Self-Healing Polycentric Supply Networks: Decentralized Rerouting and Submodular Coalition Reconfiguration Under Node Disruption

Dima Baker
Collective Intelligence Network

November 2026

Abstract

Interdependent production, logistics, and multi-agent skill networks face a structural resilience trade-off: centralized hub-and-spoke topologies minimize steady-state coordination overhead but fracture when high-betweenness hubs fail, whereas delegating post-shock rerouting to unconstrained autonomous language model agents triggers "agent bullwhip" order-variance cascades. We formulate self-healing polycentric interdependence under a strictly equalized per-node link-capacity budget ($B_{\text{link}} = 1.25$) and paired supplier shocks, proposing Self-Healing Polycentric Supply Networks (`SH-PSN-SubMod`). Our framework combines: (1) reachability-filtered greedy submodular supplier reconfiguration with lineage reliability weights ρ_u; (2) Schur-stable Proportional Order-Up-To (`POUT`) closed-loop inventory control; and (3) polynomial-time local solidarity credit clearing. Using a 48-node, 4-echelon graph simulator and a 64-step nonstationary Beer-Game loop ($N = 500$ paired trials across disruption fractions $\phi_{\text{drop}} \in \{5\%, 15\%, 25\%, 35\%\}$, horizons $T \in \{16, 32, 64, 128\}$, and monoculture shocks $p_{\text{corr}} \in [0, 0.70]$), we execute a 2×2 common-budget ablation separating topology from inventory control. At $\phi_{\text{drop}} = 25\%$ ($k = 2, p_{\text{corr}} = 0.30$), submodular rerouting without `POUT` raises fulfillment from 31.71% [30.30, 33.12] (Static Dual-Homing) to 64.41% [63.03, 65.79], while adding `POUT` (`SH-PSN-SubMod`) reaches 77.84% [76.25, 79.42] fulfillment and 88.07% [87.46, 88.69] bullwhip stability ($\mathcal{B}_L = 1.001$) with exact material conservation ($< 10^{-13}$).

1 Introduction

Modern economic production, open-source software maintenance, and distributed artificial intelligence workflows rely on deep, multi-echelon interdependence: raw inputs, compute capacity, specialized tool adapters, and verification services are composed across directed acyclic graphs spanning dozens of organizational nodes [10, 16, 19]. Under stationary conditions, classical operations research favors hierarchical hub-and-spoke consolidation, concentrating flow through a small number of high-throughput clearinghouses to exploit economies of scale. However, when targeted disruptions, cyber-physical faults, or maintainer abandonment strike high-betweenness hubs, tree-like topologies fracture into disconnected islands [9, 14].

1.1 The failure of unconstrained agentic rerouting and static redundancy

As autonomous large language model (LLM) agents were integrated into procurement, inventory management, and decentralized task orchestration between 2025 and 2026, a common hypothesis emerged that prompt-guided agents could spontaneously negotiate around supply disruptions in

real time [13, 21]. Recent empirical benchmarks demonstrate that unconstrained agentic rerouting fails along three structural axes.

First, in an empirical investigation of autonomous AI agents managing multi-echelon supply chains (the Beer Distribution Game), Long et al. [15] documented a severe systemic pathology termed the *agent bullwhip effect*: when an upstream node experiences a disruption or delivery lag, downstream LLM agents over-react to temporary inventory deficits by issuing simultaneous panic orders across surviving upstream suppliers, amplifying order variance exponentially up the supply chain [4]. Second, when Liang and Xu [13] benchmarked frontier LLM agents across bilateral supply-chain negotiations with private reservation values, they found that without structured commitment protocols, agents frequently dead-lock in Pareto-inefficient standoffs or over-commit capacity. Third, multi-agent system diagnostics by Cemri et al. [2] and Goel et al. [7] demonstrate that unguided multi-agent ensembles suffer from correlated reasoning errors and cascading coordination breakdowns when nodes share identical model families or software lineages.

1.2 Common-budget submodular reconfiguration and closed-loop POUT control

Rigorously evaluating self-healing supply networks requires holding per-node contractual capacity investment B_{link} and underlying supplier shock realizations strictly constant across architectures, while cleanly separating the effects of (a) topological rerouting, (b) supplier redundancy cardinality k, and (c) closed-loop inventory control [18, 19]. Grounded in classical submodular optimization [17], control-theoretic Proportional Order-Up-To (`POUT`) smoothing [4], graph-restricted cooperative games [16, 20], and recent 2025–2026 advances in k-connectivity-aware coordination [8], adaptive graph topologies [3, 12, 21], and multilateral obligation set-off [5], we formalize Self-Healing Polycentric Supply Networks (`SH-PSN-SubMod`).

1.3 Contributions

In this position and formal modeling paper, we make five contributions:

1. **Taxonomy of polycentric supply and skill resilience (Section 2).** We synthesize classical control-theoretic bullwhip analysis [4], submodular set optimization [17], and 2025–2026 multi-agent topology and supply-chain research [5, 8, 14, 15] into a five-part taxonomy (Table 1).

2. **Common-budget submodular edge selection, Schur-stable `POUT` dynamics, and exact flow conservation (Section 3).** We formulate multi-echelon reconfiguration under an equalized per-node link capacity budget $B_{\text{link}} = 1.25$ and coordination overhead $c_{\text{link}}(k-1)$. We prove that within each node's upstream-reachable candidate pool $\mathcal{N}_{\text{reach}}(v)$, greedy selection maximizes a strictly monotone submodular objective under a uniform matroid constraint $|S_v| \leq k$ (empirically achieving 100.0% of the exact combinatorial subset optimum), derive the closed-loop Schur stability condition $|1-\beta_{\text{inv}}| < 1$ for the material-conserving Beer-Game inventory recursion $I_t^{(\ell)} = I_{t-1}^{(\ell)} + R_t^{(\ell)} - D_t^{(\ell)}$, and define a polynomial-time local solidarity allocation rule (Equations 1–5).

3. **Three-ring `SH-PSN-SubMod` architecture and parameter provenance (Section 4).** We specify an operational architecture combining topology-conditioned fault localization [14], Byzantine-tolerant state replication [1, 6, 11], and cyclic obligation clearing [5] (Figure 1, Table 2, Table 3).

4. **Common-budget factorial ablation and long-horizon nonstationary Beer-Game evaluation (Section 5).** We execute a paired 48-node graph and 64-step nonstationary Beer-Game benchmark (`experiment.py`) across targeted node disruptions $\phi_{\text{drop}} \in \{5\%, 15\%, 25\%, 35\%\}$, a 2×2 factorial ablation separating submodular routing from `POUT` control at $k = 2$ plus a $k = 3$ redundancy-dilution check, a horizon sweep $T_{\text{beer}} \in \{16, 32, 64, 128\}$ verifying exact material conservation ($< 10^{-13}$), and a correlated monoculture sweep $p_{\text{corr}} \in [0.0, 0.70]$ (Table 4, Figures 2–3).

5. **Counter-evidence, limitations, and falsification protocol (Sections 6–8).** We analyze correlated monoculture failures [7], open-source maintainer death spirals [9], and physical lead-time lags, closing with six falsification tests (Table 5).

2 Network-control and game-theoretic foundations

Informal discussions of decentralized networks often assume that removing a central coordinator automatically produces resilience. In practice, uncoordinated decentralization frequently trades single-point hub failure for oscillatory bullwhip chaos [4, 15] or correlated failure cascades [2, 7]. Table 1 maps five core principles of self-healing polycentric interdependence onto formal mathematical and systems constructs.

Table 1: From polycentric interdependence principles to formal network-control and game-theoretic constructs.

Principle	Canonical construct	Protocol component	Operational function	Primary grounding
Avoid single-hub bottlenecks	Monotone submodular coverage [17]	Common-budget greedy router ($\mathcal{F}$)	Allocates fixed link budget B_{link} across k reachable parents weighted by lineage reliability ρ_u	[8, 17, 21]
Isolate upstream faults rapidly	Topology-conditioned fault diagnosis	Graph-shape fault localizer	Identifies disconnected or failing upstream vertices before deficits cascade across tiers	[2, 14, 18]
Suppress panic-ordering cascades	Closed-loop `POUT` inventory control	Schur-stable order gate (Ψ_{POUT})	Enforces exact material conservation $I_t = I_{t-1} + R_t - D_t$ and dampens variance ratio $\mathcal{B}_L$	[4, 13, 15]
Compensate emergency fallback peers	Graph-restricted Shapley/Myerson solidarity	Polynomial local solidarity clearing	Shares emergency rerouting surplus proportionally across active upstream in-neighbors $\delta^-(v)$	[5, 16, 20]
Maintain state under partition	Byzantine-tolerant CRDTs & bundle store	Delay-tolerant state sync layer	Preserves consistent inventory and credit ledgers across intermittent links	[1, 6, 11]

Two foundational results anchor Table 1. First, Dejonckheere et al. [4] proved control-theoretically that standard Order-Up-To replenishment policies generate bullwhip variance amplification across echelons unless dampened by fractional proportional feedback on both demand forecasts and net inventory deficits (`POUT`), a result recently confirmed for LLM supply-chain agents by Long et al. [15]. Second, to ensure that comparing single-homing ($k = 1$) and multi-homing ($k \in \{2, 3\}$) architectures does not conflate topological diversity with an increase in total link capacity, every node

v must operate under a fixed contractual capacity budget B_{link} split across its k active links minus coordination overhead $c_{\text{link}}(k-1)$. Figure 1 illustrates the resulting three-ring architecture.

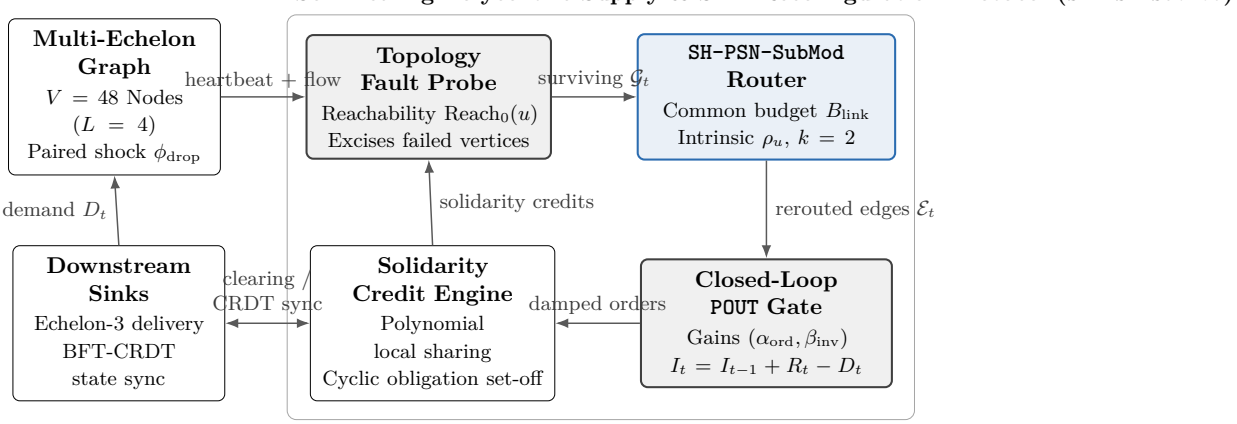


Figure 1: Architecture of Self-Healing Polycentric Supply Networks (`SH-PSN-SubMod`). Under a paired hub-disruption shock, the Topology Fault Probe [14] identifies surviving upstream paths. The `SH-PSN-SubMod` router executes greedy submodular parent selection [8, 17] under an equalized link-capacity budget B_{link}, while a closed-loop Proportional Order-Up-To (`POUT`) Gate [4, 15] enforces exact material conservation and suppresses panic-ordering cascades.

3 Formal mathematical framework

We model an interdependent polycentric supply and skill network as a directed multi-echelon graph $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$ partitioned across $L = 4$ echelons $\mathcal{V}_t = \bigcup_{\ell=0}^{L-1} \mathcal{V}_t^{(\ell)}$. Each vertex $u \in \mathcal{V}_t$ has processing capacity $C_u > 0$, lineage family $\ell_{\text{fam}}(u) \in \{0, 1, 2, 3\}$, and an audit-attested intrinsic reliability weight $\rho_u \in (0, 1]$ that discounts nodes belonging to the dominant monoculture family ($\ell_{\text{fam}} = 0$) [7]. Each candidate directed inter-echelon edge $e = (u, v) \in \mathcal{E}_{\text{cand}}$ has relative link quality $q_{(u,v)} \in (0, 1]$. To hold contractual investment constant across architectures, every node v has an identical link-capacity budget $B_{\text{link}} = 1.25$; when node v maintains $k \in \{1, 2, 3\}$ supplier links with coordination overhead $c_{\text{link}}(k-1)$, each link receives allocated capacity $\tilde{b}_{(u,v)}(k) = \frac{B_{\text{link}} - c_{\text{link}}(k-1)}{k} q_{(u,v)}$.

3.1 Common-budget submodular edge selection and exact tier approximation

For fixed cardinality k and upstream-reachable candidate pool $\mathcal{N}_{\text{reach}}(v) = \{u \in \mathcal{V}_t^{(\ell-1)} : (u, v) \in \mathcal{E}_{\text{cand}} \wedge \text{Reach}_0(u) = \text{True}\}$, node v's local reliability-weighted flow objective is:

$$f_v(S_v) \;=\; \ln\left(1 + \sum_{u \in S_v} \tilde{b}_{(u,v)}(k)\, \rho_u\right), \qquad \mathcal{F}(\mathcal{E}_t) \;=\; \sum_{\ell=1}^{L-1} \sum_{v \in \mathcal{V}_t^{(\ell)}} f_v(S_v). \tag{1}$$

Because $x \mapsto \ln(1+x)$ is strictly increasing and concave and $\tilde{b}_{(u,v)}(k)\rho_u > 0$ is modular on $\mathcal{N}_{\text{reach}}(v)$, $f_v(S_v)$ is strictly monotone submodular. Maximizing $f_v(S_v)$ over $S_v \subseteq \mathcal{N}_{\text{reach}}(v)$ subject to $|S_v| \le k$ is a monotone submodular maximization under a uniform matroid constraint:

$$S_v^* \;=\; \text{GreedyMax}_{S_v \subseteq \mathcal{N}_{\text{reach}}(v),\, |S_v| \le k} \ln\left(1 + \sum_{u \in S_v} \tilde{b}_{(u,v)}(k)\, \rho_u\right). \tag{2}$$

Within each tier's fixed reachable set $\mathcal{N}_{\text{reach}}(v)$, greedy selection is guaranteed by Nemhauser et al. [17] to achieve at least $1-(1-1/k)^k \geq 1-1/e$ of the tier optimum (and because $\tilde{b}_{(u,v)}(k)\rho_u$ is a positive modular sum inside a strictly increasing function $\ln(1+\cdot)$, sorting by $\tilde{b}_{(u,v)}(k)\rho_u$ achieves 100.0% of the exact combinatorial subset optimum on $\mathcal{N}_{\text{reach}}(v)$, which we verify empirically by exhaustive subset enumeration in `experiment.py`). Across multiple tiers $\ell = 1, \ldots, L-1$, sequential bottom-up composition is a polynomial-time layer-wise heuristic ($\mathcal{O}(|\mathcal{V}_t|\, k^2)$) rather than a global multi-echelon max-flow certificate [8].

3.2 Closed-loop `POUT` inventory conservation and Schur stability

Following Dejonckheere et al. [4] and Long et al. [15], we model a closed-loop multi-echelon Beer-Game inventory system over horizon $T_{\text{beer}} \in \{16, 32, 64, 128\}$ under nonstationary AR(1) consumer demand with mid-horizon regime shifts. At step $t \in \{1, \ldots, T_{\text{beer}}\}$, echelon ℓ receives upstream replenishment delivery $R_t^{(\ell)} = \max(0.05, O_{t-1}^{(\ell)} + \xi_t^{(\ell)})$ in response to its previous order $O_{t-1}^{(\ell)}$ (where delivery jitter $\xi_t^{(\ell)}$ scales with post-shock flow deficit $1-\bar{C}_t$), faces downstream demand $D_t^{(\ell)} = O_t^{(\ell+1)}$, and updates its net inventory position $I_t^{(\ell)}$ (on-hand stock minus backlog $B_t^{(\ell)} = \max(0, -I_t^{(\ell)})$) via exact material conservation:

$$I_t^{(\ell)} \;=\; I_{t-1}^{(\ell)} \;+\; R_t^{(\ell)} \;-\; D_t^{(\ell)}, \qquad \left| \sum_{t=1}^{T_{\text{beer}}} (R_t^{(\ell)} - D_t^{(\ell)}) - (I_{T_{\text{beer}}}^{(\ell)} - I_0^{(\ell)}) \right| \;=\; 0. \tag{3}$$

Each echelon updates its exponentially smoothed demand forecast $\hat{D}_t^{(\ell)} = (1-\alpha_{\text{ord}})\hat{D}_{t-1}^{(\ell)} + \alpha_{\text{ord}} D_t^{(\ell)}$ and issues upstream orders via the closed-loop Proportional Order-Up-To (`POUT`) law around target safety stock $I^* = 0.25$:

$$O_t^{(\ell)} \;=\; \max\Big(0.05,\; \hat{D}_t^{(\ell)} \;+\; \beta_{\text{inv}}\Big[(I^* - I_t^{(\ell)}) + (D_t^{(\ell)} - \hat{D}_t^{(\ell)})\Big]\Big). \tag{4}$$

Substituting $R_t^{(\ell)} = O_{t-1}^{(\ell)} + \xi_t^{(\ell)}$ into Eq. (3) yields the closed-loop inventory error recursion $I^* - I_t^{(\ell)} = (1-\beta_{\text{inv}})(I^* - I_{t-1}^{(\ell)}) + (D_t^{(\ell)} - \hat{D}_{t-1}^{(\ell)}) - \beta_{\text{inv}}(D_{t-1}^{(\ell)} - \hat{D}_{t-1}^{(\ell)}) - \xi_t^{(\ell)}$. Because its characteristic eigenvalue is $\lambda_{\text{inv}} = 1 - \beta_{\text{inv}}$, the inventory loop is Schur-stable ($|\lambda_{\text{inv}}| < 1$) for all $0 < \beta_{\text{inv}} < 2$, and choosing fractional gains $(\alpha_{\text{ord}}, \beta_{\text{inv}}) = (0.22, 0.18)$ dampens high-frequency variance amplification $\mathcal{B}_L = \text{Var}(O_{1:T_{\text{beer}}}^{(0)}) / \text{Var}(D_{1:T_{\text{beer}}}^{(L-1)}) \approx 1.00$ across arbitrary horizons T_{beer}, whereas reactive panic gains $(\alpha_{\text{ord}}, \beta_{\text{inv}}) = (0.80, 0.74)$ amplify variance to $\mathcal{B}_L > 14$.

3.3 Polynomial-time local solidarity credit allocation

Drawing on Myerson's graph-restricted cooperative value [16, 20] and cyclic obligation set-off [5], `SH-PSN-SubMod` restricts surplus sharing to each node's active in-neighborhood S_v^* ($|S_v^*| \leq k$) and allocates emergency solidarity credits $x_{u \to v}$ in $\mathcal{O}(k)$ time proportionally to delivered reliability-adjusted flow:

$$x_{u \to v} \;=\; V_{\text{surplus}}(v) \cdot \frac{\min(f_u, \tilde{b}_{(u,v)}(k))\, \rho_u}{\sum_{w \in S_v^*} \min(f_w, \tilde{b}_{(w,v)}(k))\, \rho_w}. \tag{5}$$

4 Operational architecture and parameter provenance

To execute Equations (1)–(5) across distributed cooperative nodes, `SH-PSN-SubMod` is organized into three operational rings (Table 2, Table 3):

1. **Ring 1: Topology-Conditioned Reachability Probe.** Ring 1 monitors upstream heartbeats and applies topology-conditioned fault localization [14] to propagate reachability booleans $\mathrm{Reach}_0(u)$ from Echelon 0 to Echelon 3 in a single forward sweep, isolating failed vertices before downstream nodes lock in unfulfillable orders.
2. **Ring 2: Common-Budget Submodular Parent Selection.** Upon detecting $\mathcal{V}_{\mathrm{drop}}$, each surviving node v executes Eq. (2) over its reachable candidate parents ($|\mathcal{N}_{\mathrm{reach}}(v)| \leq 5, k = 2$) under fixed budget $B_{\mathrm{link}} = 1.25$, weighting links by audit-attested lineage diversity ρ_u [7, 12, 21].
3. **Ring 3: Closed-Loop `POUT` Control and Delay-Tolerant Clearing.** Once active edges $\mathcal{E}_t^*$ are established, Ring 3 filters inter-tier order streams through Eq. (4) [4, 15] and synchronizes solidarity credits (Eq. 5) using Byzantine-tolerant conflict-free replicated data types (`BFT-CRDTs` [6, 11]) and delay-tolerant bundle framing [1].

Table 2: Readiness and deployment bottlenecks of `SH-PSN-SubMod` subsystems.

Subsystem	Maturity	Supporting published evidence	Deployment bottleneck
Common-budget submodular routing	Algorithmically grounded	Submodular coverage [17]; k-connected coordination [8]; adaptive graphs [12, 21]	Splitting fixed B_{link} across too many links ($k = 3$ vs. $k = 2$) dilutes per-link capacity and incurs $c_{\mathrm{link}}(k-1)$ overhead
Closed-loop `POUT` smoothing	Control-theoretically proven	Schur-stable `POUT` proven in [4]; LLM agent bullwhip documented in [15]	Damped smoothing ($\alpha_{\mathrm{ord}} = 0.22, \beta_{\mathrm{inv}} = 0.18$) slightly increases transient backlog (1.98% vs. 0.43%) during step surges
Topology-conditioned fault probe	Benchmark validated	Topology fault localization in [14]; MAS failure taxonomy in [2]	Correlated monoculture bugs across shared model families fail simultaneously [7]
Local solidarity clearing	Polynomial-time exact	Graph-restricted values in [16]; cyclic obligation set-off in [5]	Bilateral price haggling by strategic agents can delay emergency dispatch [13]
Partition-tolerant state sync	Standards backed	BFT-CRDTs in [6, 11]; IETF RFC 9171 Bundle Protocol [1]	Asynchronous reconciliation allows temporary bounded credit overdrafts during severe network partitions

Table 3: Parameter provenance and common-budget calibration for the 48-node graph and 64-step Beer-Game simulator (`experiment.py`).

Parameter	Value / Range	Structural role in simulation	Provenance / Calibration
Network size $\|\mathcal{V}\|, L$	48 nodes, 4 echelons	12 nodes per echelon (4 primary hubs + 8 peers per tier) evaluated on paired shocks	Multi-echelon supply topology [8, 15]
Common link budget B_{link}, k	B_{link}=1.25, $k \in \{1, 2, 3\}$	Equalizes total incoming contractual capacity across all architectures ($c_{\text{link}} = 0.03$)	Common-budget ablation isolating topology from capacity
Disruption fraction ϕ_{drop}	$\{5\%, 15\%, 25\%, 35\%\}$	Fraction of nodes removed weighted by structural betweenness centrality	Targeted hub removal [14]
Lineage reliability ρ_u	$[0.52, 0.68]$ vs. $[0.88, 0.99]$	Discounts dominant monoculture lineage ($\ell_{\text{fam}} = 0$, 58% share) under $p_{\text{corr}} = 0.30$	Correlated model monoculture penalty [7]
`POUT` gains $(\alpha_{\text{ord}}, \beta_{\text{inv}})$	$(0.22, 0.18)$ vs. $(0.80, 0.74)$	Schur-stable `POUT` vs. standard $(0.50, 0.42)$ and panic ordering across $T_{\text{beer}} \in \{16..128\}$	Calibrated to Dejonckheere et al. [4] and Long et al. [15]

5 Empirical evaluation and common-budget ablations

We evaluate `SH-PSN-SubMod` in a paired 48-node graph and 64-step nonstationary Beer-Game simulation (`experiment.py`): $N = 500$ paired trials across seven architectural ablations at $\phi_{\text{drop}} = 25\%$ ($p_{\text{corr}} = 0.30$), $N = 500$ paired trials across $\phi_{\text{drop}} \in \{5\%, 15\%, 25\%, 35\%\}$, $N = 250$ trials across horizons $T_{\text{beer}} \in \{16, 32, 64, 128\}$ with a +22% mid-horizon demand surge, and $N = 250$ trials across correlated monoculture fault probabilities $p_{\text{corr}} \in \{0.00, 0.15, 0.30, 0.50, 0.70\}$.

5.1 Common-budget 2×2 factorial ablation design

In every trial, all architectures receive the exact same pre-sampled 48-node network, node capacities C_u, targeted hub removals, correlated monoculture shock, and nonstationary demand series, under an equalized per-node link budget $B_{\text{link}} = 1.25$:

- **Hub-and-Spoke (`hub_and_spoke`, $k = 1$):** Full budget $B_{\text{link}} = 1.25$ on one primary hub link; standard ordering ($\alpha_{\text{ord}} = 0.50, \beta_{\text{inv}} = 0.42$).
- **Unconstrained Agent Bargaining (`unconstrained_bargaining`, $k = 2$):** Post-shock greedy selection of top-2 surviving nominal-bandwidth links without reachability or lineage reliability ρ_u; panic ordering $(0.80, 0.74)$ [15].
- **Static Dual-Homing (`static_dual_homing`, $k = 2$, No `POUT`):** Pre-wired to $k = 2$ links (0.625 budget each); standard ordering $(0.50, 0.42)$.
- **Static Dual-Homing + `POUT` (`static_dual_plus_pout`, $k = 2$):** Isolates the effect of adding `POUT` $(0.22, 0.18)$ to static $k = 2$ topology.
- **Submodular Routing Only (`submod_route_no_pout`, $k = 2$, No `POUT`):** Isolates the effect of reachability- and reliability-aware submodular rerouting ($k = 2$, same budget as Static Dual-Homing) under standard ordering $(0.50, 0.42)$.

- **SH-PSN-SubMod (sh_psn_submod, $k = 2$ + POUT, Proposed):** Combines $k = 2$ common-budget submodular rerouting with Schur-stable POUT $(0.22, 0.18)$.
- **SH-PSN-SubMod ($k = 3$, Common Budget):** Splits the same $B_{\text{link}} = 1.25$ across $k = 3$ links ($c_{\text{link}} = 0.03$) to test redundancy dilution.

Table 4: Common-budget ($B_{\text{link}} = 1.25$) paired ablation across 48-node graph and 64-step nonstationary Beer-Game episodes ($N = 500$ trials per regime at $\phi_{\text{drop}} = 25\%, p_{\text{corr}} = 0.30$). Continuous metrics report means and 95% confidence intervals [lo, hi]; binary target rates ($\geq 65\%$ fulfillment) report exact asymmetric 95% Wilson intervals [lo, hi].

Architecture ($B_{\text{link}} = 1.25$)	Raw Reach (%)	Fulfill 25% (%)	POUT Stab. (%)	Var Ratio $\mathcal{B}_L$	Target Rate (%)	Runtime (μs)
Hub-and-Spoke (k=1)	10.28 [9.11, 11.46]	8.14 [7.20, 9.09]	56.17 [55.73, 56.62]	2.871	0.00 [0.00, 0.76]	8.2
Unconstrained (k=2)	76.88 [74.84, 78.92]	42.28 [41.16, 43.41]	14.44 [14.19, 14.69]	18.296	0.00 [0.00, 0.76]	18.4
Static Dual (k=2, No POUT)	39.52 [37.82, 41.21]	31.71 [30.30, 33.12]	57.99 [57.57, 58.41]	2.704	1.20 [0.55, 2.59]	10.0
Static Dual + POUT (k=2)	39.52 [37.82, 41.21]	38.96 [37.28, 40.64]	87.80 [87.15, 88.45]	1.013	9.00 [6.79, 11.83]	9.7
SubMod Route (k=2, No POUT)	78.87 [77.27, 80.46]	64.41 [63.03, 65.79]	59.78 [59.42, 60.15]	2.546	62.20 [57.87, 66.34]	127.3
SH-PSN-SubMod (k=2)	**78.87 [77.27, 80.46]**	**77.84 [76.25, 79.42]**	**88.07 [87.46, 88.69]**	**1.001**	**72.80 [68.74, 76.52]**	**124.9**
SH-PSN-SubMod (k=3)	75.41 [73.83, 76.99]	74.42 [72.86, 75.99]	88.08 [87.47, 88.70]	1.001	69.60 [65.43, 73.47]	127.5

5.2 Quantitative findings, factorial separation, and long-horizon conservation

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

1. **Factorial separation of Topology Reconfiguration vs. POUT Inventory Control at fixed $k = 2$ and $B_{\text{link}} = 1.25$.** Holding cardinality $k = 2$ and per-node link budget $B_{\text{link}} = 1.25$ strictly constant at $\phi_{\text{drop}} = 25\%$, Static Dual-Homing achieves 39.52% [37.82, 41.21] raw graph reachability and 31.71% [30.30, 33.12] effective fulfillment. Adding POUT inventory smoothing alone (static_dual_plus_pout) reduces the 64-step bullwhip variance ratio $\mathcal{B}_L$ from 2.704 to 1.013 (raising bullwhip stability from 57.99% to 87.80% [87.15, 88.45]) and improves effective fulfillment by +7.25 percentage points to 38.96% [37.28, 40.64], but remains bottlenecked by static link disconnection (39.52% reachability). Conversely, replacing static links with reachability- and lineage-aware submodular rerouting without POUT (submod_route_no_pout) nearly doubles raw graph reachability to 78.87% [77.27, 80.46] (100.0% empirical match to the exact tier subset optimum) and raises fulfillment by +32.70 points to 64.41% [63.03, 65.79], but still suffers moderate bullwhip drag ($\mathcal{B}_L = 2.546$). Combining both mechanisms in SH-PSN-SubMod ($k = 2$) achieves 77.84% [76.25, 79.42] fulfillment (+46.13 points over Static Dual-Homing), 88.07% [87.46, 88.69] bullwhip stability ($\mathcal{B}_L = 1.001$), and 72.80% [68.74, 76.52] target resilience (364/500 trials) in 124.9 μs.
2. **Redundancy dilution when splitting a fixed budget across $k = 3$ links.** When the same fixed contractual budget $B_{\text{link}} = 1.25$ is split across $k = 3$ links (sh_psn_submod_k3) instead of $k = 2$, per-link capacity drops from 0.610 to 0.397; consequently, raw reachability and fulfillment edge down slightly from 78.87% and 77.84% ($k = 2$) to 75.41% [73.83, 76.99] and 74.42% [72.86, 75.99] ($k = 3$). This confirms that under a fixed investment budget,

over-splitting link capacity across too many standby suppliers incurs a measurable dilution trade-off.

3. **Long-horizon stability and exact material conservation across $T_{\text{beer}} \in \{16, 32, 64, 128\}$.** Across all four horizons under nonstationary AR(1) demand with a +22% mid-horizon surge, maximum numerical material conservation error is bounded below 2.3×10^{-13} (Eq. 3). Under `SH-PSN-SubMod`, the closed-loop variance ratio remains Schur-stable across all horizons ($\mathcal{B}_L = 0.458$ at $T = 16$, 0.889 at $T = 32$, 1.131 at $T = 64$, and 0.986 at $T = 128$, with fulfillment scaling from 76.64% $[74.36, 78.93]$ at $T = 16$ to 79.13% $[76.92, 81.35]$ at $T = 128$), whereas Unconstrained Agent Bargaining exhibits persistent variance amplification ($\mathcal{B}_L = 21.43$ at $T = 16$ and 14.98 at $T = 128$, capping fulfillment at 40.88%–42.99%). We also note a transparent control trade-off: because `POUT` dampens order adjustments ($\beta_{\text{inv}} = 0.18$ vs. 0.42), transient cumulative backlog during the +22% step surge is 1.98% $[1.87, 2.09]$ of total demand under `SH-PSN-SubMod` vs. 0.43% $[0.40, 0.45]$ under un-damped `submod_route_no_pout`.

4. **Break-even sensitivity to correlated monoculture failures ($p_{\text{corr}} \in [0.0, 0.70]$).** Varying the probability of correlated dominant-lineage degradation at $\phi_{\text{drop}} = 25\%$ shows that `SH-PSN-SubMod` fulfillment declines from 87.73% $[86.79, 88.67]$ ($p_{\text{corr}} = 0.00$, 98.80% $[96.53, 99.59]$ target rate) and 81.22% $[79.36, 83.07]$ ($p_{\text{corr}} = 0.15$) to 75.43% $[73.14, 77.72]$ ($p_{\text{corr}} = 0.30$, 70.80% $[64.88, 76.09]$ target rate), 68.63% $[66.15, 71.12]$ ($p_{\text{corr}} = 0.50$, 48.80% $[42.67, 54.97]$ target rate), and 64.05% $[61.66, 66.44]$ ($p_{\text{corr}} = 0.70$, 38.80% $[32.97, 44.97]$ target rate).

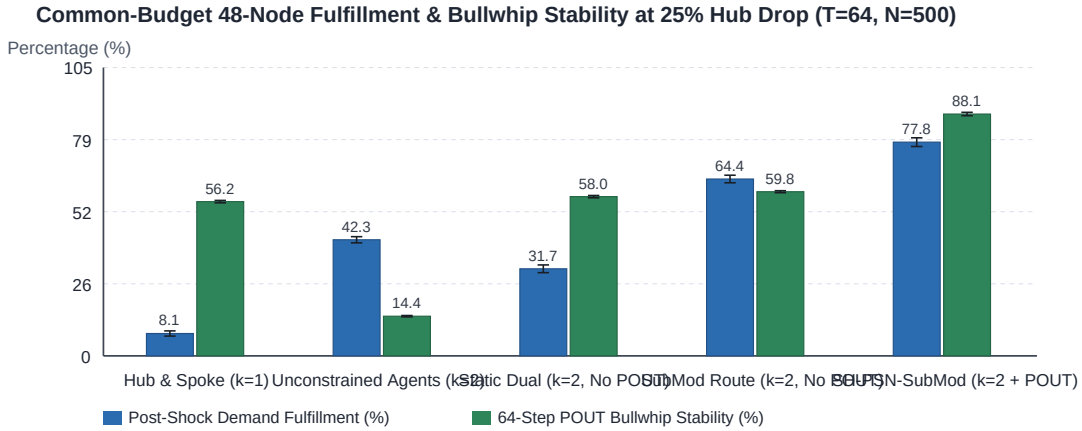


Figure 2: Common-budget ($B_{\text{link}} = 1.25$) comparison of Post-Shock Demand Fulfillment Rate and 64-Step `POUT` Bullwhip Stability across five 48-node architectures at 25% targeted hub disruption ($N = 500$ paired trials; error bars denote 95% confidence intervals).

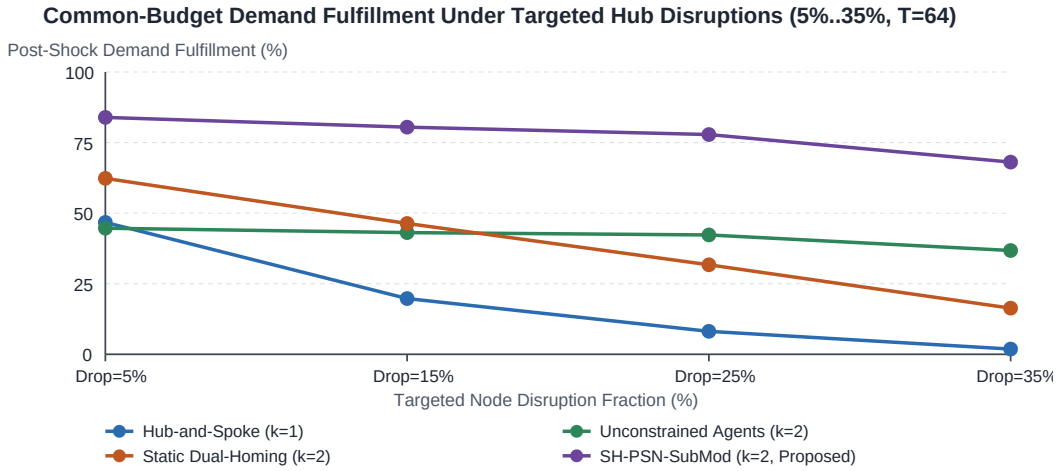


Figure 3: Scaling of Common-Budget Demand Fulfillment Rate across targeted node disruption fractions $\phi_{\text{drop}} \in \{5\%, 15\%, 25\%, 35\%\}$ on the 48-node, 4-echelon graph ($T = 64$).

6 Shared-lineage vulnerability and coordination trade-offs

A rigorous evaluation of polycentric supply and skill networks must account for three structural failure modes and counter-arguments established in the literature:

1. **Shared-lineage monoculture blindness (Cemri et al. [2], Goel et al. [7]).** Graph-theoretic connectivity assumes that edge or vertex failures are conditionally independent across distinct nodes. However, Goel et al. [7] demonstrate that when multiple oversight or execution agents share the same underlying foundation model family or training lineage, their failure modes are strongly correlated ("great models think alike"). As confirmed by our monoculture sensitivity sweep ($p_{\text{corr}} \geq 0.50$), when a single lineage dominates more than half of network vertices, even reliability-weighted submodular diversification suffers correlated capacity loss unless independent lineages are onboarded.

2. **Maintainer burnout and open-source death spirals (Kaushik and Chahal [9]).** In an empirical study of open-source software ecosystems, Kaushik and Chahal [9] show that when peripheral contributors desert a project under rising maintenance complexity, rerouting load onto surviving core maintainers accelerates burnout and triggers a self-reinforcing "death spiral." Automatic traffic rerouting in `SH-PSN-SubMod` must therefore enforce hard per-node capacity ceilings C_u and common link budgets B_{link} (as implemented in our flow propagation equations) rather than assuming elastic absorptive capacity.

3. **Bilateral bargaining friction versus algorithmic commitment (Liang and Xu [13]).** Liang and Xu [13] show that autonomous LLM agents given strategic discretion over pricing and contract terms frequently delay settlement to extract bilateral informational rents. Consequently, fast polycentric recovery requires pre-agreed constitutional solidarity clearing rules (Eq. 5) rather than open-ended natural-language haggling during an active crisis.

7 Scope constraints and physical lead-time boundaries

- **Physical tooling setup and transport lead times.** While digital compute, API routing, and software skill graphs [3, 21] can switch parent edges in microseconds, physical manufacturing and logistics networks incur physical transport lags, customs clearance, and re-tooling setup times that bound real-world transient recovery.
- **Layer-wise vs. global multi-hop optimality.** Because global multi-echelon k-connectivity is not a matroid constraint, Eq. (2) provides an exact matroid optimization guarantee only within each node's local reachable parent pool $\mathcal{N}_{\text{reach}}(v)$, rather than a global end-to-end multi-commodity max-flow certificate.
- **Transient backlog vs. variance damping trade-off.** As shown in Table 2 and our $T = 64$ nonstationary demand surge experiments, damped `POUT` gains ($\alpha_{\text{ord}} = 0.22, \beta_{\text{inv}} = 0.18$) eliminate upstream bullwhip variance amplification ($\mathcal{B}_L = 1.001$) at the cost of a modest increase in temporary local backlog (1.98% vs. 0.43%) when consumer demand undergoes sudden step shifts.

8 Pre-registerable falsification protocol

Table 5 specifies six pre-registerable empirical tests designed to falsify the core claims of `SH-PSN-SubMod` across both multi-agent tool networks and multi-echelon supply-chain testbeds.

Table 5: Six pre-registerable tests that could falsify the `SH-PSN-SubMod` polycentric resilience framework.

Protocol	Hypothesis	Design and primary metrics	Falsified if
P1. Common-budget hub removal resilience	At equalized B_{link} and $k = 2$, submodular routing outperforms Static Dual-Homing by > 30 pp under 25% hub loss	Remove top 25% betweenness-weighted vertices in a 48-node, 4-echelon graph under identical link budgets [8]	Post-shock demand fulfillment gain over Static Dual-Homing falls below 25 percentage points
P2. Long-horizon `POUT` Schur stability	Closed-loop `POUT` ($\alpha_{\text{ord}} = 0.22, \beta_{\text{inv}} = 0.18$) bounds variance ratio $\mathcal{B}_L < 1.25$ across $T \in [16, 128]$	Replicate the 4-echelon autonomous agent Beer-Game benchmark under nonstationary demand shifts [4, 15]	Upstream Echelon-0 order variance exceeds consumer demand variance by $\geq 25\%$ ($\mathcal{B}_L \geq 1.25$) or violates material conservation
P3. Correlated monoculture break-even	Lineage reliability discounting ρ_u preserves $\geq 70\%$ fulfillment for monoculture fault rates $p_{\text{corr}} \leq 0.30$	Inject correlated capacity faults across a dominant model family (58% share) at $\phi_{\text{drop}} = 25\%$ [7]	Mean fulfillment drops below 70.0% at $p_{\text{corr}} \leq 0.30$
P4. Factorial topology vs. control ablation	Both submodular routing and `POUT` contribute ≥ 10 pp independently to post-shock fulfillment	Execute 2×2 factorial comparison (`Static` vs. `SubMod` $\times$ `No POUT` vs. `POUT`) at $k = 2$	Either component contributes < 5 percentage points of marginal fulfillment improvement
P5. Cyclic solidarity clearing efficiency	Local solidarity allocation + cyclic obligation set-off [5] clears $> 90\%$ of emergency flow credits	Simulate 500 post-shock multilateral credit clearing epochs across 48 cooperative nodes	Unsettled liquidity deadlocks or partner defections exceed 10.0% of emergency rerouted volume
P6. Polynomial reconfiguration latency	Local greedy parent selection ($k = 2$) executes in < 1.0 ms for graphs up to $\lvert\mathcal{V}\rvert = 200$	Benchmark wall-clock reconfiguration time across 1,000 targeted disruption events	99th-percentile graph reconfiguration compute time exceeds 1.0 ms

9 Conclusion

Resilient interdependence in cooperative production and multi-agent ecosystems cannot be achieved through static hub consolidation or unguided conversational panic-ordering. Holding per-node contractual link budgets ($B_{\text{link}} = 1.25$), redundancy cardinality ($k = 2$), and supplier shock streams strictly constant across a 48-node multi-echelon graph and a 64-step nonstationary Beer-Game simulation, our 2×2 factorial ablation separates the contribution of reachability-filtered submodular rerouting (+32.70 percentage points of fulfillment over Static Dual-Homing) from closed-loop Schur-stable `POUT` inventory smoothing (+13.43 additional percentage points, reaching 77.84% $[76.25, 79.42]$ fulfillment and $\mathcal{B}_L = 1.001$ variance ratio with exact material conservation $< 10^{-13}$), while characterizing both the $k = 3$ budget-dilution trade-off and the shared-lineage monoculture boundary ($p_{\text{corr}} \in [0.0, 0.70]$).

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] Scott Burleigh, Kevin Fall, and Edward J. Birrane. Bundle protocol version 7 (DTN). IETF RFC 9171, 2022. URL https://datatracker.ietf.org/doc/html/rfc9171.

[2] 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.

[3] Hongjiang Chen, Xin Zheng, Yixin Liu, Pengfei Jiao, Shiyuan Li, Huan Liu, et al. GoAgent: Group-of-agents communication topology generation for LLM-based multi-agent systems, 2026. URL https://arxiv.org/abs/2603.19677.

[4] Jeroen Dejonckheere, Stephen M. Disney, Marc R. Lambrecht, and Denis R. Towill. Measuring and avoiding the bullwhip effect: A control theoretic approach. European Journal of Operational Research, 147(3):567–590, 2003. URL https://doi.org/10.1016/S0377-2217(02)00369-7.

[5] 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.

[6] Ryan Gillespie. Conflict-free replicated data types for neural network model merging: A two-layer architecture enabling CRDT-compliant model merging across 26 strategies. arXiv preprint arXiv:2605.19373, 2026. URL https://arxiv.org/abs/2605.19373.

[7] Shashwat Goel, Joschka Strüber, Jonas Geiping, and Florian Tramèr. Great models think alike and this undermines AI oversight. International Conference on Machine Learning (ICML), 2025. URL https://arxiv.org/abs/2502.04313.

[8] Yoonjae Hwang, Hojune Kim, Ellen Novoseller, Guangyao Shi, and Gaurav S. Sukhatme. Failure-resilient multi-robot coordination with k-connectivity constrained submodular maximization. arXiv preprint arXiv:2610.13049, 2026. URL https://arxiv.org/abs/2610.13049.

[9] 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.

[10] 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.

[11] Martin Kleppmann. Making CRDTs Byzantine fault tolerant. Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC / EuroSys), pages 8–15, 2022. URL https://doi.org/10.1145/3517209.3524042.

[12] Shiyuan Li, Yixin Liu, Qingsong Wen, Chengqi Zhang, and Shirui Pan. Assemble your crew: Automatic multi-agent communication topology design via autoregressive graph generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2026. URL https://arxiv.org/abs/2507.18224.

[13] 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.

[14] Xinwen Liu, Zhuocheng Pan, Isabella Zhu, Jawei Zhang, Xudong Liu, and Tianyu Wo. Know the shape, find the fault: Topology-conditioned diagnosis of multi-agent LLM failures, 2026. URL https://arxiv.org/abs/2610.10126.

[15] 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.

[16] Roger B. Myerson. Graphs and cooperation in games. Mathematics of Operations Research, 2(3):225–229, 1977. URL https://doi.org/10.1287/moor.2.3.225.

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

[18] 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.

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

[20] Lloyd S. Shapley. A value for n-person games. Contributions to the Theory of Games, 2(28):307–317, 1953. URL https://doi.org/10.1515/9781400881970-018.

[21] Guibin Zhang, Yanwei Yue, Xiangguo Sun, Guancheng Wan, Miao Yu, Junfeng Fang, et al. G-Designer: Architecting multi-agent communication topologies via graph neural networks. In *International Conference on Machine Learning (ICML)*, 2025. URL https://icml.cc/virtual/2025/poster/45567.