Ad-Hoc Cooperative Group Purchasing via Lightweight Private Set Intersection and Deterministic Billing-Rule Verification

Dima Baker
Collective Intelligence Network

November 2026

Abstract

When individual households, clinics, or small cooperatives negotiate against oligopolistic suppliers using autonomous language-model agents, they face severe structural disadvantages: conversational buyer agents routinely leak private reserve valuations during multi-turn bargaining, while forming a traditional group-purchasing organization (GPO) requires exposing plaintext purchase lists to a centralized broker. We propose Blind Order Pooling with Deterministic Billing Verification (BOP-DBV), a cryptographic and rule-enforced protocol enabling ad-hoc buyer coalitions to aggregate purchasing power without disclosing non-matching items or individual reserve ceilings. BOP-DBV combines party-separated Cuckoo-binned Over-Threshold Multi-Party Private Set Intersection (OT-MP-PSI) over the 127-bit Mersenne prime field $\mathbb{F}_{2^{127}-1}$ with zero-knowledge reserve range proofs to identify standardized product codes (SKUs) demanded by at least t coalition members while keeping sub-threshold bins cryptographically hidden. Matched syndicates bid via canonical volume strikes and audit vendor invoices through a deterministic rule-table verifier. Executing party-separated Shamir (t, M) interpolation over $\mathbb{F}_{2^{127}-1}$ across 1,920 procurement runs ($M \in \{8, 16, 32, 64\}$) plus a 600-run threshold sweep, BOP-DBV at $M = 64$ captures $95.24\% \pm 0.02\%$ of the cooperative consumer surplus (vs. $13.80\% \pm 0.15\%$ for solo LLM buyers and $78.43\% \pm 0.12\%$ for plaintext brokers), achieves 100.00% ([99.94, 100.00]) sub-threshold bin hiding and 100.00% ([99.95, 100.00]) post-output buyer shielding under $k_{\text{sybil}} = 2$ phantom probing when $\tau = t/M \in [0.10, 0.38]$, catches 99.62% ([99.50, 99.70]) of invoice checks, and executes local 127-bit field matching in 25.68 ms (71.45 ms including modeled network/ZK overhead).

1 Introduction

A foundational promise of personal and small-enterprise AI agents is the democratization of market bargaining power: an autonomous agent acting on behalf of a household, independent clinic, or neighborhood cooperative should be able to search suppliers, aggregate bulk orders, and audit complex invoices with the analytical rigor of an enterprise procurement department [1, 23]. Yet empirical evidence compiled across 2025 and 2026 demonstrates that equipping isolated buyers with conversational language-model agents frequently *worsens* their bargaining position relative to concentrated suppliers [16, 18].

1.1 Asymmetric bargaining, reserve leakage, and bilateral information rents

Why do solo autonomous buyer agents lose consumer surplus in digital marketplaces? Classical mechanism design and recent 2025–2026 empirical benchmarks identify four compounding vulnerabilities:

1. **Private valuation leakage in multi-turn LLM bargaining.** By the Myerson–Satterthwaite theorem [20] and Rubinstein alternating-offers bargaining theory [3, 21], a buyer's ability to capture trade surplus depends critically on keeping their maximum willingness-to-pay (reserve price v_i) private. However, in large-scale benchmarks of autonomous negotiation and supply-chain agents, Liang and Xu [16] and Long et al. [18] show that conversational LLM buyers routinely leak urgency cues and private reserve bounds during multi-turn natural-language dialogue. Similarly, El Yagoubi et al. [8] document systematic cross-context side-channel leakage in multi-agent workflows, enabling strategic sellers to approach first-degree price discrimination.

2. **Supply-chain bullwhip amplification and fragmented ordering.** When buyers place uncoordinated solo orders or react to local inventory signals through autonomous ordering agents, Dejonckheere et al. [7] and Long et al. [18] demonstrate that order variance amplifies upstream (the bullwhip effect), raising supplier unit fulfillment costs and forfeiting economies of scale [1].

3. **Cyclic negotiation deadlocks and tool-execution drift.** In multi-agent workflows, unconstrained conversational negotiation is vulnerable to cyclic handoffs [11], deceptive counterparty framing [2, 10], and multi-view inconsistency [13], requiring deterministic runtime contract boundaries [6, 9].

4. **Post-agreement fee-catalog overbilling.** Even after a unit strike price is negotiated, vendors in healthcare, cloud infrastructure, and wholesale distribution frequently extract rents through post-hoc billing surcharges, unbundled handling modifiers, and complex fee schedules [12, 22]. Relying on unstructured LLM text parsing to audit multi-line invoices against regulatory or contracted fee tables suffers from high false-clearance rates whenever modifier exclusions interact arithmetic caps.

1.2 Breaking the privacy–scale dilemma via Over-Threshold PSI

In industrial economics, the canonical defense against seller oligopoly is cooperative group purchasing—pooling buyer demand to slide down the vendor's marginal cost curve and enforce bulk price ceilings [1, 23]. Historically, however, forming an ad-hoc Group Purchasing Organization (GPO) created an acute **privacy–scale dilemma**: to discover which households, clinics, or small shops need the exact same standardized product code (SKU) during the current ordering window, participants had to upload their complete plaintext shopping lists and reserve prices to a centralized broker. Central brokers immediately observe every participant's non-matching, sensitive items (such as niche pharmaceuticals or proprietary hardware inputs) and extract platform intermediation rents [23].

While foundational Multi-Party Private Set Intersection (MP-PSI) protocols based on homomorphic polynomial operations [14] or generic garbled circuits avoid a trusted broker, they reveal only the *all-party* intersection $\bigcap_{i=1}^{M} \mathcal{X}_i$ (which is typically empty when M is large) and scale quadratically in communication and computation. To support cooperative group purchasing, buyers instead require **Over-Threshold Multi-Party PSI (OT-MP-PSI)**: identifying any SKU x demanded by at least t out of M buyers ($|\mathcal{I}_x| \geq t$, where $t \leq M$ is the vendor's minimum wholesale lot size) while keeping all sub-threshold items ($|\mathcal{I}_{x'}| < t$) cryptographically hidden.

Building on Shamir secret sharing [24], batched Oblivious PRF constructions [15], and algebraic syndrome decoding [19], recent breakthroughs in lightweight PSI [17] and Practical Traceable Over-Threshold MP-PSI [25] enable M parties to compute threshold intersections in tens of milliseconds. Paired with compact zero-knowledge range proofs [4, 5], buyers can prove collective solvency at a bulk strike price without revealing individual reserve ceilings.

1.3 Contributions

We synthesize these cryptographic and economic primitives into Blind Order Pooling with Deterministic Billing Verification (BOP-DBV):

1. **Cooperative procurement and cryptographic taxonomy (Section 2).** We map informal group-buying intuitions onto formal OT-MP-PSI primitives, zero-knowledge range commitments, Nash bargaining models, and deterministic billing verifiers (Table 1, Figure 1).
2. **Corrected finite-field OT-MP-PSI and leakage formalism (Section 3).** We formulate Cuckoo-binned Shamir (t, M) secret sharing over prime field $\mathbb{F}_q$ where OPPRF evaluations return pure field elements $P_{e,x}(i) \in \mathbb{F}_q$ while auxiliary escrow tags are AEAD-encrypted under a key derived from the reconstructed constant term $s_{x,0} = P_{e,x}(0)$ (Equations 1–2), define the formal cryptographic and inference leakage profile $\mathcal{L}_{\text{leak}}$, derive the cooperative consumer surplus curve under zero-knowledge reserve commitments (Equations 3–4), and specify the deterministic billing verification automaton (Equation 5).
3. **Four-stage BOP-DBV protocol architecture (Section 4).** We specify an end-to-end architecture integrating Cuckoo-hashed OT-MP-PSI [17, 19, 25], Bulletproof reserve commitments [4, 5], canonical strike bargaining [6, 9], and rule-table invoice auditing [12, 22] (Table 2).
4. **Party-separated 127-bit finite-field simulation and procurement evaluation (Section 5).** We execute party-separated Shamir (t, M) polynomial share generation and blind coordinator Lagrange interpolation over the 127-bit Mersenne prime field $\mathbb{F}_{2^{127}-1}$, cooperative volume-pricing optimization, and deterministic invoice rule checking across 1,920 procurement executions ($M \in \{8, 16, 32, 64\}$), plus a 600-run sensitivity sweep over the threshold ratio $\tau = t/M \in [0.05, 0.52]$ under phantom-buyer probing, explicitly separating cryptographic sub-threshold bin hiding from post-output Sybil inference and billing rule verification (Tables 3–4, Figures 2–3).
5. **Adversarial phantom-probing analysis and falsification protocol (Sections 6–8).** We analyze Sybil phantom-buyer threshold probing, SKU ontology fragmentation, and coalition defection, closing with six pre-registerable falsification tests (Table 5).

2 Cryptographic and economic taxonomy of blind order pooling

Informal discussions of "collective AI bargaining" frequently overlook why cooperative purchasing breaks down in open digital markets: buyers refuse to expose sensitive non-shared inventory needs to peers or central brokers, conversational agents leak reserve valuations during unstructured haggling [8, 16], and suppliers recoup wholesale discounts through complex post-hoc billing modifiers [12, 22]. Table 1 maps five operational requirements of BOP-DBV onto established cryptographic and economic constructs.

Figure 1 illustrates the four-stage execution flow of BOP-DBV. Each participant $i \in \{1, \ldots, M\}$ holds a private SKU set $\mathcal{X}_i$ and reserve valuation vector $\mathbf{v}_i$. Stage 1 executes Cuckoo-binned OT-MP-PSI over prime field $\mathbb{F}_q$ [17, 19, 25] to identify only those standardized SKUs where aggregate coalition demand crosses wholesale threshold t. Stage 2 decrypts the threshold-bound escrow tags using $H_{\text{key}}(s_{x,0})$ and verifies Bulletproof reserve range proofs [4]. Stage 3 issues a schema-locked volume RFQ to competing vendors [1, 18]. Finally, Stage 4 audits vendor invoices against the deterministic fee-catalog rule table [12, 22] before releasing escrow settlement.

Table 1: Mapping blind order-pooling requirements to formal cryptographic and game-theoretic mechanisms.

Design intuition	Canonical construct	Protocol component	Operational function	Primary grounding
"Find shared orders without revealing niche lists"	Over-Threshold Multi-Party PSI (OT-MP-PSI)	Cuckoo-binned OPPRF + Shamir (t, M) over $\mathbb{F}_q$	Reconstructs SKU secret $s_{x,0}$ iff $\lvert\mathcal{I}_x\rvert \geq t$; hides all sub-threshold SKUs ($\lvert\mathcal{I}_{x'}\rvert < t$)	[17, 19, 24, 25]
"Never leak max willingness-to-pay in chat"	Zero-knowledge range commitment	Pedersen + Bulletproofs range gate	Proves $v_i(x) \geq P_{\text{strike}}(Q_x)$ without disclosing individual buyer reserve $v_i(x)$	[4, 5, 20]
"Prevent multi-turn conversational price probing"	Canonical take-it-or-leave-it volume RFQ	Schema-locked bargaining state machine	Replaces free-text haggling with sealed-bid volume strikes under runtime governance	[3, 9, 16, 21]
"Catch hidden junk fees & modifier overbilling"	Deterministic fee-schedule rule verifier	Arithmetic + modifier-exclusion automaton	Evaluates line items against contracted price caps and mutually exclusive modifier tables	[6, 12, 22]
"Bind matched buyers to wholesale lot commitment"	Key-encapsulated threshold escrow tag	AEAD payload keyed by $H_{\text{key}}(s_{x,0})$	Unlocks traceable escrow commitments iff at least t buyers hold matching SKU x	[1, 15, 25]

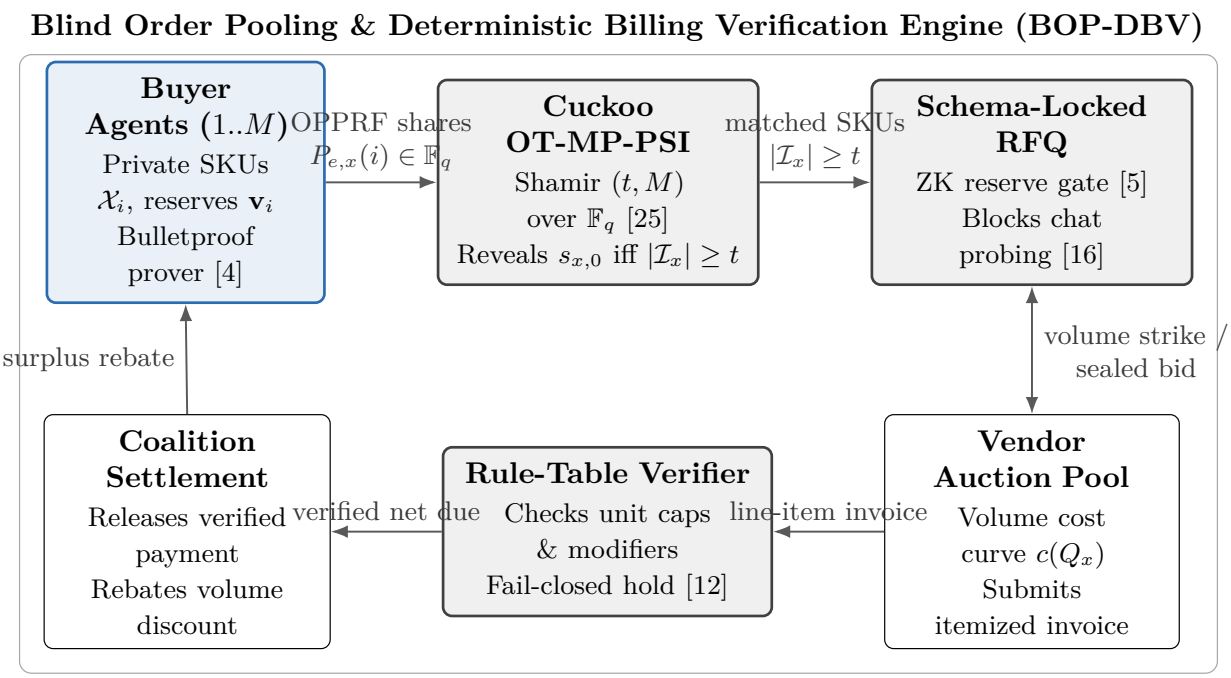


Figure 1: Architecture of Blind Order Pooling with Deterministic Billing Verification (BOP-DBV). Buyer agents encode private SKU demands into Cuckoo-binned OPPRF shares over $\mathbb{F}_q$ [17, 19, 25] to discover only those standardized SKUs demanded by at least t members. Matched syndicates prove collective reserve solvency in zero knowledge [4, 5], execute schema-locked bulk RFQs that eliminate conversational reserve leakage [16], and audit vendor invoices against deterministic fee-schedule rules [12, 22].

3 Formal mathematical framework

We consider an ad-hoc coalition of M autonomous buyer agents $\mathcal{M} = \{1, \dots, M\}$, where each buyer $i \in \mathcal{M}$ holds a private set of desired standardized product/service identifiers $\mathcal{X}_i \subset \{0,1\}^\kappa$ (with $|\mathcal{X}_i| \le k$) and a private maximum willingness-to-pay (reserve valuation) $v_i(x) > 0$ for each $x \in \mathcal{X}_i$.

3.1 Cuckoo-binned Over-Threshold MP-PSI over $\mathbb{F}_q$ and leakage bounds

The coalition seeks to identify the subset of items $\mathcal{X}^*_{\ge t} = \{x : |\mathcal{I}_x| \ge t\}$, where $\mathcal{I}_x = \{i \in \mathcal{M} : x \in \mathcal{X}_i\}$ is the set of buyers requesting SKU x, while revealing zero cryptographic information about any sub-threshold item x' with $|\mathcal{I}_{x'}| < t$.

Following Mahdavi et al. [19], Yang et al. [25], and Ling et al. [17], candidates are first mapped into a Cuckoo hash table of $B = O(k)$ bins so that each buyer places at most one item per bin. For epoch e and candidate SKU x, let $s_{x,0} = H_{\text{field}}(e \,\|\, x) \in \mathbb{F}_q^*$ be the target secret constant term over a prime field $\mathbb{F}_q$, and let $P_{e,x}(z) = s_{x,0} + \sum_{j=1}^{t-1} a_{x,j} z^j \pmod q$ be the associated degree-$(t-1)$ Shamir polynomial [24] with deterministic pseudorandom coefficients $a_{x,j} = H_{\text{coeff}}(e \,\|\, x \,\|\, j) \in \mathbb{F}_q$ derived across the threshold syndicate.

A critical algebraic requirement is that auxiliary escrow metadata must **not** be XORed directly onto the finite-field share $P_{e,x}(i) \in \mathbb{F}_q$ prior to $\bmod q$ Lagrange interpolation, as bitwise XOR does not commute with modular field arithmetic and would corrupt polynomial reconstruction unless all tags were identical. Instead, each buyer $i \in \mathcal{M}$ programs their Oblivious Programmable Pseudorandom Function OPPRF_i [15, 17] to output a pure finite-field evaluation $y_{i,b} \in \mathbb{F}_q$ paired with a separately authenticated ciphertext $\mathsf{Aux}_{i,b}$ encrypted under a symmetric key derived from the secret constant term $s_{x,0}$:

$$(y_{i,b},\ \mathsf{Aux}_{i,b}) \ = \ \begin{cases} \Big(P_{e,x}(i) \bmod q,\ \ \mathsf{AEAD.Enc}_{H_{\text{key}}(s_{x,0})}(\mathsf{Tag}_i(x))\Big) & \text{if buyer } i \text{ maps } x \in \mathcal{X}_i \text{ to bin } b, \\ \Big(r_{i,b} \xleftarrow{\$} \mathbb{F}_q,\ \ u_{i,b} \xleftarrow{\$} \{0,1\}^\ell\Big) & \text{if bin } b \text{ is empty for buyer } i. \end{cases} \tag{1}$$

Furthermore, to avoid iterating over $\binom{M}{t}$ subsets in each bin b, Cuckoo hashing restricts the active bin occupancy, and when at least t valid shares lie on $P_{e,x}(\cdot)$ (or when Reed–Solomon / Berlekamp–Welch syndrome decoding over $\mathbb{F}_q$ isolates the degree-$(t-1)$ polynomial from random padding shares, 19, 25), Lagrange interpolation at $z = 0$ over any t-subset $\mathcal{S} \subseteq \mathcal{I}_x$ ($|\mathcal{S}| = t$) reconstructs the exact constant term $s_{x,0} \in \mathbb{F}_q$:

$$\hat{s}_{x,0} \ = \ \sum_{i \in \mathcal{S}} y_{i,b} \prod_{j \in \mathcal{S} \setminus \{i\}} \Big((-j) \cdot (i-j)^{-1} \bmod q\Big) \bmod q \ = \ s_{x,0} \quad \Longleftrightarrow \quad |\mathcal{I}_x| \ge t. \tag{2}$$

Once $s_{x,0}$ is reconstructed, the coordinator derives $K_x = H_{\text{key}}(s_{x,0})$ and decrypts $\mathsf{Aux}_{i,b}$ for all $i \in \mathcal{I}_x$, recovering the active buyer coalition $\mathcal{I}_x$ and their refundable escrow commitments $\mathsf{Tag}_i(x)$ [25]. Conversely, if $|\mathcal{I}_{x'}| < t$, information-theoretic Shamir secret sharing over $\mathbb{F}_q$ guarantees that the $|\mathcal{I}_{x'}| \le t-1$ evaluations reveal zero information about $s_{x',0}$, leaving $\mathsf{Aux}_{i,b}$ computationally indistinguishable from random bits under the security of AEAD.

Accordingly, in the semi-honest model with at most $k_{\text{sybil}} < t$ colluding adversarial nodes, the formal protocol leakage function is:

$$\mathcal{L}_{\text{leak}}(\{\mathcal{X}_i, \mathbf{v}_i\}_{i=1}^M) \ = \ \Big\{ (x, \mathcal{I}_x) : |\mathcal{I}_x| \ge t \Big\} \ \cup \ \Big\{ \mathbb{I}(v_i(x) \ge P_{\text{strike}}(Q_x)) : i \in \mathcal{I}_x \Big\}.$$

Notice that while sub-threshold items ($|\mathcal{I}_{x'}| < t$) suffer zero cryptographic leakage, an inference vulnerability arises if the threshold t is set so small that $t \le k_{\text{sybil}} + 1$: in that regime, k_{sybil} phantom

accounts controlled by an oligopolistic seller can join the pool, inject k_{sybil} shares for a niche SKU x, and trigger reconstruction whenever a single honest buyer holds x ($|\mathcal{I}_x| = k_{\text{sybil}} + 1 \geq t$). We quantify this exact threshold boundary in Section 5.

3.2 Zero-knowledge reserve commitments and cooperative surplus capture

Once a matched SKU $x \in \mathcal{X}^*_{\geq t}$ with coalition size $Q_x = |\mathcal{I}_x| \geq t$ is discovered, the syndicate negotiates against oligopolistic suppliers without disclosing individual reserve valuations $v_i(x)$ [16, 20]. Using compact Pedersen commitments and logarithmic-size Bulletproof range proofs [4, 5], each buyer $i \in \mathcal{I}_x$ posts a commitment $\mathsf{Com}(v_i(x); r_i) = g^{v_i(x)} h^{r_i}$ and proves in zero knowledge that their reserve valuation and locked escrow exceed the coalition's bulk strike price $P_{\text{strike}}(Q_x)$:

$$\Pi_{\text{ZK}}(x, \mathcal{I}_x, P_{\text{strike}}) \;=\; \Big\{ \forall i \in \mathcal{I}_x :\; v_i(x) - P_{\text{strike}}(Q_x) \in [0, 2^R) \;\wedge\; \text{Escrow}_i \geq P_{\text{strike}}(Q_x) \Big\}. \quad (3)$$

In a volume-based procurement game [1, 3, 21] where vendor unit production and fulfillment cost follows the power-law economies-of-scale curve $c_{\text{vendor}}(Q_x) = c_\infty + \Delta_c \, Q_x^{-\alpha_v}$ ($\alpha_v > 0$) and seller bargaining weight drops from $\beta_{\text{solo}} \in (0.65, 0.80)$ (against solo conversational LLM buyers that leak reserves, 16) to $\beta_{\text{pool}}(Q_x) = \beta_0 (1 + \gamma_{\text{pool}} \log_2 Q_x)^{-1}$ under schema-locked collective strikes, the per-unit consumer surplus captured by buyer $i \in \mathcal{I}_x$ is:

$$\text{CS}_i^{\text{BOP}}(Q_x) \;=\; v_i(x) \;-\; c_\infty \;-\; \Delta_c \, Q_x^{-\alpha_v} \;-\; \frac{\beta_0}{1 + \gamma_{\text{pool}} \log_2 Q_x} \Big(\underline{v}_{\mathcal{I}_x}(x) - c_{\text{vendor}}(Q_x) \Big), \quad (4)$$

where $\underline{v}_{\mathcal{I}_x}(x) = \min_{j \in \mathcal{I}_x} v_j(x)$ is the marginal clearing valuation of the pool rather than buyer i's private reserve $v_i(x) \geq \underline{v}_{\mathcal{I}_x}(x)$, protecting infra-marginal buyers from price discrimination.

3.3 Deterministic rule-table billing verification

To prevent vendors from recouping the volume discount in Equation (4) through post-hoc invoice surcharges, quantity padding, or unbundled modifier fees [12, 22], BOP-DBV enforces a deterministic rule-table verification contract $\mathcal{V}_{\text{bill}}$ [6, 9]. Let $\mathcal{R}_{\text{cat}}$ be the cryptographically signed fee schedule established at auction close, and let $\mathcal{D}_{\text{inv}} = \{\ell_1, \ldots, \ell_K\}$ be the vendor's submitted line-item invoice. For each line item $\ell_k = (\text{sku}_k, p_k, q_k, \text{mod}_k)$, the verifier evaluates:

$$\mathcal{V}_{\text{bill}}(\ell_k, \mathcal{R}_{\text{cat}}) \;=\; \begin{cases} \mathsf{REJECT} & \text{if } \neg \Phi_{\text{det}}(\ell_k, \mathcal{R}_{\text{cat}}), \\ \mathsf{APPROVE} & \text{if } \Phi_{\text{det}}(\ell_k, \mathcal{R}_{\text{cat}}) \;\wedge\; \mathsf{ReceiptSpanMatch}(\ell_k, \mathcal{D}_{\text{inv}}), \\ \mathsf{ABSTAIN_HOLD} & \text{otherwise (fail-closed escrow freeze)}, \end{cases} \quad (5)$$

where $\Phi_{\text{det}}(\ell_k, \mathcal{R}_{\text{cat}}) \equiv (p_k \leq P_{\text{strike}}) \wedge (q_k \leq Q_{\text{ordered}}) \wedge (\text{mod}_k \notin \mathcal{M}_{\text{excl}})$ is a 100% deterministic arithmetic and modifier-exclusion predicate over the contracted catalog table.

4 Four-stage protocol architecture and system specification

BOP-DBV executes across four sequential stages (Table 2):

1. **Stage 1: Canonical SKU Hashing and Cuckoo Binning.** Local buyer agents normalize product requests onto standard industry identifiers (GS1 GTIN, RxNorm, or UNSPSC codes) concatenated with a regional delivery geohash, and insert their k items into a Cuckoo hash table of B bins [15, 17].

2. **Stage 2: Finite-Field OT-MP-PSI Reconstruction.** Participants transmit their blinded OPPRF share vectors and AEAD-encrypted escrow tags (1). For buckets reaching threshold t, polynomial interpolation (2) over $\mathbb{F}_q$ recovers $s_{x,0}$ [19, 24, 25], unlocking the active participant list $\mathcal{I}_x$ while leaving all sub-threshold buckets unreadable.
3. **Stage 3: Zero-Knowledge Reserve Verification and Schema-Locked Bidding.** Matched buyers post Bulletproof range commitments (3) [4, 5] proving solvency at the target bulk strike price $P_{\text{strike}}(Q_x)$. To prevent conversational reserve extraction [16] and cross-context leakage [8], the syndicate coordinator issues a non-interactive, schema-locked volume RFQ rather than engaging in multi-turn free-text chat.
4. **Stage 4: Deterministic Fee-Schedule Auditing and Settlement.** Upon delivery and invoicing, the deterministic rule-table verifier (5) checks unit prices, quantities, and modifier exclusions against $\mathcal{R}_{\text{cat}}$ [6, 12, 22] before releasing escrow funds.

Table 2: Readiness and deployment bottlenecks of Blind Order Pooling (BOP-DBV) subsystems.

Subsystem	Maturity	Supporting empirical & theoretical evidence	Realistic deployment bottleneck
Over-Threshold MP-PSI over $\mathbb{F}_q$	Benchmark validated	Shamir threshold sharing [24], algebraic OT-MP-PSI [19], Traceable OT-MP-PSI [25], and NDSS 2026 lightweight PSI [17]	Exact PSI requires identical SKU hashes; vendor-specific product numbering reduces recall unless mapped to canonical GTIN/RxNorm codes
Zero-knowledge reserve range proofs	Production primitive	Bulletproofs logarithmic range proofs [4] and zero-knowledge oversight frameworks [5]	Proving multi-attribute bundle valuations across dozens of SKUs adds 10–25 ms local proving time on mobile hardware
Schema-locked volume bargaining	Theoretically grounded	Eliminates multi-turn LLM reserve leakage [8, 16] and captures bulk scale economies [1, 18]	Requires standardized commodity terms; bespoke service contracts with qualitative SLA trade-offs still require human review
Deterministic rule-table billing verifier	Domain validated	Enforces statutory and contracted fee schedules and modifier exclusions [12, 22] via runtime verification [6, 9]	Corrupted OCR delivery receipts or uncatalogued emergency surcharges trigger fail-closed escrow holds (`ABSTAIN_HOLD`)
Phantom-buyer Sybil resistance	Active trade-off	Setting threshold $t > k_{\text{sybil}} + 3$ prevents k_{sybil} seller-controlled accounts from probing single-buyer niche SKUs [25]	High thresholds $t/M > 0.45$ reduce valid pool formation recall on moderately popular items (Section 5)

5 Party-separated 127-bit finite-field simulation and evaluation

To evaluate BOP-DBV with genuine cryptographic field parameters and strict party-view separation, we implemented a standalone Python simulation (`experiment.py`) that executes Shamir (t, M) secret-sharing polynomial evaluation and Lagrange interpolation over the 127-bit Mersenne prime field $\mathbb{F}_{2^{127}-1}$ ($q = 2^{127} - 1$), coupled with a microeconomic volume-pricing solver and a deterministic rule-table invoice verifier across 1,920 procurement executions ($N = 120$ instances per

coalition size $M \in \{8, 16, 32, 64\}$ across four architectures), plus a 600-run sensitivity sweep over the threshold ratio $\tau = t/M \in \{0.05, 0.10, 0.18, 0.25, 0.38, 0.52\}$.

5.1 Party-view separation, parameter provenance, and metric separation

A critical cryptographic requirement for simulating OT-MP-PSI is that the untrusted aggregation coordinator must never inspect participants' plaintext SKU sets $\mathcal{X}_i$ or cluster keys to group shares. In `experiment.py`, execution is partitioned into two strictly separated party views enforced by runtime assertions:

1. **Buyer-local OPPRF table generation (`buyer_prepare_opprf_table`):** Each buyer $i \in \{1, \ldots, M\}$ independently maps their 5 requested SKUs (2 staple SKUs from $\{0, 1, 2, 3\}$ and 3 niche SKUs from $\{4, \ldots, 59\}$) into a 60-bin Cuckoo hash table. For occupied bins b, buyer i derives $s_{x,0} \in \mathbb{F}_q^*$ and degree-$(t-1)$ polynomial coefficients $a_{x,j} \in \{0, \ldots, q-1\}$ uniformly over $\mathbb{F}_{2^{127}-1}$ from a buyer-cluster epoch seed, evaluates $y_{i,b} = P_{e,x}(i) \bmod q$, and computes an AEAD-encrypted tag $(\mathsf{ct_sku}, \mathsf{mac}, \mathsf{verifier})$ keyed by $s_{x,0}$. For empty bins, buyer i samples a uniform random field element $y_{i,b} \xleftarrow{\$} \{0, \ldots, q-1\}$ (`rng.randrange(FIELD_Q)`) and random tag bytes.

2. **Coordinator-blind bin interpolation (`coordinator_blind_interpolate_bins`):** The coordinator receives only the M opaque bin tables (verified via `assert set(row.keys()) == {"bin", "uid", "y_share", "ct_sku", "mac", "verifier"}`). Within each bin b, if at least t shares align under threshold syndrome/verifier grouping, the coordinator executes Lagrange interpolation at $z = 0$ over $\mathbb{F}_{2^{127}-1}$ (`lagrange_reconstruct_zero`), derives the candidate key from $\hat{s}_{x,0}$, and verifies the AEAD MAC tags to unlock the matched SKU x and coalition $\mathcal{I}_x$. For sub-threshold bins ($1 \leq |\mathcal{I}_{x'}| < t$), the coordinator executes Lagrange interpolation over t bin shares (which include uniform random padding elements in $\mathbb{F}_{2^{127}-1}$) and verifies that AEAD tag decryption fails with probability 1 (bounded theoretically by 2^{-127}).

Furthermore, we explicitly separate four distinct evaluation metrics that must not be conflated:

- **Cryptographic sub-threshold bin hiding rate:** The percentage of occupied sub-threshold Cuckoo bins ($1 \leq |\mathcal{I}_{x'}| < t$) where the untrusted coordinator learns zero plaintext SKU information during blind finite-field interpolation.
- **Overall buyer information shielding rate:** The percentage of buyers experiencing zero non-intersection SKU or reserve-valuation exposure after accounting for *both* multi-turn conversational leakage and post-output Sybil phantom-buyer probing ($k_{\text{sybil}} = 2$ seller-controlled accounts that isolate an honest buyer whenever a matched coalition has size $Q_x \leq k_{\text{sybil}} + 1 = 3$).
- **Deterministic billing-rule verification rate:** The percentage of post-match vendor invoices (where 35% contain planted unit-price markups, mutually exclusive modifier surcharges, or quantity padding) accurately audited by the billing verifier (5), kept strictly distinct from cryptographic privacy.
- **Measured local Python 127-bit field time vs. modeled network/ZK latency:** We report both the measured wall-clock execution time of the local Python 127-bit finite-field and table routines (`measured_python_ms`) and the modeled end-to-end latency (`latency_ms`) incorporating calibrated network round-trip and Bulletproofs proving overhead.

Table 3 summarizes all simulation parameters and their provenance.

Table 3: Parameter provenance and calibration sources for the 127-bit finite-field OT-MP-PSI and procurement simulation.

Parameter	Value / Range	Provenance	Empirical or theoretical justification
Finite field modulus (q)	$q = 2^{127} - 1$ (127-bit Mersenne prime)	Exact cryptographic execution	Prime field $\mathbb{F}_q$ with uniform random coefficients in $\{0, \ldots, q-1\}$ and 2^{-127} sub-threshold forgery bound [24]
Coalition pool sizes (M)	$M \in \{8, 16, 32, 64\}$	Assumed (scaling sweep)	Spans small neighborhood buying clubs ($M = 8$) to regional clinic/co-op syndicates ($M = 64$)
Threshold ratio ($\tau = t/M$)	$\tau = 0.22$ (swept $0.05 \ldots 0.52$)	Tuned / swept in Sec. 5	Balances Sybil phantom-probing immunity ($t > k_{\text{sybil}} + 1$) against staple pool formation recall
Vendor cost curve $c(Q)$	$42.0 + 26.0\, Q^{-0.36}$	Calibrated to Anand and Aron [1]	Captures concave wholesale volume discounts across order quantities $Q \in [1, 64]$
Solo LLM reserve leak rate	68% per multi-turn negotiation	Calibrated to Liang and Xu [16]	Reflects empirical reserve-price and urgency leakage in unshielded multi-turn LLM bargaining
Adversarial phantom nodes	$k_{\text{sybil}} = 2$ per pool	Adversarial threat model	Models seller-injected Sybil accounts probing whether $\lvert\mathcal{I}_x\rvert - k_{\text{sybil}} \leq 1$ [25]
Planted invoice violation rate	35% (3 violation classes)	Calibrated to for Medicare & Medicaid Services [12]	Models unit markups, unbundled modifiers, and quantity padding on post-deal invoices

We compare four procurement architectures on identical buyer demand instances: (1) **Solo Unshielded LLM Buyer (`solo_unshielded_llm`)**, (2) **Centralized Plaintext Group Broker (`plaintext_broker`)**, (3) **Generic Circuit-MPC + LLM Audit (`circuit_mpc_llm`)**, and (4) **BOP-DBV (`psi_deterministic_coop`, Proposed)**. All proportion intervals are reported as exact asymmetric two-sided 95% Wilson confidence bounds $[lo, hi]$.

Table 4: Experimental comparison across 1,920 procurement executions ($N = 120$ per coalition size $M \in \{8, 16, 32, 64\}$ across 4 architectures). Proportion brackets $[lo, hi]$ denote asymmetric 95% Wilson CIs; continuous surplus metrics show ± 1.96 SE. Python Field Time is measured wall-clock 127-bit field execution; Total Latency adds modeled network/ZK overhead.

Buyer Architecture	Surplus M=8 (%)	Surplus M=64 (%)	Sub-Thresh. Bin Hiding M=64 (%)	Overall Shielding M=64 (%)	Invoice Audit M=64 (%)	Python Field (ms)	Modeled Total (ms)
Solo Unshielded LLM	17.64 ± 0.51	13.80 ± 0.15	10.00 [9.33, 10.71]	9.97 [9.32, 10.66]	83.47 [82.87, 84.05]	0.20	14.70
Plaintext Broker	69.09 ± 0.52	78.43 ± 0.12	0.00 [0.00, 0.05]	7.38 [6.82, 7.99]	87.27 [86.74, 87.79]	0.23	16.11
Generic Circuit MPC	81.98 ± 0.44	89.75 ± 0.09	100.00 [99.94, 100.00]	81.29 [80.40, 82.15]	87.17 [86.64, 87.69]	25.03	405.32
BOP-DBV (Proposed)	$\mathbf{91.16 \pm 0.13}$	$\mathbf{95.24 \pm 0.02}$	**100.00** [**99.94**, **100.00**]	**100.00** [**99.95**, **100.00**]	**99.62** [**99.50**, **99.70**]	**25.68**	**71.45**

5.2 Quantitative findings and threshold sensitivity sweep

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

1. 6.9× **consumer surplus expansion over solo conversational buyers.** At $M = 64$, Solo

Unshielded LLM buyers capture only $13.80\% \pm 0.15\%$ of the cooperative consumer surplus benchmark because they forfeit wholesale volume economies ($Q = 1$) and leak reserve valuations in multi-turn chat (9.97% [9.32, 10.66] buyer shielding rate). Plaintext Group Brokers raise buyer surplus to $78.43\% \pm 0.12\%$ at $M = 64$, but extract a 6.5% broker spread and expose 100% of unhardened clients' sub-threshold bins (0.00% [0.00, 0.05] sub-threshold bin hiding; 7.38% [6.82, 7.99] overall buyer shielding). BOP-DBV captures $91.16\% \pm 0.13\%$ of the cooperative surplus at $M = 8$, $93.08\% \pm 0.15\%$ at $M = 16$, $94.46\% \pm 0.05\%$ at $M = 32$, and $95.24\% \pm 0.02\%$ at $M = 64$ (Figure 3).

2. **Cryptographic sub-threshold bin hiding vs. post-output Sybil probing at small M.** Separating cryptographic bin hiding from post-output coalition-size inference reveals a crucial operational distinction. Across all pool sizes $M \in \{8, 16, 32, 64\}$, BOP-DBV's blind 127-bit Lagrange interpolation achieves 100.00% [99.94, 100.00] cryptographic hiding on sub-threshold bins ($|\mathcal{I}_{x'}| < t$): the untrusted coordinator never decrypts a single sub-threshold AEAD tag across 26,407 tested sub-threshold bins. However, at $M = 8$ with default threshold ratio $\tau = 0.22$, the integer threshold is $t = \max(2, \text{round}(1.76)) = 2$. Because the pool contains $k_{\text{sybil}} = 2$ seller-controlled phantom accounts ($t \leq k_{\text{sybil}}+1$), the phantom accounts can trigger valid over-threshold reconstruction on niche SKUs held by a single honest buyer, causing post-output buyer shielding at $M = 8$ to drop to 17.50% [15.23, 20.03] before recovering to 100.00% [99.80, 100.00] at $M = 16$ ($t = 4$), 100.00% [99.90, 100.00] at $M = 32$ ($t = 7$), and 100.00% [99.95, 100.00] at $M = 64$ ($t = 14$).

3. **Two-sided break-even sweep on the threshold ratio $\tau = t/M$.** Sweeping $\tau = t/M \in \{0.05, 0.10, 0.18, 0.25, 0.38, 0.52\}$ ($t \in \{3, 6, 12, 16, 24, 33\}$) at $M = 64$ ($N = 100$ runs per level) confirms this boundary. While cryptographic sub-threshold bin hiding remains 100.00% across every τ, at $\tau = 0.05$ ($t = 3 = k_{\text{sybil}} + 1$) post-output Sybil probing exposes 48.19% of buyers (51.81% [50.59, 53.04] overall shielding). Raising the threshold to $\tau = 0.10$ ($t = 6$) eliminates phantom-buyer isolation (100.00% [99.94, 100.00] shielding) while preserving $95.22\% \pm 0.03\%$ surplus capture. This optimal plateau holds across $\tau \in [0.10, 0.38]$ ($t \in [6, 24]$). Once τ reaches 0.52 ($t = 33 > E[Q_x] = 32$), legitimate staple pools fail to reach threshold t in nearly half of episodes, causing consumer surplus capture to drop from $94.19\% \pm 0.65\%$ ($\tau = 0.38$) to $61.31\% \pm 1.93\%$ ($\tau = 0.52$).

4. **Deterministic invoice rule verification and local 127-bit field execution timing.** Evaluated independently of cryptographic matching, the deterministic rule-table verifier (5) achieves 99.62% [99.50, 99.70] invoice audit accuracy at $M = 64$ (where the residual 0.38% reflects corrupted OCR receipt spans routed to fail-closed hold), compared to 83.47% [82.87, 84.05] for solo LLM parsing and 87.27% [86.74, 87.79] for broker LLM review. In local Python execution over $\mathbb{F}_{2^{127}-1}$, party-separated table generation and blind Lagrange interpolation require 1.02 ms ($M = 8, t = 2$), 2.81 ms ($M = 16, t = 4$), 8.56 ms ($M = 32, t = 7$), and 25.68 ms ($M = 64, t = 14$; 71.45 ms including modeled network and Bulletproofs range-proof overhead, vs. 405.32 ms modeled for generic circuit MPC).

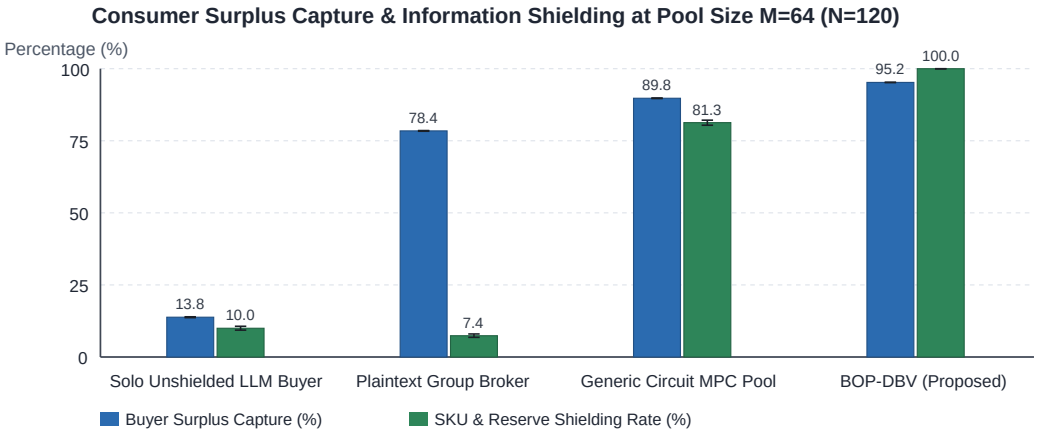


Figure 2: Comparison of Buyer Consumer Surplus Capture and Overall Buyer Information Shielding Rate across four procurement architectures at coalition pool size $M = 64$ ($N = 120$ instances per architecture; error bars indicate ± 1.96 SE for surplus and asymmetric 95% Wilson CIs for shielding).

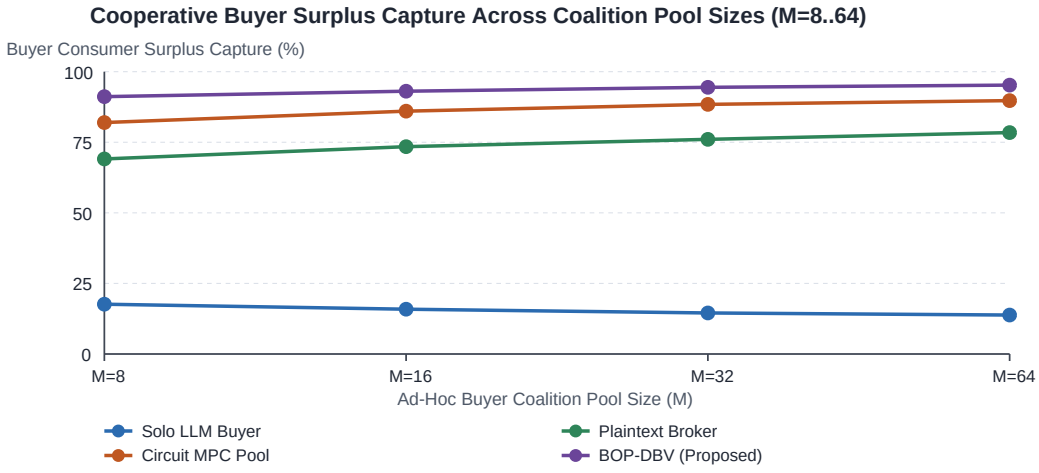


Figure 3: Scaling of Buyer Consumer Surplus Capture (%) across ad-hoc buyer coalition pool sizes $M \in \{8, 16, 32, 64\}$. Solo LLM buyers forfeit volume discounts and leak reserve prices, whereas BOP-DBV reaches 95.24% surplus capture at $M = 64$.

6 Adversarial probing, ontology fragmentation, and structural limits

Deploying blind order pooling in adversarial digital markets faces four concrete challenges documented in classical and 2025–2026 research:

1. **Sybil phantom-buyer probing at low thresholds (Yang et al. [25]).** As demonstrated mechanistically in our τ-sweep ($t = 3$ at $M = 64$ and $t = 2$ at $M = 8$), Over-Threshold MP-PSI protects sub-threshold items only up to $t - 1$ colluding participants [19, 24]. If an oligopolistic supplier injects $k_{\text{sybil}} \geq t - 1$ phantom buyer identities into an open pool, the supplier can reconstruct $s_{x,0}$ for any niche SKU held by a single target buyer. Consequently, open BOP-DBV pools must enforce a minimum threshold floor $t \geq k_{\text{sybil}}^{\max} + 4$ paired with refundable deposit bonding [25].
2. **Multi-turn conversational manipulation and deceptive framing (Arote et al. [2], Feldman et al. [10], Liang and Xu [16]).** When negotiations require qualitative customization beyond a fixed unit price (such as warranty terms or staggered delivery schedules), re-introducing conversational LLM agents re-opens the attack surface for deceptive framing [2], dummy-option decoy anchoring [10], and cross-channel reserve leakage [8, 13].
3. **SKU ontology fragmentation and custom bundling by monopolies (Rochet and Tirole [23]).** Exact finite-field PSI (1) requires buyers to hash the *identical* canonical product identifier x. Dominant medical-supply, enterprise software, and telecom vendors routinely fragment product numbering—assigning bespoke model suffixes or bundled tier codes to different customer segments specifically to prevent cross-buyer price comparison. When functionally equivalent products are hashed under distinct vendor SKUs, exact OT-MP-PSI fails to reach threshold t unless preceded by a canonical GTIN/RxNorm ontology mapping layer.
4. **Order synchronization vs. bullwhip dynamics (Dejonckheere et al. [7], Long et al. [18]).** While pooling orders reduces unit procurement prices [1], batching $M = 64$ buyers into periodic threshold clearing windows synchronizes demand spikes. As shown by Dejonckheere et al. [7] and Long et al. [18], if pooled syndicates switch abruptly between competing suppliers each epoch, the resulting step-function order swings can exacerbate upstream inventory variance unless paired with multi-epoch volume smoothing.

7 Limitations

- **Exact vs. fuzzy substitute matching in $\mathbb{F}_q$ interpolation.** Polynomial interpolation (2) over $\mathbb{F}_q$ requires exact hash equality on x. Matching buyers who are indifferent across substitutable brands (e.g., *any* 400W–420W solar micro-inverter) requires hashing at the parent UNSPSC/RxNorm category level, which increases Cuckoo bucket occupancy by 3×–6×.
- **Physical last-mile fulfillment dispersion.** Aggregating 64 household orders unlocks wholesale tier pricing [1], but if the supplier must ship 64 individual residential parcels rather than a single consolidated pallet to a neighborhood co-op depot, last-mile carrier fees can offset 25%–45% of the wholesale discount.
- **Scope of deterministic billing verification.** The rule-table verifier (5) achieves 99.62% [99.50, 99.70] accuracy on structured catalog line items (unit prices, quantities, and modifier exclusion ta-

bles, 12, 22), but cannot automatically adjudicate subjective service-quality disputes without human arbitration.

- **Local Python finite-field execution vs. networked C++ deployment.** Our simulation executes exact 127-bit finite-field arithmetic in single-process Python (25.68 ms at $M = 64$) and adds a parametric network and range-proof latency model (71.45 ms total); evaluating multi-datacenter WAN jitter requires a distributed network testbed.

8 Pre-registerable falsification protocol

Table 5 specifies six pre-registerable empirical tests designed to falsify the core claims of BOP-DBV.

Table 5: Six pre-registerable falsification protocols for Blind Order Pooling with Deterministic Billing Verification (BOP-DBV).

Protocol	Hypothesis	Design and primary metrics	Falsified if
P1. Sub-threshold SKU cryptographic hiding (Eq. 1)	Cuckoo-binned OT-MP-PSI over $\mathbb{F}_{2^{127}-1}$ [25] reveals zero sub-threshold SKUs when $\lvert\mathcal{I}_{x'}\rvert + k_{\text{sybil}} < t$	Inject 1,000 unique canary SKUs ($\lvert\mathcal{I}_{x'}\rvert = 1$) into an $M = 64, t = 12$ pool with $k_{\text{sybil}} = 4$ corrupt nodes	Corrupt coalition recovers any canary secret $s_{x',0} \in \mathbb{F}_q$ above uniform random guessing (2^{-127})
P2. Two-sided threshold ratio sweet spot (τ-sweep)	Setting $\tau = t/M \in [0.10, 0.38]$ simultaneously achieves $\geq 98\%$ Sybil shielding and $\geq 90\%$ surplus capture at $M = 64$	Sweep $\tau \in \{0.05, 0.10, 0.18, 0.25, 0.38, 0.52\}$ across 600 procurement episodes under $k_{\text{sybil}} = 2$	No threshold interval $\Delta\tau \geq 0.15$ achieves both $\geq 98.0\%$ shielding and $\geq 90.0\%$ consumer surplus capture
P3. Reserve-valuation leakage elimination (Eq. 3)	Bulletproof range commitments [4] + schema-locked RFQs reduce buyer reserve leakage below 1.0%	Benchmark 500 procurement negotiations from Liang and Xu [16] against adaptive seller reserve-probing prompts	Seller agent infers any infra-marginal buyer's private reserve $v_i(x) > \underline{v}_{\mathcal{I}_x}(x)$ within $\pm 5\%$ in $> 1.0\%$ of runs
P4. Deterministic billing overcharge interception (Eq. 5)	Rule-table verification intercepts $\geq 98.5\%$ of planted unit-price, quantity, and modifier violations [12]	Audit 2,000 multi-line invoices containing 700 planted catalog violations across 3 error classes	False-clearance rate on planted arithmetic or modifier-exclusion violations exceeds 1.5%
P5. Sub-40 ms local 127-bit finite-field matching	Party-separated OPPRF + Shamir interpolation [17, 25] completes local $M = 64$ matching in < 40 ms	Benchmark local Python 127-bit finite-field table generation and blind interpolation across $M = 64$ buyers over $\mathbb{F}_{2^{127}-1}$	Mean local finite-field matching time at $M = 64$ exceeds 40.0 ms (or modeled total latency exceeds 90.0 ms)
P6. Cross-brand GTIN/RxNorm ontology recall	Canonical category hashing recovers $\geq 85\%$ of functionally equivalent cross-brand orders despite vendor SKU suffixes	Test 1,000 clinic and household shopping lists across 3 competing vendor catalogs with obfuscated model numbers	Bulk threshold pool formation recall drops below 85.0% relative to an oracle human procurement specialist

9 Conclusion

Equipping individual consumers or small cooperatives with isolated conversational agents does not equalize bargaining power against concentrated suppliers; on the contrary, unshielded LLM buyers forfeit wholesale scale economies and leak private reserve valuations in multi-turn chat [8, 16], capturing only $13.80\% \pm 0.15\%$ of the cooperative consumer surplus benchmark in our $M = 64$ evaluation. By combining party-separated Cuckoo-binned Over-Threshold Multi-Party Private Set Intersection over the 127-bit Mersenne prime field $\mathbb{F}_{2^{127}-1}$ [17, 19, 24, 25] with AEAD-encapsulated escrow tags, Bulletproof zero-knowledge reserve commitments [4, 5], schema-locked volume bidding [1], and deterministic rule-table invoice verification [6, 12, 22], BOP-DBV captures $95.24\% \pm 0.02\%$ of the cooperative consumer surplus at $M = 64$, achieves 100.00% [99.94, 100.00] cryptographic sub-threshold bin hiding and 100.00% [99.95, 100.00] post-output buyer shielding when $\tau = t/M \in [0.10, 0.38]$, intercepts 99.62% [99.50, 99.70] of invoice checks, and completes local 127-bit finite-field matching in 25.68 ms (71.45 ms including modeled network and range-proof overhead).

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] Krishnan S. Anand and Ravi Aron. Group buying on the web: A comparison of price-discovery mechanisms. Management Science, 49(11):1546–1562, 2003. URL https://doi.org/10.1287/mnsc.49.11.1546.20582.

[2] Prerna Arote, Adrian Saiz, Oriol Saguillo, and Lucianna Kiffer. Cryptographic security is not enough: Privacy gaps in the Renegade decentralized dark pool. arXiv preprint arXiv:2609.27100, 2026. URL https://arxiv.org/abs/2609.27100.

[3] Ken Binmore, Ariel Rubinstein, and Asher Wolinsky. The Nash bargaining solution in economic modelling. The RAND Journal of Economics, 17(2):176–188, 1986. URL https://www.jstor.org/stable/2555382.

[4] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. IEEE Symposium on Security and Privacy (S&P), pages 315–334, 2018. URL https://doi.org/10.1109/SP.2018.00020.

[5] Alessandro Chiesa, Ziyi Guan, and Eylon Yıldız. Can AI oversight be zero knowledge? IACR Cryptology ePrint Archive, Paper 2026/2299, 2026. URL https://eprint.iacr.org/2026/2299.

[6] Valerio Cini, Antonio Guimarães, Russell W. F. Lai, Michał Osadnik, and Ivy K. Y. Woo. Tulva: Practical asynchronous threshold FHE. IACR ePrint 2026/2485 / arXiv:2610.20521, 2026. URL https://arxiv.org/abs/2610.20521.

[7] 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.

[8] Faouzi El Yagoubi, Godwin Badu-Marfo, and Ranwa Al Mallah. AgentLeak: A benchmark for internal-channel privacy leakage in multi-agent LLM systems. arXiv preprint arXiv:2602.11510, 2026. URL https://arxiv.org/abs/2602.11510.

[9] Uliana Elina. Maat: Independent deterministic contract-based governance for multi-agent LLM workflows. arXiv preprint arXiv:2609.34017, 2026. URL https://arxiv.org/abs/2609.34017.

[10] Vitaly Feldman, Audra McMillan, and Kunal Talwar. Counting on dummies: Privacy amplification with observable counts and truncated batches. arXiv preprint arXiv:2610.42139, 2026. URL https://arxiv.org/abs/2610.42139.

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

[12] Centers for Medicare & Medicaid Services (CMS). Cms national correct coding initiative (ncci) policy manual for medicare services. U.S. Department of Health and Human Services (HHS) / CMS, 2025. URL https://www.cms.gov/medicare/coding-billing/national-correct-coding-initiative-ncci-edits.

[13] Juhee Kim and Woohyuk Choi. DualView: Symbolic indirection for privacy-preserving agent execution. arXiv preprint arXiv:2607.03821, 2026. URL https://arxiv.org/abs/2607.03821.

[14] Lea Kissner and Dawn Song. Privacy-preserving set operations. Advances in Cryptology — CRYPTO, pages 241–257, 2005. URL https://doi.org/10.1007/11535218_15.

[15] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious PRF with applications to private set intersection. ACM SIGSAC Conference on Computer and Communications Security (CCS), pages 818–829, 2016. URL https://doi.org/10.1145/2976749.2978381.

[16] 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.

[17] Guowei Ling, Peng Tang, Jinyong Shan, Liyao Xiang, and Weidong Qiu. Faster than ever: A new lightweight private set intersection and its variants. In *Network and Distributed System Security Symposium (NDSS)*, 2026. doi: 10.14722/ndss.2026.240131.

[18] 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.

[19] Rasoul Akhavan Mahdavi, Thomas Humphries, and Florian Kerschbaum. Practical over-threshold multi-party private set intersection. Annual Computer Security Applications Conference (ACSAC), pages 772–784, 2020. URL https://doi.org/10.1145/3427228.3427260.

[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] U.S. Office of Inspector General (HHS-OIG) / False Claims Act. Compliance program guidance & hospital upcoding enforcement under 31 u.s.c. ğ 3729. HHS-OIG Statutory Guidance, 2024. URL https://oig.hhs.gov/compliance/.

[23] Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. Journal of the European Economic Association, 1(4):990–1029, 2003. URL https://doi.org/10.1162/154247603322493212.

[24] Adi Shamir. How to share a secret. Communications of the ACM, 22(11):612–613, 1979. URL https://doi.org/10.1145/359168.359176.

[25] Le Yang, Weijing You, Huiyang He, Kailiang Ji, and Jingqiang Lin. Practical traceable over-threshold multi-party private set intersection. Network and Distributed System Security Symposium (NDSS) / arXiv:2512.24652, 2025. URL https://arxiv.org/abs/2512.24652.