HectoVault: A Framework for Adversarial-Secure Deal Optimization

Technical Whitepaper v3.0

Abstract

HectoVault presents a framework for adversarial-secure optimization in commodity trading using Secure Multi-Party Computation (MPC). We enable multiple parties to discover optimal trade configurations without revealing confidential inputs. Acknowledging fundamental constraints in both technology and mechanism design, we propose a pragmatic approach focused on high-value, complexity-bounded scenarios where privacy preservation justifies computational overhead. Our economic model abandons traditional fee-on-surplus approaches in favor of subscription-based access with outcome-independent pricing, eliminating incentives for strategic manipulation. This paper presents our technical architecture, explicitly acknowledges scalability limitations, and defines the narrow but valuable market segments where the framework provides genuine utility.

1. Introduction: Reframing the Problem

The commodity trading paradox—optimal outcomes require information sharing, yet sharing erodes competitive advantage—cannot be fully "solved" by any technical system. Instead, we identify specific scenarios where:

- The value of privacy preservation exceeds computational costs.
- The number of participants is naturally bounded.
- Strategic behavior can be constrained through mechanism design.
- Time sensitivity permits multi-minute computation cycles.

Rather than claiming broad applicability, we focus on these constrained but valuable use cases.

2. Technical Architecture

2.1 MPC Protocol Selection

We implement variants of the BGW/GMW protocol families optimized for specific trade structures:

- **Honest-majority settings:** Information-theoretic security with lower overhead.
- **Dishonest-majority settings:** Computational security with higher resilience.
- **Trade-off:** Accept higher latency for stronger adversarial models when required.

Critical Acknowledgment: The Simplex algorithm's data-dependent branching makes efficient secure implementation challenging. We employ:

- Approximation algorithms with bounded iterations for large problems.
- Exact solutions only for problems with <1000 variables.
- Pre-computation phases to amortize costs where possible.

2.2 Computation Classes

We explicitly limit supported computations to maintain practical performance:

- **Class A: Simple Linear Allocation (seconds to minutes)**
  - Basic supply-demand matching
  - Linear objective, linear constraints
  - No integer requirements
- **Class B: Constrained Integer Programs (minutes to hours)**
  - Discrete unit trades
  - Limited integer variables (<100)
  - Simplified constraint structures
- **Class C: Multi-Objective Optimization (hours)**
  - Generate 3-5 Pareto-optimal points maximum
  - Focus on extremal solutions rather than full frontier
  - Acknowledge this provides guidance, not comprehensive analysis

2.3 Scalability Reality

Hard Constraint: Our architecture supports a maximum of **20 participants** due to $O(n^2)$ communication complexity.
This is not a bug—it defines our market. We explicitly target:

- Regional commodity pools
- Specialized product markets
- Consortium-based trading groups
- High-value, low-volume segments

For larger markets, we offer hierarchical approaches where sub-groups optimize locally before inter-group coordination.

3. Economic Model: Avoiding the Baseline Trap

3.1 Why Fee-on-Surplus Fails

Any model comparing "optimized" vs "baseline" outcomes invites manipulation: parties can collude to establish poor baselines, creating an artificial surplus to be shared. The platform becomes an unwitting accomplice to this collusion. We reject this entire approach.

3.2 Subscription-Based Access Model

Our economic model is built on transparency and neutrality:

- **Fixed Subscription Tiers:**
  - Based on computation complexity classes (A, B, C).
  - Monthly/annual pricing with no dependence on computed "value."
- **Usage-Based Components:**
  - Computation time (node-hours).
  - Data volume processed.
  - Priority queue access.
- **Outcome-Independent Pricing:**
  - Parties pay regardless of optimization results.
  - This removes the incentive to manipulate inputs and aligns the platform as neutral infrastructure.

3.3 Value Proposition

Subscribers pay for a secure computational capability, not magical value creation. The value is in:

- Privacy-preserving computation infrastructure.
- Cryptographic guarantees of confidentiality.
- Audit trails for regulatory compliance.
- Elimination of information leakage risk.

4. Security Architecture

4.1 Adversarial Model

We assume sophisticated adversaries who control up to 49% of computation nodes, coordinate across parties, and possess deep market knowledge. We explicitly **do not** claim resistance against:

- Majority collusion
- Nation-state adversaries
- Zero-day protocol vulnerabilities
- Social engineering of all participants

4.2 Economic Security

- **Performance Bonds:** Sized at 2-5x maximum monthly revenue per node, held in escrow. Forfeiture requires cryptographic proof and governance review.
- **Reputation System:** Historical reliability tracking, public disclosure of violations, and exclusion from high-value computations. We make no claim of "trustlessness"—

reputation matters.

4.3 Practical Security Measures

Beyond cryptography, we mandate:

- Operational security audits.
- Background checks for node operators.
- Insurance requirements for high-value computations.
- Clear legal agreements with jurisdiction clauses.

5. Use Case Analysis

5.1 Where HectoVault Works

- **Regional Energy Trading Pools (5-15 participants):** Daily/weekly optimization cycles where privacy is crucial for competitive positioning.
- **Specialized Commodity Consortiums (10-20 members):** Markets for rare earth elements, pharmaceutical precursors, etc., where trust exists but requires verification.
- **Strategic Reserve Allocations (3-10 entities):** Government and emergency response coordination where privacy is paramount over speed.

5.2 Where HectoVault Fails

- **Global Spot Markets & Commodity Exchanges:** Too many participants, real-time requirements, and low margins cannot support computation costs.
- **High-Frequency Trading:** Latency requirements are impossible to meet.

6. Implementation Roadmap

- **Phase 1: Foundation (Current):** Class A computations, 5-10 participants, single-region operations.
- **Phase 2: Expansion (12-18 months):** Class B computations, hierarchical coordination, multi-region support.
- **Phase 3: Maturity (24+ months):** Limited Class C support, specialized hardware acceleration, hybrid protocols.

Non-Goals: We explicitly will NOT pursue real-time trading, support for >20 direct participants, a general-purpose optimization platform, or cryptocurrency/DeFi integration.

7. Performance Engineering

7.1 Realistic Benchmarks

We provide bounds rather than specific numbers, as performance depends on network topology, problem structure, and security parameters.

- **Computation overhead:** 100-1000x vs. plaintext
- **Communication rounds:** O(depth of circuit)
- **Network requirements:** All-to-all connectivity

7.2 Optimization Strategies

To achieve practical performance, we use extensive pre-computation, problem-specific protocol selection, and approximate algorithms for large instances. We do not claim breakthroughs in core MPC efficiency but focus on competent implementation of existing techniques.

8. Governance and Operations

- **Technical Governance:** A conservative approach to protocol updates, requiring external audits and backward compatibility. Incidents are met with 24-hour disclosure and coordinated patching.
- **Business Governance:** A transparent pricing committee, published rate cards, and defined dispute resolution procedures with binding arbitration.

9. Competitive Analysis

- **vs. Trusted Third Parties (Clearing Houses):** HectoVault provides computation without data disclosure.
- **vs. Bilateral Negotiations:** HectoVault discovers multi-party optimizations impossible in direct deals.
- **vs. Centralized Platforms:** HectoVault offers privacy preservation as its key differentiator.

10. Risk Assessment

- **Technical Risks:** Protocol vulnerabilities, side-channel attacks, and the long-term threat of quantum computing.
- **Business Risks:** Limited market size, adoption barriers due to complexity, and competition from "good enough" solutions.
- **Operational Risks:** Node operator collusion, adverse regulatory changes, and technology obsolescence.

11. Conclusion

HectoVault provides a specialized tool for a specific class of problems. We do not claim to

revolutionize commodity trading. Instead, we offer:

- Proven cryptographic techniques applied to a commercial problem.
- An honest assessment of limitations and appropriate use cases.
- Pragmatic engineering choices favoring reliability over novelty.
- A sustainable business model that avoids perverse incentives.

For the narrow set of scenarios where privacy truly matters, participant count is limited, and time permits careful computation, HectoVault provides genuine value. Outside these constraints, simpler solutions are more appropriate. Success is not measured by broad adoption but by effective service to the specific markets where our approach makes sense.

References

[1] Ben-Or, M., Goldwasser, S., and Wigderson, A. "Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation." STOC, 1988.
[2] Goldreich, O., Micali, S., and Wigderson, A. "How to Play any Mental Game." STOC, 1987.
[3] Bogetoft, P., et al. "Secure Multiparty Computation Goes Live." Financial Cryptography, 2009.
[4] Damgård, I., and Nielsen, J. "Scalable and Unconditionally Secure Multiparty Computation." CRYPTO, 2007.
[5] Keller, M. "MP-SPDZ: A Versatile Framework for Multi-Party Computation." ACM CCS, 2020.
[6] Mohassel, P., and Zhang, Y. "SecureML: A System for Scalable Privacy-Preserving Machine Learning." IEEE S&P, 2017.
[7] Demmler, D., Schneider, T., and Zohner, M. "ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation." NDSS, 2015.
[8] NIST. "Security Requirements for Cryptographic Modules." FIPS Publication 140-2, 2001.
[9] Naor, M., Pinkas, B., and Sumner, R. "Privacy Preserving Auctions and Mechanism Design." ACM EC, 1999.
[10] Meadows, C. "Formal Methods for Cryptographic Protocol Analysis: Emerging Issues and Trends." IEEE Journal on Selected Areas in Communications, 2003.

This whitepaper reflects our current technical understanding and business strategy. We reserve the right to modify our approach based on operational experience and market feedback.