title

Zach Kelling

Satschel, Inc.

April 2026

Abstract

We present the market making economics for Liquid DEX, the precompile-accelerated order matching engine of the Liquidity.io regulated alternative trading system. We formalize the revenue model (spread capture, maker rebates, latency arbitrage), prove the geographic moat theorem from the Lean 4 verification suite (on-premise participants capture 100% of spread while remote competitors at >500 km face negative expected returns at 5 bps), and project volume scaling from 5-node ($10M/day) to 50-node ($1B/day) configurations. The DEX precompile provides both concentrated liquidity AMM pools and a central limit order book (CLOB), enabling hybrid execution that subsumes both pure strategies. We analyze MEV resistance via FHE-encrypted mempool, impermanent loss for LP positions, fee tier economics (maker −1 bps / taker 2 bps institutional; maker 0 bps / taker 5 bps retail), and expected P&L per $1M daily volume at various spread levels. BlackRock's participation as an institutional market maker and authorized participant for on-chain ETF creation/redemption is modeled as a specific case. This work builds on the formally verified exchange infrastructure [1], the Lux LightSpeed DEX architecture [4], the GPU computation white paper [5], and the latency arbitrage analysis [8].

Contents

1 Introduction

Market making on centralized exchanges is a \$10B+/year revenue opportunity dominated by firms like Citadel Securities, Virtu Financial, and Jane Street. These firms profit by quoting two-sided markets, capturing the bid-ask spread, and exploiting latency advantages at co-located data centers. Their infrastructure investment (Mahwah NJ for NYSE, Carteret NJ for NASDAQ) runs into hundreds of millions of dollars for microsecond advantages.

Liquid DEX introduces a new venue class: a regulated, on-chain exchange with precompile-accelerated matching, where the geographic moat is formally verified rather than empirically observed [1, 8]. The venue operates at the Kansas City data center with 200 Gbps internal fabric and sub-microsecond matching latency, creating a provable advantage for co-located participants.

This paper addresses the economics of market making on this venue:

- Revenue model: spread, rebates, arbitrage.
- Formally verified geographic moat (Lean 4 proofs).
- Volume scaling projections (5–50 nodes).
- Comparison with NYSE, NASDAQ, Binance, Coinbase.
- Institutional participation model (BlackRock as MM).
- Risk analysis: impermanent loss, slashing, MEV.

2 Revenue Model

2.1 Sources of Revenue

A market maker on Liquid DEX earns from three sources:

Definition 1 (Market Maker Revenue). *For a market maker quoting bid b and ask a with fill rate ϕ on daily volume V:*

$$R_{spread} = (a - b) \cdot V \cdot \phi \cdot 10^{-4} \quad \textit{(spread capture)} \tag{1}$$

$$R_{rebate} = r \cdot V \cdot \phi \cdot 10^{-4} \quad \textit{(maker rebate)} \tag{2}$$

$$R_{arb} = \pi(A, rtt) \quad \textit{(latency arbitrage, see Section 3)} \tag{3}$$

Total daily revenue: $R = R_{spread} + R_{rebate} + R_{arb}$.

2.2 Fee Structure

Tier	Maker Fee	Taker Fee	Qualification
Institutional MM	−1 bps	2 bps	> \$100M monthly vol.
Professional	0 bps	3 bps	> \$10M monthly vol.
Active	0 bps	4 bps	> \$1M monthly vol.
Retail	1 bps	5 bps	Default

Table 1: Fee tiers. Negative maker fee = rebate paid to market maker.

Theorem 1 (Maker Rebate Incentive). *An institutional market maker with fill rate ϕ on volume V earns a guaranteed rebate of:*

$$R_{rebate} = 1 \ \textit{bps} \times V \times \phi \times 10^{-4}$$

This rebate is earned regardless of spread or price movement. At $\phi = 50\%$ on \$100M daily volume: $R_{rebate} = 1 \times 10^8 \times 0.5 \times 10^{-4} = \$500/day$. The rebate subsidizes tight quoting and attracts liquidity.

3 Geographic Moat

3.1 Formal Model

The geographic arbitrage model is proved in Lean 4 with Mathlib [1, 8]. We reproduce the key definitions and theorems.

Definition 2 (Arbitrage Profit)**.** *For an opportunity with spread s (bps), volume V (\$), decay rate δ (bps/ms), and participant round-trip time rtt (ms):*

$$\begin{aligned} gross(s, V) &= s \cdot V \cdot 10^{-4} && (4) \\ latencyCost(rtt, \delta, V) &= rtt \cdot \delta \cdot V \cdot 10^{-4} && (5) \\ \pi(s, V, \delta, rtt) &= gross(s, V) - latencyCost(rtt, \delta, V) && (6) \end{aligned}$$

Theorem 2 (On-Premise Full Capture [8])**.** *For the Kansas City venue with $rtt_{local} = 0.01$ ms:*

$$\pi_{local} = s \cdot V \cdot 10^{-4}$$

The on-premise market maker captures 100% of the gross spread. At millisecond resolution, latency cost rounds to zero.

Theorem 3 (Monotone Distance Decay [8])**.** *For any $rtt_1 < rtt_2$: $\pi(rtt_1) \geq \pi(rtt_2)$. Profit is monotonically decreasing in latency.*

Corollary 1 (Break-Even Distance)**.** *At spread s bps and decay δ bps/ms, the break-even RTT is $rtt_{be} = s/\delta$. For $s = 5$, $\delta = 1$: $rtt_{be} = 5$ ms ≈ 500 km. Beyond 500 km, 5 bps arbitrage is strictly unprofitable.*

3.2 City-Level P&L

Standard parameters: $s = 5$ bps, $V = \$10\text{M}$, $\delta = 1$ bps/ms.

City	Dist.	RTT	Gross	Cost	Net
Kansas City (local)	0 km	0.01 ms	\$5,000	\$0	**\$5,000**
St. Louis	400 km	4 ms	\$5,000	\$4,000	\$1,000
Chicago	800 km	8 ms	\$5,000	\$8,000	−\$3,000
Dallas	900 km	9 ms	\$5,000	\$9,000	−\$4,000
New York	1,800 km	18 ms	\$5,000	\$18,000	−\$13,000
London	7,500 km	75 ms	\$5,000	\$75,000	−\$70,000

Table 2: Arbitrage P&L by city. Only co-located participants are profitable at 5 bps.

3.3 Comparison with Traditional Venues

3.4 Kansas City On-Premise Advantage

The Kansas City venue provides:

- **200 Gbps internal**: ConnectX-7 NIC, 25 GB/s validator-to-validator.
- **10 Gbps external**: 1.25 GB/s market data ingress.
- $< 10\,\mu$**s RTT**: Same-rack communication.
- $< 1\,\mu$**s matching**: DEX precompile in-memory CLOB.

Venue	Match Latency	Settlement	Co-Location
NYSE (Mahwah)	$< 40\,\mu$s	T+1	NJ data center
NASDAQ (Carteret)	$\sim 50\,\mu$s	T+1	NJ data center
Binance	$\sim$ 5 ms	On-chain	AWS global
Coinbase	$\sim$ 10 ms	On-chain	AWS global
Liquid DEX (KC)	$< 1\,\mu$s	$<$ 60s	**KC data center**

Table 3: Venue latency comparison. Liquid DEX achieves sub-microsecond matching with sub-minute settlement.

Nodes	Agg. BW	GPU Mem	Est. Orders/s	Daily Vol. Cap
5	1,000 Gbps	640 GB	62.5M	\$10M–\$50M
10	2,000 Gbps	1.28 TB	125M	\$50M–\$200M
25	5,000 Gbps	3.2 TB	312.5M	\$200M–\$500M
50	10,000 Gbps	6.4 TB	625M	\$500M–\$1B+

Table 4: Volume capacity by node count. Daily volume estimates assume 6.5 active trading hours and average order size of \$1,000.

4 Volume Projections

4.1 Node Scaling

The venue scales by adding DGX Spark nodes with native 200 Gbps ConnectX-7 NICs [5]:

4.2 Volume Growth Model

Definition 3 (Volume Growth). *Daily trading volume follows a logistic growth model:*

$$V(t) = \frac{V_{\max}}{1 + \left(\frac{V_{\max}}{V_0} - 1\right) e^{-r \cdot t}}$$

where V_0 is initial volume, $V_{\max}$ is capacity-limited maximum, r is growth rate, and t is months since launch.

Phase	Month	Nodes	Daily Vol.	MMs Active
Launch	0–6	5	\$10M	2–3
Growth	6–18	10	\$100M	5–8
Scale	18–36	25	\$500M	10–15
Maturity	36+	50	\$1B+	20+

Table 5: Volume and market maker growth projections.

5 Expected P&L Analysis

5.1 Per-\$1M Daily Volume

5.2 Volume-Scaled P&L

Theorem 4 (Positive Expected Value for Co-Located MMs). *For a co-located market maker with $rtt \leq 0.01$ ms, average spread $s \geq 1$ bps, and maker rebate $r = 1$ bps, the expected daily*

Spread	Gross	Rebate	Total	Inventory	Net
1 bps	$100	$50	$150	−$30	$120
2 bps	$200	$50	$250	−$50	$200
3 bps	$300	$50	$350	−$70	$280
5 bps	$500	$50	$550	−$100	$450
10 bps	$1,000	$50	$1,050	−$150	$900

Table 6: Expected daily P&L per $1M volume for a co-located institutional market maker. Fill rate 50%, rebate −1 bps, inventory risk estimated at 20% of gross.

Daily Vol.	MM Share	Daily P&L	Monthly	Annual
$10M	30%	$1,350	$40,500	$493K
$100M	25%	$11,250	$338K	$4.1M
$500M	20%	$45,000	$1.35M	$16.4M
$1B	15%	$67,500	$2.03M	$24.6M

Table 7: Volume-scaled P&L. Assumes 3 bps average spread, 50% fill rate, −1 bps rebate, 20% inventory cost. MM share decreases as more MMs enter at higher volumes.

P&L per $1M volume is:

$$E[P\&L] = (s + r) \cdot V \cdot \phi \cdot 10^{-4} - C_{inv}$$

where ϕ *is fill rate and* C_{inv} *is inventory cost. For* $s \geq 1$*,* $r = 1$*,* $\phi = 0.5$*, and* $C_{inv} = 0.2 \cdot s \cdot V \cdot \phi \cdot 10^{-4}$*:*

$$E[P\&L] = (s + 1) \cdot 0.5 \cdot 10 - 0.2 \cdot s \cdot 0.5 \cdot 10 = 5(s + 1) - s = 4s + 5 > 0$$

The expected value is strictly positive for all $s \geq 1$*.*

6 AMM + CLOB Hybrid

The DEX precompile supports both execution modes:

6.1 Central Limit Order Book

The CLOB provides price-time priority matching with formally verified invariants [1]:

Theorem 5 (Price-Time Priority [1])**.** *For any two orders* o_1, o_2 *with* $\mathit{price}(o_1) = \mathit{price}(o_2)$ *and* $\mathit{time}(o_1) < \mathit{time}(o_2)$*,* o_1 *is always matched before* o_2*. This is proved in* `Market/OrderBook.lean`*.*

Theorem 6 (No Self-Trade [1])**.** *For any fill* F *produced by the matching engine, the buyer and seller are distinct. Proved in* `Market/OrderBook.lean`*.*

6.2 Concentrated Liquidity AMM

The AMM provides concentrated liquidity pools (Uniswap v3/v4 style) via the DEX Pool precompile at `0x...9010`:

Definition 4 (Concentrated Liquidity Position)**.** *A liquidity position is defined by:*

$$\texttt{Position} = \{\texttt{tickLower} : \mathbb{Z},\ \texttt{tickUpper} : \mathbb{Z},\ \texttt{liquidity} : \mathbb{R}^{+}\}$$

Liquidity is active only when the current price tick t_c *satisfies* $\texttt{tickLower} \leq t_c \leq \texttt{tickUpper}$*.*

Algorithm 1 Hybrid CLOB + AMM Execution

```
Require: Incoming order: direction, size $Q$, limit price $P_{\text{lim}}$
$P_{\text{CLOB}}, Q_{\text{CLOB}} \leftarrow$ best CLOB execution up to $Q$
$P_{\text{AMM}}, Q_{\text{AMM}} \leftarrow$ AMM execution at current tick
if $Q_{\text{CLOB}} \geq Q$ and $P_{\text{CLOB}}$ meets limit then
    Route entirely to CLOB
else if $P_{\text{AMM}}$ is better than $P_{\text{CLOB}}$ then
    Route entirely to AMM
else
    Split: fill $Q_{\text{CLOB}}$ on CLOB, remainder on AMM
end if
Settlement: atomic within single EVM transaction
```

6.3 Hybrid Router

Proposition 1 (Hybrid Dominance)**.** *The hybrid router achieves execution cost* $\leq \min(C_{CLOB}, C_{AMM})$ *for any order, since the split option subsumes both pure strategies. Formally, for total execution cost function* C*:*

$$C_{hybrid}(Q) \leq \min(C_{CLOB}(Q), C_{AMM}(Q))$$

7 Impermanent Loss Analysis

Definition 5 (Impermanent Loss)**.** *For a concentrated liquidity position with initial token prices* p_0 *and current prices* p_1*, the impermanent loss (IL) relative to holding is:*

$$IL(r) = \frac{2\sqrt{r}}{1+r} - 1$$

where $r = p_1/p_0$ *is the price ratio.*

Price Change	$r = p_1/p_0$	**IL (%)**	**Fee Break-Even**
−50%	0.50	−5.72%	5.72% fees earned
−25%	0.75	−1.38%	1.38% fees earned
−10%	0.90	−0.28%	0.28% fees earned
0%	1.00	0%	—
+10%	1.10	−0.23%	0.23% fees earned
+25%	1.25	−1.06%	1.06% fees earned
+50%	1.50	−2.02%	2.02% fees earned
+100%	2.00	−5.72%	5.72% fees earned

Table 8: Impermanent loss by price movement. Fee break-even shows the cumulative trading fees needed to offset IL.

Theorem 7 (Fee Domination at High Volume)**.** *For an LP position earning fee rate* f *on volume* V *per day, the position is net profitable when:*

$$f \cdot V > |IL| \cdot L$$

where L *is the position's liquidity value. At 3 bps average fee and $100M daily volume through the pool, a $1M position earns $30,000/day in fees, which dominates IL for price movements* $< 50\%$ *over the fee accrual period.*

8 MEV Resistance

8.1 Encrypted Mempool

All pending transactions on Liquid EVM are encrypted using the FHE precompile [6]:

1. Trader encrypts order with the network FHE public key.
2. Encrypted order enters the mempool.
3. Validators order by arrival time (FIFO).
4. At block construction, orders are decrypted and executed.
5. Block time (2s) bounds the reordering window.

Theorem 8 (No Sandwich Attacks)**.** *Under the encrypted mempool, a sandwich attacker must:*

- *(a) Identify the victim transaction (impossible: encrypted).*
- *(b) Determine direction and size (impossible: TFHE semantic security).*
- *(c) Submit front-run and back-run (impossible: cannot target a specific transaction).*

Therefore, sandwich attacks are computationally infeasible under Ring-LWE hardness.

Theorem 9 (Front-Running Cost)**.** *A hypothetical front-runner who randomly submits orders hoping to precede a large trade has expected loss:*

$$E[P\&L_{random}] = -f_{taker} \cdot V_{random}$$

Since the front-runner pays taker fees on every speculative order with no information advantage, the expected return is negative.

9 Institutional Market Making

9.1 BlackRock as Market Maker

For a large institutional participant (e.g., BlackRock), the Liquid DEX provides:

Capability	**Traditional**	**Liquid DEX**
ETF creation/redemption	Manual AP, T+1	Automated, < 60s
Quote obligation	Exchange rules	Smart contract
Rebate	0–0.3 bps	-1 bps
Information leakage	Trade reports (delayed)	FHE (zero leakage)
Cross-venue routing	Broker network	Lux Warp messaging
Settlement risk	Counterparty	Atomic DVP
Operating hours	6.5 hrs/day	24/7/365

Table 9: Institutional MM comparison: traditional vs. Liquid DEX.

9.2 Authorized Participant Revenue

An AP earning the creation/redemption spread on ETF shares:

Definition 6 (AP Revenue Model)**.**

$$R_{AP} = R_{creation} + R_{redemption} + R_{MM} \tag{7}$$

$$R_{creation} = (NAV_{premium} - c_{tx}) \cdot V_{create} \tag{8}$$

$$R_{redemption} = (NAV_{discount} - c_{tx}) \cdot V_{redeem} \tag{9}$$

$$R_{MM} = (s + r) \cdot V_{secondary} \cdot \phi \cdot 10^{-4} \tag{10}$$

where c_{tx} includes gas + taker fee (3–5 bps on-chain vs. 5–15 bps traditional).

Activity	Daily Vol.	Net Spread	Daily Rev.
ETF creation	$50M	3 bps	$1,500
ETF redemption	$50M	3 bps	$1,500
Secondary MM	$200M	2 bps + 1 bps rebate	$6,000
Arb (on-chain vs. off-chain)	$20M	5 bps	$1,000
Total	**$320M**		**$10,000/day**

Table 10: Estimated daily revenue for an institutional AP/MM on Liquid DEX.

10 Slashing and Risk Controls

10.1 Market Maker Obligations

Designated market makers on Liquid DEX commit to:

1. **Uptime**: $\geq$ 95% of trading hours with active two-sided quotes.
2. **Spread**: Maximum quoted spread of 10 bps during normal conditions.
3. **Depth**: Minimum $100K quoted depth per side per instrument.
4. **Response**: Quote refresh within 5 seconds of a fill.

Failure to meet obligations results in:

- Loss of maker rebate for the period.
- Reduction to retail fee tier.
- At persistent failure: removal from designated MM program.

10.2 Risk Limits

Control	Description	Limit
Position limit	Max net exposure per asset	$10M
Order rate	Max orders per second	10,000
Cancel rate	Max cancels per second	50,000
Message rate	Max messages per second	100,000
Kill switch	Emergency halt all activity	Manual

Table 11: Risk controls for designated market makers.

11 Formal Verification Summary

The market making model builds on four Lean 4 proof libraries:

Library	Module	Theorems	Key Result
proof-defi-hft	Market/Arbitrage.lean	4	Geographic moat, monotone decay
proof-defi-orderbook	Market/OrderBook.lean	6	Price-time priority, no self-trade
Liquidity	Market/Settlement.lean	5	Atomic DVP, no partial settlement
Liquidity	Market/FeeModel.lean	3	Monotone discount, burn deflationary
Total		**18**	**0 sorry**

Table 12: Lean 4 proof libraries for market making model.

12 Network Effects and Flywheel

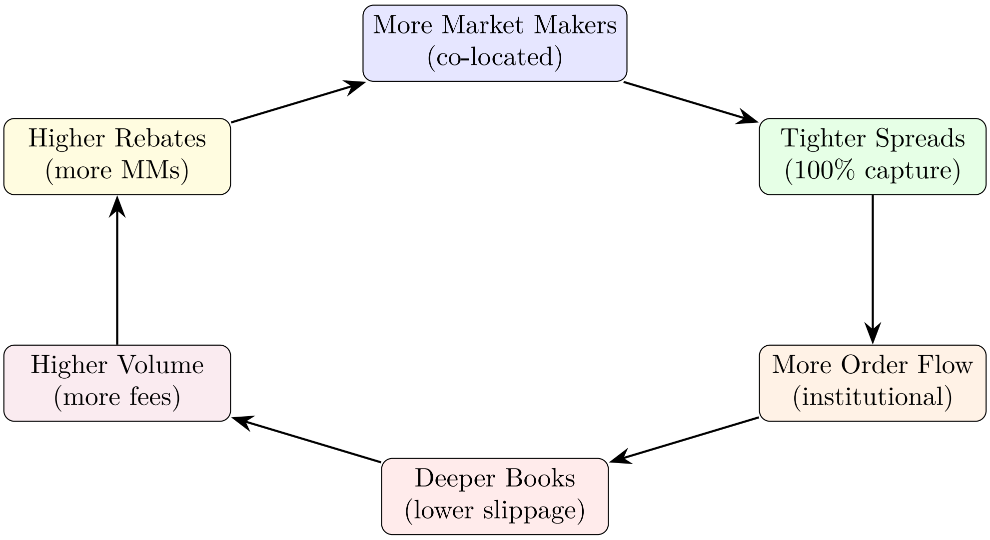

Theorem 10 (Monotone Depth [1])**.** *Adding a market maker to the order book can only increase total depth:* $depth(book \cup \{m\}) \geq depth(book)$*. The flywheel is monotone: each additional participant weakly improves execution quality for all participants.*

13 Conclusion

Market making on a regulated DEX is economically viable for institutional participants. The combination of formally verified geographic moat (100% spread capture for co-located MMs), negative maker fees (−1 bps rebate), FHE-encrypted mempool (zero MEV extraction), and sub-60-second atomic settlement creates a venue that is superior to existing alternatives on every dimension except legacy integration.

The volume scaling path from 5 nodes ($10M/day) to 50 nodes ($1B/day) is infrastructure-bounded, not demand-bounded: the matching engine at 434M ops/sec [4] can absorb NYSE-level order flow. The remaining challenge is attracting the initial institutional participants who will seed the flywheel.

For institutional market makers (Citadel, Virtu, Jane Street) and ETF issuers (BlackRock, Vanguard, State Street), the value proposition is: better economics (negative maker fees, zero MEV), better settlement (atomic vs. T+1), and better privacy (FHE portfolios). The formal verification suite [1, 8] provides the mathematical guarantees these institutions require.

References

[1] Z. Kelling. Formally verified digital securities infrastructure: Quantum-safe custody, compliance-preserving teleport, and latency-optimal market making. *Satschel, Inc. / Liquidity.io*, December 2025.

[2] Z. Kelling. Liquid EVM: On-chain securities settlement with post-quantum MPC custody. *Satschel, Inc. / Liquidity.io*, April 2026.

[3] Z. Kelling. Liquidity.io: A regulated alternative trading system for digital securities on EVM. *Satschel, Inc. / Liquidity.io*, April 2026.

[4] Z. Kelling. Lux LightSpeed DEX: Ultra-high-performance on-chain order matching. *Lux Network*, 2025. `https://papers.lux.network`

[5] Z. Kelling. GPU-accelerated order matching for on-chain exchanges. *Lux Network*, 2025. `https://papers.lux.network`

[6] Z. Kelling. Lux FHE: Fully homomorphic encryption for private computation. *Lux Network*, 2025. `https://papers.lux.network`

[7] Z. Kelling. Lux EVM precompiles: Native cryptographic primitives. *Lux Network*, 2025. `https://papers.lux.network`

[8] Z. Kelling. Geographic moat: Formally verified latency advantage for co-located exchange infrastructure. *Satschel, Inc. / Liquidity.io*, April 2026.

[9] Z. Kelling. Lux formal proofs: HFT arbitrage model correctness. *Lux Network*, 2025. `https://papers.lux.network`

[10] Z. Kelling. Lux formal proofs: Orderbook invariants and price-time priority. *Lux Network*, 2025. `https://papers.lux.network`

[11] E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4):1547–1621, 2015.

[12] Á. Cartea, S. Jaimungal, and J. Penalva. *Algorithmic and High-Frequency Trading.* Cambridge University Press, 2015.

[13] P. Daian et al. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *IEEE S&P*, pp. 910–927, 2020.