Discussion of

"Who Loses and Why? Market Microstructure of Perpetual Futures"

by Jia, Moallemi, Wang, and Zeng.

Daniele Bianchi
Queen Mary, University of London

OMI-SBS Conference on FinTech and Blockchain Economics

Why care? (1) Hyperliquid is *the* DEX perp market

Perps are now the dominant crypto instrument (~\$62T notional in 2025, $>3\times$ spot). On-chain, the market is *one* venue:

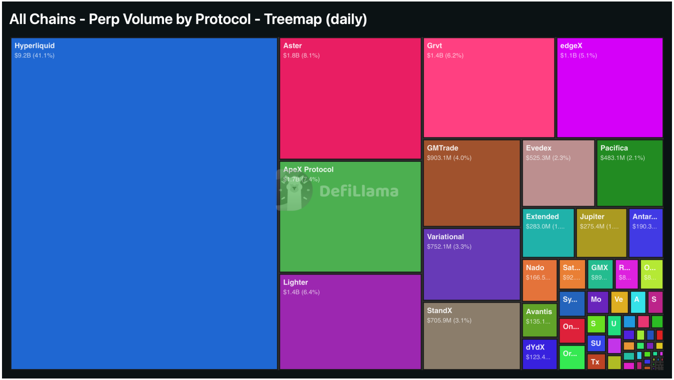

Daily perp volume by protocol, all chains. Hyperliquid is **41%** of all decentralized perp volume — larger than the next ten venues combined. *Source*: DefiLlama.

Why care? (2) The on-chain share is rising

And it is winning share from *centralized* exchanges — the part of the market we usually *cannot* see at the wallet level:

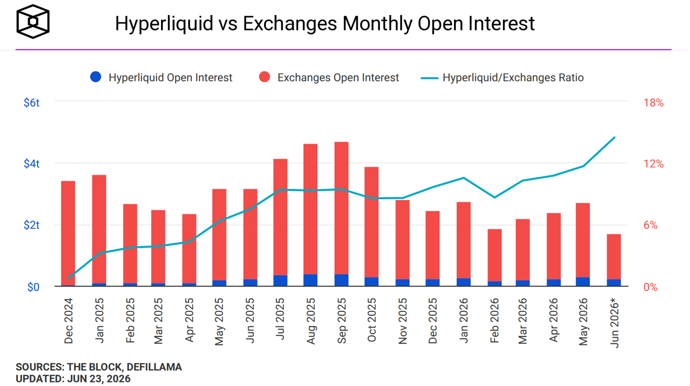

The Hyperliquid/exchange OI ratio has roughly tripled, to ∼14%. *Source*: The Block, DefiLlama.

This paper

Setting: the complete on-chain record of Hyperliquid — every order, fill, and liquidation, under a stable wallet ID. ∼\$1.95T notional, 386,211 wallets, Mar–Nov 2025.

This paper

Setting: the complete on-chain record of Hyperliquid — every order, fill, and liquidation, under a stable wallet ID. $\sim$\$1.95T notional, 386,211 wallets, Mar–Nov 2025.

Question: in a zero-sum market, *who loses, how much, and through which channels?*

This paper

Setting: the complete on-chain record of Hyperliquid — every order, fill, and liquidation, under a stable wallet ID. ∼$1.95T notional, 386,211 wallets, Mar–Nov 2025.

Question: in a zero-sum market, *who loses, how much, and through which channels?*

Three steps:

- ↪ **Classify** wallets from on-chain behavior ⇒ market makers / retail / algos.
- ↪ **Account** for each group's dollar P&L (complete-market, sums to zero).
- ↪ **Explain** the losses via four channels: fees, execution cost, post-trade reversion, liquidations.

The headline: retail finances the venue

After fees, **retail loses ~$731M**; market makers and algos are net positive.

Retail P&L, M_{120} accounting	**$M**
Voluntary trading — taker	−141
Voluntary trading — maker	+63
Forced liquidations (net of ADL)	−152
Fees	−502
Total, after fees	−731

↪ A clean, transparent decomposition.

↪ My comments are about *what is doing the work*.

My comments:

Comments/observations:

#1: Participants classification is the load-bearing wall.

#2: The retail “slippage premium” — mechanical or behavioral?

#3: Why does retail lose? Immediacy vs. information.

+ *minor comments at the end.*

Comment #1: Participants classification.

Comment #1a: Every dollar passes through the labels

The paper is conditioned on three labels (MM, retails, algos) from *behavioral thresholds* — no on-chain ground truth.

Comment #1a: Every dollar passes through the labels

The paper is conditioned on three labels (MM, retails, algos) from *behavioral thresholds* — no on-chain ground truth.

The retail signal is quite good. `FrontendMarket` orders route through Hyperliquid's own UI — an MM or bot has no reason to use it; hard to spoof *toward* retail. But it is just **0.03%** of order events (footnote 13) ...

Comment #1a: Every dollar passes through the labels

The paper is conditioned on three labels (MM, retails, algos) from *behavioral thresholds* — no on-chain ground truth.

The retail signal is quite good. `FrontendMarket` orders route through Hyperliquid's own UI — an MM or bot has no reason to use it; hard to spoof *toward* retail. But it is just **0.03%** of order events (footnote 13) ...

. . . and it labels only *part* of the bucket:

- ↪ **Affirmative**: `retail_a/b/c` — carry a real frontend share.
- ↪ **Weak**: `retail_d` (low-activity, *no* frontend) + the unclassified `default` residual.

Comment #1a: Every dollar passes through the labels

The paper is conditioned on three labels (MM, retails, algos) from *behavioral thresholds* — no on-chain ground truth.

The retail signal is quite good. `FrontendMarket` orders route through Hyperliquid's own UI — an MM or bot has no reason to use it; hard to spoof *toward* retail. But it is just **0.03%** of order events (footnote 13) ...

. . . and it labels only *part* of the bucket:

↪ **Affirmative**: `retail_a/b/c` — carry a real frontend share.

↪ **Weak**: `retail_d` (low-activity, *no* frontend) + the unclassified `default` residual.

↪ It would be useful to see **headline P&L for the affirmative core (`retail_a/b/c`) beside full retail.** *If the core still finances the venue, the result is robust to the thin bucket.*

Comment #1b: Can "Algo" really be 2.3% of volume?

In a venue running ∼200k orders/sec at sub-second latency, **2.3% algorithmic volume** (Table 1) seems implausibly small.

Comment #1b: Can “Algo” really be 2.3% of volume?

In a venue running $\sim$200k orders/sec at sub-second latency, **2.3% algorithmic volume** (Table 1) seems implausibly small.

The algos are in the 'MM' bucket. 'MM' is defined by heavy post-only (`Alo`) quoting — but quoting, updating, and cancelling at that rate is an algorithm. Ask: 'MM' is really algorithmic liquidity provision, and most algos sit inside it.

Comment #1b: Can “Algo” really be 2.3% of volume?

In a venue running $\sim$200k orders/sec at sub-second latency, **2.3% algorithmic volume** (Table 1) seems implausibly small.

The algos are in the 'MM' bucket. 'MM' is defined by heavy post-only (`Alo`) quoting — but quoting, updating, and cancelling at that rate is an algorithm. Ask: 'MM' is really algorithmic liquidity provision, and most algos sit inside it.

Implication: The 2.3% measures *where the line was drawn*, not how much algo trading occurs.

↪ “MM earns, retail loses” may partly be *automated* earns, *manual* loses.

↪ So the labels are not retail / market-maker / algo, but really **algorithmic liquidity providers** / **frontend-&-residual takers** / non-MM automated takers.

Comment #2: The retail “slippage premium” — mechanical or behavioral?

Comment #2: Split the execution cost

Retail pays **2.10 bps** slippage vs. **0.19** for MMs — $\sim 11\times$.
Alarming? Split the cost (Eq. 4):

bps of taker notional	**MM**	**Algo**	**Ret**
Half-spread *(asset / book)*	0.37	0.38	**0.33**
Slippage *(trader-controlled)*	0.19	1.20	**2.10**

Comment #2: Split the execution cost

Retail pays **2.10 bps** slippage vs. **0.19** for MMs — $\sim 11\times$.
Alarming? Split the cost (Eq. 4):

bps of taker notional	**MM**	**Algo**	**Ret**
Half-spread *(asset / book)*	0.37	0.38	**0.33**
Slippage *(trader-controlled)*	0.19	1.20	**2.10**

- $\hookrightarrow$ **Half-spread** — the cost of reaching the best quote — is the *asset* component. **Common across groups (retail even lowest)**.
- $\hookrightarrow$ **Slippage** — walking the book past that quote — is *trader-controlled*: it scales with your own size vs. depth.

Comment #2: Split the execution cost

Retail pays **2.10 bps** slippage vs. **0.19** for MMs — $\sim 11\times$.
Alarming? Split the cost (Eq. 4):

bps of taker notional	MM	Algo	Ret
Half-spread *(asset / book)*	0.37	0.38	**0.33**
Slippage *(trader-controlled)*	0.19	1.20	**2.10**

↪ **Half-spread** — the cost of reaching the best quote — is the *asset* component. **Common across groups (retail even lowest)**.

↪ **Slippage** — walking the book past that quote — is *trader-controlled*: it scales with your own size vs. depth.

The "premium" lives entirely in the part the trader controls.

Comment #2: Low depth elasticity — skill, or objective?

Why does retail pay more slippage?

$\hookrightarrow$ Because retail does not shrink its orders when the book thins: market makers scale order size to displayed depth **2.5**$\times$ more than retail (Table 7).

Comment #2: Low depth elasticity — skill, or objective?

Why does retail pay more slippage?

$\hookrightarrow$ Because retail does not shrink its orders when the book thins: market makers scale order size to displayed depth **2.5**$\times$ more than retail (Table 7).

The paper reads this as retail *failing* to condition on the book. But high vs. low elasticity may just reflect *different objectives*, not skill:

Comment #2: Low depth elasticity — skill, or objective?

Why does retail pay more slippage?

↪ Because retail does not shrink its orders when the book thins: market makers scale order size to displayed depth **2.5**× more than retail (Table 7).

The paper reads this as retail *failing* to condition on the book. But high vs. low elasticity may just reflect *different objectives*, not skill:

Market maker

↪ Sizing into available liquidity *is* the strategy.

Retail

↪ Size set by conviction; slippage is the rational price of immediacy.

Comment #2: Low depth elasticity — skill, or objective?

Why does retail pay more slippage?

$\hookrightarrow$ Because retail does not shrink its orders when the book thins: market makers scale order size to displayed depth **2.5**$\times$ more than retail (Table 7).

The paper reads this as retail *failing* to condition on the book. But high vs. low elasticity may just reflect *different objectives*, not skill:

Market maker

$\hookrightarrow$ Sizing into available liquidity *is* the strategy.

Retail

$\hookrightarrow$ Size set by conviction; slippage is the rational price of immediacy.

Ask: do longer-holding wallets show low elasticity *regardless* of group? That separates objective from skill.

Comment #3: Why does retail lose?
Immediacy vs. information.

Comment #3a: What is new here?

"Retail loses to sophisticated players" is the *prior*, not the finding:

- ↪ Documented in equities and assumed in crypto.
- ↪ Crypto perp retail are leveraged, self-selected, possibly sophisticated (e.g., cash and carry) — losing across *all* counterparties is the striking part.

And the flow mechanics are *familiar*: MM/algo contrarian, supplying liquidity to trend-chasers — standard inventory/liquidity provision.

Comment #3a: What is new here?

"Retail loses to sophisticated players" is the *prior*, not the finding:

- ↪ Documented in equities and assumed in crypto.
- ↪ Crypto perp retail are leveraged, self-selected, possibly sophisticated (e.g., cash and carry) — losing across *all* counterparties is the striking part.

And the flow mechanics are *familiar*: MM/algo contrarian, supplying liquidity to trend-chasers — standard inventory/liquidity provision.

Two things are genuinely new — and under-sold:

- ↪ The **channel decomposition** on a zero-sum ledger.
- ↪ **Why** retail loses: immediacy, not information — via loss-realizing closes (next slide).

Comment #3b: Immediacy, not information

Two ways to lose money trading. Which is retail?

- $\hookrightarrow$ **Information**: you are *wrong about prices* — buy before they fall. A prediction failure.
- $\hookrightarrow$ **Immediacy**: you pay to trade *now* — spread + slippage + fees. The price of speed, even when right.

Comment #3b: Immediacy, not information

Two ways to lose money trading. Which is retail?

- ↪ **Information**: you are *wrong about prices* — buy before they fall. A prediction failure.
- ↪ **Immediacy**: you pay to trade *now* — spread + slippage + fees. The price of speed, even when right.

The paper's own evidence points to immediacy. Reversion is governed by trade purpose, not group identity (p. 28):

- ↪ Loss-realizing closes (LRC) revert in every group
- ↪ but only for retail does it revert hard enough: **+2.66** (1 min) → **−1.58 bps** (2 h).

Comment #3b: Immediacy, not information

Two ways to lose money trading. Which is retail?

- ↪ **Information**: you are *wrong about prices* — buy before they fall. A prediction failure.
- ↪ **Immediacy**: you pay to trade *now* — spread + slippage + fees. The price of speed, even when right.

The paper's own evidence points to immediacy. Reversion is governed by trade purpose, not group identity (p. 28):

- ↪ Loss-realizing closes (LRC) revert in every group
- ↪ but only for retail does it revert hard enough: **+2.66** (1 min) → **−1.58 bps** (2 h).

Takeaway: Retail pays the spread to close *now*, paying for speed when it can least afford to. Is it because they are unskilled or they have to?

Minor comments

- **SVAR placement.** Results sit in the appendix; the main-text home is missing (broken ?? refs, e.g. pp. 22, 39).
- **"Disagreement lives in openings"** is true on *sign* (Table 11) — but the economic action is in retail's LRC. Worth stating both.
- **Distribution of retail losses.** \$731M over 377k wallets is $\sim$\$1.9k each, but surely skewed — whales vs. the mass changes the welfare read.

Summary

A terrific, novel dataset and a clean decomposition.

I learned a lot — looking forward to the next version!