

How to Serve Your Sandwich? MEV Attacks in Private L2 Mempools

Jan Gorzny
Zircuit

Joint work with **K. Gogol, M. Schneider, C. Tessone**

*The 5th International Workshop on Cryptoasset Analytics (CAAW)
Co-located with Financial Cryptography and Data Security 2026*

March 6, 2026 | St. Kitts



Why sandwich attacks may not generalize to rollups

Ethereum L1	Major Rollups
Public mempool	Private mempool
Builder market	Single centralized sequencer
Atomic bundles	No bundle guarantees
Deterministic MEV	Probabilistic ordering

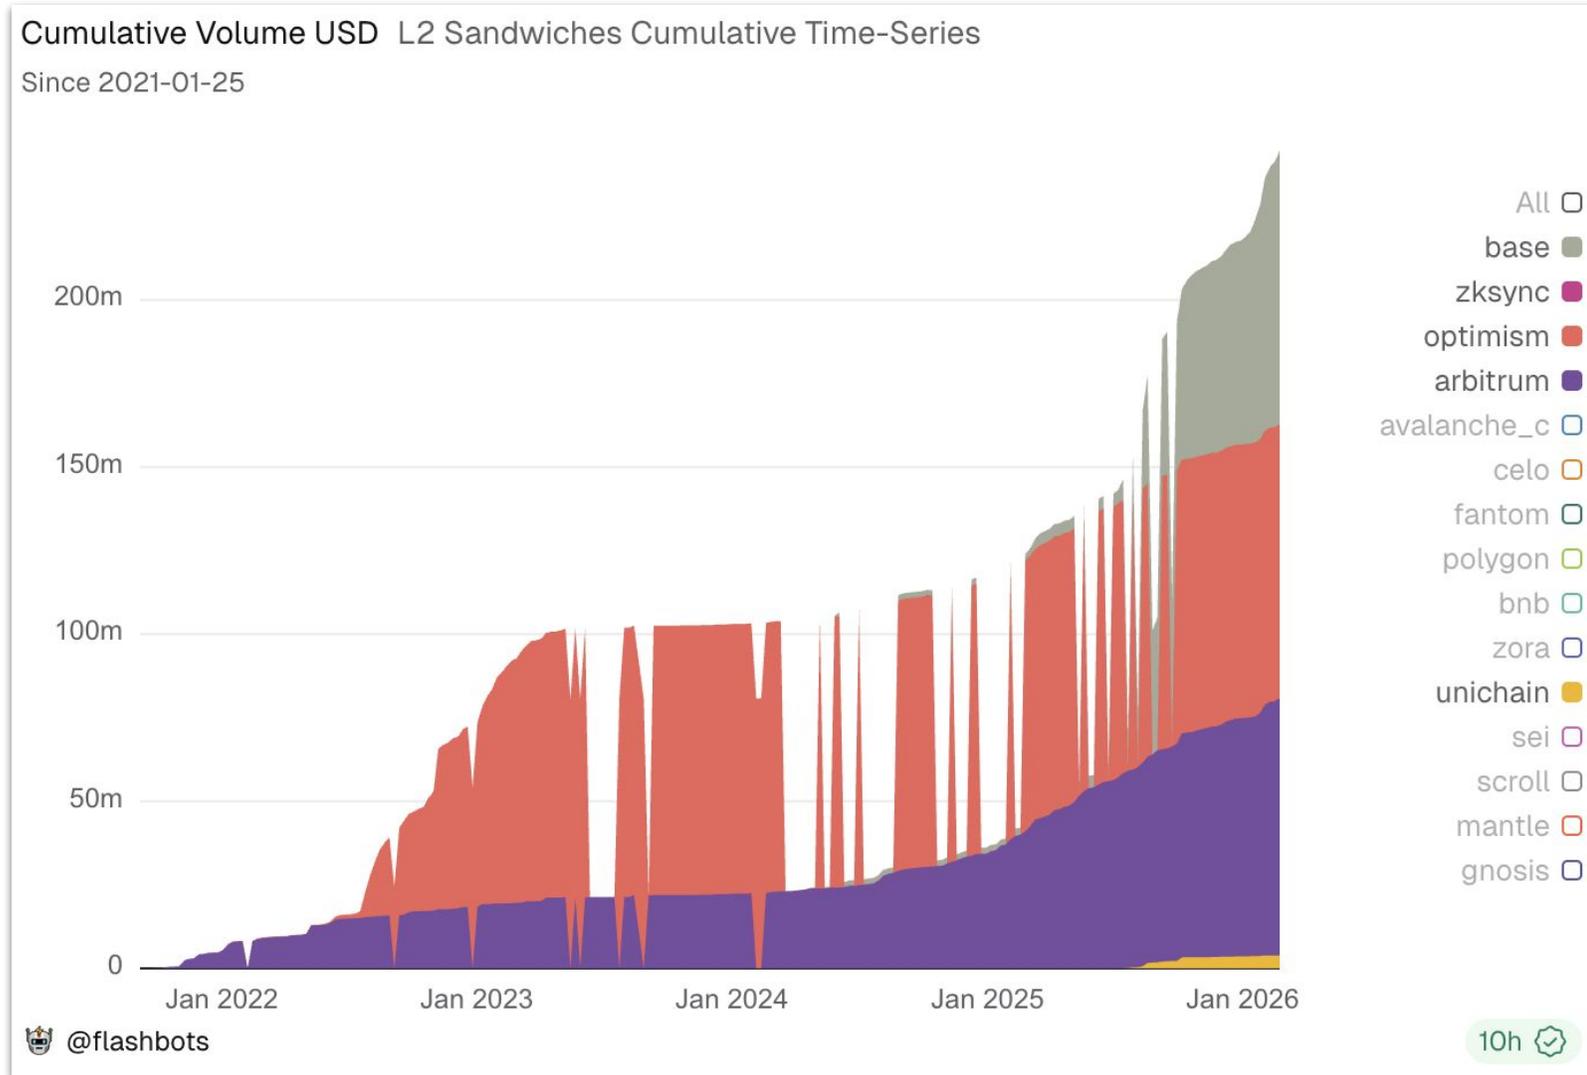
L2 architecture fundamentally changes execution guarantees

Research Question

Early reports claim sandwiches exist on L2.

We ask:

Are sandwich attacks economically profitable and technically feasible under private sequencing?



Related Work

L1 MEV:

- Flash Boys 2.0
- Qin et al. - Sandwich attack efficiency on Ethereum

L2 MEV:

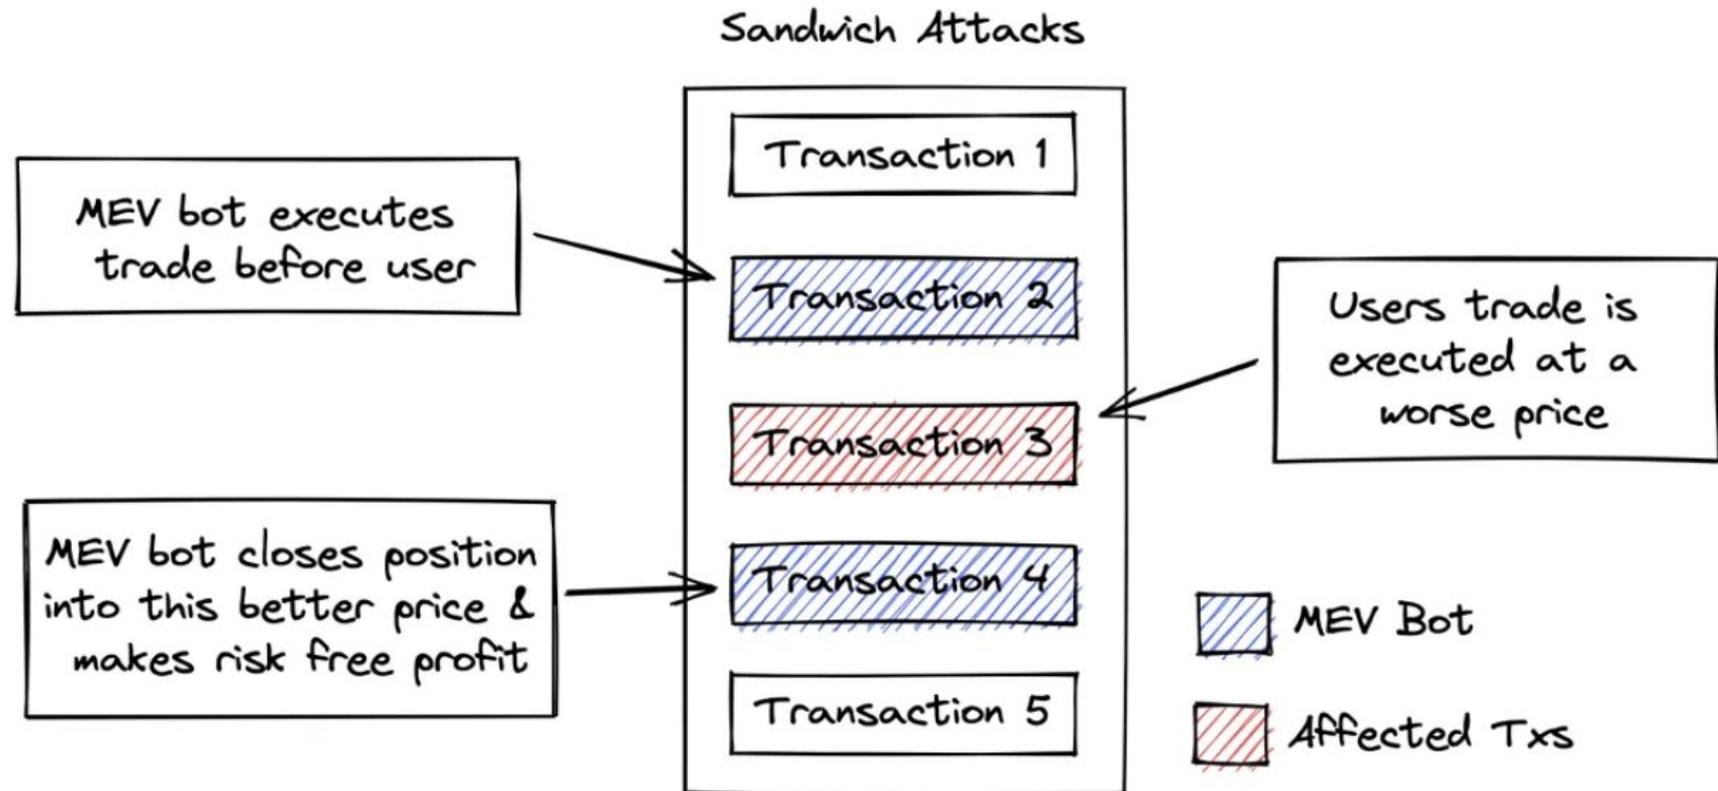
- Torres et al. - Quantification of MEV on L2s
- Oz et al. Gogol et al. - Non-atomic arbitrage MEV

No prior work jointly models economics + execution feasibility on L2.

What makes a sandwich attack work?

1. Observability of victim trade
2. Ordering control
3. Atomic co-inclusion of
 - front running tx
 - victim tx
 - back running tx

Atomicity is essential.



Execution model shift on rollups

Ethereum L1	Major Rollups
Public mempool	Private mempool
Builder market	Single centralized sequencer
Atomic bundles	No bundle guarantees
Deterministic MEV	Probabilistic ordering

Execution of sandwich attack on rollups becomes probabilistic, not atomic.

Execution Feasibility under Private Sequencing

Why same (or close)-block sandwich execution is hard?

Constraints:

- $\text{nonce}(b) > \text{nonce}(f)$
- $\text{tip}(f) \geq \text{tip}(v) \geq \text{tip}(b)$
- same-block inclusion required

Two strategies:

1. **Delay** (FCFS) → fails on fast chains
2. Sequential submissions via **intermediary contract** → no atomic inclusion guarantee

Co-inclusion probability on rollups

On rollups:

- Block time: **~200–500 ms**
- Sequencer batching window: **~300–800 ms**
- Network jitter: **tens of ms**
- Ordering rule: **FCFS vs tip-priority**

$$p(\text{arrival}) \approx \exp\left(-\frac{\Delta^2}{2\sigma^2}\right) \quad p(\text{co-inc}) \approx \left(1 - \frac{\Delta}{T_s}\right) p_{\text{policy}}(\text{priority}) \exp\left(-\frac{\Delta^2}{2\sigma^2}\right)$$

On fast rollups, even small timing deviations break execution.

Empirical implication: Co-inclusion probability \approx 5%–20%

Methodology and Data

The 5-step detection pipeline:

1. Remove atomic **arbitrage**
2. **Trading pool DEX** grouping
3. **Actor** attribution (EOA vs router)
4. **Ordering** constraint
5. **Direction** constraint (buy, sell)

Crucial improvement in methodology over prior work:

- Avoid ERC20 transfer heuristics

Data:

- 5 major rollups
- Jan–Sep 2025
- Swap-level events



Do attacker trades match victim swap?

Theory predicts:

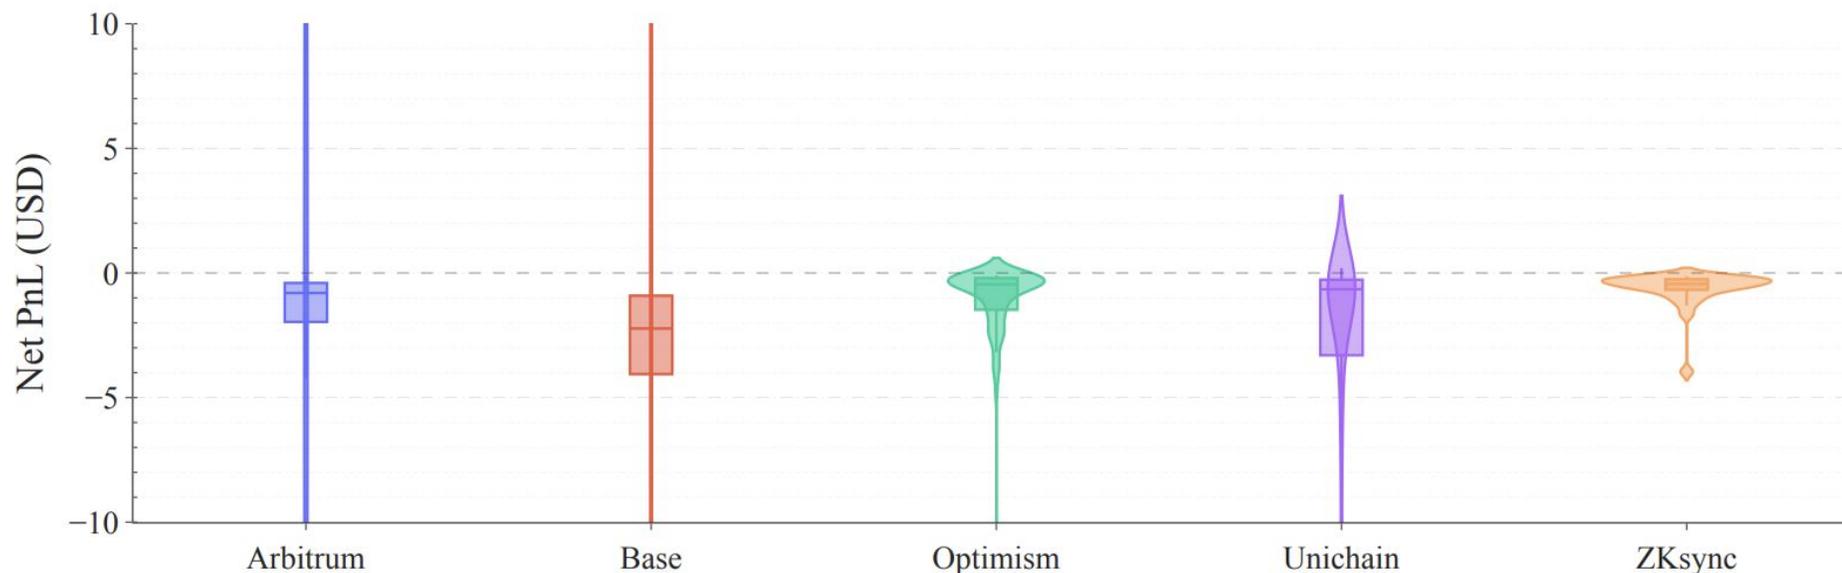
- Size of Front \approx
 $\frac{1}{2}$ Victim swap
- Back \approx Front

Empirical:

- Weak correlations
- Only 5–25% strong signature cases
(difference in size below 10%)

Chain	Strong Sig.	Median b/f	IQR b/f	Victim \leftrightarrow Front	Victim \leftrightarrow Back	Front \leftrightarrow Back
Arbitrum	24.6%	1.035	[0.816, 2.672]	0.426	0.409	0.627
Base	14.3%	0.967	[0.549, 1.583]	0.276	0.375	0.352
Optimism	5.1%	0.728	[0.282, 1.526]	0.032	0.592	0.082
Unichain	7.3%	0.825	[0.358, 1.913]	0.162	0.182	0.398
ZKsync	7.4%	0.674	[0.392, 0.785]	0.081	0.567	0.354

Median net profitability is negative



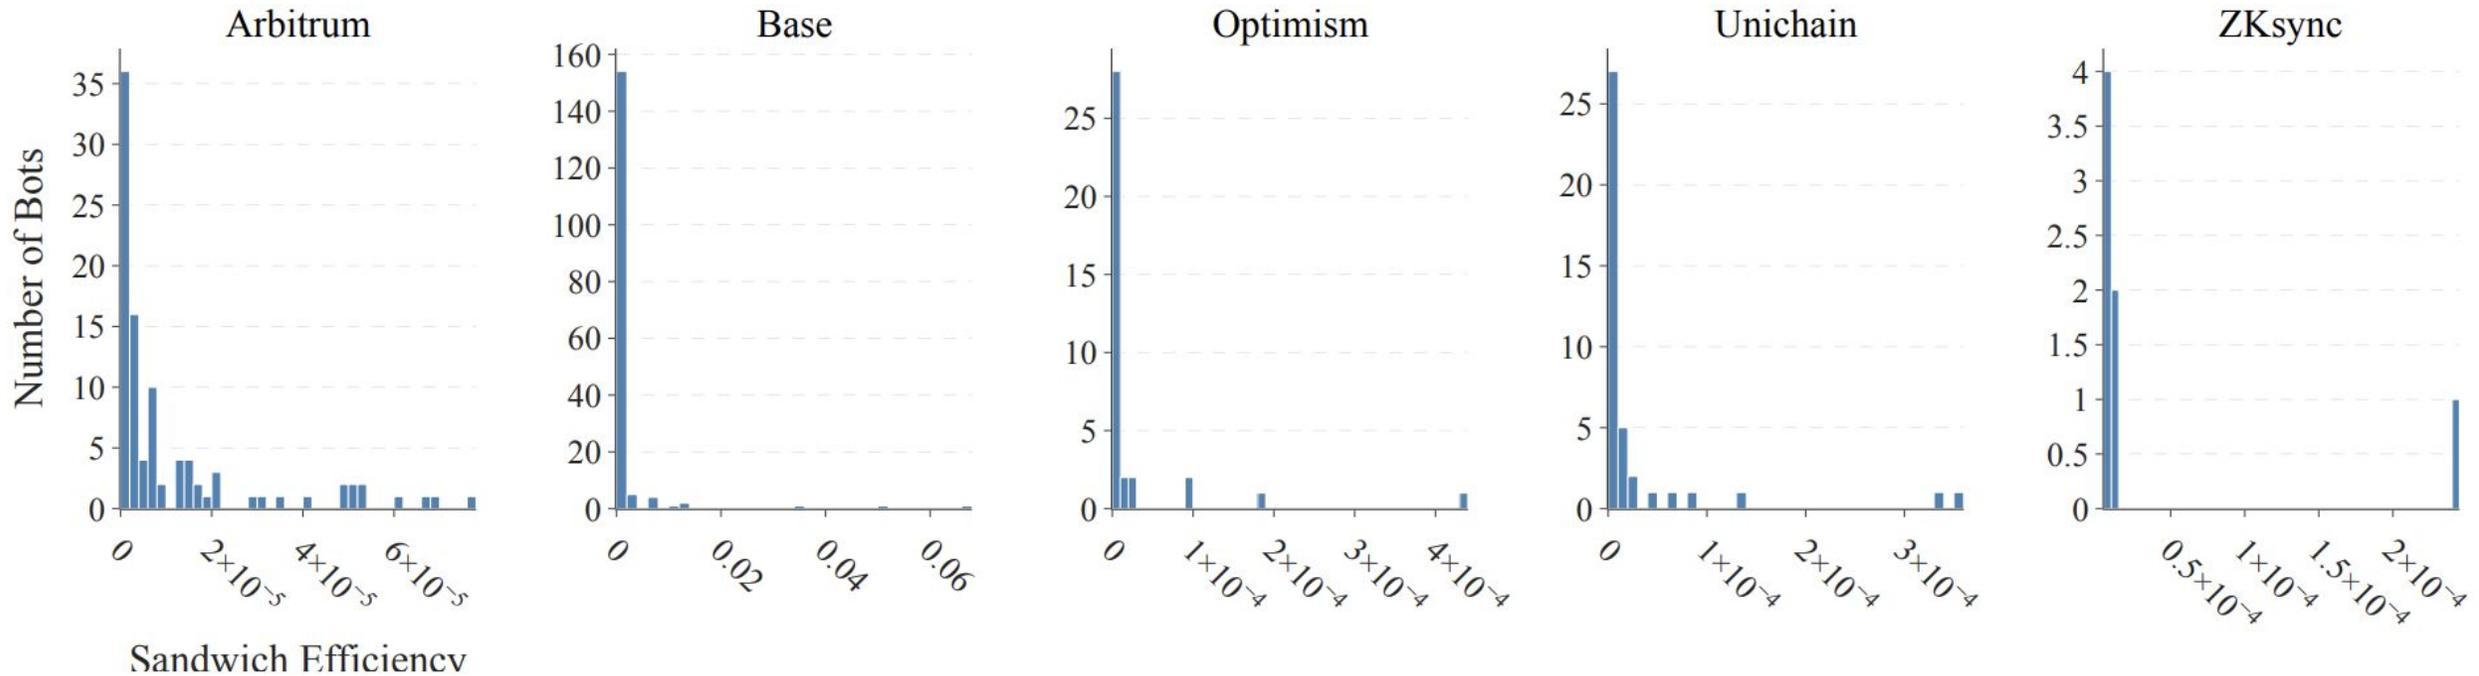
Violin distribution of estimated net profit in USD

Median profits of identified sandwich patterns are negative

Most bots are not dedicated sandwichers

We analyze all originators of sandwich patterns

Sandwich efficiency is defined as the share of sandwich patterns among a bot's total transactions.



- Few concentrated actors
- Sandwich efficiency often $< 10^{-3}$
- Most activity is general arbitrage (based on analysis of other bot transactions)

Main Takeaways

1. Execution feasibility of sandwich attacks on rollups sharply limits expected value from the attack
2. Sandwich attacks are rare and unprofitable on private L2s.
~95% of sandwich-like triples fail consistency checks.
Median profitability is negative.
3. **Private mempools materially change MEV equilibria.**

Rollups do not eliminate MEV

They restructure it

**Private sequencing suppresses
sandwich attacks**

When could sandwiches return?

If rollups introduce:

- Decentralized builders/sequencers
- Public mempools
- Order-flow auctions
- Pre-confirmation markets

then:

The feasibility frontier of MEV Sandwich attacks may shift back toward L1 conditions



Access the full paper here

Thank you

Jan Gorzny
Zircuit