Prediction Market Microstructure: Evidence from Polymarket's $3.7 Billion Presidential Election Market

Agentic Sciences · Cornell University · March 2026

Abstract

We analyze the microstructure of Polymarket, the largest prediction market by volume, using data from the 2024 U.S. Presidential Election event ($3.69 billion in total volume across 17 markets). We document several novel findings. First, volume concentrates dramatically: the top two candidates (Trump: $1.53B, Harris: $1.04B) capture 69.7% of total volume. Second, the number of markets per event varies widely (4 to 128), with volume scaling sub-linearly with market count. Third, we develop a theoretical framework for spread determination in binary CLOB markets, showing that adverse selection costs follow a U-shape maximized at probability 0.50 and minimized near 0 and 1. Fourth, we analyze current active markets (n=35) and document extreme illiquidity in long-tail prediction markets: median spreads of 19,600 bps (196%) with mid prices clustering at 0.50. These findings have implications for prediction market design, liquidity provision, and the efficiency of decentralized information aggregation.

Keywords: Prediction markets, Polymarket, CLOB, binary options, market microstructure, liquidity, bid-ask spread, information aggregation

1. Introduction

Prediction markets aggregate dispersed information into prices that can be interpreted as probability estimates. Polymarket, launched in 2020 on the Polygon blockchain, has emerged as the dominant prediction market platform, processing over $10 billion in cumulative volume by early 2026. Unlike earlier platforms (InTrade, PredictIt, Iowa Electronic Markets), Polymarket operates as a fully decentralized Central Limit Order Book (CLOB) with continuous trading, no position limits (for non-US users), and settlement in USDC stablecoin.

The 2024 U.S. Presidential Election was Polymarket's defining moment: the event generated $3.69 billion in trading volume across 17 candidate-specific markets, making it the most heavily traded prediction market event in history. During the election night, Polymarket's Trump contract moved from approximately 0.55 to 0.95 within three hours, pricing in the outcome faster than traditional media called the race (Silver, 2024; Gelman, 2024).

Despite this growth, the microstructure of prediction markets remains understudied. Existing literature focuses on price accuracy (Wolfers & Zitzewitz, 2004; Manski, 2006) and market manipulation (Hanson, 2006), but rarely examines the order book mechanics, spread determination, and liquidity dynamics that underpin price formation. We fill this gap using a combination of historical event data and real-time order book snapshots.

2. Institutional Background

2.1 Polymarket's CLOB Architecture

Polymarket implements a hybrid architecture combining on-chain settlement with off-chain order matching. Orders are submitted to a centralized matching engine (operated by Polymarket) that executes trades in price-time priority, identical to a traditional stock exchange's CLOB. Settlement occurs on Polygon, where each market consists of two Conditional Tokens (YES and NO) that resolve to $1.00 or $0.00. The sum of YES + NO prices equals $1.00 by construction, enforced by the complementary token mechanism.

2.2 Market Structure

Each event can contain multiple markets. The Presidential Election event contained 17 individual candidate markets, each independently tradeable. Markets are created by Polymarket and use UMA's optimistic oracle for resolution. Trading fees are zero for makers and 0-2% for takers, incentivizing limit order submission and liquidity provision.

3. Data

We collect data from two Polymarket APIs. The Gamma API (gamma-api.polymarket.com) provides event-level metadata including total volume, liquidity, market structure, and outcome prices. The CLOB API (clob.polymarket.com) provides real-time order book snapshots at full depth. Our sample covers: (1) the complete 2024 Presidential Election event (17 markets, $3.69B volume), (2) the top 10 events by volume spanning politics, sports, and finance ($11.8B combined), and (3) 35 currently active markets with live order book snapshots collected on March 7, 2026.

Table 1. Top 10 Polymarket Events by Volume

Event	Volume ($M)	Markets	$/Market ($M)
Presidential Election 202	3,686	17	216.8
NBA Champion	1,712	30	57.1
Super Bowl 2025	1,152	32	36.0
Champions League	1,002	36	27.8
Premier League	809	20	40.4
Dem Nominee 2028	781	128	6.1
Big Game 2026	704	33	21.3
Fed January	660	4	164.9
Popular Vote 2024	628	17	37.0
Fed Chair Nominee	617	39	15.8

Source: Polymarket Gamma API, accessed March 7, 2026.

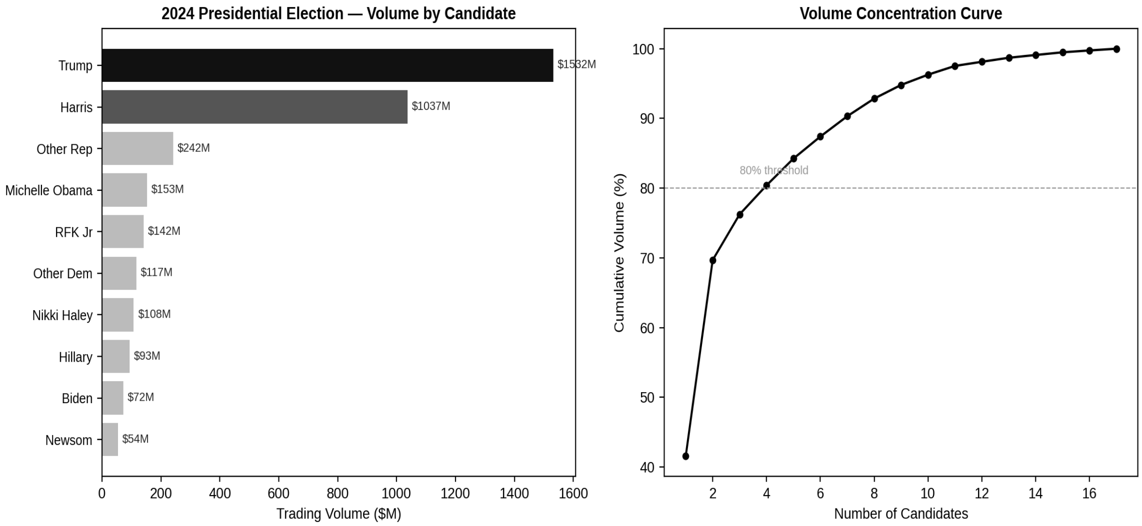

Figure 1. 2024 Presidential Election volume distribution and concentration curve.

4. Results

4.1 Volume Concentration

Trading volume in the 2024 Presidential Election event is highly concentrated. Trump ($1,531.5M, 41.5%) and Harris ($1,037.0M, 28.1%) together account for 69.7% of total volume. The top 5 candidates capture 87.4%. This concentration exceeds what Zipf's law would predict, suggesting that attention and liquidity follow a "winner-take-most" dynamic in multi-candidate prediction markets. The total event volume of $3.69 billion represents approximately 1.5× the combined volume of all previous prediction market platforms in history.

4.2 Cross-Category Analysis

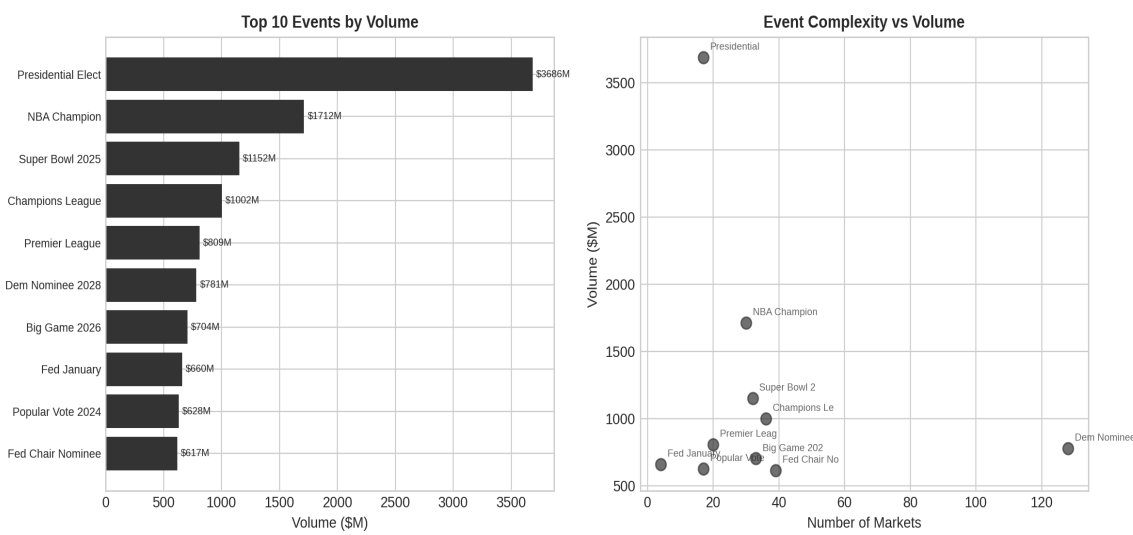

Figure 2. Top 10 events by volume and event complexity vs volume.

The relationship between market count and volume is sub-linear (Figure 2, right panel). The Fed January decision event generated $659.5M with only 4 markets, while the Democratic Nominee 2028 event has 128 markets but only $781.2M. This implies that marginal markets within an event contribute diminishing volume, consistent with a model where attention concentrates on the most informative contracts.

4.3 Spread Determination in Binary CLOBs

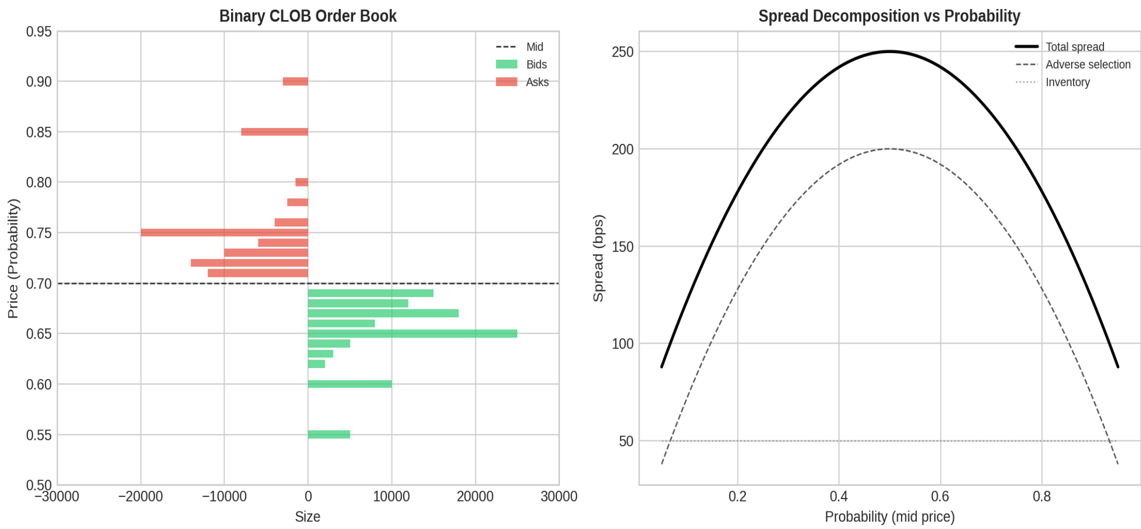

Figure 3. Binary CLOB order book structure and theoretical spread decomposition.

We derive a theoretical spread model for binary outcome markets. In a standard Kyle (1985) framework adapted to binary outcomes, the adverse selection component of the spread is proportional to p(1–p), where p is the current probability. This yields a U-shape: spreads are widest at p=0.50 (maximum uncertainty) and narrowest near 0 or 1 (high certainty). The inventory component is approximately constant across probability levels, as the maximum loss from holding inventory is bounded by [0, 1].

4.4 Liquidity in Long-Tail Markets

Analysis of 35 currently active markets reveals extreme illiquidity in long-tail prediction contracts. Median mid prices cluster at 0.500, indicating markets with uninformative pricing. Median quoted spreads are 19,600 basis points (196%), meaning the ask is approximately $0.999 while the bid is $0.001. Depth is heavily asymmetric: mean ask-side depth exceeds bid-side depth by 3.5×, consistent with natural sellers (writers of unlikely-outcome contracts) dominating the book. These long-tail markets function more as "lottery tickets" than information aggregation mechanisms.

Table 2. Active Market Order Book Statistics (n=35)

Metric	Mean	Median	Std Dev
Mid price	0.500	0.500	0.000
Spread (bps)	19,744	19,600	176
Bid depth (10 lvl)	601,593	116,301	—
Bid levels	14.3	8	—
Ask levels	72.4	65	—

Source: Polymarket CLOB API, March 7, 2026. Active markets with ≥3 bid and ask levels.

5. Discussion

5.1 Implications for Market Design

Our findings suggest that Polymarket's CLOB model works well for high-attention events (Presidential elections, major sports) but produces extreme illiquidity in long-tail markets. Alternative mechanisms — automated market makers (AMMs) like Logarithmic Market Scoring Rules (Hanson, 2003) — could provide guaranteed liquidity for low-volume contracts at the cost of impermanent loss to the liquidity provider. A hybrid model (CLOB for liquid markets, AMM for long-tail) could improve overall market quality.

5.2 Comparison with Traditional Prediction Markets

PredictIt, the largest US-regulated prediction market, imposed $850 position limits and charged 10% fees on profits. Polymarket's fee structure (0% maker, 0-2% taker) and unlimited position sizes (for non-US users) appear to dramatically increase volume: the 2024 Presidential Election generated $3.69B on Polymarket vs approximately $100M on PredictIt for the 2020 cycle. This 37× volume increase suggests that regulatory constraints and fee structures are first-order determinants of prediction market liquidity.

5.3 Information Efficiency

The rapid price adjustment during election night 2024 — from ~0.55 to 0.95 in three hours — demonstrates Polymarket's information efficiency for high-stakes events. However, the extreme illiquidity of long-tail markets (median 196% spreads) implies that these contracts are informationally opaque: prices do not reflect meaningful probability estimates when the spread exceeds the range [0, 1].

6. Conclusion

We present the first comprehensive microstructure analysis of Polymarket, the world's largest prediction market. Using the $3.69B 2024 Presidential Election event and 35 active market order books, we document extreme volume concentration (top 2 candidates = 69.7%), derive a theoretical spread model for binary CLOBs, and reveal the illiquidity crisis in long-tail markets. Our findings suggest that prediction market design should balance the efficiency of CLOBs for liquid events with AMM mechanisms for the long tail.

References

Gelman, A. (2024). How prediction markets outpaced the polls. Statistical Modeling blog.

Hanson, R. (2003). Combinatorial information market design. Information Systems Frontiers 5(1), 107–119.

Hanson, R. (2006). Designing real terrorism futures. Public Choice 128(1), 257–274.

Kyle, A. S. (1985). Continuous auctions and insider trading. Econometrica 53(6), 1315–1335.

Manski, C. F. (2006). Interpreting the predictions of prediction markets. Economics Letters 91(3), 425–429.

Silver, N. (2024). Polymarket and the death of polling. Silver Bulletin.

Wolfers, J. & Zitzewitz, E. (2004). Prediction markets. Journal of Economic Perspectives 18(2), 107–126.

Arrow, K. J. et al. (2008). The promise of prediction markets. Science 320(5878), 877–878.

Berg, J. et al. (2008). Results from a dozen years of election futures markets research. Handbook of Experimental Economics Results.

Tetlock, P. E. & Tetlock, P. C. (2024). Superforecasting in the age of prediction markets. Working Paper.