From Bets to Narratives: Understanding Whale Traders' Behavior in the 2024 US Election Prediction Market

Anonymous Author(s)

ABSTRACT

Prediction markets aggregate collective beliefs through trading, offering a unique lens into how information, sentiment, and capital interact in real time. While traditional polling outlets such as FiveThirtyEight (538) struggled to capture the dynamics of the 2024 U.S. Presidential Election, decentralized platforms like Polymarket provided more adaptive and timely forecasts. In this study, we take Polymarket's election market as a case to investigate how large traders (`Whales`) shape both market behavior and public discourse. By integrating on-chain transaction data with user-generated comments, we characterize trading patterns and linguistic signals across user groups. Our analysis reveals that (1) high ROI (Return on Investment) users often exhibited early insight by purchasing *Harris Yes* shares before key campaign events, with a subset achieving ROI > 10 through concentrated bets (a pattern we term "small trial, lucky strike"); (2) high PnL (Profit and Loss) `Whales` (|PnL| > 10k) exhibit distinct trading frequencies and automation tendencies compared with regular users; (3) cross-market analyses support the Efficient Market Hypothesis, suggesting limited transferability of abnormal returns, while Kyle model estimation further indicates that whale traders exert leading influence on market price formation; and (4) opinion mining reveals that `winning whales` rely more on conviction-driven and insider-like reasoning, whereas `losing whales` emphasize external signals such as polls and debates. Together, these findings uncover the behavioral and informational mechanisms underlying modern prediction markets, contributing to the broader understanding of collective intelligence, financial microdynamics, and online political forecasting.

CCS CONCEPTS

• **Information systems → Collaborative and social computing systems and tools**; **Web mining**.

KEYWORDS

Polymarket, U.S. 2024 Presidential Election, Predict Market

ACM Reference Format:

Anonymous Author(s). 2018. From Bets to Narratives: Understanding Whale Traders' Behavior in the 2024 US Election Prediction Market. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/1122445.1122456

Woodstock '18, June 03–05, 2018, Woodstock, NY

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00

https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

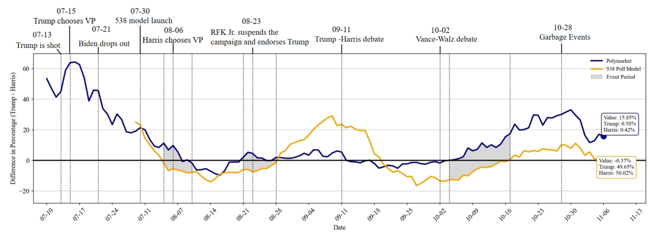

Figure 1: FiveThirtyEight prediction results and Polymarket odds difference. The final prediction result is that Polymarket wins.

Prediction markets are exchange-like platforms where participants trade contracts whose payoffs depend on uncertain future events [41]. As early as the 16th century, Italians were already betting on papal elections in the streets of Rome, and the fluctuating odds were interpreted as signals of shifting political power [21, 40]. Similar speculative venues emerged in 18th-century London coffeehouses and 19th-century New York billiard halls, where market odds were often treated as informal forecasts of political or economic outcomes [32–34]. In the 20th century, prediction markets were largely replaced by scientific polling methods represented by George Gallup [11], yet the academic fascination with collective intelligence never faded. The Iowa Electronic Markets (IEM) [15], founded by the University of Iowa in 1988, became a milestone in demonstrating that real-money political futures contracts frequently matched or surpassed opinion polls in forecasting accuracy [5, 7, 37] Some studies report that in U.S. presidential elections since 1988, prediction markets were closer to the eventual vote shares about 74% of the time compared to polls [5].

Recent years have witnessed a revival of prediction markets, driven by blockchain technology (especially, smart contract) and decentralized finance (DeFi) [25, 44, 45]. Among them, Polymarket, built on the Polygon blockchain [23], has become one of the most active decentralized prediction markets. It allows users to trade binary outcomes such as "*Will candidate Trump win the 2024 U.S. Presidential Election?*"[1] with tokenized shares reflecting the crowd's belief in each outcome. Each event can correspond to an independent market, where traders buy *Yes* or *No* shares using the `USDC` stablecoin. The trading price, ranging between 0 and 1, reflects the market's real-time consensus on the probability of an event's occurrence rather than any intrinsic value. Unlike traditional betting exchanges, Polymarket has no centralized bookmaker or fixed odds. All transactions are executed through a Continuous Limit Order Book (CLOB) implemented by smart contracts, ensuring price discovery purely from user interactions. Participants can freely enter or exit positions at any time before the outcome is decided.

[1]https://polymarket.com/event/presidential-election-winner-2024

During the 2024 U.S. Presidential Election, Polymarket gained widespread public attention for its forecasting accuracy. Presidential Election, particularly because they appeared to outperform traditional polling aggregators such as FiveThirtyEight (538). As illustrated in Figure 1, while FiveThirtyEight (538) consistently favored Harris in the final prediction, Polymarket traders dynamically updated their expectations in response to key campaign events—most notably, the Trump shooting and the presidential debate. These fluctuations captured in real time the evolving public sentiment and informational flow, ultimately converging on the correct outcome (Trump win). Interestingly, both systems exhibited correlated directional trends after major events, suggesting that prediction markets and opinion polls may share informational inputs, yet differ fundamentally in responsiveness and volatility.

However, such pronounced volatility also raises a deeper question: **who drives these sharp market movements**? In the cryptocurrency community, large traders (known as `Whales`) are capable of shifting prices through concentrated trades and liquidity provision. In this paper, we focus on these `whale` users, using the 2024 U.S. Presidential Election on Polymarket as a case to explore how they bet and how their opinions and narratives differ from other participants. By integrating transaction-level blockchain data with user-generated comments, we bridge financial behavior and social expression to examine the interplay of belief, information, and emotion in decentralized prediction markets. Based on this motivation, we articulate the following research questions:

(1) **RQ1.** What are the characteristics of the most representative users (`Whale` user) in the 2024 U.S. Presidential Election market?
(2) **RQ2.** Do `Whale` influence market prices through informational advantages? Can such dynamics be explained by the Kyle model?
(3) **RQ3.** Does the Efficient Market Hypothesis (EMH) hold in Polymarket? Specifically, do `Whale` who profit in the election market also consistently succeed in other event categories?
(4) **RQ4.** How do users' opinions and narratives differ across groups with varying market performance? In particular, do losing `Whale` and winning `Whale` exhibit distinctive attention toward political events (e.g., polls, debates) in their discourse?

Through these questions, our study provides a multi-faceted view of how information, belief, and capital co-evolve in modern prediction markets. We aim to deepen our understanding of collective intelligence and market-based forecasting in the age of decentralized platforms.

2 BACKGROUND

In this section, we introduce some background information about Polymarket as a basis for subsequent research.

2.1 Polymarket Prediction Market

Prediction markets aggregate dispersed information by allowing participants to trade contracts whose payoffs depend on future events. Market prices reflect collective beliefs, translating subjective expectations into quantitative probabilities and embodying the "wisdom of crowds."

Polymarket, a leading Web3 prediction market on the Polygon blockchain, combines DeFi primitives with transparent forecasting.

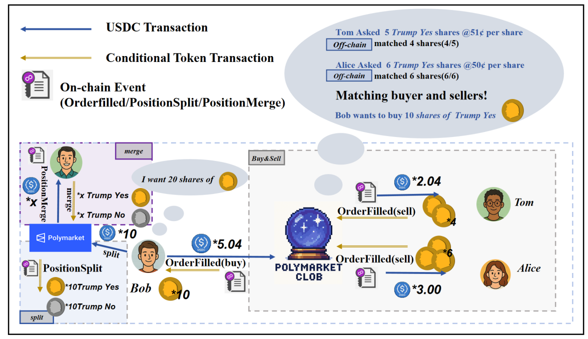

Figure 2: The CLOB Matching Mechanism on Polymarket and Users' `Buy`, `Sell`, `Split`, and `Merge` Actions

Users trade on politics (e.g., *Will Trump win the 2024 U.S. Presidential Election?*), macroeconomics, sports, or entertainment, with all markets denominated in USDC stablecoin. Each market can represent a binary proposition (*Yes* or *No*). The *Yes* price (0–1 USD) corresponds to the market-implied probability. For example, if the contract is priced at $0.51, the market collectively implies a 51% likelihood of Trump's victory. Upon resolution, the correct outcome token `redeems` $1, while the other expires worthless. It is worth mentioning that participants can freely enter or exit positions at any time before the outcome is decided—for instance, purchasing *Harris Yes* contracts when her chances were only 10% and selling later at $0.5 to secure profit without waiting for the election result.

2.2 Polymarket's Decentralized Architecture

Polymarket's markets are created via the Gnosis **Conditional Tokens Framework (CTF)** [39], trading uses a hybrid **Central Limit Order Book (CLOB)**, and outcomes are finalized on-chain through **Universal Market Access (UMA)** Optimistic Oracle with community dispute resolution.

Polymarket primarily operates on a **Central Limit Order Book** (CLOB) mechanism, where buy and sell orders are matched.

As illustrated in Figure 2, Alice and Tom place limit sell orders for 6 shares at $0.50 and 5 shares at $0.51. Bob submits a market buy order for 10 shares, which executes 6 and 4 shares against Alice's and Tom's orders, triggering three `OrderFilled` events. Bob can then `split` 10 USDC into 10 *Trump Yes* and 10 *Trump No* tokens, sell unwanted positions, or `merge` equal quantities of opposite shares back into USDC. Through *buy*, *sell*, `split`, and `merge`, users can flexibly construct trading strategies on Polymarket.

The **Conditional Tokens Framework** (CTF) proposed by Gnosis enables on-chain representation of arbitrary event outcomes. In Polymarket, a major event such as *Presidential Election Winner 2024* can be decomposed into multiple candidate-specific markets, such as *will-kamala-harris-win-the-2024-us-presidential-election.* Each market issues binary outcome tokens—*Yes* and *No*—implemented as unique ERC-1155 conditional tokens. These tokens are generated through the protocol's `split` operation by locking collateral (typically USDC) and can later be recombined into the underlying collateral via `merge`. Orders are first matched off-chain in the CLOB, and only after a match is found are the token transfers and USDC settlements executed on-chain.

3 DATASET

As we discussed in Section 2, Polymarket employs a CLOB architecture, where user orders are matched off-chain and settled on-chain on the Polygon network. Each transaction invokes a Polymarket smart contract, generating event logs that can be parsed through a Web3 node. These on-chain transactions are publicly recorded and can be fully traced through blockchain nodes or data APIs, allowing researchers and developers to reconstruct user activities, verify market outcomes, and analyze trading behaviors with high transparency and verifiability. Our data consists of two parts: **Polygon blockchain** and **Polymarket Gamma API**.

- **Polygon blockchain:** We obtained on-chain transaction data from the Polygon network using the Rust-based open-source tool **cryo**[2], connected to a locally deployed Polygon Erigon3 archive node. We filtered and decoded the event logs emitted by Polymarket's smart contracts, including `OrderFilled` events (corresponding to `Buy` and `Sell`) and `PositionSplit`/`PositionMerge` events (corresponding to `split` and `merge`). Each transaction was uniquely identified by two key parameters: the `event_id`, representing the ERC-1155 token associated with market outcomes, and the `conditionId`, representing the underlying market condition in `split`/`merge` operations. After parsing and normalization, we extracted a total of 33,263,782 `OrderFilled` transactions and 6,778,728 `PositionSplit`/`PositionMerge` transactions. The resulting dataset was further structured into standardized fields (e.g., address, price, and timestamp), as detailed in Table 6 in the Appendix (Section A).
- **Polymarket's Gamma API** offers a structured, read-only interface for accessing market-level data on the platform. Each market is uniquely identified by a `slug` and is associated with a set of `outcomes` representing possible event results, along with their corresponding `outcomePrices`, which reflect market-implied probabilities. The `closed` flag indicates whether the market has concluded, while the `conditionId` uniquely references the corresponding on-chain data for the event on the Polygon blockchain. The API supports retrieval by user `hash`, enabling us to obtain supplemental data on users' cross-market activity and comments based on the underlying on-chain information. In this study, we utilized the Gamma API to collect all markets participated in by users who placed bets on the US Election 2024 event for Trump and Harris, yielding a total of 18,203 markets. From these, we selected 16,646 markets that had `closed=true` as of September 24, 2025, for subsequent analysis of cross-market user behavior. In addition, we used the Gamma API to retrieve 210,753 comments under the *Presidential Election Winner 2024* event, providing further insights into user sentiment and interactions during the election period.

4 GROUPED USER PROFILING (RQ1)

In this section, we investigate the characteristics of the most representative users in the 2024 U.S. Presidential Election market. We first define **Profit and Loss (PnL)** and **Return on Investment (ROI)** , and then define `Whale` users. We then perform a group analysis on 230,190 users and discuss the distribution of users and the operations of `Whales`.

[2] https://github.com/paradigmxyz/cryo

> Findings 1. *(1) We found a behavior that can be described as "small trials, lucky bets", where high-ROI users often showed early insight by buying Harris Yes contracts at very low prices and selling them at higher prices; (2)* `Whale` *users ($|PnL| > 10k$) exhibited significantly different trading behaviors, compared to the rest of the population.*

4.1 PnL, ROI and `Whale` Definitions

In this subsection, we give the definition for **Profit and Loss (PnL)** and **Return on Investment (ROI)** as follows.

- **Profit and Loss (PnL):** A user's PnL consists of three main components: (1) **`Buy/Sell` net cash flow**: the difference between the USDC earned from sell transactions and the USDC (U) spent on buy transactions; (2) **`Split/Merge` net cash flow**: splitting consumes an equal amount of USDC, while merging returns an equal amount of USDC (e.g., *Trump Yes*, *Trump No*); (3) **`Redeemable` positions**: the value of remaining redeemable positions (e.g., *Trump Yes*, *Harris No*) at market resolution, where 1 token = 1 USDC. The PnL of user u is computed as:

$$\mathrm{PnL}_u = \underbrace{\Big(\sum_{t\in\mathrm{Sell}_u} U_t - \sum_{t\in\mathrm{Buy}_u} U_t\Big)}_{\texttt{Buy/Sell}} + \underbrace{\Big(\sum_{t\in\mathrm{Merge}_u} U_t - \sum_{t\in\mathrm{Split}_u} U_t\Big)}_{\texttt{Split/Merge}} + \underbrace{\Big(\sum_{m\in\mathrm{Resolution}_u} \mathrm{Position}_{u,m}\Big)}_{\texttt{Redeemable}}. \tag{1}$$

  By analyzing user-level PnL in this representative market, we can capture the distribution of profits and losses and stratify users into distinct groups.
- **Return on Investment (ROI)** measures the efficiency of capital usage, defined as:

$$\mathrm{ROI}_u = \frac{\mathrm{PnL}_u}{\mathrm{Investment}_u}, \tag{2}$$

  where, $\mathrm{Investment}_u$ denotes the total capital committed by user u, calculated as the sum of all `Buy` expenditures and Split costs. `Buy` operations involve direct purchase of outcome tokens, while `Split` operations lock an equivalent amount of USDC to create token pairs. `Sell` and `Merge` operations return USDC and are therefore excluded from the investment calculation. Incorporating ROI further accounts for users' initial capital investment, enabling a more multidimensional characterization of user performance.
- **`Whale`**: Based on PnL definition, we define users with PnL > 10k referred to as `Winning Whales`; users with PnL<-10k referred to as `Losing Whales`. It is worth mentioning that the definition of Whales may vary in different studies. We choose \$10k as the threshold because these users account for more than 80% of PnL. In this paper, we use the PnL definition for the subsequent analysis.

4.2 User Distribution

Based on the definition on Section 4.1, we analyze user performance by jointly considering Investment, ROI and PnL.

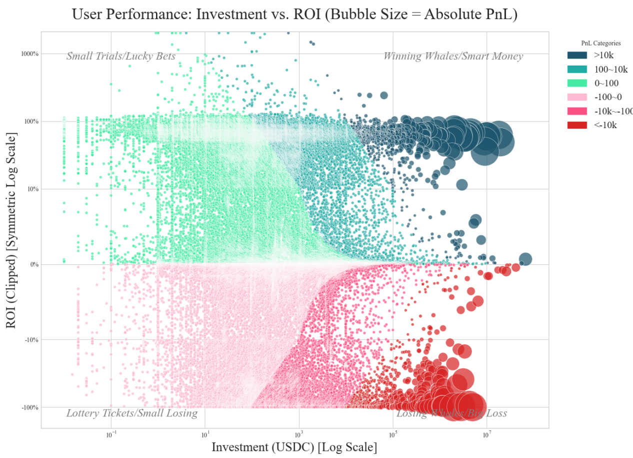

Figure 3: User performance scatter chart, where the x-axis represents Investment, the y-axis represents ROI, and the circle size indicates PnL. Each user is a point in the chart.

Figure 3 illustrates user performance with respect to investment size, ROI, and absolute PnL. The majority of participants placed wagers ranging from $1 to $100k. Notably, users achieving high ROI are mostly small-scale bettors; however, their absolute PnL remains relatively modest, indicating that high percentage returns do not necessarily correspond to large monetary gains. Conversely, traders with exceptionally large investments (>$100k) display highly polarized outcomes: either doubling their gains or incurring substantial losses. Despite representing only 0.35% of the total user base, these `winning whales` account for an aggregate PnL of 110,789,752.45, comprising 88.33% of all positive PnL in the market (see Table 1). On the other hand, `losing whales`, constitute merely 0.49% of users but are responsible for 84.42% of the total negative PnL.

Given that the total PnL across all 230,190 users approximately sums to zero, Polymarket operates as a zero-sum market. Among users with total PnL below $10000, losses generally exceed equivalent gains, reinforcing the casino-like characteristics of the platform. In this zero-sum context, smaller-scale participants (total PnL below $10000) collectively realized limited profits and substantial losses, whereas high-stakes winners captured the majority of funds from the market. Table 1 provides a detailed breakdown of user counts and aggregate PnL across different profit and loss intervals, highlighting the extreme contribution of a small fraction of users to the overall market outcome.

4.3 High ROI Bettors

In Figure 3, we observed some high ROI users. In this subsection, we analyze users achieving ROI above 100.

Our analysis indicates that users achieving ROI above 100% especially those whose ROI above 1000%, are closely associated with early purchases at relatively low prices in the *Harris Yes* market. We identified two notable users who relied almost exclusively on the *Harris Yes* market. The first, **@ShitposterIsBack**, accumulated *Harris Yes* positions between July 4–12, 2024, executing 88 buy trades totaling $65,270.59. In total, 429,741.31 shares were purchased at

Table 1: Distribution of Users and PnL across Profit and Loss Intervals in the Market.

PnL Interval	User Count	% within Group	% of Total Users	PnL Sum	% of Group PnL
Positive PnL (Total Users = 87,236, Total PnL = 125,424,480.67)					
0–100	70,814	81.18%	30.76%	1,150,456.01	0.92%
100–10,000	15,626	17.91%	6.79%	13,484,272.21	10.75%
>10,000	796	0.91%	0.35%	110,789,752.45	88.33%
Negative PnL (Total Users = 142,954, Total PnL = -125,424,480.70)					
-100–0	123,048	86.08%	53.45%	-1,397,840.62	1.11%
-10,000– -100	18,789	13.14%	8.16%	-18,138,624.78	14.46%
< -10,000	1,117	0.78%	0.49%	-105,888,015.29	84.42%

Table 2: User Group Analysis Summary for the Harris Yes Market (ROI > 0).

Group	User Count	Avg Buy Price	Avg Sell Price	Harris Net PnL / Total PnL Ratio (%)
ROI 0–1	82,997	0.4034	0.4168	-8.36
ROI 1–10	3,280	0.0879	0.2972	28.08
ROI >10	23	0.0166	0.3443	100.00

a weighted average price of 0.1519 USDC and later liquidated between August 10–12 at 0.5161 USDC, yielding a profit of $156,513.94 and making this user the top-earning `whale` from the event. No other `whales` deployed comparable early capital, and this user did not participate in subsequent events.

The second case, **@lasso**, illustrates an extreme high-ROI outcome: investing only $600 for 39,999.996 shares at 0.0150 USDC and selling at 0.3400 USDC, the user realized $13,000 profit (ROI 21.67), showing that small-scale bettors could achieve outsized returns with early entry.

Table 2 summarizes *Harris Yes* user group analysis. Users with ROI 0–1 contribute minimally to net PnL (-8.36%), ROI 1–10 users contribute positively (28.08%), and the 23 users with ROI >10 account for 100% of net PnL in this extreme category. These outliers, above the 1000% ROI line in Figure 3, illustrate that extraordinary returns stemmed from early, small-scale investments. Overall, high ROI indicates effective timing and strategy but does not guarantee large absolute gains without sufficient investment scale.

4.4 Whale User

In this subsection, we analyze whether users with extreme profit or loss (`Whale`) exhibit distinctive behavioral and transactional characteristics, including substantially larger investments, higher trading frequency, and a markedly greater reliance on automated trading systems.

4.4.1 Investment. As shown in Table 3, users in the PnL >10k group (`Winning Whales`) have substantially higher median investments compared to other PnL groups. Conversely, users in the PnL <-10k group (`Losing Whales`) also tend to have higher investments compared to most small PnL users, although their PnL outcomes are negative. This pattern is consistent with the distribution depicted in the boxplot (Figure 4), illustrating the clear difference in transaction amounts between `Winning Whales`, `Losing Whales`, and others .

4.4.2 Transaction Frequency. In this subsection, we calculated the total trading volume (i.e., number of transactions) of users across different PnL intervals as a measure of trading frequency.

Table 3: Comparison of Investment Amounts Across PnL Groups. `Winning Whales` (Pnl >10k) exhibit significantly higher investment than all other groups, consistent with the boxplot shown in Figure 4.

PnL Group	User Count	Median Investment	Mann-Whitney U	p-value
`Losing Whales` $(-\infty, -10{,}000]$	1,117	32,990.18	535,931.50	<0.001
$(-10{,}000, -100]$	18,789	557.00	14,662,587.00	0.0
$(-100, 0)$	123,048	310.87	97,103,029.50	0.0
$[0, 100)$	70,814	74.00	55,972,914.50	0.0
$[100, 10{,}000)$	15,626	749.72	12,106,504.50	0.0
`Winning Whales` $[10{,}000, \infty)$	796	50,765.97	-	-

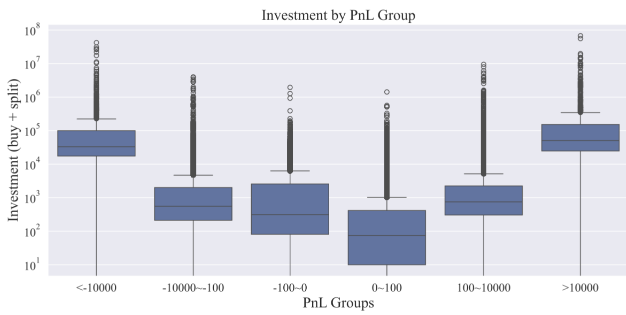

Figure 4: The median and distribution range of total investment are significantly higher in the `Winning Whales` group (PnL > 10k) compared to all other PnL groups, while the investment of the `Losing Whales` group (PnL < -10k) is also higher than the remaining PnL groups.

The analysis shows that high-frequency trading (more than 500 transactions) is concentrated among users with extreme PnL. Specifically, among `Losing Whales`, approximately 10.8% are high-frequency traders (95%, CI: 8.6%–12.9%), while among `Winning Whales`, the proportion is about 10.6% (95%, CI: 8.6%–12.9%). In contrast, users with moderate gains or losses exhibit much lower high-frequency trading rates. For example, the (-10k,-100] group has a proportion of 0.4% (95%, CI: 0.3%–0.5%), and the [0,100) gain group only 0.3% (95%, CI: 0.2%–0.3%). These findings indicate that extreme PnL users, both on the winning and losing side, tend to engage in high-frequency trading, whereas typical users predominantly trade at low frequency. The symmetry between winning and losing `whales` suggests that extreme PnL, rather than gain or loss direction, is the main driver of high-frequency trading behavior.

4.4.3 Automated Trading. In our analysis, we observed that some users exhibit multiple trades recorded at the exact same timestamp. Such patterns can arise for two distinct reasons: (i) the *Polymarket matching mechanism*, where trades matched through the CLOB are simultaneously submitted on-chain within a single transaction, resulting in multiple fills recorded under the same `transaction_hash`; or (ii) *automated transaction systems*, where users employ bots or scripts to place trades in a highly synchronized manner.

We identify automated trading by leveraging the `transaction_hash` field: multiple trades with the same timestamp and hash are considered internal matches, whereas trades at the same timestamp but across different hashes indicate *Automated transaction* usage.

Applying this method across user groups, high-frequency automation is concentrated among extreme PnL users: 52.7% of *losing whales* (589/1,117) and 49.6% of `Winning Whales` (395/796) exhibit automated trading, while moderate PnL users show substantially lower rates (11.7–13.6%). Two-proportion z-tests confirm these differences are highly significant ($p < 0.001$) relative to typical users, whereas `Winning Whales` vs. `Losing Whales` show no significant difference ($z = 1.34$, $p = 0.18$). This indicates that automation amplifies exposure to extreme outcomes on Polymarket, independent of profit direction.

5 KYLE MODEL ANALYSIS (RQ2)

> FINDINGS 2. *(1) Other users show negative λ, acting primarily as post-event price followers. (2)* `Winning Whales` *exhibit short-term positive λ following high-impact events, reflecting active trading that directly influences prices rather than merely following trends.*

In this section, we use the Kyle Model [24] to analyze how `Whales` influence the market.

5.1 Kyle Theory

Kyle's model studies how informed traders impact prices in a market with asymmetric information. The core insight is that price changes ΔP relate linearly to net order flow Q, with the slope coefficient λ (*price impact coefficient*) quantifying the influence of a unit trade on price:

$$\Delta P_t = \lambda Q_t + \epsilon_t, \tag{3}$$

where $\Delta P_t = P_{t+1} - P_t$ represents the price change following a trade, and Q_t is the net trade volume within the period (positive for net buying pressure, negative for net selling).

5.2 λ Estimation in Polymarket

In Polymarket, we estimate λ for different user groups and markets based on observed transaction data.

5.2.1 Methods. Our estimation process consists of four steps:

(1) **User Group Classification**: Users are categorized according to realized PnL: `Winning Whales`, `Losing Whales`, and others.
(2) **Market Selection**: We focus on the *Trump* and *Harris* presidential election markets, including both *Yes* and *No* positions.
(3) **Price and Net Trade Volume**: For each trade, we define

$$P_t = \begin{cases} \text{price}, & \text{if trading the } \textit{Yes} \text{ position,} \\ 1 - \text{price}, & \text{if trading the } \textit{No} \text{ position,} \end{cases} \tag{4}$$

$$\Delta P_t = P_{t+1} - P_t, \tag{5}$$

$$Q_t = \text{sign} \times \frac{\text{value}}{\text{price}}, \tag{6}$$

where sign = +1 for trades that push the price upward (buying *Yes* or selling *No*), and sign = −1 for trades that push the price downward (selling *Yes* or buying *No*). Because the sum of the *Yes* and *No* token prices typically fluctuates around 1 USDC, we use 1 − price to represent the *No* side price without loss of generality.

(4) **Time Window and Regression**: All transactions are chronologically sorted and aggregated into non-overlapping 12-hour windows. Within each window, we run an OLS regression without intercept:

$$\Delta P_t = \lambda Q_t + \epsilon_t, \tag{7}$$

where the estimated coefficient λ represents the group-level price impact in that time window.

5.2.2 Per-Capita Regression. To further control for disparities in group sizes (e.g., the number of `Whales` vs. others), we employ a *per-capita* λ estimation procedure. Within each 12-hour window, we perform an individual-level regression for each trader using Equation 7. For each user group g (`Winning whales`, `Winning whales`, others), we then compute the mean of all individual λ_i values to obtain the per-capita price impact using Equation 7: This normalization ensures that differences in $\bar{\lambda}$ reflect behavioral or informational heterogeneity, rather than disparities in participant counts.

By adopting this *per-capita* normalization, we eliminate the confounding effect of unequal group sizes between `whales` and other traders, allowing for a more accurate comparison of intrinsic price impact intensity across user segments and market domains.

5.2.3 Interpretation of λ.

Sign of λ. The sign of λ reflects a user group's trading direction relative to the market. Positive λ ($\lambda > 0$) indicates trend-following trades that push prices with market flow, while negative λ ($\lambda < 0$) indicates contrarian trades against market movements.

Magnitude of λ. Larger $|\lambda|$ signals stronger price sensitivity to trades, reflecting high market influence, event-driven activity, or low liquidity amplifying price changes.

Temporal dynamics. Increasing $|\lambda|$ denotes stronger price impact and active or informed trading; decreasing $|\lambda|$ indicates market stabilization and deeper liquidity.

5.3 Event-Driven Liquidity Shifts Analysis

In this subsection, we combine major events, analyze λ, and study Event-Driven Liquidity Shifts. In Figure 5, we do λ analysis on *Harris Markets*. The time points marked by green dashed lines (*2024-07-13*, *2024-08-06*, *2024-09-11*, *2024-10-02*, *2024-11-05*)—corresponding to the *Trump was shot*, *Harris chooses VP*, *Trump-Harris debate*, *Vance-Walz debate*, and *Election Day*, respectively—generally exhibit such dynamics, except for the performance of `Winning Whales` during the *2024-09-11* presidential debate.

Prior to major events, market liquidity is typically thin, with shallow order books amplifying individual trades' price impacts and resulting in larger absolute values of λ. During event windows, participation surges on both the buy and sell sides, deepening liquidity and driving λ toward zero. Consequently, event periods are characterized by sharp pre-event fluctuations in λ followed by stabilization.

However, the persistently negative λ among ordinary users reflects their limited influence and low profitability. After the Trump shooting incident (2024-07-13)—perceived as pro-Trump—their trading volume in Harris markets spiked (e.g., Harris Buy +105.9%, Harris No +358.7%), yet λ remained negative. These users intensified trading in already liquid zones without moving prices, acting as post-event followers rather than price leaders. In contrast, following the September 11 presidential debate, `winning whales` exhibited a short-term positive λ, signaling temporary price leadership. Their trading volume expanded drastically (e.g., Harris Yes/No Buy +486.7% / +4450.1%), enhancing overall liquidity while influencing prices. This transient positive λ thus reflects informed trading rather than illiquidity.

Overall, the contrasting λ dynamics confirm that large players with informational advantages—`winning whales`—actively shape prices during high-impact events, consistent with the Kyle model. Informed traders strategically move prices based on superior information, whereas less informed participants primarily react to these movements, reinforcing the market's information-absorption efficiency.

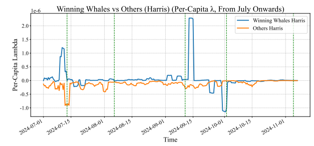

(a) Winning Whales vs Others (Harris)

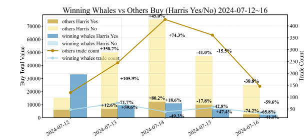

(b) Winning Whales vs Others Buy Transactions Fluctuation

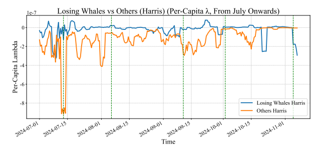

(c) `Losing Whales` vs Others (Harris)

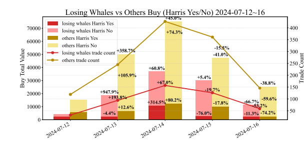

(d) `Losing Whales` vs Others Buy Transactions Fluctuation

Figure 5: Comparison of Whale and Other User Buy Behavior in Harris Markets (Around July 13, 2024).

6 EMH ANALYSIS (RQ3)

> FINDINGS 3. *(1) Users' performance in Elections events is only related to Trumps events and shows no correlation with other events. (2)* `Winning Whales` *outperform in Elections/Trump, while* `losing whales` *slightly outperform in some non-political categories, with limited generalization. (3) Overall,* `Whales` *profitability is domain-specific and not systematically transferable, consistent with EMH.*

In this section, we discuss on Efficient Market Hypothesis (EMH), which posits that participants cannot systematically earn persistent excess returns. We focus on whale-level profitability to see if it exhibits cross-category persistence, indicating transferable predictability and potential arbitrage [35].

We used Gamma API to obtain various events, excluding *Presidential Election Winner 2024*, totaling 43,265. We further divided these markets into 12 different group[3]. The classification criteria are in the appendix. Further, we standardized 43,265 events into Polymarket's taxonomy to build a user–category PnL matrix, then applied Spearman correlations for cross-category consistency

[3]The classification criteria are in the Appendix (Section C).

and Mann–Whitney U tests to compare PnL distributions between groups.

Table 4: Cross-Category Correlation and Winning vs. Losing Whales Comparison Significance. The p-values *, **, and * correspond to < 0.05, 0.01, and 0.001, respectively.**

Category	Correlation	Winner Mean	Loser Mean
Elections	/	165,052.43***	-104,003.69***
Trump	0.2378***	85,537.36***	-2,778.20***
Earnings	0.1107	3,074.68	1,226.67
Crypto	0.0707	1,864.66	4,631.45
Culture	0.0517	7,035.70	9,074.98
Politics	0.0470	61,710.15	17,237.64
Mentions	0.0363	1,526.75	2,859.09
World	-0.0059	10,082.23	3,315.47
Sports	-0.0087	-2,225.73*	6,170.80*
Economy	-0.0188	9,532.23	13,650.25
Geopolitics	-0.0457	5,203.50	14,934.40
Tech	-0.1044*	-9,533.33***	16,392.43***

6.1 Event Category Distribution

Figure 7 in Appendix illustrates `Whales`' total and average PnL across categories. We find that political domains (e.g., *Elections*/*Politics*) attract more whale participation and feature larger PnL magnitudes; meanwhile, average PnL varies across categories, suggesting domain expertise and strategic preferences. Interestingly, aside from setbacks in Trump-related events, `Losing Whales` from the *Presidential Election Winner 2024* market actually realize positive PnL to varying degrees in other categories, both in terms of total and average PnL. Conversely, previously highly successful `winning whales` show similar patterns, and in some cases—such as *Tech* and *Sports* events—even register negative PnL.

This suggests that, at the preliminary data level, profitable `whales` in the *Presidential Election Winner 2024* market **do not consistently maintain a performance advantage across other markets**.

6.2 Cross-Category Profit Relationship

To examine the consistency of user performance across different categories, we compute Spearman rank correlations between *Elections* PnL and PnL in other categories, restricting to users with non-zero PnL in both categories to reduce sparse-noise effects. Table 4 summarizes the correlations and their significance: correlations are particularly strong and highly significant for US election–related events; for example, *Trump* and *Elections* are positively correlated ($r \approx 0.238$, $p < 10^{-11}$). Across non-election-related categories, correlations are generally weak and non-significant, with a small but significant negative correlation in *Tech* ($r \approx -0.104$, $p < 0.05$). These findings suggest that user performance consistency is largely limited to US election–related events, indicating minimal cross-category transferability outside election markets, which aligns with EMH.

6.3 `Whales` Comparison

To compare the performance differences between distinct `whales` groups, we classify users in the *Presidential Election Winner 2024* market based on their total PnL into *winning whales* (net-positive) and *losing whales* (net-negative).

For each of the 12 event categories, we compute the PnL of each `whale` and rank users within the category. We then apply two-sided Mann–Whitney U tests on these rank distributions to determine whether winning and losing `whales` exhibit systematic differences in performance across categories.

Table 4 presents mean PnL by category and significance levels. We found: (1) `Winning Whales` significantly outperform in *Elections*/*Trump* markets (e.g., $p < 0.001$), consistent with domain expertise and strategic advantages in specific political contexts. (2) `Losing Whales` slightly outperform in non-political categories such as *Tech* and *Sports*, but effects are limited and do not generalize across categories.

Together with the correlation analysis, the rank-based distribution, and the Mann–Whitney U comparisons, these results indicate that `whales`' profitability is largely domain-specific rather than systematically transferable across categories. While `winning whales` outperform in specific markets (e.g., *Elections*/*Trump*), their success does not generalize to unrelated domains.

7 OPINION AND NARRATIVE MINING (RQ4)

In this section, we conduct opinion mining on user-generated comments associated with the 2024 U.S. Presidential Election market. Using the official Polymarket Gamma API, we collected a total of 210,753 user comments related to this event. Building on our previous analysis in Section 4, we categorize users into five distinct groups: `Winning Whale`, `Losing Whale`, `Winning User`, `Losing User` and `Others`[4]. Table 5 summarizes their descriptive statistics. Overall, we observe that only a small fraction of traders who placed bets actively participated in the comment discussions (i.e., 5563/230190 = 0.024). Notably, `Winning Whales` contributed relatively few comments, whereas `Losing Whales` exhibited higher engagement both in comment count and comment length.

To characterize linguistic differences across user groups, we adopt the **log-odds ratio with informative Dirichlet prior** [26] to identify group-specific bigrams that are statistically distinctive from those used by other groups. Before estimation, we apply standard preprocessing steps, including tokenization, lemmatization, and stopword removal using **spaCy** [1], and we further construct bigrams to capture multiword expressions. Additionally, we employ a profanity lexicon to remove obscene or offensive words[5].

Figure 6 presents group-specific word clouds weighted by the log-odds z-scores. First, the comment corpus contains a substantial amount of emotional and evaluative expressions. Except for a few well-known traders (e.g., the "`French whale`", whose comments are shown in Appendix D), very few posts convey genuinely informative or trade-relevant content. Across multiple thresholds and resampling checks, we observe a consistent lexical divergence between `Whale` groups. In summary, the linguistic patterns reveal two distinct communication regimes among `whale` traders.

[4]Others represent users who did not participate in the betting but commented
[5]https://github.com/snguyenthanh/better_profanity

(a) Winning whales (b) Losing whales (c) Winning users (d) Losing users (e) Others

Figure 6: Group-specific word cloud, sized by positive log-odds z-scores.

Table 5: Group-wise corpus statistics.

Group	Users	Total tokens	Avg. tokens/user	Total msgs	Avg. msgs/user
Winning Whales	160	47,389	296.2	2,420	15.1
Losing Whales	176	203,468	1,156.1	10,026	57.0
Winning Users	2,947	713,533	242.1	35,162	11.9
Losing Users	2,280	762,976	334.6	36,045	15.8
Others	7,163	2,112,164	294.9	127,100	17.7

`Winning Whales` tend to employ self-referential, meme-driven language infused with irony, humor, and dominance signaling (typified by expressions such as *stay_mad*, *seethe_cope*, and *touch_grass*). Such usage suggests a discourse style rooted in internal conviction and group identity rather than external information sources. In contrast, `Losing Whales` frequently reference *polls*, *debates*, and political figures or slogans such as *maga*, *chump*, and *brandon*, reflecting a reasoning mode anchored in mainstream narratives and external validation.

> Findings 4. `Whales` *who resist prevailing media frames and rely on private beliefs tend to exhibit superior market performance in U.S. 2024 Election prediction.*

8 RELATED WORK

8.1 Prediction Markets and On-Chain Platforms

Classical work shows that small, well-incentivized real-money prediction markets can aggregate dispersed information and often rival polls in forecast accuracy [5, 42]. Overviews clarify when market prices can be read as probabilities and what frictions bias them [43]. Mechanism-design advances—e.g., Hanson's LMSR and combinatorial markets—explain how automated market makers sustain liquidity and express dependencies among outcomes [17, 18, 28].

With the advent of blockchain, Decentralized Prediction Markets (DPMs) such as Augur, Omen/Gnosis, and Polymarket port these designs on-chain, exposing rich, high-frequency order-book data and enabling cross-market constraints [29, 31, 39]. Recent systematizations and case studies document both opportunities and pathologies unique to DPMs (e.g., oracle resolution, liquidity fees, and dependency mispricing) [3, 20]. These literatures [5, 17, 29, 36, 42] jointly motivate our **RQ1** (profiling users).

8.2 Informed Trading and Market Efficiency

Kyle's model predicts a linear relation between price changes and net order flow, with the price-impact coefficient λ summarizing market depth and the extent of information revelation [24]. Empirical tools estimate the informativeness of trades and decompose quote revisions [19]. No-dynamic-arbitrage constraints further link the shape and decay of impact [12], while execution studies document concave impact at scale [2]. These frameworks have been adapted beyond equities to test whether "`whales`" disproportionately move prices. We build on these to test **RQ2**: whether `whale` exhibit higher λ consistent with informational advantages.

Further, Efficient Market Hypothesis (EMH) posits that abnormal, systematic excess returns are difficult to sustain [8, 9]. Yet anomalies (e.g., momentum) and investor underperformance suggest bounded rationality [4, 22]. In prediction markets, long-run studies generally find high accuracy, with mixed evidence on exploitable inefficiencies and manipulation [5, 6]. We evaluate **RQ3** by testing whether `whales` who prospered in the 2024 U.S. election market also succeed across others categories.

8.3 Opinion and Narrative Mining

Recent studies on stance detection and political text analysis highlight how linguistic cues reveal ideological or informational divides in online discourse [10, 14, 27, 46]. Following this line, we analyze whether user groups differ in their opinions and attention toward political events. We adopt the *log-odds with informative Dirichlet prior* method [26] in **RQ4**, because it provides interpretable group-specific word salience and reliable uncertainty estimates under sparse data. Alternative methods such as TF-IDF [30], BERTopic [16], or inverse-regression estimators [13, 38] are also considered for robustness and alignment with prior work in web mining.

9 CONCLUSION

In this paper, we examined the co-evolution of information, belief, and capital in the 2024 U.S. Presidential Election prediction market on Polymarket. By integrating blockchain transactions with large-scale comment analysis, we offered a multi-dimensional perspective on market dynamics, emphasizing the role of large traders (`whales`). Our findings reveal that high-ROI users often possessed early informational advantages, entering positions ahead of major events and securing exceptional gains through concentrated bets. `Whales` displayed distinctive behavioral signatures—such as higher trading frequencies and automation tendencies—setting them apart from regular participants. Despite these disparities, cross-category consistency tests supported the Efficient Market Hypothesis, indicating that sustained abnormal profits are rare. Finally, linguistic analysis uncovered a divergence in reasoning styles: `winning whales` exhibited conviction-driven language, whereas `losing whales` relied on external cues like debates and polls. Together, these insights illuminate how individual cognition and market structure jointly shape predictive efficiency in decentralized environments.

REFERENCES

[1] Explosion AI. 2017. spacy: Industrial-strength natural language processing in python.

[2] Robert Almgren, Chee Thum, Emmanuel Hauptmann, and Hong Li. 2005. Direct estimation of equity market impact. *Risk* 18, 7 (2005), 58–62.

[3] Hamed Amini, Maxim Bichuch, and Zachary Feinstein. 2023. Decentralized Prediction Markets and Sports Books. *SSRN Electronic Journal* (2023). https://api.semanticscholar.org/CorpusID:259951199

[4] Brad M Barber and Terrance Odean. 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The journal of Finance* 55, 2 (2000), 773–806.

[5] Joyce E Berg, Forrest D Nelson, and Thomas A Rietz. 2008. Prediction market accuracy in the long run. *International Journal of Forecasting* 24, 2 (2008), 285–300.

[6] Joyce E. Berg and Thomas A. Rietz. 2014. Market Design, Manipulation, and Accuracy in Political Prediction Markets: Lessons from the Iowa Electronic Markets. *PS: Political Science & Politics* 47 (2014), 293 – 296. https://api.semanticscholar.org/CorpusID:154673614

[7] Robert S Erikson and Christopher Wlezien. 2012. Markets vs. polls as election predictors: An historical assessment. *Electoral Studies* 31, 3 (2012), 532–539.

[8] Eugene F Fama. 1970. Efficient capital markets: A review of theory and empirical work. *The journal of Finance* 25, 2 (1970), 383–417.

[9] Eugene F. Fama. 1991. Efficient Capital Markets: II. *Journal of Finance* 46 (1991), 1575–1617. https://api.semanticscholar.org/CorpusID:264619767

[10] Yi Fang, Luo Si, Naveen Somasundaram, and Zhengtao Yu. 2012. Mining contrastive opinions on political texts using cross-perspective topic model. In *Web Search and Data Mining*. https://api.semanticscholar.org/CorpusID:16954033

[11] George Gallup and Saul Forbes Rae. 1940. The pulse of democracy: the public-opinion poll and how it works. (1940).

[12] Jim Gatheral. 2010. No-dynamic-arbitrage and market impact. *Quantitative finance* 10, 7 (2010), 749–759.

[13] Matthew Gentzkow, Jesse M Shapiro, and Matt Taddy. 2019. Measuring group differences in high-dimensional choices: method and application to congressional speech. *Econometrica* 87, 4 (2019), 1307–1340.

[14] Parush Gera and Tempestt J. Neal. 2025. Deep Learning in Stance Detection: A Survey. *Comput. Surveys* (2025). https://api.semanticscholar.org/CorpusID:279310559

[15] Paul Gomme. 2003. Iowa electronic markets. *Economic commentary* (2003), 1. https://api.semanticscholar.org/CorpusID:154124861

[16] Maarten R. Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *ArXiv* abs/2203.05794 (2022). https://api.semanticscholar.org/CorpusID:247411231

[17] Robin Hanson. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1 (2003), 107–119.

[18] Robin Hanson. 2007. Logarithmic markets coring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets* 1, 1 (2007), 3–15.

[19] Joel Hasbrouck. 1991. Measuring the information content of stock trades. *The Journal of Finance* 46, 1 (1991), 179–207.

[20] Richard Hobeck, Christopher Klinkmüller, Herath Mudiyanselage Nelanga Dilum Bandara, Ingo Weber, and Wil M.P. van der Aalst. 2021. Process Mining on Blockchain Data: A Case Study of Augur. In *International Conference on Business Process Management*. https://api.semanticscholar.org/CorpusID:237393224

[21] John M Hunt. 2015. Betting on the Papal Election in Sixteenth-Century Rome. (2015).

[22] Narasimhan Jegadeesh and Sheridan Titman. 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of finance* 48, 1 (1993), 65–91.

[23] Jaynti Kanani, Sandeep Nailwal, Mihailo Bjelic, and Anurag Arjun. 2021. Polygon Whitepaper. https://polygon.technology/papers/pol-whitepaper. Accessed: YYYY-MM-DD.

[24] Albert S Kyle. 1985. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society* (1985), 1315–1335.

[25] Bhabendu Kumar Mohanta, Soumyashree S Panda, and Debasish Jena. 2018. An overview of smart contract and use cases in blockchain technology. In *2018 9th international conference on computing, communication and networking technologies (ICCCNT)*. IEEE, 1–4.

[26] Burt L Monroe, Michael P Colaresi, and Kevin M Quinn. 2008. Fightin'words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis* 16, 4 (2008), 372–403.

[27] Stefan Neumann, Yinhao Dong, and Pan Peng. 2024. Sublinear-Time Opinion Estimation in the Friedkin–Johnsen Model. *Proceedings of the ACM Web Conference 2024* (2024). https://api.semanticscholar.org/CorpusID:269362796

[28] Abraham Othman, David M Pennock, Daniel M Reeves, and Tuomas Sandholm. 2013. A practical liquidity-sensitive automated market maker. *ACM Transactions on Economics and Computation (TEAC)* 1, 3 (2013), 1–25.

[29] Jack Peterson, Joseph Krug, Micah Zoltu, Austin K Williams, and Stephanie Alexander. 2015. Augur: a decentralized oracle and prediction market platform. *arXiv preprint arXiv:1501.01042* (2015).

[30] Juan Enrique Ramos. 2003. Using TF-IDF to Determine Word Relevance in Document Queries. https://api.semanticscholar.org/CorpusID:14638345

[31] Pavel Řežábek. 2024. Analysis of Prediction Markets in Crypto: Investigating Convergence in Time, Volatility, and Biases in Polymarket. (2024).

[32] Paul W Rhode and Koleman Strumpf. 2008. *Historical Political Futures Markets: An International Perspective.* Technical Report. National Bureau of Economic Research.

[33] Paul W Rhode and Koleman Strumpf. 2013. The long history of political betting markets: An international perspective. (2013).

[34] Paul W Rhode and Koleman S Strumpf. 2004. Historical presidential betting markets. *Journal of Economic Perspectives* 18, 2 (2004), 127–142.

[35] Oriol Saguillo, Vahid Ghafouri, Lucianna Kiffer, and Guillermo Suarez-Tangil. 2025. Unravelling the Probabilistic Forest: Arbitrage in Prediction Markets. https://arxiv.org/abs/2508.03474v1. arXiv:2508.03474v1 [cs.CR],.

[36] Oriol Saguillo, Vahid Ghafouri, Lucianna Kiffer, and Guillermo Suarez-Tangil. 2025. Unravelling the Probabilistic Forest: Arbitrage in Prediction Markets. *arXiv preprint arXiv:2508.03474* (2025).

[37] Gary Stix. 2008. When markets beat the polls. *Scientific American* 298, 3 (2008), 38–45.

[38] Matt Taddy. 2013. Multinomial inverse regression for text analysis. *J. Amer. Statist. Assoc.* 108, 503 (2013), 755–770.

[39] Gnosis Team. 2017. *Gnosis Whitepaper: A Prediction Market Platform Overview.* Technical Report. Gnosis. https://www.allcryptowhitepapers.com/wp-content/uploads/2018/05/Gnosis.pdf

[40] Leighton Vaughan Williams and David Paton. 2015. Forecasting the Outcome of Closed-Door Decisions: Evidence from 500 Years of Betting on Papal Conclaves. *Journal of Forecasting* 34, 5 (2015), 391–404.

[41] Justin Wolfers and Eric Zitzewitz. 2004. Prediction Markets. *Stanford Graduate School of Business Research Paper Series* (2004). https://api.semanticscholar.org/CorpusID:10803317

[42] Justin Wolfers and Eric Zitzewitz. 2004. Prediction markets. *Journal of economic perspectives* 18, 2 (2004), 107–126.

[43] Justin Wolfers and Eric Zitzewitz. 2006. *Interpreting Prediction Market Prices as Probabilities.* Technical Report 12200. National Bureau of Economic Research, Cambridge, MA. https://doi.org/10.3386/w12200

[44] Dirk A Zetzsche, Douglas W Arner, and Ross P Buckley. 2020. Decentralized finance. *Journal of Financial Regulation* 6, 2 (2020), 172–203.

[45] Zibin Zheng, Shaoan Xie, Hong-Ning Dai, Xiangping Chen, and Huaimin Wang. 2018. Blockchain challenges and opportunities: A survey. *International journal of web and grid services* 14, 4 (2018), 352–375.

[46] Minghui Zou, Ronghui Guo, Sai Zhang, Xiaowang Zhang, and Zhiyong Feng. 2025. TAT: Improving Stance Detection on Social Media through Thought Alignment with LLMs. *Companion Proceedings of the ACM on Web Conference 2025* (2025). https://api.semanticscholar.org/CorpusID:278822747

APPENDIX

A DATASET DETAILS

A.1 Data Field

Table 6: Detailed Description of Polymarket Transaction Event Fields: This table lists all the key fields extracted from on-chain Polymarket smart contract events, including their meaning and how they are used to identify and parse user actions. The dataset covers `OrderFilled` events (buy/sell) and `PositionSplit`/`PositionMerge` events (split/merge).

Field	Description
address	The Polymarket wallet address of the user who executed the transaction.
event_id	The ERC-1155 token ID corresponding to the market outcome. Used for distinguishing **buy** and **sell** events.
conditionId	The unique identifier of the underlying condition in the prediction market. Used for distinguishing **split** and **merge** events.
action	The parsed user action: **buy**, **sell**, **split**, **merge**, or **redeem**.
value	The total transaction amount (in USDC).
price	The average execution price of the token in the transaction.
time	The timestamp in Coordinated Universal Time (UTC).

Table 6 provides a detailed description of all key fields extracted from on-chain Polymarket smart contract events. These fields are used to identify and parse user actions, covering both `OrderFilled` events (buy/sell) and `PositionSplit` or `PositionMerge` events (split/merge).

A.2 Data Overview

Our study collected on-chain transaction data from the Polygon blockchain using an Erigon archive node. The dataset covers blocks from 54,000,000 (February 27, 2024) to 64,000,000 (November 7, 2024), comprising 1,166,714 blocks in total. This limited range was selected to allow our analysis methods to run on a standard workstation,

while still encompassing all on-chain transactions related to the 2024 U.S. Presidential Election event on Polymarket. The block range can be freely adjusted to include more data for comprehensive feature analysis or to shorten processing time.

Table 7: Overview of the dataset filtered from the Polymarket event "Presidential Election Winner 2024". Buy/sell statistics are computed via the event_id field, while split/merge statistics are computed via the conditionId field.

Attribute	Value
Block range	54,000,000 to 64,000,000
Number of related blocks	1,166,714
Number of unique users	230,190
Time span	2024-03-01 to 2024-11-06
Total number of transactions	7,946,647
Trump-related total volume (buy/sell)	$1,629,781,310.22
Harris-related total volume (buy/sell)	$886,592,556.08
Trump-related transactions (buy/sell)	5,105,869
Harris-related transactions (buy/sell)	2,811,789
Trump unique users (split/merge)	1,306
Trump split: total volume / count	$55,879,444.05 / 11,887
Trump split: avg per user / per tx	$50,251.30 / $4,700.89
Trump merge: total volume / count	$105,769,320.84 / 7,300
Trump merge: avg per user / per tx	$201,849.85 / $14,488.95
Harris unique users (split/merge)	866
Harris split: total volume / count	$21,630,726.52 / 3,027
Harris split: avg per user / per tx	$30,725.46 / $7,145.93
Harris merge: total volume / count	$71,816,856.54 / 6,775
Harris merge: avg per user / per tx	$223,728.53 / $10,600.27

B WHALES' PNL IN DIFFERENT CATEGORIES

This figure shows the distribution of whales' total and average PnL across categories. The "user count" section on the left shows the distribution of the number of the two types of whales across different categories of Polymarket events. The two plots above and below the points corresponding to each category represent the total PnL and the average PnL of the two types of whales within that category, respectively.

C EVENT CLASSIFICATION STANDARDS

First, we classify each trade's `slug` using a priority-based keyword scheme, matching as many categories as possible (e.g., **US Election Winner 2024** slugs to *Elections*, Trump-related slugs to *Trump*), and use **Gemini-2.5-Pro** to categorize the remaining slugs that cannot be reliably matched by keywords into one of twelve main categories consistent with Polymarket's taxonomy (*Politics, Sports, Crypto, Earnings, Geopolitics, Tech, Culture, World, Economy, Trump, Elections, Mentions*). After classification, we build a user–category PnL matrix, in which each cell records a user's aggregated (total) PnL within that category. This matrix supports distributional analysis and serves as the unified basis for correlation and winning–losing whales comparisons.

D POLYMARKET FRENCH WHALE COMMENTS

This user uses polling deviations, neighbor effects, and historical overperformance data for quantitative analysis to arbitrage in key swing states and national markets.

His comments reflect both information sensitivity and strategic behavior, revealing the trading style and psychology of a typical information-advantage whale.

- **Michie** (2024-10-18 09:14:18): Nationally: In the 2020 election, RCP's average polls 2-3 weeks before showed Biden leading by 8.9 points. The final result was Biden winning by 4.5 points, meaning Trump outperformed the polls by 4.4 points.
- **Michie** (2024-10-18 09:12:52): Trump overperformance: PA 4.4 - MI 4.4 - WI 5.4
- **Michie** (2024-10-18 09:11:52): Actual election results: PA 1.2 - MI 2.8 - WI 0.7
- **Michie** (2024-10-18 09:08:27): Spread of surveys on RealClearPolitics on 10/17/2020: PA 5.6 - MI 7.2 - WI 6.1
- **Michie** (2024-10-18 09:26:43): So Trump should outperform again this year, all pollsters and media know it... and I just make money. I don't know what is better actually: being greedy or a manipulator LOL
- **Michie** (2024-10-18 09:25:14): 1. Redfield & Wilton (September 2024): Own vote: Harris 48%, Trump 46%. Neighbors: Trump 47%, Harris 44% (Neighbor effect: Trump +5).
  2. New York Times/Siena (mid-September 2024, Pennsylvania): Own vote: Harris 49%, Trump 46%. Neighbors: Trump 47%, Harris 44% (Neighbor effect: Trump +6).
  3. Fox News/Politico (late September 2024, Pennsylvania): Own vote: Harris 49%, Trump 46%. Neighbors: Trump 47%, Harris 44% (Neighbor effect: Trump +6).
- **Michie** (2024-10-18 09:19:44): All is around 5% – same thing in 2024 measured by the "neighbor vote effect" both nationally and in PA.
- **Michie** (2024-10-12 12:31:20): The "Shy Trump Voter" Effect: This "neighbor voting" pattern, where Trump consistently scores higher in perceived neighbor support, is a well-documented "shy Trump voter" phenomenon. This effect suggests that Trump supporters may be reluctant to declare their intent directly, particularly given the current polarized environment.
- **Michie** (2024-10-12 12:30:02): Here's a compelling case for why Trump could outperform current polls by as much as 5 points,

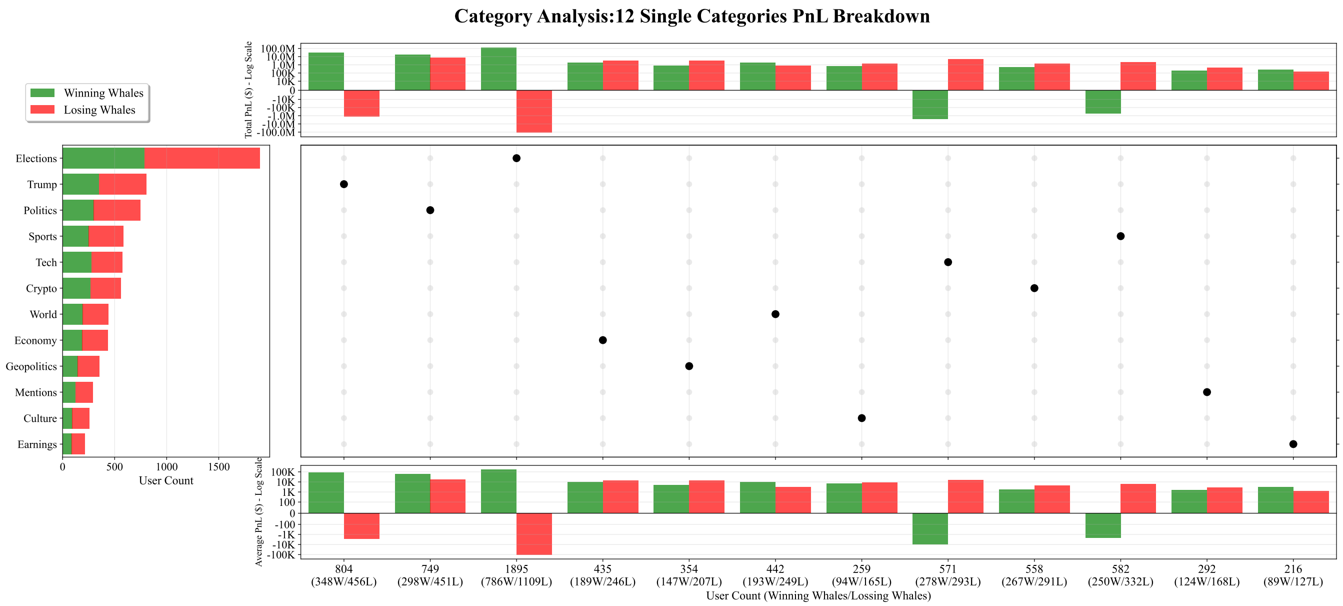

Figure 7: Distribution of whales' total and average PnL across categories

both in Pennsylvania and nationally:

1. Recent Polls Show a Pattern: In Pennsylvania, two recent polls reveal the "neighbor effect":

• NYT/Siena College (mid-September 2024): Respondents favor Harris 49% to Trump 46% for their own vote, but when asked about their neighbors' voting intentions, Trump leads 47% to Harris's 44%.

• Fox News/Politico (Sept. 20-24, 2024): Similar results, with Harris leading 49%-46% personally, but Trump leads 47%-44% among "neighbors."

Nationally, this trend also appears:

• Redfield & Wilton (September 2024): Respondents personally favor Harris 48%-46%, yet think their neighbors lean 47% for Trump vs. 44% for Harris.