

Betting on Informed Trades in Prediction Markets

Maylin Chen, Andrea Li, Elizabeth Ma

I. Introduction

Prediction markets have unrivaled potential for quantitative strategies, with tremendous volume of over \$63.5 billion dollars annually.¹ Some estimates show that annual volume will grow to over \$1 trillion dollars by 2030.² Moreover, according to a 2026 study by researchers at Columbia Law School and the University of Haifa, there was over \$143 million dollars in anomalous profit since 2024 on Polymarket alone.³ For instance, regarding the Iran ceasefire, at least 50 newly created Polymarket accounts spent and won thousands betting that a ceasefire would occur, just minutes before Trump announced one.⁴ This pattern is reflected in the profit distribution on Polymarket. According to a 2026 Wall Street Journal investigation, just 0.1% of Polymarket accounts hold 67% of the profits, compared to the vast majority of accounts which are unprofitable.⁵

Thus, we hypothesize that most prediction market users base their decisions on recently-announced public information, which will have usually been reflected in the market by the time of action. This delay provides an opportunity to profit. If we can identify these anomalous bets and match those bets proportionally, we can capitalize on the inside information these bettors have and buy before the markets catch up, letting us profit on the price difference.

There are two main prediction markets: Kalshi and Polymarket. Kalshi is US-based and significantly more regulated, as it requires ID verification for every account. Kalshi has also been known to track down and ban users for insider trading. For instance, it suspended several Congressional candidates for trading in their own election markets, even though some of the bets were less than \$100 dollars.⁶ In contrast, Polymarket doesn't inherently link your account to your ID, and moves money using crypto, which is a lot less traceable. The liquidity in Polymarket is also much deeper overall and prices change less with very large bets, which is an important property to allow us to enter positions at favorable prices. However, Polymarket's lack of regulation has gotten it banned in over 20 countries, including the US.⁷ There is theoretically US access to Polymarket, but it's locked behind a waitlisted app account that's linked to your ID and no longer uses crypto.⁸ It is also illegal to trade on Polymarket using a VPN.

Despite the legal ambiguity, we have chosen to focus on Polymarket, as its looser regulatory environment remains in effect in several jurisdictions, most relevantly Canada, where our algorithm can be legally deployed.

¹ Ryan Butler. "Prediction Market Volume Quadrupled in Past 2 Years, Report Finds." *Yahoo Finance*, 13 Mar. 2026, finance.yahoo.com/news/prediction-market-volume-quadrupled-past-190400921.html.

² Giangulio, Davis. "Prediction Markets Will Grow to \$1 Trillion by 2030, Bernstein Estimates." *CNBC*, 14 Apr. 2026, www.cnbc.com/2026/04/14/prediction-markets-will-grow-to-1-trillion-by-2030-bernstein-says.html.

³ Jack Newsham. "'Informed' traders on Polymarket netted \$143 million in 'anomalous' profit since 2024, researchers find." *Business Insider*, 28 Mar. 2026, www.businessinsider.com/polymarket-informed-insider-trades-study-2026-3.

⁴ Smith, David. "Newly Created Polymarket Accounts Win Big on Well-timed Iran Ceasefire Bets." *The Guardian*, 24 Apr. 2026, www.theguardian.com/business/2026/apr/08/polymarket-trump-us-iran-ceasefire.

⁵ Mehta, Neil, et al. "Why Almost Everyone Loses—Except a Few Sharks—on Prediction Markets." *The Wall Street Journal*, 3 May 2026, www.wsj.com/finance/investing/polymarket-kalshi-betting-profits-prediction-markets-eb23ac11.

⁶ "Kalshi Fines and Suspends 3 Congressional Candidates for Betting on Their Own Elections." *PBS News*, 23 Apr. 2026, www.pbs.org/newshour/politics/kalshi-fines-and-suspends-3-congressional-candidates-for-betting-on-their-own-elections.

⁷ "Geographic Restrictions - Polymarket Documentation." *Polymarket Documentation*, docs.polymarket.com/api-reference/geoblock.

⁸ Mayowa, Sofoluwe. "Is Polymarket Legal in the U.S. in 2026?" *Gambling Insider*, 29 Jan. 2026, www.gamblinginsider.com/in-depth/106291/is-polymarket-legal-in-the-us.

II. Strategy

Our strategy focuses on event-based contracts and options markets on Polymarket that are sensitive to high-stakes government actions, such as regulatory shifts, legislative votes, federal contract awards. These environments seem to be particularly prone to informed trading spikes ahead of public announcements. Thus, our strategy is to identify anomalous bids in political events that meet our predefined thresholds and then replicate them by matching both the quantity and price of the original trades.

To assess the feasibility of implementing this approach in practice, we make two key assumptions when evaluating our trading strategy. First, we assume that we are able to buy contracts at the same price as the original trader placing the large long-shot bet. This is reasonable because most of these markets have deep enough liquidity to minimize price changes from these large bets. Second, we assume that Polymarket is able to fill our orders. This is also reasonable given that political events are widely followed and heavily traded, which should be able to support large volume orders.

A. Data Filter Selection Analysis

To obtain the dataset, we used the Polymarket Gamma API to request the max limit of 10,000 events from 2023 to 2026. We looked for events with political keywords and labels such as “geopolitics” and “elections.” For each political event, we retrieved bid histories to identify trades meeting our selection criteria. When a qualifying bid was found, we stored key attributes in a JSON file, including event ID, liquidity pool, and resolution status.

To further narrow down our results, we explored three filters: market-implied probability, fixed floor, and relative to order book depth. The market-implied probability filter enforces a maximum threshold of 20%, concentrating on longshot outcomes where informed trading is more likely to stand out and maximize returns. The fixed floor filter checks for a minimum bid amount of \$1,000 that filters out retail noise since most individual trades are well under \$500. This threshold is held constant across all markets. The relative to order book depth requires the bet to be some meaningful fraction of the total liquidity (eg. $\geq 2\%$), helping isolate true anomalies rather than simply large absolute bets. However, this filter required historical liquidity snapshots at the time of execution, which we were unable to obtain, so it was not implemented. After applying the maximum implied probability and fixed floor filters, we identified 5,102 unusual bets. Excluding the 0–1% implied probability bucket reduced the sample to 4,691 contracts within the signal filter.

We excluded the 0-1% range for three reasons. First, there could be price staleness: an 82% win rate in this bucket for events such as “Will Trump be inaugurated,” suggests these outcomes were not driven by informed trading. Instead, these cases probably reflect markets failing to immediately update already-obvious outcomes, allowing opportunistic traders to exploit mispriced probabilities. Second, extreme illiquidity at sub-1% prices means that even moderately sized orders can move prices substantially, making such trades difficult to replicate in practice. Third, this segment behaves more like a retail “lottery ticket” zone, where participants

place highly speculative bets on extreme longshots rather than expressing informational advantage.




Figure 1. Time Before Market Close (0-5% Implied Probability, 4 Weeks to Event Resolution). We further tested for price staleness by analyzing the timing of trades relative to event resolution and measuring corresponding win rates. If staleness were the primary driver, we would expect the strongest edge immediately after market open, when prices are least likely to have fully incorporated available information. However, we do not observe such a pattern, nor any meaningful early-window edge. Across the first four weeks after market open, performance remains essentially flat, with no evidence of systematic staleness.

However, stronger edges appear closer to resolution. The 0-1% bucket exhibits a distinct pattern closer to event resolution, with an anomalously high edge of +72.7% at 1-4 weeks before close, suggesting prices stay stuck near zero on long-dated markets even as resolution approaches. In contrast, the 2-5% range shows a gradual increase in edge as the event approaches closure, potentially reflecting late-stage bandwagoning or delayed reaction to informed trading signals. Given its distinct behavior and unclear cause of behavior, we exclude the 0-1% category from further analysis. We instead focus on the 1-5% range, where the edge is more modest, consistent across timing windows, and plausibly actionable.




Figure 2. Event-Level Unusual Bet Volume Share by Probability Silo (1-5% Implied Probability). The relative size of the dots represents the relative value of the bet. Overall, anomalous bets constituted just 0.02%, 0.04%, and 0.03% of total volume share in the 0-1%, 1-3%, and 3-5% implied probability silos respectively. Aside from a few outliers, this is reflected in the individual event level. These statistics are unable to confirm price staleness because these anomalous bets still constitute a rather consistently small percentage of the total bet volume in their probability silos. One flaw to this analysis is that we're unable to recover the total liquidity of the event pool at the time of the anomalous bet. Instead, we are using the final liquidity pool at the time of event resolution. It may still be the case that these bets constitute an abnormally large percentage of the event pool at the time they were placed.




Figure 3. Unusual Bet Volume as Share of Event Liquidity/Volume by Category. This figure shows that our assumptions are reasonable, with the average event liquidity in the political sector in the hundreds of millions of dollars. For example, the elections category has liquidity pools above \$2,500 million which is more than sufficient to fill our long-shot orders. Additionally, unusual bet volume is extremely small, less than 1% of total liquidity across all categories, so markets are unlikely to experience significant price changes after a trader places a large long-shot bet. This allows us to replicate their trade at similar prices and expected value.

Table 1. Threshold Scan (Trade-Level).

min_usd \ max_prob	1-5%	1-10%	1-15%	1-20%
1,000	0.1740	0.1245	0.1122	0.0988
2,500	0.1820	0.1240	0.0848	0.0648
5,000	0.1996	0.0887	0.0398	0.0192
10,000	NaN	-0.0059	-0.0549	-0.0542

This table identifies where the signal is strongest as the filtering criteria are varied. We conducted this analysis to evaluate whether our chosen filters are optimal. Each cell reports the edge (realized - implied probability) for a given combination of max implied probability and minimum bet amount. We ultimately select a **minimum bet amount of \$1,000**, as it produces the largest edge across the 1–10%, 1–15%, and 1–20% probability ranges while maintaining a sufficiently large sample size. This threshold also allows for direct comparison with the broader strategy discussed later, which uses the 1–20% filter. Using the same analysis, we choose a **10% maximum implied probability threshold**. Although the 1–5% range generates a larger edge, it contains substantially fewer trades and would make the strategy overly dependent on a small number of events. The 1–10% range provides a favorable balance between edge magnitude and trading frequency, resulting in a more robust and scalable signal.

Table 2. Bucket Statistics (1–20%, USD ≥ \$1,000, Trade-Level).

Bucket	n	Mean price	Realised rate	95% CI	Edge	t-stat	p-value	Sig.
1–3%	619	0.0204	0.1922	[0.1631, 0.2252]	0.1719	30.261	0.0	True
3–5%	400	0.0402	0.2175	[0.1799, 0.2605]	0.1773	18.059	0.0	True
5–10%	958	0.0754	0.1472	[0.1262, 0.1710]	0.0717	8.407	0.0	True
10–15%	1,312	0.1249	0.2188	[0.1972, 0.2419]	0.0938	10.281	0.0	True
15–20%	1,402	0.1746	0.2418	[0.2201, 0.2649]	0.0672	6.623	0.0	True

We broke down our results by implied probability bucket and tested whether large trades within each range resolve successfully more often than market prices would suggest. The variable n denotes the number of observations in each bucket. Mean price represents the average implied probability of the trades, while realized rate is the observed outcome frequency. For example, if 18 out of 100 events resolve positively, the realized rate is 18%. Edge is defined as *realized rate - the average implied probability*. The t-statistic measures the strength of this edge relative to the underlying variance, while the p-value measures the probability of observing an edge at least this large under the null hypothesis of no true edge. Throughout this analysis, we define statistical significance as $p < 0.05$.

This results shows that every implied probability bucket from 1-20% is significant, suggesting the presence of either systematic mispricing or informed trading. The largest edges occur in the 1–3% and 3–5% buckets, both near 17 percentage points. We also include the 5–10% range in our strategy despite its smaller edge, as it substantially increases the number of available trading opportunities. Interestingly, many of the largest events in the 1–10% range are Fed rate decisions, while the 10–20% range is dominated by the Polish election. This suggests that some of the observed edges may be driven by event- or category-specific effects.

B. Profitability Summary

Our goal was to generate a positive return after accounting for Polymarket's takers fees. Fees are calculated using the following formula: $fee = C \times feeRate \times p \times (1 - p)$, except for geopolitical markets, which are fee free.⁹ For the political markets examined in this study, the fee rate is 0.04. As discussed earlier, we selected a 10% maximum implied probability threshold and a \$1,000 minimum bet size because this combination produced a strong edge while retaining a sufficiently large sample. Applying these constraints reduced the dataset to 1,977 unusual bets. The resulting win rate was 17.6%, which is expected given that the strategy targets low-probability outcomes. However, successful trades generate large payouts that more than offset the high loss rate. By replicating these long-shot trades, we invested a total of \$4.7 million and achieved a 366.1% return on investment (ROI). The strategy produced a Sharpe ratio of 0.3, which is also expected given the highly volatile nature of the returns.

Breaking profitability results into max implied probability buckets of 1–3%, 3–5%, and 5–10% reveals a clear pattern. Win rates are relatively similar across buckets (19.2%, 21.8%, and 14.7%), but ROI declines sharply as implied probability increases (824.9%, 480.7%, and 83.6%), driven by their outsized payouts when they do resolve successfully. Additionally, this nearly tenfold spread in returns highlights a strong asymmetry between implied probabilities.

C. Market Breakdowns

Table 3. Political Sector Analysis (1–10%, USD ≥ \$1,000, Event-Level).

Category	Events	Implied	Realised	Edge	p-value	Sig
Elections	97	6.0%	20.6%	+14.6%	0.0000	✓
Executive / Legislative	53	6.2%	9.4%	+3.3%	0.1595	
Fed / Monetary Policy	27	4.4%	40.7%	+36.3%	0.0000	✓
Geopolitical	39	7.1%	5.1%	-1.9%	0.6812	
Other	35	6.5%	5.7%	-0.8%	0.5777	
Regulatory / Legal	22	6.2%	63.6%	+57.4%	0.0000	✓

We examined which categories generated the largest edge within the 1-10% implied probability range. Fed/Monetary Policy, Regulatory/Legal, and Elections emerge as the only statistically significant categories, with Fed/Monetary Policy and Regulatory/Legal both exhibiting edges above 35%. The key question is whether these results are driven by informed trading or by market mispricing, which is important for assessing the true risk and longevity of the strategy. To investigate this, we examine several of the highest-performing events within each category.

⁹ "Fees - Polymarket Documentation." *Polymarket Documentation*, docs.polymarket.com/trading/fees.

Table 4. Top Events for the Fed/Monetary Policy Category (1-10% Implied Probability)

Rank	Event Name	Trades	Outcome
1	No change in Fed interest rates after 2024 September meeting	29	YES
2	Fed decreases interest rates by 50+ bps after May 2025 meeting	16	NO

For the Fed/Monetary Policy category, roughly half the events ultimately resolved positively, consistent with the category's 48.5% realized probability. Additionally, the category contains 27 distinct events, suggesting that the observed edge is consistent across many events rather than concentrated in a few cases. It seems unlikely that the observed edge is driven by informed trading from one of the 12 FOMC members, suggesting that much of the edge may instead reflect market mispricing. Unlike other categories, Federal Reserve decisions occur within a structured and highly scrutinized process with a limited set of possible outcomes. As a result, traders may be better able to identify situations where market probabilities understate the true likelihood of an outcome.

Table 5. Top Events for the Regulatory/Legal Category (1-10% Implied Probability)

Rank	Event Name	Trades	Outcome
1	TikTok banned in the US before May 2025?	62	NO
2	Will Biden pardon SBF? (Sam Bankman-Fried)	35	YES

The Regulatory/Legal category is dominated by two large events, after which the next-largest event contains only six trades. At the event level, this category exhibits a realized probability of 63.5%. Removing the two largest events reduces the realized probability only slightly, to 62.0%, indicating that the results are not driven solely by the two outliers. This suggests that the elevated edge is broadly distributed across the category.

From a common-sense perspective, Regulatory/Legal events are prone to information-based trading. Individuals who are within or adjacent to government institutions may have access to information about regulatory decisions before it becomes public. While this analysis cannot directly distinguish informed trading from market mispricing, such information advantages could help explain why the realized probability in this category is unusually high, even relative to categories such as Fed/Monetary Policy.

Table 6. Top Events for the Elections Category (1-10% Implied Probability)

Rank	Event Name	Trades	Outcome
1	Will Biden resign before the election?	29	NO
2	Will Rafał Trzaskowski be the next President of Poland?	24	YES

The Elections category contains 459 qualifying trades, while the largest individual event accounts for only 29 trades. This suggests that the observed edge is broadly distributed across many events rather than being driven by a handful of outliers. Similar to the Regulatory/Legal category, elections are a plausible setting for information-based trading, as individuals close to campaigns, governments, or political figures may become aware of important developments before they are publicly announced. For example, someone close to President Biden would likely know about a decision to resign before it became public knowledge. Additionally, because the Elections category includes both U.S. and international markets, we compare these groups separately to assess whether the observed edge is concentrated in regions where information advantages may be more pronounced.

Table 7. Top Events for the Elections Category (1-10% Implied Probability, Event-Level)

Country	Events	Win Rate	Implied	Edge	T-stat	Sig
Australia	3	66.7%	5.4%	+61.3%	4.714	✓
Brazil	1	0.0%	2.9%	-2.9%	-0.173	
Canada	2	0.0%	6.0%	-6.0%	-0.357	
France	1	100%	6.9%	+93.1%	3.683	✓
Germany	3	0.0%	4.2%	-4.2%	-0.361	
Other	6	0.0%	6.6%	-6.6%	-0.653	
Poland	2	50.0%	5.1%	+44.9%	2.896	✓
Romania	5	20.0%	7.1%	+12.9%	1.128	
Syria	1	0.0%	3.7%	-3.7%	-0.195	
UK	4	25.0%	6.4%	+18.6%	1.525	
US	69	18.8%	6.1%	+12.8%	4.434	✓

We broke the Elections category down by geography to determine whether certain countries exhibit stronger predictive signals or greater market mispricing. The table shows that Australia, France, Poland, and the United States all have statistically significant edges. However, with the exception of the United States, the number of events in each country is too small to support meaningful country-level analysis.

To address this limitation, we grouped events into U.S. and Foreign categories. In the U.S., there were 69 events with a realized win rate of 18.8% and an edge of +12.8%. In foreign markets, there were 28 events with a realized win rate of 21.4% and an edge of +15.5%. These results suggest that the signal is not unique to U.S. elections and persists across international markets as well. While the slightly higher edge in foreign markets may indicate greater mispricing or larger information asymmetries, the difference is relatively small.

Table 8. Profitability Summary for Edge Categories per Event (1–10%, Event-Level).

Category	Events	Win%	Invested	P&L	ROI	Sharpe
Elections	97	20.6%	\$1.31 M	\$2.82 M	216.0%	0.18
Fed / Monetary Policy	27	40.7%	\$0.903 M	\$13.4 M	1487.2%	0.71
Regulatory / Legal	22	63.6%	\$0.539 M	\$2.22 M	411.9%	0.51
TOTAL	146	26.3%	\$2.75 M	\$18.5 M	671.9%	0.42

If we constructed a strategy using only the three categories that exhibited a statistically significant edge under the 1–10% implied probability filter and \$1,000 minimum bet threshold, the strategy would have generated a 671.9% ROI. Breaking the results down into the 1–3%, 3–5%, and 5–10% implied probability buckets reveals that all three achieved win rates of roughly 30%, despite producing substantially different ROIs. This difference is driven by the larger payouts associated with lower-probability contracts.

III. Benchmarks

We evaluated our strategy against two comparable strategies:

- **“Majority”**: For the events with the anomalous trades, bet an equal amount on the option with the majority implied probability at the time. This majority option will always be the opposite choice. We assume the opposing bid price equals one minus the original bid price.
- **“Sports/Crypto”**: Sports and crypto are the 2 largest non-political markets from our dataset. We implemented our exact same strategy and data filtering approach on the sports and crypto markets to see if insider information edge is unique to politics. We also run the “Majority” strategy on these markets for further comparison.

Table 9. Backtested Performance of “Majority” Strategy.

Performance Metrics		Return Metrics	
Number of Bets	5,012	Total Returns	\$12, 646, 284.59
Total Invested	\$12, 285, 301.85	Total Profit	\$360, 982.64
Average Win Rate	91.73%	Overall ROI	2.94%
Sharpe Ratio (per-bet returns)	0.062	Average ROI per Bet	1.99%
		Median ROI per Bet	9.89%

The “Majority” strategy has a very low ROI of 2.94%, and that’s before taker fees, which would make the returns negligible or cause us to incur a small loss.

Table 10. Comparison Against Our Original Strategy.

<u>Strategy</u>	<u>Win Rate</u>	<u>ROI</u>	<u>Amount Invested</u>
“Majority” Strategy	91.73%	2.94%	\$12.4M
Our Strategy (1–20%)	20.7%	157.6%	\$11.9M
Our Strategy (1–10%)	17.6%	366.1%	\$4.7M

Despite its significantly higher win rate, the “Majority Strategy” significantly underperforms ours. This is to be expected: majority implied probability options are priced higher. When we win bets, returns are modest because the value of our bids don’t increase significantly. However, if we lose, we lose a relatively large amount of our investment. While we would expect the win rates of the two strategies to sum up to 100% since the strategies take opposite sides of the same contracts, the “Majority” strategy was tested on all anomalous bids, while our strategy excluded anomalous bids with implied probabilities of 0-1%, causing a discrepancy.

Table 11. Performance of Original Strategy in Sports Markets.

	Our strategy	Generic (majority)
Trades	1,559	1,559
Win rate	9.3%	90.7%
Total return	-93.9%	+12.5%
Profit factor	0.67×	1.43×
Sharpe	-0.10	0.12
Max drawdown	-100.7%	-2.6%
Final value	\$6,084	\$112,521

Using the same data filters, we implemented both our original strategy and the generic “Majority” strategy. We lost basically all of our initial investment using our original strategy, and only achieved a return of 12.5% using the generic strategy. We hypothesize this is because it’s harder to attain insider information on sports just because much of the information is publicly available on the media. Furthermore, insider information is less likely to yield concrete results since there can be multiple factors at play.

Table 12. Performance of our strategy in Crypto markets.

	Our strategy	Generic (majority)
Trades	1,284	1,284
Win rate	13.9%	86.1%
Total return	+44.5%	-5.5%
Profit factor	1.20×	0.85×
Sharpe	0.05	-0.05
Max drawdown	-46.2%	-11.9%
Final value	\$144,517	\$94,459

For the crypto market, we actually made a 44.5% return using our strategy. We hypothesize that we got these results because there is more insider information in the crypto market especially since these events do overlap with the political market, such as when Trump dubbed himself a “crypto-president” and supported the bitcoin reserve.¹⁰ The loss on investments for the generic strategy is likely due to the inherent volatility in cryptocurrencies.

¹⁰ Di Salvo, Mat. “Donald Trump and Bitcoin: From ‘Not a Fan’ to Crypto President.” *Decrypt*, 5 Oct. 2025, decrypt.co/240580/donald-trump-bitcoin-evolution-timeline-president.

IV. Risks

A. Crowding

Several websites and communities, including Polywhaler and Reddit, have begun tracking and aggregating unusual Polymarket trades. As a result, other traders can potentially identify and replicate the same signals, leading to crowding and reduced profitability. To evaluate whether this effect is present, we performed a year-by-year analysis from 2023–2026.

Table 13. Per-Year Breakdown Analysis (1–10%, USD $\geq$ \$1,000, Event-Level)

Year	Events	Trades	Win Rate	Implied	Edge	Sig
2023	4	11	0.0%	4.5%	-4.5%	
2024	183	1069	18.0%	6.2%	+11.8%	✓
2025	92	881	25.0%	6.0%	+19.0%	✓
2026	2	16	50%	7.5%	+42.5%	

The table shows that the majority of qualifying trades occurred in 2024 and 2025, which are also the only years with a statistically significant edge. To check whether our criteria were too restrictive, we analyzed the number of long-shot trades in the 1–20% implied probability range, which also varied considerably across years: 22 trades in 2023, 2,804 trades in 2024, 2,173 trades in 2025, and 103 trades in 2026. The low count in 2026 is partially explained by the dataset ending on February 17, leaving only about six weeks of resolved events. However, the number of qualifying trades remains notably lower than in previous years at the same point in time. The only two events that met our 1–10% implied probability criteria in 2026 were “Maduro out by January 31, 2026?” and “Will Trump deport 250,000–500,000 people?”.

There are too few trades in 2023 and 2026 to draw meaningful conclusions. In 2024, the strategy generated an edge of 11.8% across the largest number of trades. In 2025, the number of trades fell by 17.6%, while the number of events declined by roughly 50%. This pattern may suggest that activity became increasingly concentrated in a smaller set of high-profile events, consistent with a market containing a few informed traders alongside a larger number of bandwagoners. Such behavior could be indicative of crowding, as traders increasingly focus on the same opportunities. This interpretation is also consistent with the higher 2025 edge of 19.0%, as more participants may have been concentrating their trades on a limited number of events with strong signals.

B. False Positives

A single wealthy trader or meme-driven narrative can produce a volume spike without possessing any informational edge. In February 2026, for instance, a single trader bet nearly \$100,000 that the Trump administration will confirm the existence of alien life by the end of this

year. Another wager appeared for nearly double that amount.¹¹ With each passing month, it becomes less likely that these bettors were acting based on inside information. To mitigate these false signals, a modified strategy could be to wait for multiple anomalous bets to be made before considering the activity as a signal. This threshold will allow us to detect coordinated or widespread abnormal behavior, like the 50+ accounts betting on the Iran ceasefire, while filtering out isolated cases like the alien believers.

C. Increasing Regulation

Some anomalous bets could reveal information regarding sensitive US military operations, such as the US military intervention in Venezuela, which has raised growing national security concerns.¹² In response, Congress has been increasing scrutiny into prediction markets, pressuring markets to tighten regulations or fear being banned. US states have also started to individually ban Kalshi and Polymarket, such as Minnesota on 5/19/26, over concerns about the increased accessibility to gambling.¹³ Internationally, the trend continues as well, with Spain being the most recent country to ban Polymarket on 5/26/26.¹⁴ While these developments introduce increasing regulatory constraints, the strategy can be adapted to follow legal guidelines and implemented on Kalshi or other US-based markets.

¹¹ Landymore, Frank. "Anonymous Kalshi User Makes Huge Bet That the White House Is About to Confirm Alien Life or Technology." *Yahoo Finance*, 28 Feb. 2026, finance.yahoo.com/news/anonymous-kalshi-user-makes-huge-131500695.html.

¹² "A \$400,000 Payout After Maduro's Capture Put Prediction Markets in the Spotlight. Here's How They Work." *PBS News*, 12 Jan. 2026, www.pbs.org/newshour/nation/a-400000-payout-after-maduros-capture-put-prediction-markets-in-the-spotlight-heres-how-they-work.

¹³ Allyn, Bobby. "Minnesota Becomes First State to Ban Prediction Markets." *NPR*, 19 May 2026, www.npr.org/2026/05/19/nx-s1-5821265/minnesota-ban-prediction-markets.

¹⁴ Kassam, Ashifa. "Spain Blocks Access to Polymarket and Kalshi as It Launches Gambling Licence Investigation." *The Guardian*, 26 May 2026, www.theguardian.com/world/2026/may/26/spain-blocks-access-polymarket-kalshi-gambling-licence-investigation.

V. Future Work

Equal-weighting

In Section II, our strategy assumes that anomalous bets should be replicated proportionally, based on the idea that larger bets may reflect stronger conviction or an informational advantage. To evaluate whether this sizing assumption is justified, we compared three portfolio construction methods: equal-weighting, square-root weighting, and full dollar-weighting. Under equal-weighting, the same amount of capital is allocated to every qualifying unusual bet regardless of the original position size. Square-root weighting represents an intermediate approach in which larger bets receive more capital, but less than proportionally to their size. Full dollar-weighting directly matches the size of the original anomalous bet, which is the strategy implemented in Section II.

Initial experiments suggested that equal-weighting outperformed both square-root weighting and full dollar-weighting. Under the event-level settlement procedure used in our preliminary analysis, equal-weighting produced a 208.5% ROI on the full 1–20% signal set and a 423.2% ROI on the stricter 1–10% signal set compared to the lower returns on square-root weighting (407.7 %) and full dollar-weighting (366.1%), which suggests that larger anomalous bets did not contain substantially more predictive information than smaller ones.




Figure 4. Distribution of Unusual Bet Sizes. Most unusual bets are clustered near the \$1,000 minimum threshold, while a relatively small number of observations form a long upper tail of much larger bets. This skewed distribution suggests that a small number of very large bets account for a disproportionate share of total volume.




Figure 5. Bet Size and Predictive Performance. Examines how predictive performance varies with bet size under the event-level signal definition.

Figure 5 shows the almost inverse relationship between bet size and performance. Win rates decline from roughly 23% in the smallest bet-size quintile to 17% in the largest quintile, while equal-weight ROI falls from 280% to 168%. A logistic regression controlling for price, squared price, and event volume finds a statistically significant negative relationship between bet size and success probability (odds ratio = 0.80, $p = 0.002$). Together, these results contradict our original assumption that larger anomalous bets contain more predictive information than smaller ones. This could be due to differences in trader wealth, where what constitutes a “large” bet varies across individuals.

Wallet and Category Filtering

Rather than treating all political markets equally, we can expand upon our previous analysis of performance at category levels to focus our strategy on markets that consistently exhibit anomalous activity. Furthermore, wallet-level analysis may reveal traders who repeatedly place profitable abnormal bets across unrelated markets. Thus, future versions of the strategy should prioritize filtering for historically strong categories, events, and wallets, while using an equal-weighted portfolio.