

Political Bias in Decentralized Prediction Markets: Evidence from Trump-Related Contracts on Polymarket

Simone Skovgaard* Noah Wenneberg Junge[†]
Bulat Ibragimov[‡]

July 2026

Abstract

This paper investigates whether political bias influences pricing dynamics in decentralized online prediction markets. Using Polymarket, a blockchain-based prediction market platform, we simulate a rule-based Dollar-Cost Averaging strategy that systematically bets on anti-Trump outcomes across 218 Trump-related binary markets over a 287-day observation period from July 2025 to May 2026. The anti-Trump strategy is benchmarked against 10,000 Monte Carlo simulations of neutral directional trading, holding market selection, trade timing, and portfolio construction constant. The anti-Trump strategy ranks at the 100th percentile of the benchmark distribution, generating a cumulative return of +\$1,006 over the full 287-day period against −\$1,028 for the mirrored pro-Trump strategy. An OLS trend regression yields a slope of +\$14.07 per day ($R^2 = 0.662$, $p < 0.001$). The effect is not explained by the favourite–longshot bias and is distributed broadly across market topics rather than driven by any single theme.

*University of Copenhagen, Department of Computer Science. Email: wnf255@alumni.ku.dk

[†]University of Copenhagen, Department of Computer Science. Email: qxk266@alumni.ku.dk

[‡]Corresponding author. Associate Professor, Department of Computer Science, University of Copenhagen. Email: bulat@di.ku.dk

The findings are consistent with systematic overpricing of pro-Trump outcomes, likely driven by politically motivated retail participants on a cryptocurrency-native platform.

JEL Classification: G14, G41, D72, D84, G17

Keywords: prediction markets, political bias, market efficiency, behavioural finance, Polymarket, cryptocurrency

Contents

1	Introduction	5
1.1	Motivation	5
1.2	Hypotheses and Research Question	7
1.3	Methodology Overview	8
1.4	Contributions	8
1.5	Report Structure	9
2	Literature Review	9
2.1	Prediction Markets	9
2.2	Online and Decentralized Prediction Markets	10
2.3	Market Efficiency and Information Aggregation	11
2.4	Political Polarization and Ideological Bias	12
2.5	Theoretical Motivation	14
3	Method	15
3.1	Research Design	15
3.2	Market Identification and Selection	16
3.3	Data Collection	17
3.4	Simulation Architecture	18
3.5	Simulation Framework	21
3.6	Limitations of the Simulation Framework	22
3.7	Return and Performance Measurement	25
3.8	Statistical Analysis	27
4	Results	29
4.1	Dataset Description	29
4.2	Profit and Loss Analysis	30
4.3	Dual Equity Curve Analysis	33
4.4	Market Size Analysis	35
4.5	Statistical Analysis	36
4.6	S&P 500 Correlation	40
5	Discussion	42
5.1	Why Mean-Return Tests Were Weak	42
5.2	Relative Mispricing Rather than Absolute Profitability	43
5.3	Alpha, Beta, and the Source of Returns	44

5.4	Interpretation of Unrealized Returns	45
5.5	A Question of Who Was Trading	46
5.6	What Prediction Markets Aggregate	47
5.7	Limits of Informational Efficiency in Polarized Environments .	48
5.8	Distinguishing Political Mispricing from Favourite-Longshot Bias	50
5.9	Robustness to Topic Concentration	51
6	Conclusion	52

1 Introduction

Prediction markets have been considered powerful mechanisms for aggregating dispersed information about future events. By allowing participants to trade contracts where the payoff depends on the realization of specific outcomes, these markets transform beliefs into prices. The prices are commonly interpreted as probabilistic forecasts. In theory, competitive trading incentivizes participants to incorporate all available information into prices, making prediction markets efficient forecasting tools.

In recent years, decentralized blockchain-based prediction markets have emerged as an extension of this idea. One of the most prominent platforms is Polymarket, a cryptocurrency-based prediction market launched in 2020 [22]. Polymarket operates on a blockchain infrastructure and allows users to trade on the outcomes of real-world events, including elections, geopolitical developments, economic indicators and public policy decisions. Trading on Polymarket occurs in USDC, which is a cryptocurrency stable coin pegged to the US dollar [4], approximately 1:1 equivalent with the U.S. dollar. For readability, monetary values are expressed in USD throughout this paper unless otherwise specified. Market participants buy and sell shares corresponding to a specific outcome of an event, where prices range between 0 and 1 USD and reflect the market-implied probability of that outcome occurring. Unlike traditional regulated prediction markets, Polymarket operates in a decentralized environment and facilitates trading through stablecoins on blockchain networks.

Polymarket enables open participation and real-time trading on politically sensitive events [22]. Due to its cryptocurrency-based infrastructure and global accessibility, Polymarket provides a unique setting in which retail investors, politically motivated participants and speculative traders interact within the same marketplace.

1.1 Motivation

As mentioned previously, prediction markets are considered efficient mechanisms for aggregating dispersed information, meaning that prices should theoretically reflect the collective expectation regarding a future outcome, therefore making it difficult to achieve persistent abnormal returns. However,

increasing political polarization and the emergence of decentralized betting platforms raises questions regarding whether politically sensitive markets remain informationally efficient in polarized climates.

This paper is motivated by the possibility that market participants on cryptocurrency-based platforms may not be politically neutral. Polymarket operates as a cryptocurrency-based prediction market. Users are required to engage with blockchain infrastructure, cryptocurrency wallets and decentralized finance tools. We hypothesize that cryptocurrency ownership may have become increasingly associated with support for Donald Trump during the 2024 U.S. presidential election cycle. A poll conducted by Fairleigh Dickinson University found that cryptocurrency owners were disproportionately likely to support Trump, estimating that approximately one in seven U.S. voters owned cryptocurrency [36]. If this association holds among active Polymarket traders, crypto-based platforms may attract a politically distinct subgroup of participants, though we cannot verify this directly from our data.

The 2024 election also revealed unusually concentrated trading behavior on Polymarket. The presidential election market closed with a value of nearly \$3.7 billion US dollars in trading volume. Historically, this was the largest prediction market observed in terms of volume [21]. During the campaign, Trump contracts frequently traded at higher implied probabilities than many polling averages. Reports also revealed that a small number of accounts placed approximately \$30 million in pro-Trump bets, increasing Trump’s implied probability to almost 60% [12]. Subsequent investigations showed that several of these accounts were controlled by the same trader, demonstrating how concentrated participation may influence prices in these markets.

The association between prediction market platforms and pro-Trump networks was strengthened in 2025 when Donald Trump Jr. joined the advisory boards of both Polymarket [37] and Kalshi. At the same time, his investment firm, 1789 Capital, invested in Polymarket [27]. These affiliations do not imply direct market manipulation, but they may attract politically aligned traders and increase the likelihood of sentiment-driven pricing distortions.

These findings support our intuition of a possibility for prices in prediction markets to be affected by politically motivated trading activity, particularly if pro-Trump sentiment contributes to persistent overvaluation of Trump-

favorable contracts.

1.2 Hypotheses and Research Question

The study examines the following research question:

Does systematic betting on anti-Trump outcomes in Polymarket reveal persistent political bias in market pricing?

To evaluate this question, this paper implements a rule-based trading strategy that systematically bets against Trump-favorable outcomes across selected political markets. If politically aligned traders disproportionately favor pro-Trump positions, anti-Trump contracts may become undervalued and generate positive abnormal returns.

The study formulates the following hypotheses.

Null Hypothesis H_0 : Market Efficiency

A systematic anti-Trump betting strategy yields zero abnormal returns over time. Any observed profits or losses are attributable to random market fluctuations.

Under this hypothesis, Polymarket prices efficiently incorporate available information and no persistent politically driven mispricing exists.

Alternative Hypothesis H_1 : Pro-Trump Market Bias

A systematic anti-Trump betting strategy yields persistent positive abnormal returns.

If supported, this would suggest that pro-Trump outcomes are systematically overpriced relative to their true probabilities, potentially reflecting ideological overconfidence or politically motivated trading behavior among market participants.

1.3 Methodology Overview

To examine the proposed hypotheses, we adopt an empirical approach grounded in automated data collection and rule-based simulation. Price data for Trump-related binary markets has been scraped daily from Polymarket’s Gamma API and persisted in an Elasticsearch backend. A two-stage filter restricts the sample to binary Yes/No markets discovered through either the `trump` tag (included unconditionally) or the `politics` tag (included only when the question text matches a configurable Trump keyword set). On this filtered universe we simulate a Dollar-Cost Averaging (DCA) strategy that places a fixed daily bet on the anti-Trump outcome of every tracked market [32]. Due to the binary structure of the contracts, the corresponding pro-Trump prices and metadata is also retained and analyzed by taking the opposite side of each simulated trade. To evaluate whether the observed performance reflects politically driven pricing distortions rather than random directional exposure, the strategy is compared against a neutral benchmark that executes identical trades in the same markets and at the same times, but randomly selects between YES and NO positions. The neutral benchmark is simulated repeatedly using Monte Carlo methods to estimate the expected return distribution under random trading behavior.

The resulting equity curve is then evaluated using a one-sample t -test on daily returns, as well as a linear regression on cumulative profit and loss to detect a statistically significant trend.

1.4 Contributions

This paper makes two contributions to the literature on prediction-market efficiency. First, we develop a fully replicable simulation framework for testing ideological pricing bias in real-money prediction markets. The framework is deterministic, requires no market-specific signals and is released as an open-source Docker Compose application together with a versioned seed dataset, so that any subsequent researcher may reproduce or extend our results on equivalent data. Second, we provide novel empirical evidence on whether a systematic anti-Trump betting rule generates abnormal returns on Polymarket during the post-inauguration period of the second Trump administration. Because the strategy contains no information advantage by construction, any persistent deviation from zero return is informative about

the structural pricing of politically charged contracts and therefore about the ideological efficiency of the market.

1.5 Report Structure

The remainder of this report is organized as follows. Section 2 reviews the relevant literature on prediction markets, market efficiency and political bias in financial decision-making. Section 3 describes the methodology, including the data collection pipeline, the market identification and filtering procedure, the DCA simulation framework and the statistical tests employed. Section 4 presents the empirical results, with a description of the dataset and proceeding through the primary statistical analysis and the dual equity-curve comparison. Section 5 discusses the findings in the context of the existing literature and addresses limitations and potential confounders. Section 6 concludes by summarizing the answers to the research question and hypotheses.

2 Literature Review

2.1 Prediction Markets

Wolfers and Zitzewitz describe prediction markets as markets in which “participants trade contracts whose payoff depends on unknown future events,” where prices reflect the market’s expectation of the likelihood of those events occurring ([39] p. 109). Prediction markets function as information aggregation mechanisms where traders incorporate dispersed knowledge into market prices through trading activity [9]. In such markets, prices can be interpreted as consensus probability estimates reflecting the collective beliefs of market participants [9].

There are several different types of contracts that can be used in prediction markets. One common structure is the winner-take-all contract, where a trader pays a price p for a contract that pays a fixed amount if a specific event occurs. Under the assumption of risk neutrality, the price of the contract represents the market’s estimated probability that the event will occur [39]. Other contract structures include index contracts, where the payoff depends on a numerical outcome such as vote share, and spread contracts, where traders bet on whether an outcome exceeds a specified threshold.

Prediction markets have been applied across a wide range of domains. One of the first recognized prediction markets was the Iowa Electronic Markets [1], which allowed participants to trade contracts based on election outcomes. Empirical analyses of the Iowa Electronic Markets show that the markets have consistently produced highly accurate election forecasts, often outperforming traditional polling methods and achieving very small average forecast errors in presidential elections [2].

Because of their ability to aggregate diverse information and incentives for truthful revelation, prediction markets have been increasingly studied as forecasting tools in economics, political science and business forecasting. Their potential applications include predicting political outcomes, financial indicators, technological developments and corporate performance.

2.2 Online and Decentralized Prediction Markets

Prediction markets have increasingly moved to digital platforms that allow participants from around the world to trade contracts related to future events. With the development of blockchain technology and decentralized finance (DeFi), prediction markets have begun to operate in decentralized environments where market transactions are recorded on distributed ledgers rather than controlled by a central authority. Blockchain technology provides an immutable record of transactions and allows digital assets to be exchanged directly between users without intermediaries [3].

In blockchain-based prediction markets, trades are often executed through smart contracts that automatically enforce the rules of the market and distribute payouts once an outcome has been verified. This allows prediction markets to operate in a trustless environment where participants do not need to rely on a central platform operator to manage trades or settle contracts [3].

Decentralized prediction markets offer several advantages compared to traditional online platforms. The use of distributed ledgers provides transparency and reduces the risk of manipulation, while cryptocurrencies allow global participation without traditional financial intermediaries. In addition, decentralized oracle systems can be used to determine the final outcome of events by aggregating information from multiple participants, reducing re-

liance on a single authority to resolve market outcomes [3]. Recent research has also examined the efficiency of cryptocurrency markets, which form the economic infrastructure underlying many decentralized prediction markets. Evidence suggests that these markets may evolve towards greater efficiency over time as liquidity and participation increase [15].

These advancements have led to the emergence of blockchain-based prediction market platforms such as Augur, Gnosis, Kalshi and Polymarket, which allow users to trade on the probability of real-world events using cryptocurrencies. These platforms represent a new generation of prediction markets that combine financial incentives, distributed computing infrastructure and collective intelligence to forecast uncertain outcomes.

2.3 Market Efficiency and Information Aggregation

Financial markets are often viewed as mechanisms that aggregate dispersed information from many individuals into market prices. One of the earliest theoretical explanations of this process was provided by Hayek, who argued that economic knowledge is widely dispersed across individuals and that prices serve as signals that communicate information about relative scarcity and preferences. In this framework, the price system allows individuals to coordinate their actions without requiring any single participant to possess complete information about the entire economy [10].

Hayek emphasized that the fundamental economic problem is not merely the allocation of given resources, but rather the utilization of knowledge that is distributed among many individuals. Because each participant possesses only partial knowledge of economic conditions, decentralized market mechanisms allow this information to be incorporated into prices through the interactions of many market participants [10].

Building on this idea, the Efficient Market Hypothesis (EMH) formalizes the concept that market prices reflect available information. According to Fama, an efficient market is one in which asset prices “fully reflect” all available information. Under this assumption, it becomes difficult for traders to consistently achieve abnormal returns using publicly available information, since prices already incorporate the relevant information known to market participants [7].

The efficient markets framework suggests that competition among traders quickly incorporates new information into prices, causing asset prices to adjust rapidly when new information becomes available. As a result, price changes in efficient markets are often modeled as unpredictable, since future price movements depend only on new information that cannot be anticipated in advance [7].

Later research has highlighted important limitations of perfectly efficient markets. Grossman and Stiglitz argue that perfectly informationally efficient markets cannot exist because traders require incentives to acquire and process costly information. If prices fully reflected all available information, rational traders would have no incentive to invest resources in obtaining new information, which would ultimately prevent markets from remaining fully efficient [8].

This insight suggests that real-world markets may exhibit some degree of informational inefficiency, as prices may not always fully incorporate all available information. As a result, market prices should be interpreted as reflecting the aggregated beliefs and information of market participants, while still allowing the possibility that behavioral biases or informational frictions may lead to temporary price distortions.

These theoretical perspectives provide an important foundation for understanding prediction markets. If prediction markets function as efficient information aggregation mechanisms, market prices should reflect the collective beliefs of traders regarding the probability of future events. At the same time, behavioral biases or motivated reasoning among traders may lead to deviations from fully efficient pricing, particularly in politically sensitive markets.

2.4 Political Polarization and Ideological Bias

Research examines how political polarization shapes and influences belief formation, particularly through motivated reasoning. Across the literature reviewed, a consistent finding is that individuals do not process information in a purely objective manner but rather that they interpret evidence in ways that reinforce prior beliefs and align with their political or social identity.

One of the key explanations is that individuals rely on identity-protective cognition, where beliefs signal group membership rather than reflect objective truth [13]. Motivated reasoning is not a result of low cognitive ability or a reliance on heuristics, it is identity-protective and can be stronger among individuals with higher cognitive reflection. Kahan’s [13] study finds that individuals who are better at analytical thinking are often more capable of interpreting evidence in a way that supports their group identity, meaning that increased reasoning ability can actually intensify polarization rather than reduce it. Motivated reasoning is shown to be largely symmetric across political ideologies, with both liberals and conservatives displaying similar tendencies to align beliefs with identity [13].

At the individual level, Taber & Lodge [34] demonstrate that motivated reasoning operates through several cognitive mechanisms. Individuals exhibit confirmation bias, actively seeking out information that supports their prior belief, as well as a disconfirmation bias, where they devote more cognitive effort to counter-arguing opposing evidence. These processes lead to attitude polarization, even when individuals are exposed to balanced arguments, their views become more extreme rather than converging [34]. This suggests that information does not correct beliefs but instead reinforces them.

These individual-level biases are reinforced by the broader political environment. Druckman, Peterson & Slothuus [5] show that elite polarization increases motivated reasoning by making partisan identity more salient. In politically polarized environments, individuals rely more heavily on party cues and evaluate arguments on whether they align with their political group rather than their objective strength. As a result, individuals may ignore substantively stronger argument if they come from an opposing party [5]. This leads to greater ideological division as political identity overrides evidence-based evaluation.

Together this literature suggests that beliefs in politically charged environments are systematically biased. These findings have direct implications for prediction markets such as Polymarket. When individuals participate in financial environments such as prediction markets, these biases may not disappear but instead become embedded in trading behavior and ultimately market prices.

Recent empirical work supports this mechanism. Evidence from political betting markets shows that trader behavior and market outcomes are influenced by partisan beliefs rather than purely objective information. Becker, McGurk and LeGrasso [16] demonstrate that partisan sentiment significantly predicts asset returns, particularly for Republican-linked contracts, indicating that market prices can be driven by ideological optimism rather than unbiased expectations. This suggests that biased information processing extends beyond individual beliefs and directly affects price formation.

At the same time, deviations from efficiency are not solely driven by behavioral factors. Structural features of prediction markets also contribute to systematic distortions. Rothschild [24] highlights that both polling data and prediction markets contain inherent biases, including the well-documented favorite-longshot bias, where probabilities are systematically misrepresented. These distortions may arise even when individuals hold accurate beliefs, due to market frictions such as liquidity constraints and transaction costs.

Finally, evidence from Polymarket further illustrates the interaction between ideological bias and financial incentives. Saama [25] finds that partisan traders exhibit greater directional position persistence than neutral traders following major electoral events, indicating the presence of belief-driven trading. However, these traders also partially adjust their positions in response to new information, suggesting that financial incentives impose some discipline on behavior.

This suggests that financial incentives may reduce the strength of ideological bias, but do not fully eliminate it as persistent belief-driven trading can lead to statistically significant deviations from market efficiency.

2.5 Theoretical Motivation

The literature review suggests that the informational efficiency of prediction markets depends on the availability of information, as well as the composition of traders interpreting this information. Prediction markets are efficient as they aggregate dispersed knowledge through competitive trading. However, research on motivated reasoning and partisan cognition strongly suggests that individuals may interpret political information through ideological priors rather than only probabilistic evaluation. If participation within a

market becomes concentrated among politically aligned individuals, there is a possibility of these biases becoming embedded in market prices. This is particularly relevant in decentralized cryptocurrency-based prediction markets where participation requires engagement with cryptocurrency infrastructure, and therefore may attract a politically distinct community. Given these conditions, politically driven mispricing may emerge even though participants have access to the same public information. This theoretical perspective is the motivation for the empirical analysis in this paper.

3 Method

This study employs an empirical simulation design that combines automated data collection with a deterministic rule-based trading strategy to investigate whether political bias manifests as persistent pricing distortions in Polymarket’s Trump-related binary markets. A binary market is a contract where the outcome of an event has exactly two options - either Yes or No. This means there is no middle ground and there can only be winning or losing bets. The methodology is structured around five components: research design, market identification and selection, data collection, simulation framework and return measurement and statistical analysis.

3.1 Research Design

This study takes a quantitative empirical approach. The underlying idea is straightforward: if politically motivated traders are systematically overpricing pro-Trump outcomes on Polymarket, then consistently betting on the opposite side should generate positive returns over time. Under the Efficient Market Hypothesis, a strategy with no informational advantage should produce returns indistinguishable from zero, so any persistent deviation from that baseline is interpreted as evidence of ideologically driven mispricing.

The strategy is kept deliberately simple. It uses no market-specific information, makes no attempt to time entries, and applies the same rule uniformly across every tracked market. This is important for interpretation. If a mechanical, rule-based strategy with no informational edge consistently produces abnormal returns, the explanation has to lie in the structure of the prices themselves rather than in any skill on the part of the trader.

To make sure the results are not simply an artifact of which markets were selected or how the portfolio is constructed, the study also implements a neutral benchmark strategy. The benchmark executes the same trades in the same markets at the same times, but randomly assigns each position to either YES or NO uniformly. Repeating this 10,000 times produces a distribution of outcomes that a direction-agnostic trader could expect. The performance of the anti-Trump strategy is then compared both against zero and against this distribution, which makes it possible to isolate the effect of the directional bet from everything else.

3.2 Market Identification and Selection

Markets are sourced from Polymarket’s Gamma API (<https://gamma-api.polymarket.com>), which provides real-time metadata and pricing for all active markets on the platform. To ensure that only politically relevant markets are included in the simulation, a two-stage filtering procedure is applied.

In the first stage, markets are identified by tag. All markets carrying the **trump** tag on Polymarket are included unconditionally, as this tag is applied by the platform to markets specifically framed around Donald Trump or his administration. Markets tagged **politics** are subjected to an additional keyword filter in the second stage: the market question must contain at least one of the following terms, matched case-insensitively: *trump*, *donald trump*, *djt*, *maga*, *potus*, or *president trump*.

In addition to topical relevance, two structural filters are applied. First, only binary Yes/No markets are included. Examples of such market questions include “Will Trump implement 25% tariffs on Canadian goods?”, “Will Trump fire Federal Reserve Chair Powell?”, and “Will Trump issue a pardon for [specified individual]?” Second, only markets that are active and non-closed at the time of collection are included, ensuring that the portfolio reflects live tradeable contracts.

Applying these filters, the system tracked approximately 218 markets simultaneously over the observation period, with the count varying as markets resolved and new ones were added.

It is important to note that the simulated strategy does not uniformly bet Yes across all markets. While the strategy is defined as consistently taking the anti-Trump position, the directional bet (Yes or No) depends on how each individual market question is framed. For example, a market such as “*Will Trump implement 25% tariffs on Canadian goods?*” warrants a No bet as the anti-Trump outcome, whereas a market such as “*Will Trump be impeached by June 2026?*” warrants a Yes bet, since a Yes resolution in that market would represent a negative outcome for Trump. The anti-Trump direction for each market is therefore determined manually on a market-by-market basis prior to simulation, with each market assigned a directional stance that reflects alignment with Trump’s political loss. This manual classification is a deliberate methodological choice. It ensures that the portfolio genuinely represents the behavior of an ideologically motivated anti-Trump trader.

3.3 Data Collection

The dataset is built from two sources, one historical and one collected in real time.

For the retrospective period covering 19 July 2025 to 25 January 2026, daily prices were retrieved from Polymarket’s CLOB API after the fact. The API returns one closing price per market per calendar day, but only for the YES side of each contract. The NO-side price is therefore inferred as $p_{\text{NO},t} = 1 - p_{\text{YES},t}$, which means the spread between the two sides is treated as zero for all retrospective data. This is a simplification that we discuss further in the limitations section.

For the prospective period, we collected daily price snapshots from Polymarket’s Gamma API using a purpose-built data collection application. Each collection run queried the API for all events tagged as `trump` or `politics`, applied the filters described in the previous section and saved a snapshot for each qualifying market. The Gamma API returns both the YES and NO prices directly, so the real bid-ask spread is captured. Each snapshot record contains the following fields: `timestamp_utc`, `market_id`, `question`, `yes_price`, `no_price`, `yes_cents`, `no_cents`, `spread`, `volumeNum`, `liquidityNum`, `active`, `closed`, and `market_slug`. The `yes_price` field represents the market-implied probability that the YES outcome occurs and is the primary

price variable used throughout the simulation.

To avoid duplicate records, each snapshot is given a unique identifier of the form `{timestamp_utc}|{market_id}`. Re-running the collection at any prior timestamp therefore produces no new data. All snapshots are stored in an Elasticsearch database and exported to a versioned Excel file that is committed to the project repository after each collection run, making the full dataset recoverable from any point in the Git history.

The observation window runs from 19 July 2025 to 1 May 2026, covering 287 days in total. The first 191 days make up the retrospective period, based on CLOB historical prices. The remaining 96 days form the prospective period from 26 January to 1 May 2026. The first 27 days of this period initially had irregular live coverage while the pipeline was being set up and were later back-filled using CLOB daily prices to give complete daily coverage across the full prospective window. Consistent Gamma API collection runs from 22 February onward. The number of tracked markets grew throughout the observation period as new Trump-related contracts were listed on Polymarket, settling at 217 active markets at termination.

3.4 Simulation Architecture

To support the data collection and simulation, we built a self-contained full-stack application that handles market discovery, snapshot storage, simulated trading and data export. The system runs as a Docker Compose deployment with three services (Figure 1): an Elasticsearch 8.11 cluster that stores all data, a Python FastAPI backend that contains the collection and simulation logic, and a React/Vite frontend used for market discovery and manual configuration. The full codebase is publicly available at <https://github.com/NoahJunge/polymarkettracker>.

Data Layer

All data is stored across seven Elasticsearch indices, each with a deterministic document identifier to make the pipeline idempotent. The `markets` index holds discovered market metadata, including the question text, slug, outcome list, volume and liquidity. The `snapshots_wide` index is the append-only price history; each document is keyed by `{timestamp_utc}|{market_id}`,




Figure 1: Three-service Docker Compose deployment. The backend is the only component that communicates with the Polymarket APIs; all persistent state lives in Elasticsearch.

which means re-running the collection on the same time window cannot produce duplicate records. The `tracked_markets` index records which markets are included in the simulation, along with the manually assigned `pro_outcome` field (Yes or No) that records which outcome represents the pro-Trump side of the contract. The anti-Trump DCA subscription is then placed on the opposite outcome. The `paper_trades` index holds simulated trade records; positions are recomputed by aggregating trades per market and side, with realized profit and loss calculated on a first-in-first-out basis. The `dca_subscriptions` index stores one configuration document per active DCA subscription, including the market, side, fixed quantity and a `last_executed_date` field that prevents more than one trade being recorded per market per day. The `settings` and `alerts` indices hold runtime configuration and user-defined price notifications respectively.

Reproducible Seed Pipeline

The application exports the full dataset to an Excel workbook (`seed.xlsx`) after each collection run. The workbook contains all tracked markets, price snapshots, simulated trades and DCA subscriptions and is automatically committed and pushed to the project’s GitHub repository. The consequence is that the Git history doubles as a versioned data archive: any commit can be checked out to recover the exact dataset at that point in time. When the application is started against an empty database, it imports the seed file automatically, so anyone who clones the repository begins from an identical state. In addition to the Gamma API snapshots, all CLOB historical prices retrieved during the retrospective backfill are included in the seed, giving a complete record of both data sources in a single file.

Market Discovery and Stance Assignment

The web dashboard serves one purpose that cannot be automated: deciding which markets belong in the simulation and which side of each contract represents a bet against Trump. The discovery page lists every market the backend has found but not yet included in the simulation. We reviewed each candidate manually and decided whether it was relevant to the study. For each included market, we also had to identify whether YES or NO was the anti-Trump outcome. Most Trump-related markets are framed affirmatively, for example “Will Trump implement tariffs?”, where NO is the anti-Trump bet. But a meaningful number are framed the other way around, for example “Will Trump be impeached?”, where YES is the anti-Trump bet. Since this cannot be reliably determined from the question text alone, we classified each of the markets manually. This reduces scalability but ensures that every position in the simulation genuinely reflects an anti-Trump directional bet.

Collection Pipeline

Each collection run proceeds in five steps: discover and update market meta-data from the Gamma API, store a fresh price snapshot for every tracked market, check any active price alerts, execute the daily DCA trade for each

subscription that has not yet traded on the current calendar day, and export and push the updated seed file to GitHub. Every step is idempotent, so restarting the application mid-run does not produce inconsistencies. Collection was carried out manually on a daily basis throughout the observation window.

3.5 Simulation Framework

Dollar-Cost Averaging (DCA) is an investment strategy in which a fixed amount is invested at regular intervals regardless of the current price [32]. Rather than attempting to time the market, an investor using DCA simply buys the same amount on a set schedule and accepts whatever price the market offers at that moment. We chose this approach for three reasons.

First, DCA is a natural model of an ideologically motivated retail investor who responds to ongoing political events by repeatedly purchasing contracts aligned with their beliefs, rather than making a single timed entry. Second, fixing the trade direction, frequency and quantity in advance across all markets eliminates timing discretion and market-specific signals, so any return pattern that emerges can only be explained by structural properties of the prices themselves. Third, applying DCA simultaneously across approximately 200 markets produces a diversified return series that reduces the influence of any single market outcome on the overall result.

The strategy works as follows. For each tracked market, we create a DCA subscription that places one simulated trade per calendar day, always on the anti-Trump side of the contract, for a fixed single quantity. The entry price for each trade is taken from the most recent price snapshot at the time of execution. More than one trade cannot be executed per market day to avoid duplication.

To extend the return series as far back as possible, the simulation includes a backfilling procedure. When a DCA subscription is created, one simulated trade is generated retroactively for each historical snapshot day available in the database for that market. This extends the series back to the earliest available price record for each market, though it carries the simplifying assumption that sufficient liquidity was available at historical prices to execute trades at those levels. We discuss this further in the limitations section.

When a market resolves, the corresponding DCA subscription is cancelled automatically. If the market resolves in favour of the YES outcome, the `yes_price` converges to 1.0 USDC, representing a full payout on any YES position. If it resolves in favour of the NO outcome, the `yes_price` converges to 0.0 USDC, representing a total loss on a YES position and a full payout on a NO position. These final prices come directly from the last snapshot recorded before market closure was detected.

The neutral benchmark is constructed alongside the anti-Trump strategy using the identical DCA procedure, across the same markets, on the same calendar days and in the same fixed quantities. The only difference is that instead of always betting anti-Trump, each market is randomly assigned to either YES or NO with equal probability at the start of each simulation run. This process is repeated 10,000 times to produce a distribution of outcomes that a direction-agnostic trader could expect. Because every other feature of the trading process is held constant, any persistent gap between the anti-Trump strategy and this distribution can be attributed to the directional bet itself rather than to portfolio construction effects.

3.6 Limitations of the Simulation Framework

All trades in the simulation are simulated rather than executed in live markets. No actual capital is deployed on Polymarket and no real contracts are purchased. This is intentional: it allows the strategy to be applied uniformly across all tracked markets without requiring substantial capital. But it also introduces several limitations that matter when interpreting the results.

The most significant is that the simulation does not account for price impact. In a live prediction market, each trade moves the price in the direction of the bet. A participant systematically buying YES contracts across hundreds of markets would cause YES prices to rise and NO prices to fall through their own activity alone, partially creating the very pricing pattern the study aims to detect. The paper trading design avoids this circularity entirely, since simulated trades are recorded against observed prices without affecting them. That said, the fixed daily quantities used in the simulation are deliberately modest. Given the liquidity levels observed across the tracked markets, real trades of equivalent size would be unlikely to generate meaningful price im-

pact in the more actively traded markets. The concern applies mainly to smaller, less liquid markets, where even modest order sizes can shift prices noticeably.

Relatedly, the simulation assumes that each trade can be executed at the observed price without slippage, queue effects or liquidity constraints. In practice, the bid-ask spread and available liquidity in each market would determine the actual execution price, which may deviate from the mid-price snapshot recorded by the collection pipeline. This is most likely to matter in smaller or less actively traded markets, where spreads are wider and depth is thinner.

The backfilling procedure assumes that historical snapshot prices were tradeable at the time they were recorded. Liquidity conditions at past time points are unknown and cannot be reconstructed from price snapshots alone.

Mixed realized and unrealized positions

Of the 218 markets tracked across the experiment, 97 (approximately 44%) had resolved by the experiment’s end date of 1 May 2026, contributing realized profit and loss at the final binary settlement price. The remaining 121 markets (approximately 56%) were still active at that date, contributing unrealized profit and loss marked to the prevailing market price. Because unrealized positions depend on current market sentiment rather than confirmed outcomes, the portfolio value at any given point reflects a mixture of settled results and probabilistic valuations. This introduces noise into the return series that would not be present in a fully resolved portfolio. The final P&L for both strategies will only be confirmed once the remaining open markets resolve, and outcomes could move in either direction.

Retrospective data quality and the no-spread assumption

The retrospective period relies on price history fetched post-hoc from the Polymarket CLOB API, which exposes only the YES token price. The NO-side price is therefore computed as $p_{\text{NO}} = 1 - p_{\text{YES}}$ exactly, implying a zero bid-ask spread for all retrospective NO-side positions. In contrast,

the prospective data sourced from the Gamma API stores both sides independently, capturing the real spread. This means transaction costs are systematically understated in the retrospective period, and any direct comparison of execution quality across the two periods should account for this difference.

Market selection

Markets are included in the dataset only if they were active at the time data collection commenced. Markets that opened and fully resolved before collection began are not represented, which introduces a form of selection bias. The strategy is evaluated on contracts that remained open long enough to be captured, not on the full universe of Trump-related markets active during the sample period.

Cross-market correlation

The 218 tracked markets are not statistically independent. Many contracts concern the same underlying political events, so a single news event can move multiple markets in the same direction on the same day. This matters most for the per-market breakdown: treating each market's individual P&L as an independent observation would overstate the effective sample size, since correlated markets share information rather than adding independent evidence. For the primary t-test the concern is less acute, because that test is applied to 96 daily portfolio returns, one aggregate figure per day, which already absorbs cross-market correlation into the portfolio's observed variance. The Ljung-Box test confirmed no significant autocorrelation across those days, which is the independence condition the t-test actually requires. Readers should nonetheless be cautious when interpreting the per-market results as if each market were an independent trial.

The neutral benchmark and its assumptions

The Monte Carlo benchmark assigns each market a random direction at the start of each simulation and holds it constant throughout. This is well-suited to isolating the effect of directional choice, but it represents a stylized model of a neutral trader. A real direction-agnostic participant might adjust positions over time or concentrate capital differently. The benchmark therefore measures the value of a specific type of neutral strategy, one structurally identical to the directional strategies in all respects except direction, rather than the full space of non-directional alternatives.

Platform specificity and sample period

The results are specific to Polymarket during a single, politically active window spanning approximately ten months from July 2025 to May 2026. The platform has experienced legal contentions in regards to operating in the US. It exited the US market in 2022, relaunching in 2025 with limited operability [18]. It is federally legalized, but this is contested between states, with some states such as Tennessee and Nevada taking further legal action to block the platform [18].

Whether the pricing patterns documented here generalize to other prediction market platforms, other political contexts, or future time periods remains an open question.

These limitations do not invalidate the study’s core findings. They do mean that the results should be interpreted as evidence of structural pricing patterns in the market rather than as a demonstration of a practically executable trading strategy. The simulation serves as a controlled instrument for detecting systematic deviations from efficient pricing.

3.7 Return and Performance Measurement

Portfolio returns are computed using a mark-to-market valuation framework applied to all open positions at each observation date.

We define the contract direction for each market i as:

$$s_i \in \{\text{yes, no}\}$$

where s_i denotes whether the position in market i is held through YES or NO shares.

The mark-to-market value of a position in market i at time t is:

$$V_{i,t} = Q_{i,t} \times p_{i,t}^{(s_i)} \quad (1)$$

where $Q_{i,t}$ is the total number of shares held in market i at time t , and $p_{i,t}^{(s_i)}$ is the current price on the contract direction held.

For positions accumulated through multiple daily DCA trades, the average entry price $\bar{p}_{i,t}^{\text{entry}}$ is the volume-weighted mean of all historical entry prices for that market. Unrealized profit and loss for a position is:

$$\text{PnL}_{i,t}^{\text{unrealized}} = Q_{i,t} \times \left(p_{i,t}^{(s_i)} - \bar{p}_{i,t}^{\text{entry}} \right) \quad (2)$$

Total portfolio equity at time t is the sum of the mark-to-market values of all open positions:

$$E_t = \sum_{i \in \mathcal{M}_t} V_{i,t} \quad (3)$$

where $\mathcal{M}_t$ denotes the set of markets with open positions at time t .

The daily portfolio return is computed as the change in total profit and loss relative to the cumulative capital invested through the prior day:

$$r_t = \frac{\Delta \text{PnL}_t}{I_{t-1}} \quad (4)$$

where $\Delta \text{PnL}_t = \text{PnL}_t - \text{PnL}_{t-1}$ is the day-over-day change in total portfolio profit and loss, and I_{t-1} is the total capital deployed through day $t - 1$. Dividing by invested capital rather than portfolio equity avoids conflating daily valuation changes with the continuous injection of new capital through DCA trades.

Several portfolio-level performance metrics are computed from the daily return series. The Sharpe ratio [29] is the mean daily return divided by the standard deviation of daily returns, annualized by a factor of $\sqrt{365}$:

$$\text{Sharpe} = \frac{\bar{r}}{\bar{\sigma}} \times \sqrt{365} \quad (5)$$

Returns are not adjusted for the risk-free rate, as short-term risk-free rates are negligible relative to the fluctuations observed in a prediction market portfolio. Additional risk metrics reported include the Sortino ratio, maximum drawdown and Value at Risk at the 95% confidence level. The win rate is the proportion of tracked markets whose position is in profit at the end of the observation window, valued at the prevailing market price for markets that have not yet resolved. The profit factor is the ratio of total gross profit to total gross loss, computed on the same market-level basis.

The same framework is applied identically to the pro-Trump strategy. The only difference is that the contract direction s_i is flipped for every market, so each position takes the pro-Trump side of the corresponding contract.

3.8 Statistical Analysis

The statistical analysis tests whether the observed return patterns are consistent with market efficiency. Under the Efficient Market Hypothesis, a deterministic strategy with no informational advantage should produce returns indistinguishable from zero over time. Under H_1 , the anti-Trump strategy should show persistently positive returns, reflecting the systematic overpricing of pro-Trump outcomes by politically motivated traders.

Before applying hypothesis tests, we verified four conditions on the prospective daily return series ($T = 96$). The Shapiro-Wilk and Jarque-Bera tests rejected normality, so the t-test is supplemented with non-parametric alternatives. The Ljung-Box Q-test at lags 1, 5, and 10 found no significant autocorrelation, satisfying the independence assumption the t-test requires. The Augmented Dickey-Fuller test confirmed the series is stationary, and the Durbin-Watson statistic of 2.17 indicated no lag-1 autocorrelation.

The primary hypothesis test is a one-sample t-test on the mean daily portfolio return:

$$t = \frac{\bar{r}}{\hat{\sigma}/\sqrt{T}} \quad (6)$$

where $\bar{r}$ is the sample mean of daily returns, $\hat{\sigma}$ is the sample standard deviation and $T = 96$ is the number of prospective observation days. Statistical significance is assessed at $\alpha = 0.05$. A significantly positive t-statistic supports H_1 ; failure to reject the null corresponds to H_0 .

To account for the non-normality of returns, two robustness checks are applied alongside the t-test. A bias-corrected and accelerated bootstrap confidence interval (BCa) is constructed from 10,000 resamples of the daily return series. A Wilcoxon signed-rank test provides a non-parametric test of whether the median daily return differs from zero. Together these give a picture of significance that does not depend on distributional assumptions.

To assess whether the return pattern reflects a persistent directional trend rather than a short-lived deviation, we estimate an OLS regression of cumulative portfolio profit and loss on time:

$$\text{PnL}_t = \alpha + \beta t + \varepsilon_t \quad (7)$$

where t is the observation day index. The slope coefficient $\hat{\beta}$ captures the average daily drift in accumulated profit and its significance is assessed using heteroskedasticity-consistent standard errors (HC3). A statistically significant positive slope for the anti-Trump strategy supports a persistent structural upward trend consistent with H_1 .

The neutral benchmark is constructed using a Monte Carlo procedure that isolates the contribution of directional choice. In each of 10,000 simulation runs, every tracked market is independently assigned to either YES or NO with equal probability. The full DCA simulation is then re-run with those random directions, using the same markets, the same calendar days and the same fixed quantities as the actual strategy. This produces a distribution of outcomes that a direction-agnostic trader could expect. The percentile rank of the anti-Trump strategy within this distribution provides a direct measure of how unusual its performance is relative to a benchmark that has no directional view on any market.

4 Results

4.1 Dataset Description

This section describes the empirical dataset used in the analysis. All figures reported here are extracted directly from the Elasticsearch indices via the backend’s diagnostic endpoints; the underlying state is reproducible from the versioned `seed.xlsx` archive described in the previous chapter.

Aggregate Counts

A total of 287 Trump-related markets were tracked throughout the observation period. After filtering and manual processing, 218 of these markets had subscribed DCA trading. The results and data obtained are from the 218 markets with DCA subscriptions. As of May 1st, the date of termination, each trading strategy held a total number of 46,058 contracts.

Temporal Coverage

The prospective observation window spans 96 days from January 26, 2026 to May 1st, 2026. Coverage was not uniform across the early stages of the observation window, due to iterative development of the simulation. Any missing data during this time period was filled using the historical data from the CLOB API. The retrospective period spans from July 19 2025 to January 26 2026, in which historical data was retrieved from the CLOB API for each of the tracked markets. The following analysis covers both the retrospective period, prospective period and the full observation window (287 days).

Market Categories

The tracked markets include *tariff* and trade-policy markets, *personnel* markets concerning appointments and dismissals within the executive branch, *judicial* markets relating to nominations and rulings, *electoral* markets covering primaries and special elections and *international relations* markets concerning diplomatic outcomes, sanctions and agreements. The breadth of categories is deliberate. A large and diverse set of markets helps average out event-specific noise, as well as reduces the chance that any findings are driven by a single category.

4.2 Profit and Loss Analysis

Polymarket contracts are binary prediction markets in which the probabilities of opposing outcomes sum to one. Because of this, observing the market price for an anti-Trump outcome automatically determines the implied probability of the corresponding pro-Trump outcome. This allowed equivalent pro-Trump and anti-Trump trading strategies to be constructed without requiring separate datasets. Profit and loss will be presented for both pro and anti-Trump strategies, as well as the neutral trading strategy, for comparison. The profit and loss are analyzed for the entire observation period, as well as a separate analysis of both the retrospective period and prospective period.

The anti-Trump trading strategy produced an overall profit of USD 1,005.66 over 287 days, in which the mean daily return was 0.0279%. The total capital invested amounted to USD 30,020.92. The profit factor was 2.006, indicating that the strategy was overall profitable. The win rate from the strategy was 64.2%: 140 of the 218 tracked markets held a profitable anti-Trump position at the end of the observation window. Because 121 of those markets had not yet resolved on 1 May 2026, this figure is a mark-to-market measure rather than a count of confirmed correct outcomes. The annualized return for the anti-Trump strategy over the full observation period is 10.18%.

During the retrospective period, spanning 191 days from the 19th of July 2025 to the 26th of January 2026, the mean daily return was 0.0160%. This is significantly less than the mean daily return of the whole period as majority of the markets tracked were first opened from January onwards - hence there is minimal data for the retrospective period. The total invested during this period was USD 6,404.57, which produced an overall return of 0.98% (USD 62.67).

During the prospective period, 96 days from January 26th 2026 to May 1st 2026, the anti-Trump strategy produced a mean daily return of 0.0515%. The total invested amounted to USD 23,616.35, and produced a profit of USD 942.99 (3.99%). The prospective period was more successful in terms of profit. Trading volume often increases significantly as markets approach resolution. As a few long-standing markets were being resolved during the prospective period, this may have contributed to the difference in profits observed between the two time periods.

For the anti-trump trading strategy, both time periods produced positive returns.

The pro-Trump trading strategy produced an overall loss of USD $-1,028.30$ during the 287 day observation period, with a mean daily return of -0.0982% . The profit factor is 0.498, indicating the losses exceed the gains for this trading strategy. The annualized return on the pro-Trump strategy is -35.84% . The total capital invested for this trading strategy was USD 16,059.72.

During the retrospective period, the strategy produced a daily mean return of -0.1003% . The capital invested during this period was USD 3,678.4, producing a loss of -1.70% , or USD 62.67.

The prospective period produced a higher loss, with a daily mean return of -0.0941% . The total capital invested across this period was USD 12,381.29, with a final loss of USD -965.62 or -7.80% .

The neutral trading strategy was simulated 10,000 times. The average mean daily return from the neutral simulations was -0.0040% , with an average profit of USD 3.23. The 5th percentile of final P&L across 10,000 simulations was USD -598.08 , whereas the 95th percentile of final P&L was USD 607.51.

Below is a histogram of the of mean daily returns across the 10,000 neutral simulations. The daily mean return from the pro-Trump and anti-Trump from the prospective period are also shown.




Figure 2: Histogram of the mean daily returns from the neutral trading strategy, with pro- and anti-Trump mean daily returns from the prospective period

The distribution is approximately centered near zero, with most simulations producing small positive or negative mean daily returns. This is consistent with the expectation that a politically neutral strategy should not systematically generate abnormal returns. The pro-Trump strategy lies in the extreme left tail of the benchmark distribution, whereas the anti-Trump lies in the right tail of the distribution. Fewer than 0.4% of the neutral simulations achieved returns as poor as the pro-Trump strategy, in contrast the anti-Trump strategy outperformed all neutral simulations during the prospective period.

Below is the same histogram, with the mean daily returns shown for the full observation period.




Figure 3: Histogram of the mean daily returns from the neutral trading strategy, with pro- and anti- Trump mean daily returns from the full observation period

The pro-trump outcome remains far in the left tail of the benchmark distribution and corresponds to the 0.3rd percentile. Approximately 0.3% of neutral simulations performed as poorly or worse than the Pro-Trump strategy. The pro-Trump strategy performed slightly worse in the full observation period than the prospective period.

The anti-Trump strategy results still lie near the extreme right tail of the distribution and corresponds to the 99.7th percentile. Only around 0.3% of neutral benchmark simulations achieved equal or higher returns. The anti-Trump strategy still significantly outperforms the neutral trading strategy, but performed slightly better during the prospective period.

4.3 Dual Equity Curve Analysis

The dual analysis is implemented in a single backend method, `get_equity_curve_dual()`, which calls the underlying replay engine twice on the same set of trades and snapshot prices. The second invocation passes `flip_sides=True`, which causes the equity-curve replay routine to swap the YES and NO sides of every simulated trade before marking to market. Crucially, both curves are computed from a single Elasticsearch fetch, ensuring that any difference between them reflects only the side inversion and not differences in the underlying data.

Figure 4 presents the cumulative profit-and-loss trajectories of the Pro-Trump

and Anti-Trump strategies relative to the neutral Monte Carlo benchmark over the full observation period.




Figure 4: Dual Equity of pro-Trump and anti-Trump trading strategies

The dashed vertical gray line indicates the start of the prospective period, on the left of the line is the retrospective period. The gray shaded region represents the 5th–95th percentile range of cumulative returns generated by the neutral benchmark simulations, while the dashed horizontal gray line represents the median benchmark outcome. Similar to the histogram presented in Figure 2 and 3, the benchmark remains centered close to zero throughout the observation period, indicating that randomly assigning political positioning across the same markets generally produced approximately neutral returns.

During the early stages of the sample period, both strategies remain close to the neutral benchmark median. There was relatively limited divergence when fewer positions had accumulated and as some markets were further from resolution at that time. Some divergence is visible from November 2025, but both strategies fall close to the neutral 25-75th percentile. The trajectories of both trading strategies begin to separate substantially around the middle of March 2026. For the majority of April 2026, the pro-Trump strategy is significantly under-performing compared to the neutral 5th percentile, and the anti-Trump strategy is significantly outperforming compared the neutral 95th percentile.

The pro-Trump strategy exhibits persistent loss throughout the observation period, with the most significant loss on the 22nd of April of USD $-1,592.04$.

On the same day, the anti-Trump strategy exhibited its largest profit of USD 1,586.93. However, from February 27th 2026 to March 4th 2026, the anti-Trump strategy produced a loss, whereas the pro-Trump strategy profited. The largest profit of the pro-Trump strategy was from this period, on the 2nd of March, amounting to USD 174.26. The anti-Trump strategy's largest loss was on the same day, at USD 169.1.

The widening separation between the two equity curves is a notable feature of the figure. Rather than fluctuating randomly around the neutral benchmark median, the strategies display sustained directional divergence over time.




Figure 5: Distribution of P&L across individual markets for the Anti-Trump strategy

Figure 5 shows the profit or loss obtained from each individual market from the anti-Trump trading strategy. The equivalent market-level distribution for the Pro-Trump strategy is omitted for simplicity, as the strategy construction is symmetric. Profitability is observed across majority of the positions, with 140 markets obtaining a profit. However, a few markets produced gains above USD 125, over a tenth of the overall profit. From the graph, it appears that profitability is distributed across many markets, although several large winning positions contributed disproportionately to total portfolio returns.

4.4 Market Size Analysis

To understand where any mispricing comes from, all markets have been split into two groups based on their trading volume: large markets (above the

median) and small markets (below the median). If pro-Trump outcomes are only overpriced in smaller markets with few traders, the finding is less interesting because those markets are naturally less efficient. If the mispricing is also present in the biggest, most actively traded markets, it suggests that a large number of politically motivated traders are systematically willing to overpay for pro-Trump outcomes, even in markets where there is plenty of liquidity and opportunity for the price to correct itself. The following analysis is conducted on only the prospective period, as the retrospective data obtained from the CLOB API does not contain trading volume.



Figure 6: Large Markets VS Small Markets Analysis

The median market volume was USD 30,564. There is a total of 105 large markets, and 113 small markets. Large-volume markets yield an anti-Trump return on investment of +4.69%, compared to +2.64% for small-volume markets. This is notable because large markets attract more participants and more capital, which in an efficient market should make mispricings harder to sustain.

4.5 Statistical Analysis

Pre-Analysis Diagnostics

The following results are organised in the same structure as the preceding profit and loss analysis. Findings are reported separately for the retrospective period, the prospective period and the full observation window. This allows the statistical evidence to be read alongside the return figures already presented.

Before applying the hypothesis tests, we verified that the conditions for valid inference hold on the prospective daily return series ($T = 96$). The Shapiro-Wilk test yielded $W = 0.925$ ($p < 0.001$) and the Jarque-Bera statistic was $JB = 122.62$ ($p < 0.001$), both rejecting the assumption of normality [28, 11]. Because of this, we supplement the t-test with non-parametric alternatives throughout. The Ljung-Box Q-test at lags 1 and 5 found no significant autocorrelation ($p = 0.197$ and $p = 0.332$ respectively), and at lag 10 the result is marginal ($p = 0.049$) [14]. The Durbin-Watson statistic of 2.17 confirms no lag-1 serial correlation. The Augmented Dickey-Fuller test confirmed the series is stationary [26]. Together, these diagnostics support the use of the t-test as the primary tool, with bootstrap and Wilcoxon procedures as robustness checks.

Primary Hypothesis Test

The primary test is a one-sample t-test on the mean daily return of the anti-Trump strategy [33]. We apply this test to each of the three periods to show that the directional pattern is consistent across the full observation window.

Retrospective period ($T = 191$)

During the retrospective period, spanning 191 days from July 19, 2025 to January 25, 2026, the anti-Trump strategy produced a mean daily return of $\bar{r} = +0.016\%$ with a standard deviation of $\hat{\sigma} = 0.434\%$. Applying the test statistic:

$$t = \frac{\bar{r}}{\hat{\sigma}/\sqrt{T}} = \frac{+0.000160}{0.00434/\sqrt{191}} = +0.51, \quad p = 0.612 \quad (8)$$

The result is not statistically significant at $\alpha = 0.05$. The direction is positive, but returns are small in the retrospective period which is consistent with the profit and loss analysis, which showed that most market activity was concentrated in the later prospective window.

Prospective period ($T = 96$)

During the prospective period, spanning 96 days from January 26, 2026 to May 1, 2026, the mean daily return was $\bar{r} = +0.052\%$ with a standard

deviation of $\hat{\sigma} = 0.848\%$. The test statistic is:

$$t = \frac{+0.000515}{0.00848/\sqrt{96}} = +0.59, \quad p = 0.554 \quad (9)$$

This does not meet the significance threshold. The direction is positive and larger than in the retrospective period, but with 96 observations and substantial daily volatility, there is limited power to detect a moderate underlying drift. Non-rejection here may reflect the relatively short sample period and high volatility of daily returns, rather than the absence of a directional pattern.

Full observation window ($T = 286$)

Over the full 287-day window, the mean daily return was $\bar{r} = +0.028\%$ with a standard deviation of $\hat{\sigma} = 0.604\%$. The test statistic is:

$$t = \frac{+0.000279}{0.00604/\sqrt{286}} = +0.78, \quad p = 0.436 \quad (10)$$

Again, the result is directionally positive but not individually significant. Across all three periods, the anti-Trump strategy produced positive mean daily returns, and the pro-Trump strategy produced negative ones. The directions are consistent throughout.

Robustness Checks

Two robustness checks were applied to the prospective period, where the primary test is focused. A bias-corrected and accelerated bootstrap confidence interval (BCa) was constructed from 10,000 resamples of the daily return series [6]. The resulting 95% interval included zero, consistent with the t-test. A Wilcoxon signed-rank test, which tests whether the median daily return differs from zero without assuming normality, yielded $W = 2,222$ and $p = 0.699$ [38]. Both checks are consistent with the primary result: the direction is positive across the period, but the signal is not strong enough to be individually significant over 96 days.

Trend Analysis

Because the t-test has limited power over relatively short series with high daily noise, we also estimated an OLS regression of cumulative portfolio profit

and loss on time, with HC3 heteroskedasticity-consistent standard errors [40]:

$$\text{PnL}_t = \alpha + \hat{\beta} t + \varepsilon_t \quad (11)$$

Over the prospective period, the estimated slope for the anti-Trump strategy is $\hat{\beta} = +\$14.07$ per day ($p < 0.001$, $R^2 = 0.662$). The corresponding slope for the pro-Trump strategy is $\hat{\beta} = -\$14.21$ per day ($p < 0.001$), a near-perfect mirror. Over the full observation window, the estimated slope for the anti-Trump strategy is $\hat{\beta} = +\$2.63$ per day ($p < 0.001$, $R^2 = 0.397$) and for the pro-Trump strategy $\hat{\beta} = -\$2.65$ per day ($p < 0.001$). The OLS trend is statistically significant in both periods, in both strategies, and the slopes are near-identical in magnitude and opposite in sign. This symmetry is consistent with a single structural source, the pricing of the contracts themselves, rather than random variation in one direction.

Monte Carlo Benchmark

As shown in Figures 2 and 3, the anti-Trump strategy’s mean daily return lies in the extreme right tail of the neutral benchmark distribution in both the prospective period and the full observation window. Over the prospective period, the anti-Trump strategy exceeded the return of all 10,000 neutral simulations, placing it at the 100th percentile of the benchmark with an empirical p-value of approximately 0.01. Over the full observation window, it ranked at the 99.7th percentile. The pro-Trump strategy sits at the 0.4th percentile over the prospective period and at the 0.3th percentile over the full window. In all four comparisons both strategies and both periods results in the extreme tail of the neutral distribution.

Risk-Adjusted Performance

The Sharpe ratio [29] for the anti-Trump strategy over the prospective period is:

$$\text{Sharpe} = \frac{\bar{r}}{\hat{\sigma}} \times \sqrt{365} = \frac{0.000515}{0.00848} \times \sqrt{365} = +1.16 \quad (12)$$

A Sharpe ratio above one indicates that the return earned per unit of total risk is favorable. The Sortino ratio, which penalizes only downside volatility rather than total volatility, is +1.35 [31], reflecting that the strategy’s daily volatility skews toward positive days. The corresponding figures for the pro-Trump strategy are -1.16 and -1.04 respectively. The near-identical mag-

nitudes across both strategies confirm that they are drawing from the same underlying price series, they are just from opposite sides of each contract.

4.6 S&P 500 Correlation

To assess whether the anti-Trump strategy’s returns reflect broad equity market exposure rather than Polymarket-specific pricing effects, we compared its daily returns to those of the S&P 500 index. S&P 500 daily return data was obtained via the `yfinance` Python library, which sources prices from Yahoo Finance. If the returns are driven by general risk appetite in financial markets, the correlation with the S&P 500 would be positive and statistically significant and the finding would carry less weight as evidence of political mispricing.

The analysis covers 66 trading days over which both S&P 500 returns and anti-Trump daily returns were available. Figure 7 plots both series over time.

The Pearson correlation is $r = +0.068$ ($p = 0.586$), indicating no statistically significant linear relationship between the two series. The Spearman rank correlation is similarly negligible ($\rho = +0.073$, $p = 0.560$). Regressing anti-Trump daily returns on S&P 500 returns yields a market beta of $\hat{\beta} = +0.055$ ($p = 0.584$, $R^2 = 0.005$). The S&P 500 explains less than 0.5% of the variance in the anti-Trump return series.

These results suggest that the strategy’s returns are orthogonal to broad equity market movements. Whatever is driving the directional asymmetry between pro- and anti-Trump pricing on Polymarket, it is not equity sentiment.




Figure 7: Anti-Trump and S&P 500 daily returns plotted on a dual axis over the overlapping observation window, with a rolling 20-day Pearson correlation in the lower panel.

Summary

PERIOD	STRATEGY	DAYS	EVENTS	BETS	$\bar{R}$ (%)	Σ (%)	T-STAT	B (\$/DAY)
Retrospective	Pro-Trump	191	79	10,083	-0.1003	1.0917	-1.2597	\$-0.33
Retrospective	Anti-Trump	191	79	10,083	0.0160	0.4335	0.5088	\$0.33
Prospective	Pro-Trump	96	218	35,975	-0.0941	1.5489	-0.5953	\$-10.06
Prospective	Anti-Trump	96	218	35,975	0.0515	0.8481	0.5950	\$9.82
Full Series	Pro-Trump	287	218	46,058	-0.0982	1.2621	-1.3112	\$-3.58
Full Series	Anti-Trump	287	218	46,058	0.0279	0.6039	0.7813	\$3.50

Figure 8: Table of figures from pro-Trump and anti-Trump strategies across all time periods

Figure 8 presents a summary of the performance characteristics of the Pro-Trump and Anti-Trump trading strategies across the retrospective, prospective and full observation windows. In all three periods, the anti-Trump strategy produced positive mean daily returns and the pro-Trump strategy produced negative ones. No individual t-test reaches significance at $\alpha = 0.05$, which is expected as prediction market returns are noisy, and 96 or even 287 daily observations provide limited power against a moderate underlying

drift. The evidence becomes clear when looking at persistence rather than averages. The OLS trend is highly significant in both the prospective period and the full series ($p < 0.001$ in all cases), and the Monte Carlo benchmark places both strategies in the extreme tails of a 10,000-simulation neutral distribution across every time horizon examined.

The consistency of these results show the same direction in every period, the same symmetry between the two strategies and the same extremity in the benchmark comparison. This evidence supports H_1 : pro-Trump outcomes were systematically overpriced throughout the observation window and the anti-Trump strategy captured that premium.

5 Discussion

5.1 Why Mean-Return Tests Were Weak

The one sample t-tests were not able to reject the null hypothesis in the retrospective period, prospective period or the entire observation window. At first glance, this may appear inconsistent with the strong divergence observed in the dual equity curve, as well as the mean daily returns of the pro- and anti-trump strategy compared to the Monte Carlo benchmark. We believe this difference arises as the mean-return tests evaluate a narrower statistical question than the benchmark simulation.

The daily drift in the anti-Trump strategy was relatively small compared to the day-to-day volatility of the portfolio. Across all time periods, the mean daily return remained under 0.1%, but daily standard deviations fluctuated between 0.4% – 0.8%. Short-term fluctuations dominate individual observations. Under these conditions, it is difficult to distinguish a moderate underlying effect from random variation. In this case, the absence of significance does not imply the absence of a systematic directional pattern.

Most importantly, the foundational pattern observed from the simulation accumulated over time. The anti-Trump strategy did not generate large abnormal returns on isolated days. Rather, small directional advantages accumulated persistently over time. This produced a substantial divergence in cumulative returns between the pro- and anti-Trump trading strategy over

time. This effect is clearly visible in the dual equity curves and OLS trend regressions rather than in the individual daily return observations themselves.

The Monte Carlo benchmark analysis evaluates a different hypothesis than the one-sample t-tests. The t-tests evaluates whether the average daily return differs significantly from zero in absolute terms. Whereas the benchmark simulations aim to isolate the effect of directional betting while holding market selection, trade timing and portfolio construction constant. The observations from the simulation clearly show that the anti-Trump positioning systematically outperformed random directional assignment, as well as pro-Trump positioning.

The nature of prediction markets being heavily event-driven creates volatility and abrupt repricing around relevant developments. These characteristics reduce the sensitivity of simple parametric mean-return tests. Therefore, in the context of the simulation, the benchmark-relative analysis and cumulative trend behavior provide a more informative measure of persistent abnormal returns.

5.2 Relative Mispricing Rather than Absolute Profitability

An obvious takeaway from the study is that the anti-Trump trading strategy generated positive returns in absolute terms. However, more importantly is that the strategy outperformed the neutral benchmark. The benchmark preserved the same market selection, trade timing and portfolio construction. The only difference was the randomized directional betting. The primary signal observed in the analysis is therefore relative directional asymmetry rather than standalone portfolio profitability.

The one-sample t-tests evaluate whether the average daily return of the strategy differs significantly from zero in isolation. The benchmark simulations evaluate whether anti-Trump positioning systematically performed differently.

The neutral benchmark strategy was designed to isolate the effect of political direction while controlling the broader structure of the simulation itself. The

framework applied to both strategies was identical and the only changing variable was whether a simulated trade selected the YES or NO side of the contract. The persistent out-performance relative to the benchmark suggests that directional positioning contributed materially to portfolio outcomes.

The symmetry observed between the anti-Trump and pro-Trump strategies reinforces this. The anti-Trump strategy consistently occupied the extreme upper tail of the benchmark distribution, while the pro-Trump strategy occupied the extreme lower tail. The divergence between the two trading strategies cannot be only attributed to broad market appreciation, increasing market activity or the mechanics of the DCA framework. When comparing the three trading strategies, it becomes evident that the direction of political exposure systematically influenced pricing outcomes throughout the observation period.

While the results show that the anti-Trump strategy produced a profit, we are not arguing that Trump-related prediction markets always generate large positive returns for anti-Trump positioning in absolute terms. The study provides evidence that pro-Trump outcomes were persistently priced differently from what would be expected under neutral directional assignment. The contribution of the benchmark analysis is the proof of systematic asymmetry between politically opposing positions under otherwise identical trading conditions.

Although this study does not evaluate the anti-Trump strategy as a conventional investment strategy, the magnitude of returns observed is economically meaningful. The anti-Trump strategy generated annualized returns of 10.18% over the full observation period, which is comparable to some well-known investment funds. The S&P 500 has an average annualized return of 10.51%, and 6.64% when adjusted for inflation [17]. This comparison should be interpreted cautiously given the vastly different volatility and risk characteristics of prediction-market trading.

5.3 Alpha, Beta, and the Source of Returns

In portfolio theory, excess returns come from two sources: exposure to systematic risk factors or returns that are genuinely independent of those factors[29]. A strategy that profits simply by tracking equity markets is earning beta,

not alpha. The S&P 500 correlation results in section 5.6 bear directly on this question.

The anti-Trump strategy produces a market beta of essentially zero. At $\hat{\beta} = +0.055$ with $R^2 < 0.005$, there is no meaningful relationship between its daily returns and movements in the S&P 500. This rules out the most straightforward alternative explanation: that the strategy's positive returns are a compensation for bearing equity risk or that they reflect general financial market sentiment spilling into prediction markets.

What remains is a return that is specific to the directional structure of Trump-related contract pricing on Polymarket. The +10.18% annualised return over the full observation period is in the same range as the S&P 500's long-run average of approximately 10.51%, but it is generated through a completely different mechanism: one with no equity exposure and no dependence on broader market conditions. A simple, rules-based strategy with zero informational edge and no equity risk can match the S&P 500's annualised return because the return source is structural. The prices themselves contain a directional tilt, not the strategy.

This is consistent with the interpretation we advance throughout this paper. Ideologically motivated traders willing to overpay for pro-Trump outcomes do not necessarily act in ways correlated with equity markets. Their behaviour is politically driven, not financially. The resulting mispricing is therefore invisible to equity-market risk models and the anti-Trump strategy captures it precisely because it exploits a pricing anomaly that is orthogonal to conventional risk factors.

5.4 Interpretation of Unrealized Returns

The portfolio performance presented in the results is partially composed of unrealized profit and loss. 97 markets were closed at the termination of DCA trading and contributed fully realized returns, and 121 markets remained opened, contributing unrealized returns. Reported profits therefore represent the value that would be obtained if positions were liquidated at the current market prices on May 1st.

The distinction matters as, in prediction markets, prices continue to fluctu-

ate until the market reaches resolution. The period between when a market opens and resolution can span months or even years. A position that appears profitable during the observation period may converge toward a different final payoff as additional information enters the market. The results are therefore evidence of directional pricing behavior during the sample period rather than proof of fully realizable arbitrage profits.

The unrealized nature of the returns does not eliminate the significance of the observed results. The central finding of the study is that anti-Trump positioning consistently appreciated relative to pro-Trump positioning and relative to the neutral benchmark over an extended period of time. This paper investigates whether Trump-related prediction markets show persistent directional pricing asymmetries during the observation window, not whether every individual contract ultimately settled in favor of one political side. Mark-to-market valuation provides a useful measure of how the market priced political outcomes over time and how those prices evolve as new information enters the market.

Moreover, the analysis does not fully incorporate practical trading frictions. Bid-ask spreads, slippage, liquidity constraints or transaction costs may reduce the extent to which the observed returns could be captured in practice, particularly during periods of heightened volatility. The reported profitability should therefore be interpreted as an estimate of directional market behavior under the simulated execution framework rather than as a guaranteed realizable trading strategy.

5.5 A Question of Who Was Trading

The pattern we found is not dramatic in any individual observation. Mean daily returns are small and on any given day the strategy could easily be read as noise. The direction over time is harder to dismiss. The anti-Trump strategy gained +\$942.99 over 96 days, ranked at the 100th percentile of 10,000 neutral simulations, and produced an OLS slope of +\$14.07 per day with $R^2 = 0.662$. Two-thirds of the variation in cumulative value is explained by linear time alone. A market that drifts consistently in one direction over several months is not producing random outcomes.

A persistent drift of this kind points toward something stable in the pric-

ing environment rather than a reaction to isolated events. We believe the most plausible explanation relates to the context in which our observation window sits. In 2024, pro-Trump conviction on Polymarket was rewarded. A single participant placed approximately \$30 million in pro-Trump contracts and profited by approximately \$85 million when Trump won the election [20, 23, 19]. Political belief and market outcome aligned directly. Our data begins in July 2025, nine months into the Trump administration and the contracts we tracked operate differently. They do not ask whether Trump will win an election. They ask whether specific tariffs take effect on schedule, whether executive orders survive legal challenge, whether international negotiations end on favourable terms. These outcomes are contested, reversed and shaped by forces that political sentiment alone does not predict. If participants approached these contracts with the same directional confidence that proved correct in an election, they would systematically price pro-Trump policy outcomes too high and the anti-Trump side would carry a persistent profit.

We cannot verify this at the level of individual traders. The same directional result holds across the full 287-day series, including the period from July 2025 to January 2026 before the tariff escalation and is not concentrated in any single event. The tariff period of March and April 2026 appears to have sharpened the signal, but the directional pattern was present before it began.

5.6 What Prediction Markets Aggregate

Prediction markets are often described as the most accurate collective forecasting mechanism available. The argument is that they aggregate information from many participants with diverse views, and that the resulting price reflects the combined knowledge of the crowd. In political contexts, prediction market probabilities have frequently outperformed polling and expert forecasting.

That argument holds when the crowd is genuinely diverse. A market populated by participants with shared political priors does not aggregate a wide range of information about the true probability of an outcome. It aggregates the collective confidence of one type of participant. The price reflects that group’s calibration, not a broader consensus.

Our finding is consistent with this. Polymarket operates on blockchain infrastructure, requires cryptocurrency wallets and excludes the majority of US users. The participant base making predictions on US political events is therefore self-selected and not representative of informed opinion on those events more broadly. Across 218 markets and 287 days, the pro-Trump side was priced above what outcomes justified, consistently and in both time periods we examined.

We are careful not to overstate what this shows. The observation window is nine months, and we cannot rule out that some portion of the signal reflects concentrated informed trading in specific markets or that the tariff period amplified an underlying pattern that would be smaller in a calmer policy environment. The finding may not generalize beyond this political context. The bias we document is, however, large enough, consistent enough and symmetric enough to be meaningful. Prediction markets are as good as the diversity of the participants inside them. When that diversity is limited, the price reflects a worldview rather than a probability and this distinction matters for how the output of these markets should be interpreted.

5.7 Limits of Informational Efficiency in Polarized Environments

Prediction markets have been understood as efficient information aggregation mechanisms, where prices incorporate the dispersed knowledge of many participants [39]. Under the Efficient Market Hypothesis, persistent abnormal returns should be difficult to obtain as publicly available information is rapidly incorporated into prices through competitive trading [7]. The results of this study suggest that this process may become less reliable in politically polarized environments where market participants are not ideologically neutral.

The results of the study do not appear to arise due to limited information. The markets tracked throughout the observation period revolved around highly public political developments, including tariffs, executive actions and diplomatic negotiations. Information regarding these events was broadly available to market participants. This may suggest that the observed pric-

ing differences are attributable to asymmetric interpretation of information rather than unequal access to it.

This interpretation aligns with research on motivated reasoning and identity-protective cognition. Kahan [13] argues that individuals often interpret evidence in ways that reinforce pre-existing political identity rather than purely objective evaluation. In prediction markets, these biases may become embedded directly into prices if a sufficiently large number of participants share similar political alignment. Market prices may therefore partially reflect collective political conviction in addition to probabilistic expectations.

Our findings bring to light an important limitation of informational efficiency in decentralized political markets. Prediction market theory assumes that diverse participants with differing beliefs collectively discipline prices through arbitrage and competition [39]. When participation is concentrated within a politically distinct subgroup, the diversity required for efficient aggregation weakens. In the context of Polymarket, participation requires familiarity with cryptocurrency infrastructure, which became increasingly associated with pro-Trump political alignment during the recent election cycle. If the participant pool was systematically skewed in this direction, pro-Trump outcomes may have faced persistent excess demand, leading prices to drift away from informational efficiency.

A potential criticism of the findings is that the observed asymmetry may be driven primarily by smaller markets. Thinly traded markets are generally expected to be less efficient, making persistent mispricing more likely to emerge. To examine this possibility, the prospective period sample was divided into large and small volume cohorts using the median market volume as the split threshold. The anti-Trump strategy generated positive returns in both cohorts. The return on invested capital was higher among large volume markets, where the strategy produced an ROI of 4.69%, compared to 2.64% in small volume markets. This finding is noteworthy because larger markets typically attract more participants, greater liquidity and more opportunities for informed traders to correct mispricing. As a result, it may be expected that directional asymmetries diminish as market activity increases. The persistence of the signal in larger markets is consistent with the possibility that the observed pricing dynamics reflect broader participant behavior rather than isolated inefficiencies within niche markets.

While the results do not imply that prediction markets are inefficient at aggregating dispersed information, they do suggest that informational efficiency may be conditional on the composition of market participants and the political context in which the market operates. In this context, the prices within politically polarized markets may incorporate political sentiment and identity-driven trading, rather than solely probabilistic beliefs about future outcomes.

The findings of this study raise questions for how prediction market prices should be interpreted in politically polarized climates, particularly if market participation is concentrated among ideologically aligned individuals.

5.8 Distinguishing Political Mispricing from Favourite–Longshot Bias

A natural alternative explanation for the anti-Trump strategy’s returns is the *favourite–longshot bias*, the well-documented tendency in betting and prediction markets for low-probability contracts to be overpriced and high-probability contracts to be underpriced [30, 35]. Because the anti-Trump position is on average the higher-priced side of the contracts in our sample (a mean entry price near \$0.65, against \$0.35 for the pro-Trump side), one might suspect that the strategy simply harvests this generic bias rather than any politically driven mispricing.

The structure of the mispricing across the price distribution argues against this. The favourite–longshot bias is largest at the extremes: deep longshots and heavy favourites show the biggest distortions, while contracts priced near \$0.50 show almost none. When we partition the anti-Trump positions by entry price and measure the realised mispricing in each bucket — the gap between the realised resolution frequency and the entry price on markets that resolved — the pattern runs in the opposite direction (Figure 9). The mispricing does not peak at the favourite extreme. It peaks among mid-priced, contested contracts with entry prices of roughly \$0.35–0.55, which resolved in the anti-Trump side’s favour around 65–90% of the time, far more often than their prices implied. Among the heaviest anti-Trump favourites (entry prices above \$0.90), the realised edge is only a few cents per contract. A

pattern that is weakest exactly where the favourite–longshot bias predicts its strongest effect, and largest where that bias predicts near-zero distortion, cannot be explained by favourite–longshot pricing alone. It points instead toward a directional component tied to political orientation in how these contracts were priced. We note that these bucket-level results are computed on the subset of markets that resolved within the observation window and are descriptive rather than formal hypothesis tests.




Figure 9: Anti-Trump positions bucketed by entry price. Left: realised win rate versus entry price (points above the 45° line are underpriced; labels give the number of resolved trades). Right: realised and mark-to-market return on investment by bucket. The mispricing peaks in the contested mid-range, not at the favourite extreme.

5.9 Robustness to Topic Concentration

A further concern is that the aggregate result could be driven by a single cluster of correlated markets rather than a broad pricing pattern — most plausibly the tariff and trade contracts, which experienced a sequence of adverse legal and policy developments for the administration during the sample window. To assess this, we classified all 218 tracked markets into thematic categories based on their question text and decomposed the anti-Trump strategy’s profit and loss by category (Figure 10).

The returns are distributed across many distinct themes rather than con-

centrated in one. Personnel and cabinet-tenure markets contributed roughly 28% of total profit, macroeconomic and Federal Reserve markets 27%, tariff and trade markets 18%, immigration 11%, and a range of further themes the remainder. Excluding the tariff and trade markets entirely, the strategy remains clearly profitable (approximately \$1,160 compared to \$1,410 for the full sample), confirming that no single topic drives the result. Topic-level figures in this subsection are valued at a later data pull than the portfolio results in Section 4.2, which are valued at 1 May 2026. Several markets resolved in the interim, raising the full-sample total to approximately \$1,410. The decomposition is therefore indicative of where returns are concentrated rather than directly comparable to the headline figures.

Two qualifications are worth noting. First, several themes contain near-duplicate contracts — for example, the same official appearing in multiple tenure markets, or a single quantitative threshold posed at several cutoff dates — so the number of statistically independent observations is smaller than the raw count of 218 markets suggests. Second, the strategy is not profitable in every theme. Markets concerning prospective federal equity stakes in named companies produced a combined loss of approximately \$290. This reflects both the genuine difficulty of forecasting these outcomes and the ambiguity involved in assigning a clear pro-Trump direction to them — arguably the weakest application of the classification rule in the dataset. The breadth of the effect across the remaining themes indicates that the mispricing we document is a general feature of Trump-related contracts during this period rather than an artefact of a single unlucky topic.

6 Conclusion

This paper investigated whether systematic betting on anti-Trump outcomes in Polymarket reveals persistent political bias in market pricing during the observation period from July 2025 to May 2026.

The results show that the anti-Trump trading strategy consistently outperformed both the pro-Trump strategy and the neutral benchmark. As well, the pro-Trump strategy persistently underperformed in comparison to the benchmark strategy.




Figure 10: Anti-Trump profit and loss decomposed by market topic (all positions marked to final price). The edge is broad across themes; it survives excluding tariff markets entirely, and the federal equity-stakes cluster is a net loss.

The one-sample t-tests were not able to reject the null hypothesis at conventional significance levels. However, the dual equity curve, benchmark percentile analysis and trend regressions revealed substantial directional divergence and asymmetry across the observation period.

The findings of this simulation are consistent with the hypothesis that in Trump-related contracts, the pro-Trump outcome has been systematically overvalued, thereby creating positive returns with the approach of an anti-Trump betting strategy.

This paper highlights an important limitation of informational efficiency in highly polarized political markets. Prediction markets are considered objective and efficient aggregators of dispersed information and expectations. Although, their accuracy ultimately depends on the diversity of the market participants driving prices. If participation becomes concentrated among

ideologically aligned market actors, prices may begin to incorporate political sentiment and identity-driven conviction, in addition to probabilistic belief. Understanding how participant composition shapes market behavior is becoming increasingly important as block-chain based prediction markets continue to develop in political and economic relevance.

References

- [1] Joyce E. Berg, Forrest D. Nelson, and Thomas A. Rietz. Accuracy and forecast standard error of prediction markets, 2003.
- [2] Joyce E. Berg and Thomas A. Rietz. Market design, manipulation, and accuracy in political prediction markets: Lessons from the iowa electronic markets. *PS: Political Science & Politics*, pages 293–296, 2014.
- [3] Patrick Buckley. Blockchain based prediction markets. *The Journal of Prediction Markets*, pages 73–86, 2024.
- [4] Centre Consortium. USD Coin (USDC): A Digital Dollar, 2018. Available at: <https://f.hubspotusercontent30.net/hubfs/9304636/PDF/centre-whitepaper.pdf>. Accessed: May 2026.
- [5] James N. Druckman, Erik Peterson, and Rune Slothuus. How elite partisan polarization affects public opinion formation. *American Political Science Review*, 107(1):57–79, 2013.
- [6] Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. Chapman and Hall, 1993.
- [7] Eugene F. Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, pages 383–417, 1970.
- [8] Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *The American Economic Review*, pages 393–408, 1980.
- [9] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, pages 107–119, 2003.
- [10] Friedrich A. Hayek. The use of knowledge in society. *The American Economic Review*, pages 519–530, 1945.
- [11] Carlos M. Jarque and Anil K. Bera. A test for normality of observations and regression residuals. *International Statistical Review*, 55(2):163–172, 1987.

- [12] Wall Street Journal. The mysterious trader behind massive pro-trump polymarket bets, 2024. Accessed May 2026.
- [13] Dan M. Kahan. Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, 8(4):407–424, 2013.
- [14] Greta M. Ljung and George E. P. Box. On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303, 1978.
- [15] Carmen López-Martín, Sonia Benito Muela, and Raquel Arguedas. Efficiency in cryptocurrency markets: new evidence. *Eurasian Economic Review*, pages 403–431, 2021.
- [16] Marc LeGrasso Mary Becker, Zachary McGurk. Partisan sentiment and returns from online political betting markets in the 2020 us presidential election. *Scottish Journal of Political Economy*, 2025.
- [17] J.B. Maverick. S&P 500 average returns and historical performance. *Investopedia*, May 14 2026.
- [18] Sofoluwe Mayowa. Is polymarket legal in the u.s. in 2026? *GamblingInsider*, January 2026.
- [19] Joe Nocera. How a french whale made \$85 million off trump’s win. The Free Press, November 2024.
- [20] Alexander Osipovich. A mystery \$30 million wave of pro-trump bets has moved a popular prediction market. The Wall Street Journal, October 2024.
- [21] Polymarket. 2024 presidential election market, 2024. Accessed May 2026.
- [22] Polymarket. What is Polymarket? Polymarket Help Center, 2026. Available at: <https://help.polymarket.com/en/articles/13364060-what-is-polymarket>. Accessed: May 2026.
- [23] Marco Quiroz-Gutierrez. The french whale who made massive pro-trump bets is only doing it for the money. Fortune, November 2024.
- [24] David Rothschild. Forecasting elections comparing prediction markets, polls, and their biases. *Working Paper*, 2009.

- [25] Saman Saama. Political bias in prediction markets: Evidence from the u.s. 2024 presidential elections. Master’s thesis, Master of Science in Economics and Business Administration, 2025. Master’s Thesis.
- [26] Said E. Said and David A. Dickey. Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika*, 71(3):599–607, 1984.
- [27] Manya Saini. Polymarket secures investment from trump jr-backed 1789 capital. *Reuters*, August 2025.
- [28] Samuel S. Shapiro and Martin B. Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3–4):591–611, 1965.
- [29] William F. Sharpe. The sharpe ratio. *The Journal of Portfolio Management*, 21(1):49–58, 1994.
- [30] Erik Snowberg and Justin Wolfers. Explaining the favorite–longshot bias: Is it risk-love or misweighting? *Journal of Political Economy*, 118(4):723–746, 2010.
- [31] Frank A. Sortino and Robert van der Meer. On the use of downside risk in practice. *Journal of Portfolio Management*, 17(4):35–42, 1991.
- [32] Meir Statman. A behavioral framework for dollar-cost averaging. *Journal of Portfolio Management*, 22(1):70–78, 1995.
- [33] Student. The probable error of a mean. *Biometrika*, 6(1):1–25, 1908.
- [34] Charles S. Taber and Milton Lodge. Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3):755–769, 2006.
- [35] Richard H. Thaler and William T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.
- [36] Fairleigh Dickinson University. Fdu poll finds trump’s embrace of crypto pays off, 2024. Accessed May 2026.
- [37] Wikipedia contributors. Polymarket. Wikipedia, The Free Encyclopedia, 2026. Accessed: April 2026.

- [38] Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.
- [39] Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.
- [40] Jeffrey M. Wooldridge. *Introductory Econometrics: A Modern Approach*. South-Western Cengage Learning, 5th edition, 2012.