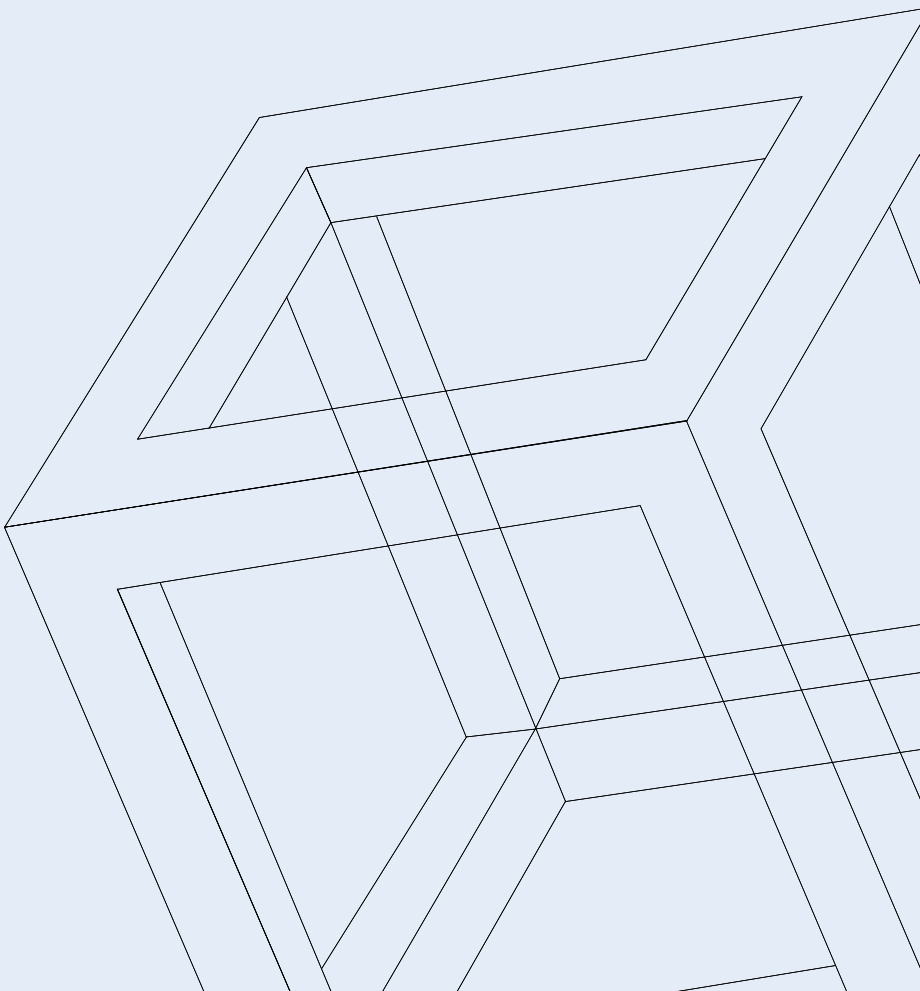
brn

Can Prediction Market Odds Predict Asset Prices?

Testing Polymarket as a Trading Signal for Crypto



May 2026

Table of Contents

1) Executive summary 05

2) Introduction 07

3) Methodology 10

4) Findings 12

5) What This Means for Institutional Allocators 15

6) Risks, Limitations, and Counterarguments 17

7) Conclusion 20

1) Executive summary

Prediction markets have emerged as one of the most-watched data sources in digital asset markets, with Polymarket alone processing over $9 billion in cumulative volume by early 2026. Yet most institutional crypto allocators have no systematic framework for using prediction market odds as a trading signal. This report fills that gap. Using 420 markets, 819 days of crypto price data, and three distinct analytical frameworks, we test whether Polymarket odds carry actionable information for crypto asset pricing.

The answer is yes, but the signal lives in a different place than most participants would expect.

What We Found

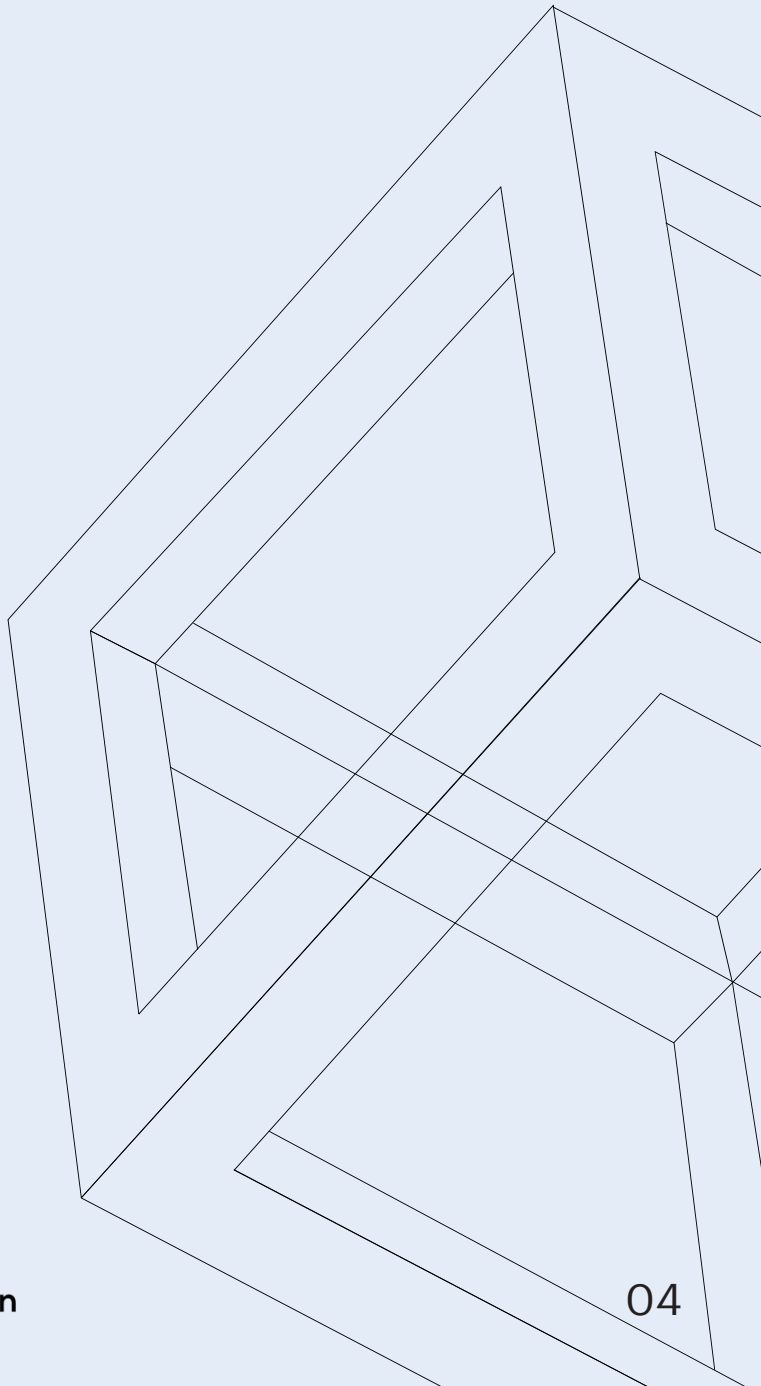
The most valuable signal is in monetary policy, not crypto-specific events. Federal Reserve rate cut odds on Polymarket Granger-cause BTC returns at a one-day lag ($F=3.18$, $p=0.076$), while BTC returns do not Granger-cause Fed odds. When rate cut probability crosses 50%, a systematic long-BTC strategy produces a Sharpe ratio of 1.12 over 11 trades. This is the single strongest and most actionable finding in the study.

Crypto-specific policy markets follow prices rather than leading them. The Bitcoin strategic reserve market, the most closely watched crypto-policy market of 2025, shows overwhelming BTC→Odds Granger causality ($F=16.6$, $p<0.001$) with no evidence of the reverse. The 2024 US presidential election market ($1.5 billion in volume) shows no lead-lag relationship with BTC in either direction. ETF approval markets show a positive but statistically insignificant same-day reaction. For allocators, this means Polymarket's crypto-specific policy markets are sentiment gauges, not leading indicators.

Price-target markets are well-calibrated but work as contrarian sell signals. Across 230 resolved BTC and ETH price-target markets, Polymarket odds achieve a Brier Score of 0.0455, well below the 0.25 coin-flip baseline. However, buying BTC when the crowd is most confident in a price target produces deeply negative returns: a Sharpe of –2.50 at the 80% threshold. The crowd is accurate on direction but consistently late, reaching peak confidence near local price tops.

What This Means for Allocators

For allocators making tactical crypto decisions, the primary Polymarket signal to monitor is Fed rate expectation markets, not crypto-native ones. For risk managers, price-target market calibration data can be used to construct probability-weighted scenarios with better accuracy than subjective estimates. For all participants, the contrarian signal from price-target markets provides a quantified framework for identifying when consensus is most likely to be wrong.



2) Introduction

2.1 Market Context

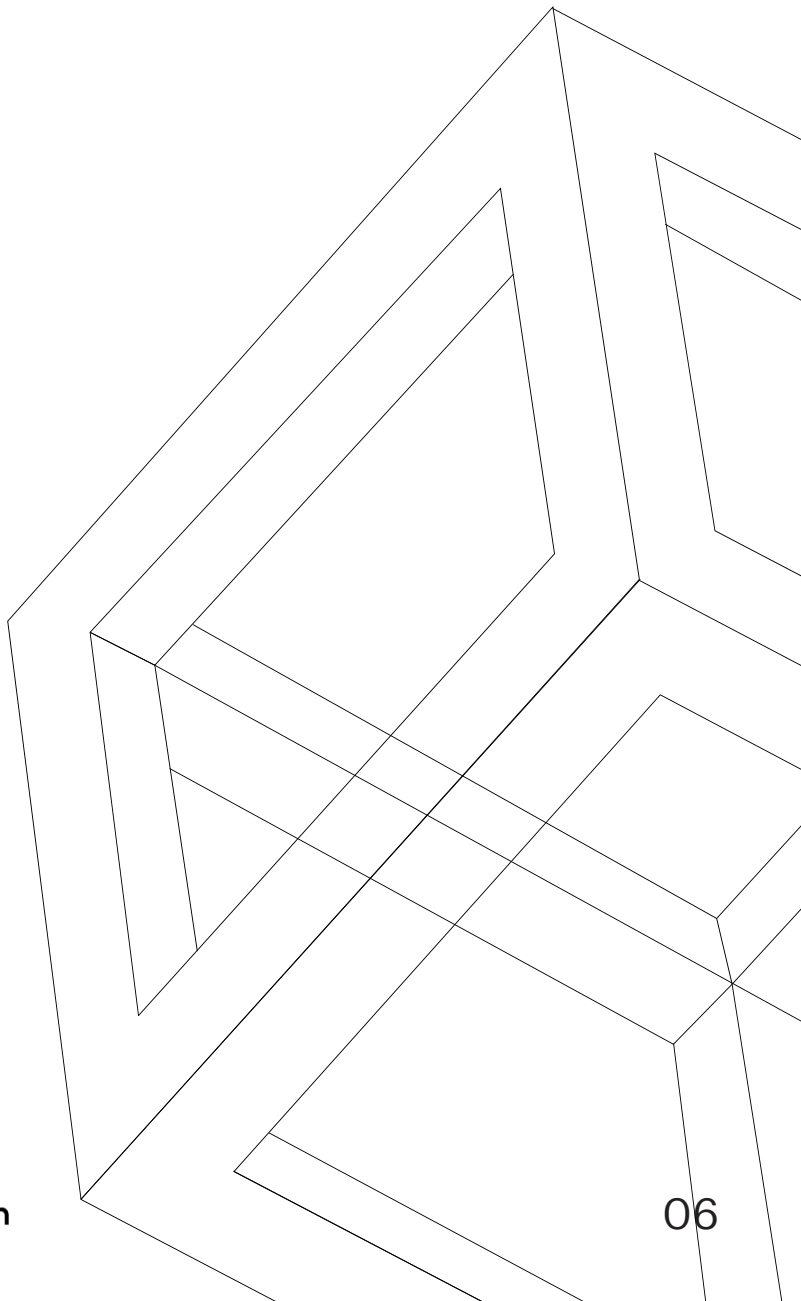
Polymarket's growth trajectory from approximately $50 million in cumulative volume before 2024 to over $9 billion by early 2026 has transformed prediction markets from a niche curiosity into a data source that institutional investors cannot afford to ignore. Crypto assets are disproportionately exposed to binary political and regulatory catalysts (the January 2024 Bitcoin ETF approval, the May 2024 Ethereum ETF approval, and the 2025 US Strategic Bitcoin Reserve executive order are the clearest examples) and prediction markets are designed precisely to price such events. Yet most institutional allocators continue to rely on options implied volatility, funding rates, and sentiment indices for event-risk gauging, leaving prediction market data largely untapped.

This study tests whether that gap represents a missed opportunity. Two conceptually distinct mechanisms could make prediction market odds informative for crypto prices. The first is an event resolution signal: policy and regulatory markets price binary outcomes such as ETF approvals and legislative milestones, and if these odds shift before the outcome is fully priced in crypto spot markets, they constitute a leading indicator. The second is a price-path signal: price-target markets aggregate crowd beliefs about whether BTC will reach a specific price by a specific date, and if these implied probabilities are well-calibrated against realized outcomes, they constitute a reliable directional signal. The two mechanisms are tested separately throughout this report.

2.2 Research Questions and Hypotheses

The study is organized around four research questions. First, do changes in Polymarket odds for crypto-relevant policy events lead or lag crypto spot price movements, and over what time horizon? Second, are Polymarket price-target market odds well-calibrated against realized outcomes? Third, how quickly is prediction market information absorbed into crypto prices? Fourth, can a systematic, rules-based strategy exploiting either mechanism generate risk-adjusted returns?

Four hypotheses follow. H1 predicts that policy market odds changes Granger-cause crypto spot price changes at one-to-two day lags, but not the reverse. H2 predicts that price-target odds are reasonably well-calibrated, within plus or minus 15 percentage points of historical frequencies, but with variation by strike distance and liquidity. H3 predicts that signal persistence is short, with information absorbed within one to three days. H4 predicts that a threshold-based trading rule generates positive returns on a small-sample basis but with Sharpe ratios below 1.0 and high event-concentration risk.



3) Methodology

3.1 Conceptual Framework

The study sits at the intersection of three bodies of literature. The prediction market efficiency literature, anchored by Arrow et al. (2008) and Manski (2006), establishes that liquid prediction markets tend to be well-calibrated but that the interpretation of market prices as probabilities requires careful attention to risk premia and marginal trader composition. The crypto market microstructure literature documents that price discovery in digital asset markets is informationally fragmented, occurring unevenly across spot, perpetuals, and options venues. This fragmentation creates windows where external signals can carry incremental information. The cross-market price discovery literature asks whether prediction markets function as a venue for informed traders who subsequently act in the underlying asset, analogous to the well-documented phenomenon of options markets leading spot equities around earnings announcements.

The testable implication is straightforward: if informed participants trade prediction markets first and crypto spot or derivatives second, we should observe odds changes leading price changes. If prediction markets simply reflect what spot markets already know, we should observe the reverse. Both possibilities are tested formally.

3.2 Data and Sample

The dataset draws on four free, publicly accessible sources. Market metadata and resolution outcomes come from the Polymarket Gamma API. Historical daily odds are reconstructed from the Polymarket CLOB API using the /prices-history endpoint with daily fidelity, achieving a 100% success rate across 420 markets. Crypto spot prices for BTC, ETH, and SOL (819 daily candles from January 2024 through March 2026) and BTC perpetual funding rates come from the Binance public API. The Fear and Greed Index (819 daily readings) comes from the Alternative.me API. All data is free and requires no authentication.

The sample is organized into two tracks. Track A covers policy and regulatory events: 10 crypto-specific markets (Bitcoin ETF, Ethereum ETF, Solana ETF, two strategic reserve markets, three US election and

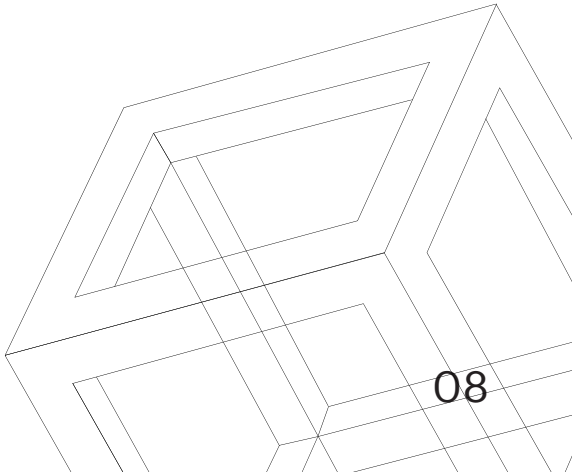
inauguration markets, and two China-unban-Bitcoin markets) plus 116 Federal Reserve rate decision markets across approximately 15 FOMC meetings. Track B covers 3,060 price-target markets for BTC and ETH, of which 2,552 are closed and 300 have daily CLOB odds history. The sample period spans January 2024 through March 2026, capturing two distinct crypto regimes (pre- and post-ETF), the 2024 US election cycle, and the 2025 regulatory shift under the new US administration.

3.3 Analytical Methods

Study 1 (Event Study): For each Track A event, we identify all single-day odds movements exceeding 10 percentage points and compute cumulative abnormal returns for BTC, ETH, and SOL in windows from t=0 through t+5 days after each shock. Abnormal return is computed as the raw log return. We supplement the event study with pairwise Granger causality tests (lags 1-5) between daily odds changes and crypto returns for the highest-volume events, controlling for BTC funding rates and the Fear and Greed Index.

Study 2 (Calibration): We group resolved Track B markets into probability bins based on the implied odds seven days before expiry and compute the observed YES resolution frequency within each bin. A well-calibrated market should lie close to the 45-degree perfect-calibration line. We compute Brier Scores across the full sample and by sub-group (BTC-only, high-volume, different time horizons). A critical data-quality step validates CLOB token mapping: we cross-reference each market's outcomePrices at resolution against the last CLOB token price, identifying and correcting 13 inverted tokens (4.9% of the sample).

Study 3 (Signal Backtest): We define threshold-based trading rules for each track. For Track A, go long BTC when a Fed rate cut market's YES odds cross above a given threshold; exit at event resolution date. For Track B, go long BTC when a price-target market's odds cross above a threshold; exit at expiry. We report average return, annualized Sharpe ratio, win rate, and number of trades across threshold levels of 50%, 60%, 70%, and 80%.



4) Findings

4.1 Event Study Results (Track A)

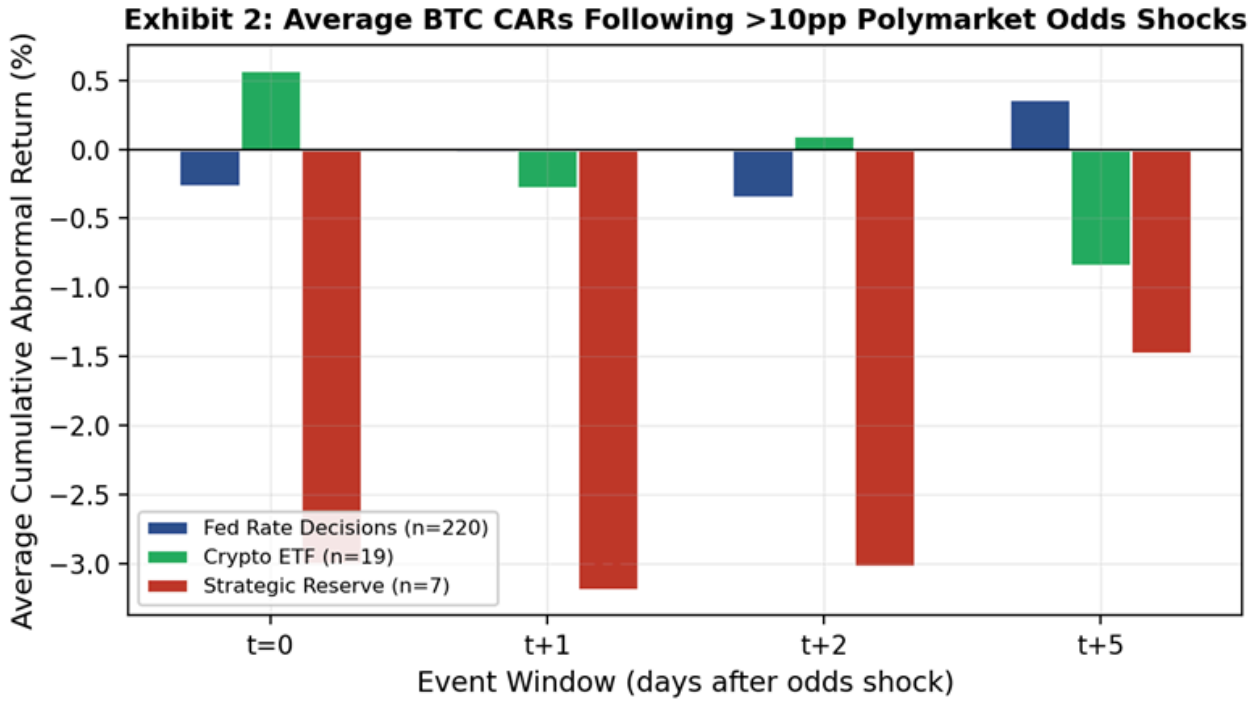
Exhibit 1 summarizes the sample coverage across both tracks. Exhibit 2 presents the average cumulative abnormal returns for BTC following large odds shocks in each event category.

Exhibit 1: Sample Coverage and Key Characteristics

Market / Event	Track	Volume	Resolution	Outcome	N Shocks	CLOB Pts	Granger
Bitcoin ETF (Jan 2024)	A	$28M	1/15/2024	YES	4	40	—
Ethereum ETF (May 2024)	A	$12M	5/31/2024	YES	15	136	—
Solana ETF (Jul 2025)	A	$1.9M	7/31/2025	YES	0	226	—
BTC Reserve – 100 Days	A	$8M	4/30/2025	YES	4	175	BTC → Odds
BTC Reserve – 2025	A	$15M	12/31/2025	Open	3	367	BTC → Odds
Trump 2024 Election	A	$1,532M	11/5/2024	YES	0	307	None
Fed Rate Decisions (116)	A	varies	various	various	220	varies	Odds → BTC*
BTC/ETH Price Targets	B	>$1M ea.	various	various	—	300	—

** Fed rate cut odds Granger-cause BTC returns at lag 1 (F=3.18, p=0.076).*
Source: Polymarket Gamma API, CLOB API, Binance.

[1] 'BTC Reserve – 2025' carries a nominal end date of 2025-12-31 but shows outcome 'Open' because, at the time of data collection (March 2026), Polymarket had not issued a final resolution. President Trump signed an executive order on March 6, 2025 establishing a Strategic Bitcoin Reserve using approximately 198,000 BTC sourced from existing government seizures. Whether this executive action satisfied the market's resolution criteria remained under review by Polymarket's resolution committee. The outcome will be updated upon final resolution.



*Source: Polymarket CLOB API, Binance. * p<0.10, ** p<0.05.*

The event study reveals a sharp divergence between macro and crypto-specific policy signals. Federal Reserve rate decision markets generate 220 odds shocks exceeding 10 percentage points across the sample period, but the average BTC cumulative abnormal return is negligible at all horizons (ranging from –0.35% at t+2 to +0.36% at t+5). None are statistically significant. This suggests that large single-day Fed odds movements are too noisy, or are absorbed too quickly, to produce a detectable multi-day price effect in the event study framework.

However, the Granger causality analysis tells a different story. For Fed rate cut markets, odds changes do Granger-cause BTC returns at a one-day lag ($F=3.18$, $p=0.076$), while BTC returns do not Granger-cause odds. This is the study's most important finding for Track A: the predictive signal from Fed markets is real, but it operates through a persistent trend channel rather than discrete shocks. When rate cut expectations are gradually building on Polymarket, BTC tends to follow with a one-day lag. This is a signal that rewards trend-following positioning, not event-driven trading.

Crypto ETF markets (Bitcoin, Ethereum, and Solana ETF approval) produce 19 shocks with a positive but statistically insignificant same-day CAR of +0.57 % (t=1.04). There is no persistent signal beyond the event day. The 2024 US presidential election, despite $1.5 billion in Polymarket volume and a widely discussed crypto narrative, shows no Granger causality with BTC in either direction at any lag tested (all p-values exceeding 0.33). The election's impact on crypto was regime-level (a structural repricing over weeks, not a signal captured by daily lead-lag analysis).

4.1.1 Strategic Reserve Deep-Dive

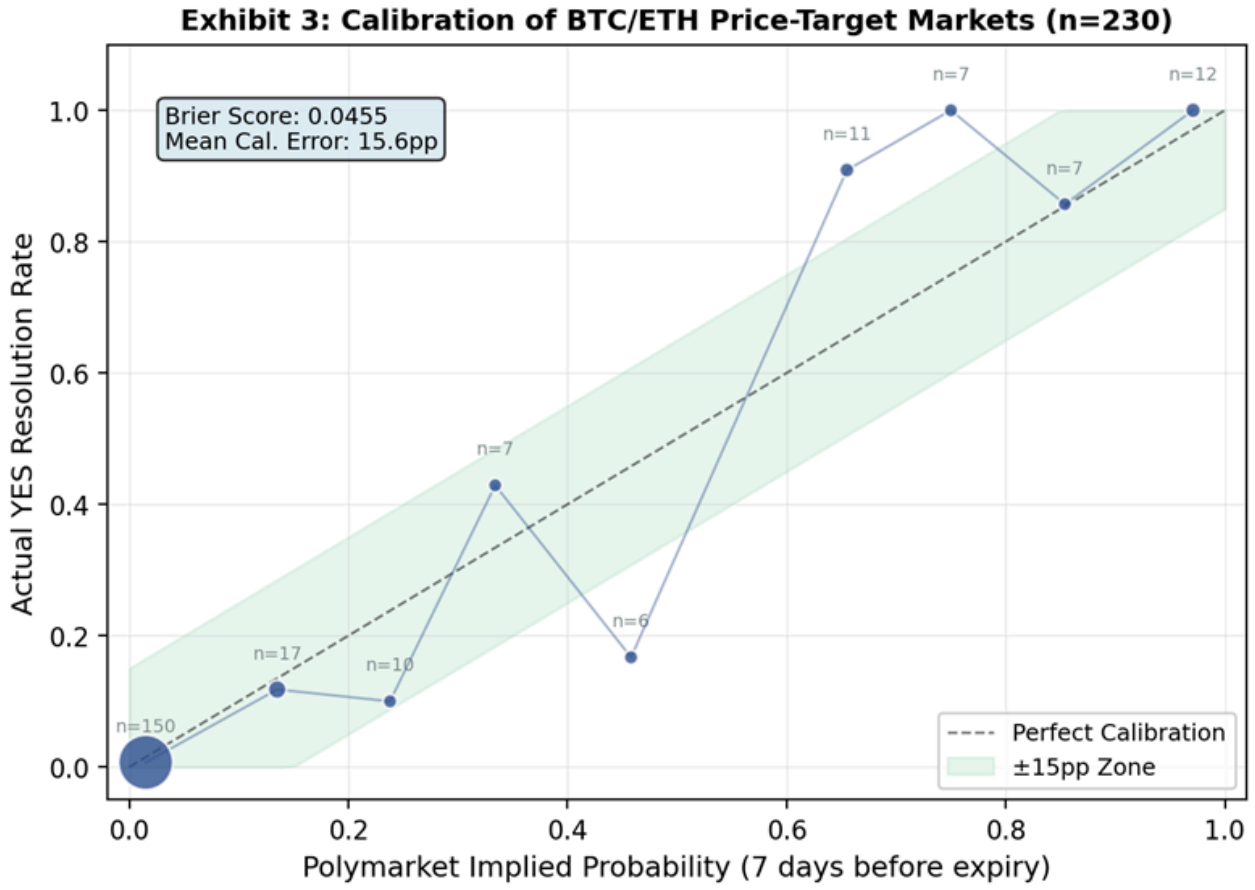
The strategic reserve markets warrant separate discussion because they produced the study's most counterintuitive preliminary result: statistically significant negative CARs at t=0 and t+1 following positive odds shocks. This appeared to suggest that higher reserve probability was bearish for BTC, the opposite of the intuitive relationship.

Investigation resolves the anomaly. Both reserve markets show overwhelming BTC→Odds Granger causality ($F=16.4–16.6$, $p<0.001$ at all tested lags) with zero evidence of Odds→BTC causality. The lagged correlation from BTC(t–1) to Odds(t) is +0.21 to +0.31, while the reverse direction is near zero. The key context is timing: during February and March 2025, BTC fell over 20% (from $100,636 to a low of $78,596) in a broad market correction, while reserve odds actually rose from 56% to 66.5% following Trump's executive order announcement on March 6. The reserve market was responding to political developments while BTC was falling for unrelated macro reasons. The negative CARs reflect coincident timing, not causation. For allocators, the implication is clear: the strategic Bitcoin reserve market is a lagging indicator of BTC sentiment, not a leading one. It has no predictive value for crypto prices.

4.2 Calibration Results (Track B)

Exhibit 3 plots the calibration of Polymarket price-target markets, comparing implied probability (seven days before expiry) against the actual YES resolution rate. A critical data-quality fix was required before this analysis: 13 of 266 markets (4.%) had inverted CLOB token IDs, where the tracked token mapped to the NO outcome rather than YES.

To understand why this matters, note that each binary Polymarket contract has two separate on-chain tokens (one redeemable for $1 if the event resolves YES, one redeemable for $1 if it resolves NO) and the API does not guarantee which token is listed first. The inversion was detected by comparing two data points at resolution: the final CLOB price for the tracked token (which converges to $1.00 for the winning side and $0.00 for the losing side) against the market's recorded YES price from the Gamma API's outcomePrices field. Where these two figures diverged (i.e., the tracked token was priced near $1.00 at resolution yet the market had resolved NO) the token was flagged as inverted and its time series replaced with (1 - price). After correction, the results are sensible and informative.



Source: Polymarket CLOB API. Token mapping validated against Gamma API outcomePrices. Bubble size proportional to sample count.

The extreme bins are well-calibrated. Markets priced below 10% odds resolved YES only 0.7% of the time (n=150), closely matching the predicted 1.5%. Markets priced above 90% resolved YES 100% of the time (n=12), matching the predicted 97.1%. The tails work as they should, which is the most practical finding for allocators who use prediction market data to construct extreme-scenario probabilities.

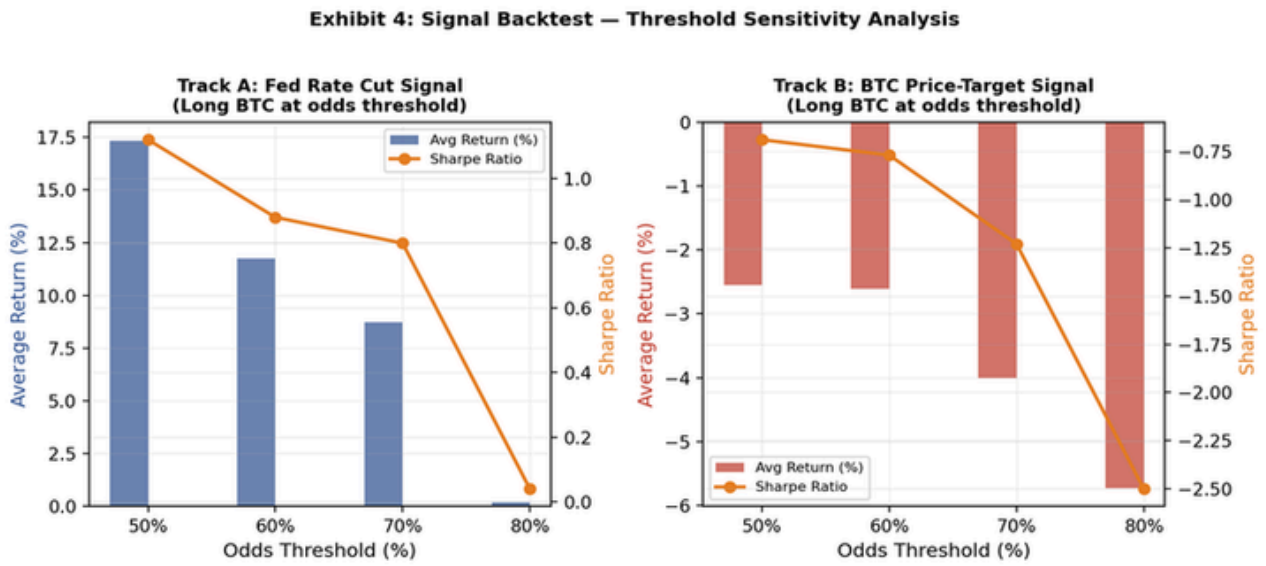
The middle bins show moderate but informative deviation. The 60–70% bin overperforms significantly: 90.9% of markets in this range resolved YES (n=11), against a predicted 65.5%. The 70–80% bin similarly overperforms at 100% actual versus 75% predicted (n=7). This pattern suggests the crowd is systematically under-confident in near-the-money BTC price targets, when Polymarket says there is a 65% chance BTC will reach a target, the historical frequency is closer to 90%. The 20–30% and 40–50% bins underperform, though sample sizes are small (10 and 6 observations respectively).

The overall Brier Score of 0.0455 is well below the 0.25 coin-flip baseline, confirming that Polymarket price-target odds contain genuine predictive information. Calibration degrades gracefully at longer horizons: the 30-day-before snapshot achieves a Brier Score of 0.0602, and the midpoint-of-market-life snapshot achieves 0.1269. BTC-only calibration (n=167, Brier 0.0514) is consistent with the full sample, confirming the pattern holds for the most liquid markets.

One methodological point deserves emphasis. The 4.9% token inversion rate is low in absolute terms but was sufficient to completely break naive calibration analysis. Any allocator or quant team attempting to build systematic signals from Polymarket data must validate CLOB token mapping before use. This is a practical barrier to entry that limits arbitrage and may partially explain why the calibration patterns we observe have not been fully priced away.

4.3 Signal Backtest

Exhibit 4 presents the threshold sensitivity analysis for both Track A and Track B trading rules.



Source: Polymarket CLOB API, Binance. Track A: Fed rate cut 25bps decrease markets. Track B: BTC reach/hit/above markets (n=148, token-mapping-corrected).

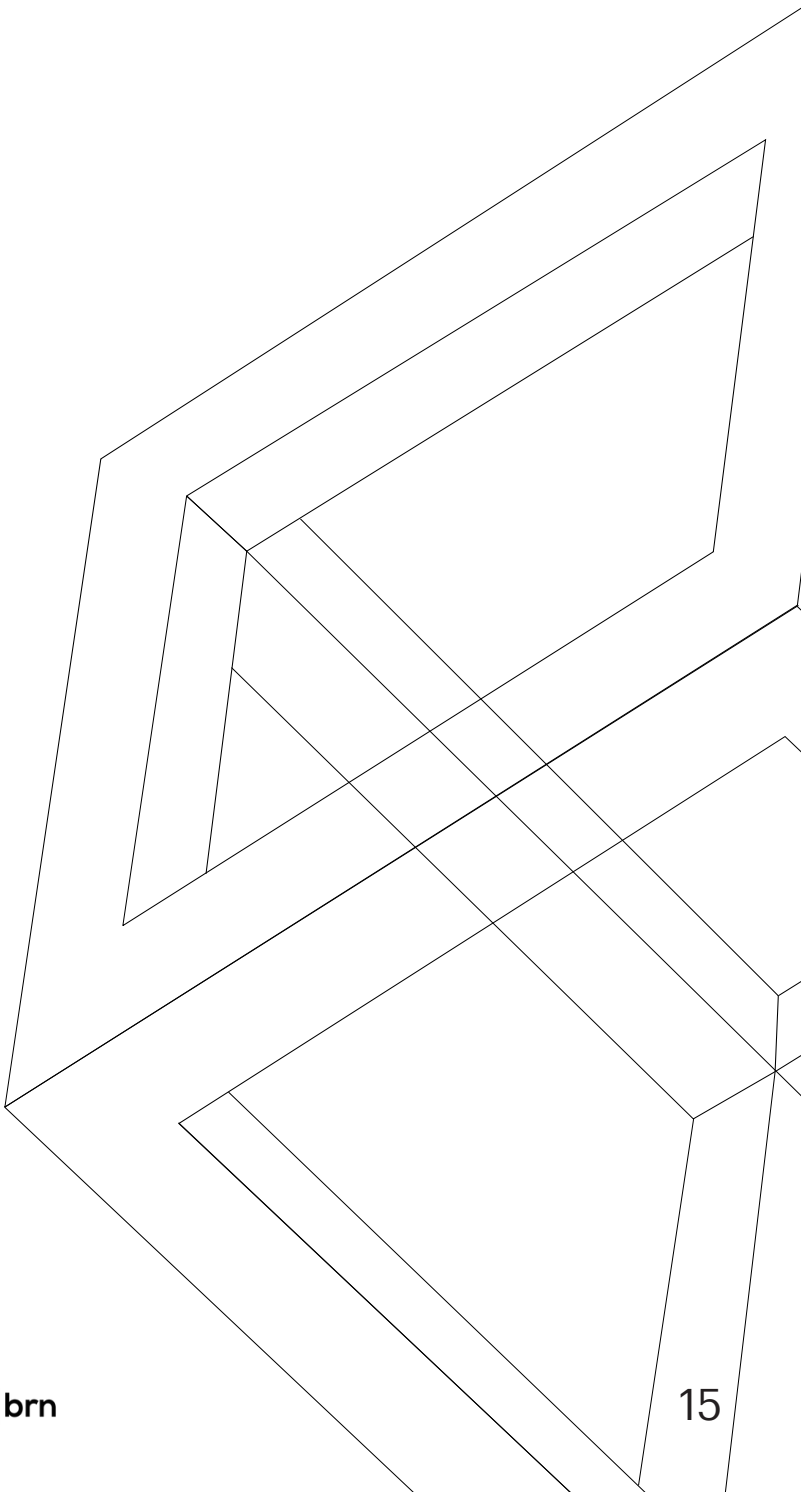
The Track A backtest produces the study's strongest risk-adjusted performance. Entering a long BTC position when Fed rate cut odds cross 50% yields an average return of +17.34% over 11 trades with a Sharpe ratio of 1.12 and a 72.7% win rate. Performance declines monotonically at higher thresholds: the 60% threshold yields a Sharpe of 0.88, the 70% threshold yields 0.80, and the 80% threshold collapses to 0.04. This pattern is consistent with the Granger causality finding: the signal's value is in early positioning when the rate cut is uncertain, not when it is priced as near-certain. By the time odds reach 80%, the information is fully absorbed into BTC prices.

Two important caveats apply. First, the 2024/2025 BTC bull market heavily influences the Track A returns; the strategy benefits from a favorable secular trend. Second, the sample of 11 trades is small, and the results should be interpreted as directional evidence rather than a statistically definitive backtest. Nevertheless, the alignment between the Granger causality result (which is model-based and independent of the backtest) and the threshold-based strategy (which is purely mechanical) reinforces the conclusion that a real signal exists.

The Track B backtest delivers the study's most surprising finding. Buying BTC when price-target market odds cross above any threshold produces negative returns at every level tested. The 50% threshold yields -2.55% average return with a Sharpe of -0.69 over 70 trades. The 80% threshold yields -5.73% with a Sharpe of -2.50 over 37 trades. The win rate drops from 34.3% at the 50% threshold to just 16.2% at 80%. This result survived the

token mapping fix with near-identical figures (only two of the 148 markets in this subset were inverted).

The contrarian interpretation is clear: price-target markets peak in confidence near local BTC highs. When the Polymarket crowd is most convinced that BTC will reach a target, the asset is most likely to underperform from that point forward. This makes high price-target odds better suited as a sell signal than a buy signal. The monotonic decline in performance across thresholds is particularly striking: it is not noise but a systematic, threshold-dependent pattern that holds across 70 independent trades.

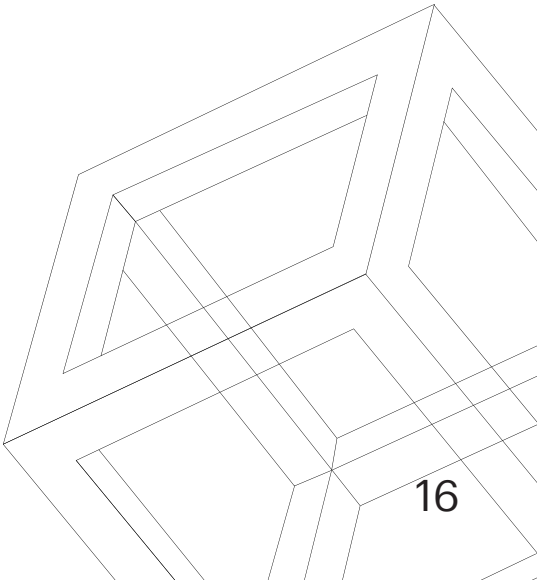


5) What This Means for Institutional Allocators

For long-only allocators: The primary Polymarket signal to incorporate into a monitoring framework is the implied probability of Fed rate cuts. When rate cut odds for the next FOMC meeting cross 50%, this has historically preceded BTC appreciation with a Sharpe of 1.12. Allocators operating on a monthly or quarterly rebalancing cycle should treat this threshold as a conviction input for overweighting BTC exposure ahead of accommodative monetary shifts. Crypto-specific policy markets (reserve, ETF, election) should not be used as leading indicators; they reflect crypto prices rather than driving them.

For tactical allocators and hedge funds: The contrarian signal from Track B price-target markets provides a quantified sell-side filter. When odds that BTC will reach a given price target exceed 70–80%, historical performance indicates a high probability of negative forward returns (average -4% to -6%, win rate between 16% and 26%). This can be used as a position-sizing reduction trigger or a stop-loss calibration input. The underlying mechanism (crowd confidence peaking near local tops) is consistent with behavioral finance findings on sentiment-driven overshoot.

For risk managers: Polymarket price-target odds are sufficiently well-calibrated (Brier Score 0.0455) to serve as probability inputs for scenario-weighted VaR models. At the extremes, calibration is tight: sub-10% odds map to sub-1% realized frequency, and above-90% odds map to 100%. Middle bins require a haircut of approximately 15 percentage points. This provides a more objective alternative to subjective probability assessments for binary regulatory or macroeconomic scenarios, with the caveat that token mapping must be validated for any programmatic implementation.



6) Risks, Limitations, and Counterarguments

Risk	Severity	Probability	Positioning Implication
Polymarket liquidity is thin and susceptible to manipulation	High	Medium	Apply a minimum $1M volume filter; weight signals by liquidity
US access restrictions limit participant pool	High	Low to Medium	Cross-validate against Kalshi where overlapping markets exist
Small crypto-specific event sample (n=10)	Medium	High	Frame as directional evidence; Fed sample (n=116) carries the quantitative weight
Regime dependency: 2024/25 bull market may not generalize	Medium	Medium	Sub-period analysis recommended; flag as caveat
Token mapping errors in CLOB API data	Medium	Confirmed	4.9% inversion rate; must validate before programmatic use
Macro controls limited to crypto-native metrics	Low to Medium	Confirmed	BTC funding rate and Fear & Greed sufficient for crypto-focused study

Acknowledged counterargument: A rational-expectations critic would argue that sophisticated participants already arbitrage prediction markets against crypto derivatives, eliminating any persistent edge. The response is that such arbitrage requires simultaneous access to Polygon-based Polymarket, centralized crypto exchanges, and derivatives venues, a multi-leg operation with execution risk, cross-chain bridging delays, and jurisdictional complexity that limits how quickly the gap closes. This microstructural friction is precisely what creates the signal window tested in Study 1. The 4.9% token mapping error rate adds a further practical barrier: without careful data validation, even well-resourced quant teams cannot reliably construct the signals documented here.

7) Conclusion

Polymarket odds carry genuine, actionable information for crypto asset pricing, but the signal mechanism is not where most market participants would look. The most reliable predictive channel runs through monetary policy expectations, not crypto-specific event outcomes. Crypto-policy markets (including the strategic Bitcoin reserve, ETF approvals, and the 2024 US presidential election) either follow BTC prices or show no statistically significant lead-lag relationship. Price-target markets, meanwhile, are well-calibrated as probability estimates but perversely effective as contrarian sell signals, with crowd confidence peaking reliably near local price tops.

These findings invert the popular narrative around prediction markets and crypto. The tool is useful, but not for the reasons most commonly assumed. Three specific recommendations follow from the evidence.

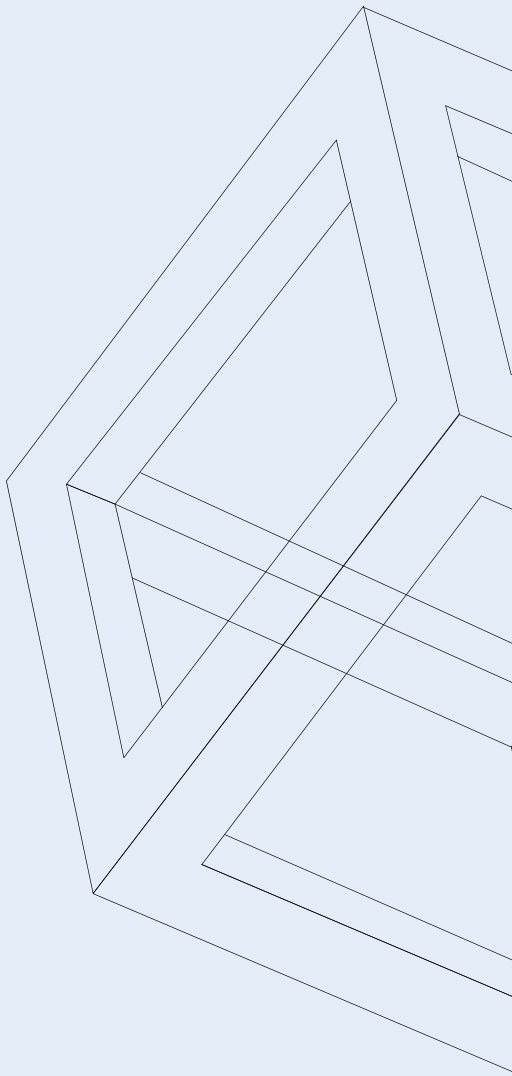
First, institutional allocators should incorporate Polymarket Fed rate cut odds into their pre-FOMC position review. When implied probability of a rate cut crosses 50%, this has historically corresponded to BTC outperformance with a Sharpe of 1.12 over the subsequent holding period. Treat a sustained move above 50% as a conviction input for directional crypto exposure.

Second, risk managers should use Polymarket price-target odds as probability inputs for scenario analysis, applying the calibration data from this study. At the tails (below 10% and above 90%), the odds can be used at face value. In the middle bins, apply a ±15 percentage point uncertainty band. This provides a more objective, market-derived alternative to subjective probability estimates for binary crypto events.

Third, tactical allocators should monitor BTC price-target market confidence as a contrarian indicator. When odds for a near-term BTC price target exceed 70%, this study's evidence suggests the crowd is reaching peak confidence near a local top. This signal should trigger a review of position sizing, stop-loss levels, and hedging activity, not a new long entry.

The next generation of crypto event risk is regulatory: stablecoin legislation, CFTC jurisdiction expansion, strategic reserve accumulation targets, and cross-border settlement frameworks. Prediction markets will price it before

traditional indicators do. The allocators who build systematic frameworks for interpreting these signals today will have a structural informational advantage as the market matures.

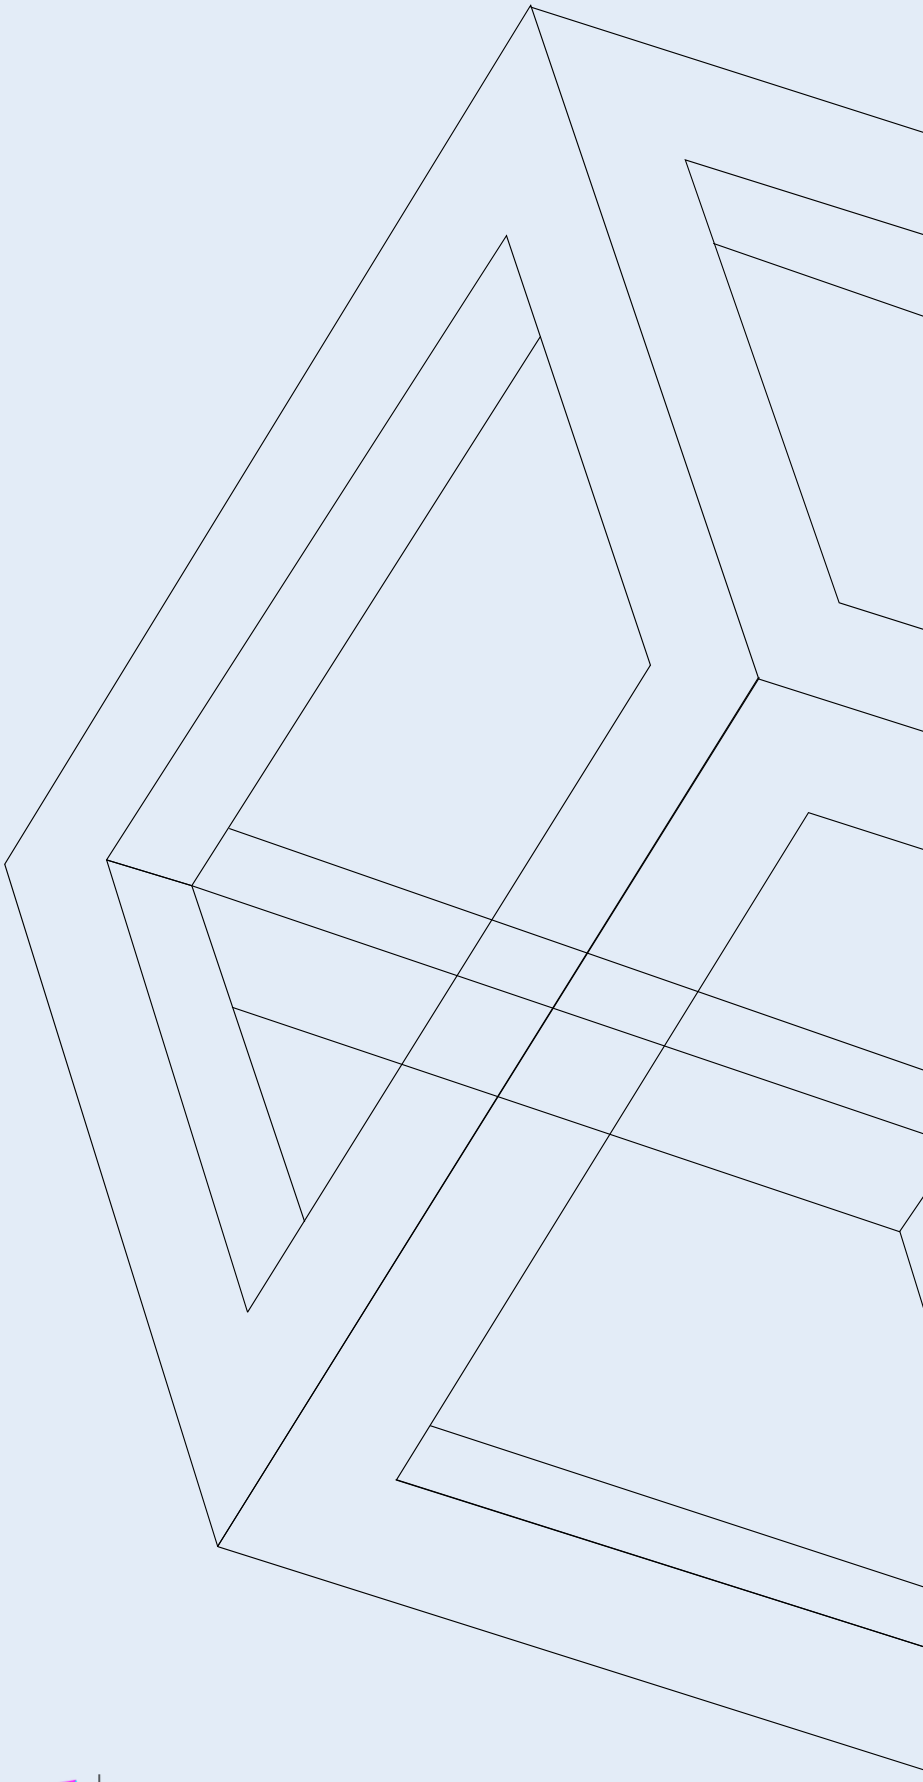


Data Sources

Polymarket Gamma API (market metadata, resolution outcomes) • Polymarket CLOB API (historical daily odds, /prices-history endpoint) • Binance Public API (BTC/ETH/SOL daily OHLCV, BTC perpetual funding rates) • Alternative.me API (Fear & Greed Index) • All data publicly accessible, no authentication required. Sample period: January 2024 – March 2026.

Methodological References

Arrow, K., Forsythe, R., et al. (2008). The Promise of Prediction Markets. Science. • Manski, C. (2006). Interpreting the Predictions of Prediction Markets. Economics Letters. • Granger, C. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. Econometrica.



brn

27 Ann Siang Road, Singapore | 069707
www.thebrn.co
hello@thebrn.co