



UNIVERSITÉ DU QUÉBEC EN OUTAOUAIS
DÉPARTEMENT DES SCIENCES ADMINISTRATIVES

WORKING PAPER NO. 8

Prediction Markets as Arrow–Debreu Securities: Evidence from Polymarket

*State Prices, No-Arbitrage Bounds and Macroeconomic News in Blockchain-Based
Contingent Claims*

Simon-Pierre Boucher

Département des sciences administratives
Université du Québec en Outaouais

simon-pierre.boucher@uqo.ca

Gatineau – Pavillon Alexandre-Taché
283, boulevard Alexandre-Taché
Gatineau, Québec, Canada J9A 1L8

October 2026

Version 1.1

Abstract

Negative-risk events on Polymarket list one binary contract for every outcome of a decision and settle exactly one of them at one dollar, so that their prices are observable analogues of Arrow–Debreu state prices. We use the five-outcome market on the June 2026 FOMC decision (146 daily closes, December 2025 to May 2026) and the seventeen-contract market on the 2024 U.S. presidential election (304 daily closes) to ask how the Arrow–Debreu pricing restrictions fare in a blockchain market. Because spanning is built into the contract design, the testable content is the law of one price for the complete set of claims, whose executable no-arbitrage bounds we derive from the platform’s minting and conversion mechanics. The sum of state prices averages 1.017 (Fed) and 1.015 (election); the overround is statistically positive but of the order of the quoted half-spreads of the low-priced legs, declines steadily towards resolution, reverts within two days, and is bracketed by the bid- and ask-side sums of the order book. It has the wrong sign to be explained by the time value of money, which places the complete set below one. The normalised prices define a risk-neutral distribution over the FOMC decision whose mean—a Polymarket analogue of the futures-implied expected policy change—moves two to five times more, in variance terms, on FOMC and employment-report days than on other days, while weekend trading carries about one third of weekday variance. Once Polymarket closes are aligned to the U.S. trading calendar and lag orders are chosen by information criteria, daily changes in this measure neither lead nor lag the ten-year Treasury yield, the S&P 500, the VIX or Bitcoin. We document how a one-day misalignment of blockchain time stamps and a minimum- p lag search generated spurious lead–lag findings in an earlier version.

Keywords: prediction markets, Arrow–Debreu securities, state prices, law of one price, monetary policy expectations, macroeconomic announcements, Polymarket

JEL Classification: G13, G14, E52, D47, C58

1. Introduction

A prediction-market contract that pays one dollar if the Federal Open Market Committee (FOMC) leaves its target range unchanged at a given meeting, and nothing otherwise, is the textbook example of a state-contingent claim. When a platform lists one such contract for every possible decision of that meeting, and when its settlement rules guarantee that exactly one of them pays, the set of contracts spans the outcome space of the decision in the sense of [Arrow \(1964\)](#) and [Debreu \(1959\)](#). Blockchain-based prediction markets have made this construction routine: Polymarket, the largest venue by volume, groups mutually exclusive contracts into “negative-risk” events whose collateral mechanics enforce that the complete set of claims is worth exactly one dollar at resolution. The resulting prices are therefore observable analogues of Arrow–Debreu state prices for a small number of macroeconomic and political state spaces, traded continuously, with a public order book and public settlement.

This paper asks what can and cannot be learned from such prices. The question has two parts. The first is a question about pricing. In an Arrow–Debreu economy the price of the complete set of state claims equals the price of a riskless bond, and each state price is the product of a discount factor and a risk-neutral probability. Because the spanning property is built into the contract design, the empirically testable content of the framework is not completeness itself but the pricing restrictions it imposes: the law of one price for the complete set, the no-arbitrage bounds implied by the mechanics of minting, converting and redeeming claims, and the interpretation of normalised prices as a risk-neutral distribution rather than as physical probabilities ([Manski, 2006](#); [Wolfers and Zitzewitz, 2006a](#)). The second is a question about information. If normalised state prices form a risk-neutral distribution over FOMC decisions, their first moment is a market-implied expected policy change comparable, in spirit, to the federal funds futures-based measures used since [Kuttner \(2001\)](#). One can then ask whether this quantity moves on scheduled macroeconomic releases, whether weekend trading—a feature absent from the futures market—carries information, and whether its daily changes lead or lag the Treasury, equity and volatility markets.

We address both parts with daily data from two Polymarket negative-risk groups: the five-outcome market on the June 2026 FOMC decision (11 December 2025 to 22 May 2026, 146 daily closes, about \$38 million of cumulative volume) and the seventeen-contract market on the winner of the 2024 U.S. presidential election (4 January to 5 November 2024, 304 daily closes, about \$3.7 billion of volume). The Fed group is the primary laboratory because its state space is small, exhaustive and economically interpretable; the election group serves to check whether the pricing regularities hold in a much larger and more heavily traded state space.

Three sets of results emerge. First, the law of one price holds for the complete set of claims up to a small and economically interpretable overround. The sum of the five Fed state prices

averages 1.017 (standard deviation 0.026) and the sum of the seventeen election prices averages 1.015 (standard deviation 0.031). Both overruns are statistically different from zero, but they are of the order of the quoted spreads, decline steadily as resolution approaches—from about 5 percent in the first month of the Fed market to 0.4 percent in the last—and revert within two days. At the order-book snapshot taken with the data, the sum of the best bids across the five Fed contracts is 0.999 and the sum of the best asks is 1.004, so that no round trip through the complete set was profitable at the quotes. We also show that the overrun has the wrong sign to be explained by the time value of money: a complete set bought and held to resolution should trade *below* one by the discount factor, which averaged 0.989 over the Fed sample, so that the executable bounds are asymmetric, and the data sit at or slightly above the upper bound rather than at the lower one. Second, the risk-neutral expected policy change implied by the Fed group moves significantly more on scheduled release days than on other days. Its absolute daily change is 1.15 basis points on the thirteen release days in the sample against 0.51 basis points on the other 132 days (variance ratio 4.9, Brown–Forsythe $p = 0.008$), and the “No change” state price moves 5.5 against 1.4 percentage points (variance ratio 14.2). The effect is concentrated on FOMC days and employment-report days; the four CPI releases in the sample produced no detectable excess movement. Weekend days, when the futures market is closed, carry about one third of the variance of weekdays. Third, once Polymarket closes are aligned to the U.S. trading calendar and the lag order of the vector autoregression is chosen by information criteria rather than by searching over lags, there is no evidence that daily changes in the Polymarket-implied policy expectation lead or lag daily changes in the ten-year Treasury yield, the S&P 500, the VIX or Bitcoin. The series are contemporaneously correlated with the Treasury yield (correlation 0.26), consistent with both responding to the same macroeconomic news within the day, but neither predicts the other.

The paper makes three contributions. The first is conceptual. The prediction-market literature has long debated whether prices are probabilities ([Manski, 2006](#); [Wolfers and Zitzewitz, 2006a](#); [Snowberg et al., 2013](#); [Ottaviani and Sørensen, 2015](#)); we show that for negative-risk groups the question can be separated into a pricing restriction that is testable (the complete set is worth one dollar within executable bounds) and an interpretation that is not (normalised prices are risk-neutral, not physical, probabilities), and we derive the bounds from the actual settlement mechanics rather than from a generic “transaction cost” argument. The second is empirical. To our knowledge this is the first study to document the term structure of the overrun in a blockchain prediction market, its relation to quoted depth, and the announcement-day behaviour of a market-implied FOMC expectation extracted from binary contracts. The third is methodological and cautionary. Polymarket’s daily price history is stamped at 00:00 UTC of the following day, so that a naive merge with U.S. closing prices misaligns the two calendars by one day in one direction only; and a lag-search procedure that selects the lag with the smallest p -value produces spurious “Granger

causality” in short samples. An earlier version of this paper reported both artefacts as findings; we document them in Appendix C because they are likely to affect other work that merges blockchain and conventional market data.

The results also delimit what prediction markets can contribute to the measurement of policy expectations. The Polymarket June 2026 market traded through four FOMC meetings, five CPI releases and five employment reports, and its implied expected policy change moved from -15.5 basis points to zero over the sample. The two largest daily moves—a fall of 11.5 percentage points in the “No change” price on 28 January 2026 and a rise of 23 percentage points on 18 March 2026—both occurred on FOMC days. These are the responses one would expect from an instrument that prices the policy path; they are not evidence of informational superiority over the futures market, and the Granger tests give none. The appropriate reading is that Polymarket offers a continuously traded, publicly settled complement to federal funds futures with a transparent and small transaction-cost wedge, not a substitute with a lead.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and the related literature. Section 3 develops the framework and states the hypotheses. Section 4 describes the data and the alignment of calendars. Section 5 presents the methods. Section 6 reports the results for the Fed market. Section 7 extends the pricing analysis to the election market and reports further checks. Section 8 discusses limitations and concludes. Appendix A reports the volatility models, Appendix B defines the variables, and Appendix C reconciles the present results with the earlier version.

2. Institutional Setting and Related Literature

2.1. Polymarket and negative-risk markets

Polymarket is a prediction-market platform on the Polygon blockchain. Each binary market issues a pair of outcome tokens, YES and NO, through a conditional-token contract: one unit of the USDC stablecoin can be split at any time into one YES and one NO token, and a YES–NO pair can be merged back into one USDC. At resolution the winning token redeems for one USDC and the losing token for nothing. Orders are matched off-chain on a central limit order book operated by the platform and settled on-chain; the order book, trade history and settlement are public. Resolution is delegated to the UMA optimistic oracle: anyone may propose an outcome by posting a bond, the proposal can be disputed during a challenge window, and disputed resolutions are decided by a vote of UMA token holders.¹

Multi-outcome questions are organised as *negative-risk* events. The outcomes of such an event

¹Polymarket documentation, <https://docs.polymarket.com> (consulted 17 October 2026). The documentation describes a proposal bond of about \$750, a two-hour challenge window, and a dispute process of four to six days.

are mutually exclusive by construction, and the platform’s adapter contract allows a NO token on any one outcome to be converted at any time into one YES token on each of the other outcomes. The conversion embodies the accounting identity that betting against one outcome is equivalent to betting on all the others. For an event with N outcomes, it implies that a trader holding one NO token on every outcome holds a claim worth $N - 1$ dollars at resolution—exactly one of the N NO tokens expires worthless—and the adapter lets this position be unwound before resolution. The event studied in this paper, “Fed decision in June?”, lists five outcomes defined on the change in the upper bound of the target federal funds range decided at the June 2026 FOMC meeting: a decrease of 50 basis points or more, a decrease of 25 basis points, no change, an increase of 25 basis points, and an increase of 50 basis points or more. These five outcomes are exhaustive, so the five YES tokens form a complete set of state claims on the decision. The same structure applies to the “Presidential Election Winner 2024” event with seventeen candidate contracts, where the last contract (“Other”) completes the state space.

Two features of the setting matter for the interpretation of prices. First, prices are quoted in USDC and USDC held as collateral earns no interest on the platform, so that the time value of money is not neutralised by the settlement design as it is in a margined futures contract. Second, the fee structure differs across markets: the metadata of the five Fed contracts report no trading fee at the time of the snapshot, whereas the platform introduced taker fees on short-dated cryptocurrency markets in early 2026 and announced an extension to other categories. We treat the fee on the markets studied as zero and discuss this assumption in Section 6.

2.2. Prices, probabilities and risk

The interpretation of prediction-market prices as probabilities is the central theme of the literature surveyed by [Wolfers and Zitzewitz \(2004\)](#), [Wolfers and Zitzewitz \(2006b\)](#) and [Snowberg et al. \(2013\)](#). [Manski \(2006\)](#) showed that, with risk-neutral traders who hold heterogeneous beliefs and face budget constraints, the equilibrium price of a binary contract need not equal the mean belief and can lie anywhere in an interval determined by the distribution of beliefs. [Wolfers and Zitzewitz \(2006a\)](#) replied that with logarithmic utility the price equals the wealth-weighted mean belief, and that for plausible preference parameters the gap between price and mean belief is small except near the extremes. [Ottaviani and Sørensen \(2015\)](#) derived the dynamics of such prices under heterogeneous beliefs and wealth effects and showed that they underreact to news, producing momentum and reversal. The empirical calibration literature finds that prices are informative but not perfectly calibrated: long-run evidence from the Iowa Electronic Markets ([Berg et al., 2008](#)), the favourite–longshot pattern documented by [Page and Clemen \(2013\)](#), and the comparison of markets and polls in [Rothschild \(2009\)](#) all point to systematic but modest deviations. [Rhode and Strumpf \(2004\)](#) show that the large historical betting markets on U.S. presidential elections behaved in much

the same way a century ago.

The present paper takes a different route to the same question. For an event whose states are correlated with aggregate wealth—a monetary-policy decision is the leading example—the price of each state claim is $\psi_j = \delta q_j$ where δ is a discount factor and q_j a risk-neutral probability. Dividing by the sum of state prices recovers q_j exactly, whatever the risk attitudes of traders, but the physical probability remains unidentified without further assumptions. [Piazzesi and Swanson \(2008\)](#) document economically large and time-varying risk premia in federal funds futures, the natural benchmark for our Fed contracts; we therefore interpret our normalised prices as risk-neutral quantities throughout and treat the recovery of physical probabilities as outside the scope of the data.

2.3. Monetary-policy expectations and macroeconomic news

A large literature uses market prices to measure monetary-policy expectations and their revisions around announcements. [Kuttner \(2001\)](#) extracts policy surprises from federal funds futures; [Gürkaynak et al. \(2005\)](#) show that FOMC statements move expectations through a second “path” factor beyond the current-rate surprise; [Bernanke and Kuttner \(2005\)](#) and [Rigobon and Sack \(2004\)](#) estimate the response of equity prices to those surprises; [Nakamura and Steinsson \(2018\)](#) document that announcements also convey information about the economy. The announcement-day variance approach, in which the arrival of scheduled information is identified by comparing the variance of price changes on release days with that on other days, goes back to [Ederington and Lee \(1993\)](#) for interest-rate futures and to [Andersen et al. \(2003\)](#) for high-frequency exchange rates, and the distinction between calendar time and trading time in the arrival of information to [French and Roll \(1986\)](#). Our announcement-day analysis of the Polymarket-implied expected policy change follows this tradition with daily data, exploiting the fact that the platform trades on weekends.

2.4. Price discovery across venues

When the same asset trades in several markets, the contribution of each venue to the common efficient price can be measured by the information share of [Hasbrouck \(1995\)](#) or the common-factor weights of [Gonzalo and Granger \(1995\)](#); [Baillie et al. \(2002\)](#) and [Yan and Zivot \(2010\)](#) discuss the relation between the two and the conditions under which either is informative. Both measures require the prices to be cointegrated with a known cointegrating vector, which is the case for the same security on two exchanges or for a futures contract and its underlying index ([Hasbrouck, 2003](#)). It is not the case for a binary contract on an FOMC decision and the ten-year Treasury yield or an equity index: the contract price is bounded in $[0, 1]$, converges to zero or one, and is related to conventional asset prices through a non-linear and time-varying mapping. We therefore

do not compute information shares and rely instead on Granger-causality tests in first differences (Granger, 1969; Sims, 1980; Lütkepohl, 2005), with the limitations discussed in Section 5. The cross-venue comparison of modern prediction markets themselves has recently been taken up by Ng et al. (2025), who compare Polymarket, Kalshi and PredictIt around the 2024 election and find that the more liquid venue leads, and by Chernov et al. (2025), who combine state-level election contracts with polls to estimate the comovement of voter preferences. Reichenbach and Walther (2025) study calibration and trader skill on Polymarket. None of these papers examines the pricing of complete state spaces or the Fed decision markets.

2.5. Microstructure of blockchain markets

The microstructure concepts we use are standard: the adverse-selection view of spreads (Glosten and Milgrom, 1985; Kyle, 1985), the serial-covariance spread estimator of Roll (1984), and the surveys of O’Hara (1995) and Madhavan (2000). Two strands of the cryptocurrency literature are directly relevant. Makarov and Schoar (2020) document large and persistent arbitrage deviations across cryptocurrency exchanges, attributable to capital controls and settlement frictions; the complete-set arbitrage on a negative-risk event is simpler, because all legs settle on one platform in one collateral, and it is therefore a cleaner test of whether blockchain markets enforce the law of one price. Cong et al. (2023) show that reported volumes on unregulated cryptocurrency exchanges are heavily inflated by wash trading; we do not use volume as an input to any test, but the point cautions against reading Polymarket’s reported volumes as measures of genuine participation. Harvey et al. (2021) provide an overview of the decentralised-finance infrastructure on which the platform is built. Experimental and theoretical work on manipulation (Hanson et al., 2006; Hanson, 2003) and the documented episodes on Polymarket itself—concentrated position-taking in the 2024 election market and a disputed oracle resolution in March 2025—are discussed in Section 8.

3. Conceptual Framework

3.1. State claims and the price of the complete set

Consider a single decision with a finite set of mutually exclusive and exhaustive outcomes $\mathcal{S} = \{s_1, \dots, s_N\}$, resolved at a known date T . A state claim on s_j pays one unit of account at T if s_j obtains and nothing otherwise; its payoff is $a_j(s) = \mathbf{1}\{s = s_j\}$. Let $p_{j,t}$ be the price of the claim at date $t < T$ and $p_t = (p_{1,t}, \dots, p_{N,t})'$.

Definition 1 (Spanning). *The market spans $\mathcal{S}$ if every payoff $c : \mathcal{S} \rightarrow \mathbb{R}$ can be replicated by a portfolio of traded claims. With one state claim per outcome the payoff matrix is the identity and spanning holds by construction.*

Spanning is therefore not an empirical property of a negative-risk group: the platform lists one claim per outcome and its settlement rules guarantee that exactly one claim pays. What the data can test is how the complete set is priced.

Proposition 1 (Absence of arbitrage). *If there is no arbitrage, there exist a strictly positive state-price vector $\psi_t \in \mathbb{R}_{++}^N$ and, equivalently, a discount factor $\delta_t \in (0, 1]$ and a probability measure $\mathbb{Q}_t$ on $\mathcal{S}$ such that $p_{j,t} = \psi_{j,t} = \delta_t q_{j,t}$ for every j . The price of the complete set is the price of a riskless claim to one unit of account at T :*

$$\Sigma_t \equiv \sum_{j=1}^N p_{j,t} = \delta_t = \frac{1}{(1 + r_t)^{\tau_t}}, \quad (1)$$

where r_t is the riskless rate and $\tau_t = T - t$ the remaining time to resolution.

Equation (1) is the Arrow–Debreu restriction tested in this paper. Its two ingredients deserve separate statements. Positivity, $p_{j,t} > 0$, is required for $\mathbb{Q}_t$ to be equivalent to the physical measure; on an order book with a minimum tick it is satisfied mechanically whenever a claim is quoted, so that it is a necessary condition without empirical content. The restriction with content is on the level of Σ_t .

3.2. Executable bounds

In a frictionless market $\Sigma_t = \delta_t$. On Polymarket the complete set can be traded in two directions with different settlement timing, which makes the no-arbitrage band asymmetric. If $\Sigma_t < \delta_t$, a trader buys one unit of each claim for Σ_t and holds the set to resolution, when exactly one claim pays one unit; the profit is $1 - \Sigma_t$ at T , or $\delta_t - \Sigma_t$ in present value, and the capital is tied up for τ_t . If $\Sigma_t > 1$, a trader splits one unit of collateral into N YES–NO pairs (one per outcome), sells the N YES tokens for Σ_t , and converts the N NO tokens through the negative-risk adapter into $N - 1$ units of collateral immediately; the round trip costs one unit and returns $\Sigma_t + N - 1$, a riskless profit of $\Sigma_t - 1$ realised at t rather than at T . Allowing for a cost c_t of executing the N legs at the quotes, the executable bounds are

$$\delta_t - c_t \leq \Sigma_t \leq 1 + c_t. \quad (2)$$

The lower bound depends on the riskless rate and on the horizon, the upper bound does not. With $r_t \approx 3.6$ percent and τ_t up to half a year, δ_t ranges from 0.98 to 1 over the Fed sample, so that the band is about two percentage points wide at the start of the sample and collapses to the width of the spread costs at the end. The transaction-cost term c_t is naturally measured by the sum across claims of half the quoted spread: for five contracts quoted at a tick of 0.001 with one-tick spreads, $c_t = 0.0025$, and we use 0.005 as a conservative value in the tables.

A positive overround, $\Sigma_t > 1$, therefore cannot be explained by the time value of money, which pushes Σ_t below one. It can be explained by the convention that the daily price is the last trade or midpoint of each contract: when four of five contracts trade at a few tenths of a cent with one-tick spreads, the sum of midpoints exceeds the sum of bids by $N/2$ ticks, and the last trades of thinly traded claims tend to occur on the ask side. It can also reflect a genuine but small premium for the convenience of the conversion mechanism, or, in the first days of a market, quotes seeded by the platform before trading begins.

3.3. The risk-neutral distribution and the expected policy change

Normalising prices by their sum removes the discount factor:

$$q_{j,t} = \frac{p_{j,t}}{\Sigma_t}, \quad \sum_{j=1}^N q_{j,t} = 1. \quad (3)$$

For the Fed group, assign to each state the policy change $x_j \in \{-50, -25, 0, 25, 50\}$ basis points, with the two open-ended states valued at ± 50 . The risk-neutral expected policy change and its dispersion are

$$\mathbb{E}_t^{\mathbb{Q}}[\Delta r] = \sum_{j=1}^N q_{j,t} x_j, \quad \text{SD}_t^{\mathbb{Q}}[\Delta r] = \left(\sum_{j=1}^N q_{j,t} (x_j - \mathbb{E}_t^{\mathbb{Q}}[\Delta r])^2 \right)^{1/2}. \quad (4)$$

$\mathbb{E}_t^{\mathbb{Q}}[\Delta r]$ is the Polymarket analogue of the futures-implied expected rate change used in the monetary-policy literature; like it, it embeds a risk premium (Piazzesi and Swanson, 2008). Under $\mathbb{Q}_t$ each $p_{j,t}$ is a martingale up to the discount factor, so that daily changes in $\mathbb{E}_t^{\mathbb{Q}}[\Delta r]$ should be unpredictable from past information and their variance should reflect the arrival of news. This is the property exploited in the announcement-day analysis.

3.4. Hypotheses

The framework yields the following hypotheses, stated so that each can be rejected by the data.

Hypothesis 1 (Law of one price for the complete set). *Σ_t lies within the executable bounds (2) on most days; deviations are small relative to quoted spreads, decline as the lower bound rises towards one, and are transitory (stationary with a half-life of a few days).*

Hypothesis 2 (Information on scheduled release days). *Absolute daily changes in $\mathbb{E}_t^{\mathbb{Q}}[\Delta r]$ and in the modal state price are larger on days of scheduled FOMC statements, CPI releases and employment reports than on other days; changes on weekend days, when conventional markets are closed, are smaller than on weekdays.*

Hypothesis 3 (Lead–lag with conventional markets). *Daily changes in $\mathbb{E}_t^Q[\Delta r]$ neither Granger-cause nor are Granger-caused by daily changes in the ten-year Treasury yield, the S&P 500, the VIX or Bitcoin, once the series are aligned on the same trading calendar.*

Hypothesis 3 is stated as a null of no predictability because it is the prediction of informational efficiency in both markets; a rejection in either direction would be the interesting finding, and we show in Appendix C how easily spurious rejections arise. A fourth hypothesis on the conditional volatility of the modal state price, which played a prominent role in the earlier version of this paper, is examined in Appendix A.

3.5. Caveats

The framework is a single-period abstraction applied to a market that trades for six months. Treating each daily close as a snapshot of the state-price vector is innocuous for the pricing restriction, which must hold at every date, but it means that the risk-neutral distribution is conditional on the information of that day and that its evolution mixes news with changes in risk premia. The finite state space is natural for an FOMC decision and for an election winner; it would not be for a continuous outcome. The “ ± 50 ” valuation of the open-ended states is a convention that matters little while those states carry a few percent of risk-neutral mass, as they do here. Finally, the mapping from the contract to the economic event depends on the resolution rules and on the oracle; the possibility of a disputed resolution adds a basis risk that we cannot measure in these two markets but discuss in Section 8.

4. Data

4.1. Sources and sample

Polymarket data were retrieved on 22 May 2026 from the platform’s public Gamma API (event and market metadata) and CLOB API (price histories and order books). For each contract we obtain the daily price history returned by the prices-history end-point at daily fidelity, which reports one price per calendar day, and for the five active Fed contracts a snapshot of the full order book at 08:51 UTC on 22 May 2026. Macroeconomic series come from FRED: the effective federal funds rate (DFF, daily), used for the discount factor, together with the CPI and unemployment series used to date releases. Daily closing prices of the SPDR S&P 500 ETF (SPY), the CBOE Volatility Index (VIX), the CBOE ten-year Treasury yield index (TNX) and Bitcoin (BTC-USD) come from Yahoo Finance. FOMC, CPI and employment-report dates are those of the official 2025–2026 calendars.

The primary sample is the negative-risk event “Fed decision in June?” (Gamma event 101772), created on 10 December 2025, a few hours after the December 2025 FOMC meeting lowered the

target range by 25 basis points (the effective rate fell from 3.89 to 3.64 percent on 11 December). The five contracts are listed in Panel A of Table 1 with their cumulative volume (\$6.5 to \$8.8 million each, \$37.8 million in total) and resting liquidity at the snapshot (\$0.2 to \$1.8 million). The price histories span 10 December 2025 to 22 May 2026; after the alignment described below the panel contains 147 days on which all five contracts have a close, out of 164 calendar days, the remaining days having no recorded daily point for at least one contract. The first day carries the quotes seeded at listing—0.15, 0.50, 0.575, 0.25 and 0.095, summing to 1.57—and is excluded from all statistics, leaving $T = 146$ daily closes, 42 of which fall on weekends.

The secondary sample is the event “Presidential Election Winner 2024” (event 903193), with seventeen candidate contracts and \$3.7 billion of cumulative volume (Panel B of Table 1). The histories span 4 January to 5 November 2024, the final close being the one recorded at 00:00 UTC on 6 November, that is, at 19:00 Eastern time on election day, before polls had closed in most states. The sample therefore ends *before* the outcome was known; it contains 304 days with a complete set of seventeen closes after exclusion of the first.

4.2. Alignment of calendars

The daily points of the CLOB price history are time-stamped at 00:00 UTC (plus a few seconds) and represent the last price of the preceding UTC day. We label each point by the UTC day it closes. This convention has three consequences that matter for the analysis. First, an FOMC statement released at 14:00 Eastern time (18:00 or 19:00 UTC) and a CPI or employment report released at 08:30 Eastern time are both contained in the Polymarket close of the same calendar day, so that “day 0” of a release is well defined. Second, the Polymarket close of day t is observed about three to four hours after the U.S. cash close of day t (16:00 Eastern); when Polymarket changes are regressed on lagged U.S. returns the lag is 27 hours, and when U.S. returns are regressed on lagged Polymarket changes it is 21 hours, so that neither direction uses information not yet available at the time of the explanatory variable. Third, Polymarket trades on weekends; for comparisons with the U.S. series we construct Polymarket changes on the U.S. trading-day grid, so that a Monday change spans Friday evening to Monday evening in both markets. Bitcoin, whose daily close on Yahoo Finance is also at 00:00 UTC, is aligned in the same way. The earlier version of this paper labelled each point by the UTC day of the time stamp, which placed every Polymarket close one day too late relative to the U.S. series and to the release dates; Appendix C quantifies the effect.

4.3. State prices and their changes

Table 2 summarises the five state prices and their daily changes over the main sample. The “No change” claim is the modal state throughout, averaging 0.61 and rising from 0.36 on 11 December

Table 1. Markets analysed: Polymarket metadata and daily closing prices

Contract	Volume	Liquidity	T	First	Mean	Min	Max	Last
<i>Panel A: Fed decision, June 2026 FOMC meeting (event 101772; volume and liquidity in USD)</i>								
25 bps decrease	6,849,280	225,586	146	0.535	0.301	0.006	0.570	0.009
25 bps increase	6,923,152	255,431	146	0.035	0.032	0.005	0.084	0.009
50+ bps decrease	8,765,442	780,011	146	0.130	0.057	0.004	0.130	0.004
No change	6,483,789	964,295	146	0.360	0.610	0.300	0.981	0.978
50+ bps increase	8,740,316	1,849,158	146	0.040	0.016	0.002	0.061	0.003
<i>Panel B: Presidential Election Winner 2024 (event 903193; volume, USD million)</i>								
Donald Trump	1,531.5	–	306	0.405	0.527	0.405	0.705	0.580
Kamala Harris	1,037.0	–	305	0.027	0.177	0.006	0.539	0.421
Other Republican Politician	241.7	–	304	0.009	0.005	0.001	0.015	0.001
Michelle Obama	153.4	–	305	0.013	0.024	0.001	0.080	0.001
Robert F. Kennedy Jr.	141.6	–	306	0.029	0.017	0.001	0.049	0.001
Other Democrat Politician	116.6	–	304	0.007	0.009	0.001	0.081	0.001
Nikki Haley	107.5	–	306	0.085	0.006	0.001	0.085	0.001
Hillary Clinton	93.3	–	305	0.008	0.002	0.001	0.025	0.001
Joe Biden	72.2	–	306	0.365	0.227	0.001	0.445	0.001
Gavin Newsom	54.2	–	306	0.050	0.012	0.001	0.074	0.001
Ron DeSantis	46.3	–	305	0.023	0.003	0.001	0.023	0.001
AOC	22.0	–	304	0.011	0.003	0.001	0.500	0.001
Vivek Ramaswamy	21.2	–	305	0.019	0.002	0.001	0.019	0.001
Elizabeth Warren	14.7	–	304	0.008	0.001	0.001	0.008	0.001
Chris Christie	14.2	–	304	0.002	0.001	0.001	0.003	0.001
Bernie Sanders	9.8	–	304	0.002	0.001	0.001	0.005	0.001
Kanye	9.2	–	304	0.007	0.001	0.001	0.007	0.001

Notes: Volume is cumulative traded notional reported by the Gamma API at the snapshot date (22 May 2026); liquidity is the resting order-book depth reported by the same end-point (not available for closed markets). T is the number of daily closes after exclusion of each contract’s first recorded point, which carries the quote seeded at listing (0.50 for most election contracts; the Fed seeds are reported in Section 4); “First” is the first regular close. The maximum of 0.500 for the AOC contract is an isolated print on 9 April 2024. Prices are in USDC per \$1 payoff. Panel B contracts are ordered by volume; the sample ends at the close of 5 November 2024 (00:00 UTC on 6 November), before results were known.

2025 to 0.978 on 22 May 2026; the “25 bps decrease” claim falls from 0.535 to 0.0085 over the same period, while the two tightening states and the large-cut state never exceed 0.13. Daily changes in the two main claims have standard deviations of about 3.3 percentage points, with excess kurtosis between 10 and 17 reflecting the concentration of movement on a few days. The augmented Dickey–Fuller statistics (Dickey and Fuller, 1979) reject a unit root in all change series at the 1 percent level; the level series are bounded and converge to zero or one by construction, so that all time-series analysis is conducted on changes.

Table 2. Summary statistics: Fed June 2026 state prices (11 December 2025 – 22 May 2026)

Contract	Level $p_{j,t}$				Daily change $\Delta p_{j,t}$ (percentage points)				
	Mean	SD	Min	Max	Mean	SD	Skew	Ex. kurt	ADF
50+ bps decrease	0.057	0.041	0.004	0.130	-0.087	0.76	-1.70	13.2	-5.35
25 bps decrease	0.301	0.195	0.006	0.570	-0.363	3.21	-2.61	14.4	-5.82
No change	0.610	0.237	0.300	0.981	0.426	3.34	1.65	16.1	-10.47
25 bps increase	0.032	0.019	0.005	0.084	-0.018	1.16	0.98	9.8	-5.76
50+ bps increase	0.016	0.015	0.002	0.061	-0.026	0.63	1.12	17.0	-7.84
State-price sum Σ_t	1.0167	0.0261	0.9605	1.1080	-0.068	1.83	-0.41	2.5	-8.71

Notes: $T = 146$ daily closes (calendar days, Polymarket trades continuously) after exclusion of the listing day. Levels are in USDC; changes are in percentage points of the \$1 payoff. ADF is the augmented Dickey–Fuller statistic (constant, AIC lag) for the daily changes; the 1% critical value is -3.48 . Prices are bounded in $[0, 1]$ and converge to 0 or 1 at resolution, so level series are not stationary by construction.

Figure 1 plots the five state prices. The dominant movements are the two FOMC days of 28 January 2026, when the “No change” price fell from 0.425 to 0.31 and the “25 bps decrease” price rose from 0.485 to 0.545, and 18 March 2026, when “No change” rose from 0.635 to 0.865 and “25 bps decrease” fell from 0.295 to 0.085. Between early April and the end of the sample the modal claim drifts upward towards one as the June meeting approaches.

4.4. Event dates

The sample contains four FOMC statements (17 December 2025; 28 January, 18 March and 6 May 2026), five CPI releases (14 January, 11 February, 11 March, 14 April and 12 May 2026) and five employment reports (9 January, 6 February, 6 March, 3 April and 8 May 2026). No Polymarket close is recorded on 14 April 2026 or on the two following days, so the April CPI release cannot be assigned a day-0 change and is dropped; the 3 April employment report, a Good Friday on which U.S. markets were closed, is assigned the first subsequent Polymarket close (4 April). The announcement-day analysis therefore uses thirteen release days.




Figure 1. State prices of the five contracts on the June 2026 FOMC decision, 11 December 2025 to 22 May 2026 (daily closes, listing day excluded). Dotted vertical lines mark the four FOMC statements in the sample (17 December, 28 January, 18 March, 6 May).

5. Methodology

5.1. Tests of the pricing restriction

For each day of the main sample we compute Σ_t from equation (1) and the deviation $\Sigma_t - 1$. The mean deviation is tested against zero with a heteroskedasticity- and autocorrelation-consistent standard error (Newey and West, 1987) with five lags, since $\Sigma_t - 1$ is serially correlated. We report the share of days on which $|\Sigma_t - 1|$ is within 0.5, 1, 2 and 5 percentage points, and the share of days on which Σ_t lies within the executable bounds (2), with the lower bound computed as

$$\delta_t = \exp(-r_t \tau_t), \quad (5)$$

where r_t is the effective federal funds rate on day t and τ_t the number of days to the resolution date divided by 365, and with $c_t = 0.005$. The dynamics of the deviation are characterised by the first-order autoregressive coefficient ϕ of $\Sigma_t - 1$, the implied half-life $-\ln 2 / \ln \phi$, and the augmented Dickey–Fuller statistic with a constant and AIC-selected lag length. Because the lower bound rises towards one as $\tau_t \rightarrow 0$, we also regress $\Sigma_t - 1$ on τ_t with HAC standard errors; a positive slope indicates that the overround shrinks as resolution approaches. The same statistics are computed for

the election group with T equal to 5 November 2024.

The order-book snapshot provides a direct check of the bounds at one point in time. For each contract we record the best bid and ask, the quoted and relative spread, the depth at the best quotes, and the notional resting within one cent of the best quote; the sum of best bids is the proceeds from selling one complete set at market and the sum of best asks the cost of buying one.

5.2. Announcement-day analysis

Let y_t denote either the daily change in the risk-neutral expected policy change, $\Delta\mathbb{E}_t^Q[\Delta r]$ in basis points, or the daily change in the “No change” state price in percentage points. Under the risk-neutral martingale property both have conditional mean close to zero, so that no estimation window is needed to define an abnormal change; the information content of a release is identified by the variance of y_t rather than by its mean, as in [Ederington and Lee \(1993\)](#). For each release type we report the change on day 0 and the cumulative changes over $[-1, +1]$ and $[-3, +3]$ calendar days for every event in the sample, and we compare the distribution of $|y_t|$ on the K release days with its distribution on the remaining days using three statistics: the ratio of means, the ratio of variances with the Brown–Forsythe test of equal variances (the Levene test centred at the median, robust to the heavy tails documented in Table 2), and the one-sided Mann–Whitney test that $|y_t|$ is stochastically larger on release days. With K between 4 and 13 these tests have limited power, which is why we report them for all releases pooled as well as by type. The same comparison between weekend and weekday observations among non-release days tests whether the market incorporates information when conventional markets are closed ([French and Roll, 1986](#)).

The earlier version of this paper reported cumulative abnormal returns for each contract with a constant-mean benchmark ([MacKinlay, 1997](#)) and a cross-sectional t -test over two or three events. We abandon that design for three reasons: simple returns on claims priced at a fraction of a cent are dominated by tick effects, the cross-sectional t -statistic with two events has one degree of freedom, and the sixty-day estimation window excluded two of the four FOMC meetings and two of the five CPI releases from the sample.

5.3. Lead–lag tests

Lead–lag relations are tested on the U.S. trading-day grid with five variables: the daily change in $\mathbb{E}_t^Q[\Delta r]$ (basis points), the change in the ten-year Treasury yield (basis points) and the log changes in SPY, VIX and BTC (percent). For a pair (x_t, y_t) the bivariate vector autoregression of order p ,

$$y_t = \alpha + \sum_{\ell=1}^p \beta_{\ell} y_{t-\ell} + \sum_{\ell=1}^p \gamma_{\ell} x_{t-\ell} + \varepsilon_t, \quad (6)$$

is estimated by least squares and the hypothesis $\gamma_1 = \dots = \gamma_p = 0$ is tested with the conventional F -statistic and with a Wald statistic using the Newey–West covariance with five lags. The lag order is chosen once, by the Akaike information criterion (Akaike, 1974) applied to the five-variable system over $p \in \{0, \dots, 5\}$; because the criterion selects $p = 0$, as do the Schwarz (Schwarz, 1978) and Hannan–Quinn criteria, the main tests use $p = 1$ and we report p -values for $p \in \{1, 2, 5\}$ as a sensitivity check. We deliberately do not select, for each pair, the lag that minimises the p -value: that procedure, used in the earlier version, is a search over five correlated tests and its nominal size is far below its actual size. The full five-variable VAR(1) is also reported equation by equation with HAC t -statistics (Lütkepohl, 2005).

Two limitations of the design are inherent in the data. Daily frequency is coarse relative to the speed at which both markets react to scheduled releases, so that a lead of minutes or hours, if it existed, would not be detectable; and the sample contains 110 trading days, so that only large predictive relations could be found. The tests are therefore informative about the absence of daily-horizon predictability, not about intraday price discovery.

5.4. Spread estimation

The implicit effective spread of Roll (1984) is estimated from the first-order serial covariance of daily price changes over a rolling twenty-day window, $\hat{S}_t = 2\sqrt{-\widehat{\text{Cov}}(\Delta p_t, \Delta p_{t-1})}$ when the covariance is negative and zero otherwise. With one price per day the estimator captures the bid–ask bounce only partially and is reported as a descriptive complement to the quoted spreads of the order-book snapshot. The Amihud illiquidity ratio (Amihud, 2002) cannot be computed because the price history does not include daily volume; the earlier version of this paper referred to it in error.

5.5. Volatility models

The conditional volatility of daily log returns on the “No change” claim is examined in Appendix A with GARCH and GJR-GARCH models (Bollerslev, 1986; Glosten et al., 1993) under Gaussian and Student- t innovations, estimated by quasi-maximum likelihood with robust standard errors, and with the heterogeneous autoregressive regression of Corsi (2009) using absolute and squared daily returns as volatility proxies. Engle’s ARCH-LM test (Engle, 1982) and Ljung–Box statistics on squared returns are reported alongside, because the presence of conditional heteroskedasticity is itself in question in a sample of 145 observations dominated by a handful of announcement days.

6. Results

6.1. The price of the complete set

Table 3 reports the tests of the pricing restriction for the Fed group (first column) and, for comparison, the election group discussed in Section 7. Over the 146 days of the main sample the sum of the five state prices averages 1.0167 with a standard deviation of 0.0261; the median is 1.0093, the minimum 0.9605 and the maximum 1.108. The mean overround of 1.67 percentage points has a HAC standard error of 0.44 and is therefore significantly different from zero ($t = 3.80$). It is, however, small relative to the quoted spreads of the first weeks and it is concentrated early in the sample: the mean of Σ_t is 1.051 (standard deviation 0.029) over the first thirty days and 1.0040 (standard deviation 0.0036) over the last thirty. On 84 percent of days Σ_t exceeds one; on 44 percent of days it is within one percentage point of one, on 65 percent within two and on 87 percent within five.

Panel C evaluates the executable bounds of equation (2). The lower bound $e^{-r_t \tau_t}$ averages 0.989 over the sample and reaches 0.981 at the start, so that Σ_t exceeds the buy-and-hold bound by 2.8 percentage points on average and is never below it. Because the overround is positive, the binding bound is the upper one: Σ_t lies within $[\delta_t - 0.005, 1.005]$ on 29 percent of days and above 1.005 on 65 percent. Taken at face value this would mean that selling the complete set and converting the NO legs was profitable on most days. The order-book snapshot in Table 4 shows why the face-value reading is wrong. On 22 May 2026 all five contracts were quoted with one-tick spreads of 0.001, the sum of the best bids was 0.999 and the sum of the best asks 1.004; the daily close recorded for that day, 1.0015, is exactly the sum of the midpoints. Neither direction of the round trip was profitable at the quotes, and the difference between the recorded Σ_t and the executable bid-side sum is the half-tick on each of five contracts. Earlier in the sample, when the Roll estimate of the effective spread on the “No change” claim alone was between one and three cents (Figure 8, Panel A), a recorded overround of two to five percentage points is of the same order as the spread costs of the five legs. We conclude that the law of one price holds for the complete set within executable bounds, and that the recorded overround is a measurement artefact of last-trade and midpoint prices in thin low-priced claims rather than an unexploited arbitrage.

Panel D confirms that deviations are transitory. The first-order autocorrelation of $\Sigma_t - 1$ is 0.72, implying a half-life of 2.1 days, and the Dickey–Fuller statistic of -2.82 rejects a unit root at the 10 percent level ($p = 0.055$) in a sample of 146 observations. The regression of the deviation on time to resolution has a slope of 0.113 per year (HAC $t = 4.87$): the overround falls by about one percentage point per month as the meeting approaches, which is the pattern expected if spreads narrow and arbitrage capital concentrates as the resolution date nears. Figure 2 plots Σ_t against the bounds and Figure 3 the distribution of the deviations.

Table 3. Price of the complete set of state claims: law-of-one-price tests

	Fed June 2026 ($N = 5$ states)	Election 2024 ($N = 17$ candidates)
<i>Panel A: distribution of $\Sigma_t = \sum_j p_{j,t}$</i>		
Observations T	146	304
Mean	1.0167	1.0146
Median	1.0093	1.0105
Standard deviation	0.0261	0.0313
Minimum	0.9605	0.9660
Maximum	1.1080	1.5150
<i>Panel B: deviation $\Sigma_t - 1$</i>		
Mean deviation (overround)	0.0167	0.0146
HAC standard error (5 lags)	0.0044	0.0020
t -statistic, $H_0: \mathbb{E}[\Sigma_t - 1] = 0$	3.80	7.42
Mean absolute deviation	0.0209	0.0162
Share of days with $\Sigma_t > 1$	83.6%	91.8%
Share within ± 0.5 pp	25.3%	12.8%
Share within ± 1 pp	43.8%	42.4%
Share within ± 2 pp	65.1%	74.0%
Share within ± 5 pp	87.0%	98.7%
<i>Panel C: discount-adjusted bounds, $e^{-r_t \tau_t} \leq \Sigma_t \leq 1$</i>		
Mean lower bound $e^{-r_t \tau_t}$	0.9890	0.9783
Mean $\Sigma_t - e^{-r_t \tau_t}$	0.0277	0.0363
Share inside bounds (± 0.5 pp)	29.5%	18.4%
Share above upper bound	65.1%	81.2%
<i>Panel D: dynamics of $\Sigma_t - 1$</i>		
AR(1) coefficient	0.721	0.060
Half-life (days)	2.1	0.2
ADF statistic (constant, AIC lag)	-2.82	-16.40
ADF p -value	0.055	0.000
Slope on time-to-resolution τ_t (years)	0.113	0.020
HAC t -statistic	4.87	3.73

Notes: Σ_t is the sum of the daily closing prices of all contracts in the neg-risk group on days when all contracts are quoted. The Fed sample excludes the listing day (10 December 2025, $\Sigma = 1.570$ on seeded quotes; including it gives mean 1.0204, SD 0.0525, maximum absolute deviation 0.570); the election sample likewise excludes its first complete day (4 January 2024, $\Sigma = 1.515$). The lower bound discounts the \$1 payoff of the complete set at the effective federal funds rate (FRED DFF) over the remaining time to resolution; the upper bound is one because a complete set can be minted and its NO legs converted immediately through the neg-risk adapter. HAC standard errors use the Newey–West estimator with 5 lags.

Table 4. Order-book snapshot, Fed June 2026 contracts (22 May 2026, 08:51 UTC)

Contract	Best	Best	Quoted	Relative	Depth at best (shares)		Depth within 1¢ (USD 000)	
	bid	ask	spread	spread (%)	Bid	Ask	Bid	Ask
50+ bps decrease	0.004	0.005	0.001	22.2	569,147	99,356	26	1
25 bps decrease	0.008	0.009	0.001	11.8	20,071	14,498	11	1
No change	0.977	0.978	0.001	0.1	22,839	20,673	748	600
25 bps increase	0.008	0.009	0.001	11.8	44,819	5,227	26	1
50+ bps increase	0.002	0.003	0.001	40.0	5,574,520	1,619,036	23	6
Sum of quotes	0.999	1.004						

Notes: Snapshot of the CLOB order book (YES tokens) retrieved with the price histories. Quoted spread is in USDC; the relative spread divides it by the midpoint. Depth within 1¢ is the notional resting within one cent of the best quote. The sum of best bids is the proceeds from selling one complete set at market; the sum of best asks is the cost of buying one.




Figure 2. Price of the complete set of five state claims, Σ_t , with the executable bounds of equation (2): the discounted value of one dollar at resolution (dashed, computed from the effective federal funds rate) and one (solid). The shaded band widens each bound by 0.5 percentage points for spread costs.




Figure 3. Distribution of the deviation $\Sigma_t - 1$ (percentage points) for the Fed group (Panel A, $T = 146$) and the election group (Panel B, $T = 304$). The dashed line in Panel A marks 0.5 percentage points, the sum of five half-ticks. The single election-group observation above 50 points is an isolated print of 0.50 on an otherwise inactive contract on 9 April 2024.

Two further remarks are in order. The positivity condition is satisfied throughout—the lowest recorded price is 0.002 for the “50+ bps increase” claim—but, as noted in Section 3, this is guaranteed by the price grid and carries no information. And the earlier version of this paper reported a mean of 1.0205, a standard deviation of 0.0527 and a maximum absolute deviation of 0.57; those figures include the listing day, on which the five seeded quotes summed to 1.57 before any trade had occurred, and the maximum deviation they reported was that artefact.

6.2. The risk-neutral distribution of the June decision

Figure 4 shows the normalised distribution $q_{j,t}$ at five dates and Figure 5 the implied expected policy change of equation (4) and its daily changes. At the start of the sample the distribution is bimodal, with 36 percent of risk-neutral mass on “No change” and 53 percent on a 25-basis-point cut; the expected policy change is -15.5 basis points and its risk-neutral standard deviation 21.7 basis points. By 22 May 2026 the distribution is concentrated on “No change” (0.978), the expected change is -0.1 basis points and the standard deviation 5.3 basis points. The expected change never exceeded -0.05 or fell below -17.4 basis points. Its path is dominated by a small number of large moves: $+6.0$ basis points on 18 March 2026 (FOMC), -1.7 basis points on 9 January 2026 (employment report), and $+1.3$ and $+1.2$ basis points on the CPI releases of 11 March and 11 February.




Figure 4. Risk-neutral distribution of the June 2026 decision, $q_{j,t} = p_{j,t}/\Sigma_t$, at five dates spanning the sample.

6.3. Scheduled releases

Table 5 examines Hypothesis 2. Panel A lists every release in the sample with the day-0 change and the cumulative changes over three- and seven-day windows, for both the expected policy change and the “No change” state price. The four FOMC statements produced day-0 changes in the modal state price of +1.5, −11.5, +23.0 and 0.0 percentage points. The January statement was read as making a June cut more likely—the “25 bps decrease” price rose by six points and the two tightening states, which had been nearly worthless, were briefly bid up to 0.054 and 0.037, so that the expected policy change actually rose slightly while its dispersion increased; the March statement removed most of the probability of a cut and moved the expected change by six basis points. The May statement, one week before the end of the sample, left the distribution unchanged. Among the data releases, the employment report of 6 March 2026 lowered the “No change” price by 12 points on the day and the CPI release of 11 February raised it by 7 points; the remaining releases moved it by less than five points.

Panel B pools the releases. On the thirteen release days the absolute change in the expected policy change averages 1.15 basis points against 0.51 on the other 132 days (ratio 2.3); the variance ratio is 4.9, the Brown–Forsythe test rejects equal variances ($p = 0.008$) and the Mann–Whitney test rejects equal distributions ($p = 0.012$). For the “No change” state price the absolute change averages 5.5 against 1.4 percentage points, the variance ratio is 14.2 and both tests reject at the




Figure 5. Upper panel: risk-neutral expected change in the upper bound of the target range at the June 2026 meeting, $\mathbb{E}_t^Q[\Delta r]$ (basis points), with FOMC statements (triangles), CPI releases (circles) and employment reports (squares). Lower panel: daily changes; release days in colour.

Table 5. Scheduled macroeconomic releases and the risk-neutral expected policy change

<i>Panel A: individual releases in the sample</i>									
Release	Date	$\mathbb{E}_t^Q[\Delta r]$ level (bps)	$\Delta \mathbb{E}^Q[\Delta r]$ (bps)			$\Delta p_{\text{No change}}$ (pp)			
			[0]	$[-1, +1]$	$[-3, +3]$	[0]	$[-1, +1]$	$[-3, +3]$	
FOMC	17 Dec 2025	-11.6	0.39	1.23	1.03	1.50	0.00	0.00	
FOMC	28 Jan 2026	-14.5	0.99	-1.06	1.08	-11.50	-9.00	-10.00	
FOMC	18 Mar 2026	-1.6	6.02	6.28	6.05	23.00	24.00	22.00	
FOMC	06 May 2026	-0.7	-0.12	-0.30	0.54	0.00	0.00	1.95	
CPI	14 Jan 2026	-14.1	-0.04	-0.40	0.13	0.00	0.00	-1.00	
CPI	11 Feb 2026	-15.5	1.24	0.19	0.34	7.00	7.00	8.00	
CPI	11 Mar 2026	-9.1	1.32	-0.48	2.50	4.00	1.50	11.00	
CPI	12 May 2026	-0.3	0.15	0.20	0.22	0.90	0.60	1.00	
Employment report	09 Jan 2026	-12.7	-1.66	-3.44	-3.27	-0.50	0.00	0.00	
Employment report	06 Feb 2026	-15.8	-0.68	-1.38	-1.92	-3.50	-11.50	-11.50	
Employment report	06 Mar 2026	-8.1	-1.02	-2.21	3.28	-12.00	-11.50	8.00	
Employment report	04 Apr 2026	-1.3	-0.96	-1.48	-0.50	6.00	4.00	6.00	
Employment report	08 May 2026	-0.4	0.31	0.45	-0.06	1.55	1.95	1.15	

<i>Panel B: absolute daily changes on release days versus all other days</i>											
Days	K	$ \Delta \mathbb{E}^Q[\Delta r] $ (bps)					$ \Delta p_{\text{No change}} $ (pp)				
		Ann.	Other	VR	p_{BF}	p_{MW}	Ann.	Other	VR	p_{BF}	p_{MW}
All scheduled releases	13	1.15	0.51	4.85	0.008	0.012	5.50	1.38	14.23	0.000	0.003
FOMC	4	1.88	0.53	10.60	0.004	0.113	9.00	1.55	30.44	0.000	0.083
CPI	4	0.69	0.56	0.51	0.908	0.352	2.98	1.72	0.92	0.576	0.176
Employment report	5	0.93	0.56	0.52	0.798	0.018	4.71	1.65	4.47	0.023	0.013
Weekend vs. weekday (no release)	41	0.37	0.58	0.33	0.113	0.861	0.91	1.59	0.34	0.053	0.970

Notes: $\mathbb{E}_t^Q[\Delta r] = \sum_j q_{j,t} x_j$ with $x = (-50, -25, 0, 25, 50)$ basis points and $q_{j,t} = p_{j,t} / \Sigma_t$; the open-ended states are valued at ± 50 bps. Day 0 is the UTC calendar day of the release (FOMC statements at 14:00 ET, CPI and employment reports at 08:30 ET), which is contained in the Polymarket close of that day (00:00 UTC). Windows cumulate daily changes; the 3 April 2026 employment report (Good Friday) is assigned the first subsequent close, and the 14 April 2026 CPI release has no close within two days and is dropped. Panel B compares absolute daily changes on the K release days with all other days ($T = 145$ daily changes, listing day excluded): “Ann.” and “Other” are means, VR is the ratio of variances, p_{BF} the Brown–Forsythe test of equal variances and p_{MW} the one-sided Mann–Whitney test that absolute changes are larger on release days. The last row compares weekend days with weekdays among non-release days. For comparison, the absolute daily change in the 10-year Treasury yield (CBOE TNX) is 3.2 bps on release days versus 3.3 bps otherwise (variance ratio 0.90, $p_{BF} = 0.966$).

1 percent level. By release type, the excess variance is concentrated on FOMC days (variance ratios of 10.6 and 30.4) and, for the state price, on employment-report days (variance ratio 4.5, Brown–Forsythe $p = 0.023$); the four CPI releases in the sample show no excess variance in either measure. The absence of a CPI effect should not be over-interpreted: four observations cannot distinguish a small effect from none, and the CPI prints of the period may simply have been close to expectations. Over the same dates the absolute daily change in the ten-year Treasury yield was 3.3 basis points on release days and 3.3 on other days, so that the Treasury market shows no release-day excess variance at the daily frequency in this particular sample either; the difference between the two markets is that the Polymarket contract prices a single meeting and therefore reacts almost exclusively to information about that meeting.

The last row of Panel B compares weekend and weekday observations among non-release days. Weekend changes in the expected policy change have a standard deviation about 0.6 times that of weekday changes (variance ratio 0.33, Brown–Forsythe $p = 0.11$) and weekend changes in the modal state price a variance ratio of 0.34 ($p = 0.053$). Weekend trading therefore carries real but reduced information flow, roughly one third of weekday variance, which is close to what [French and Roll \(1986\)](#) found for equity returns over non-trading periods and consistent with the view that most of the news relevant to the June decision arrives during business hours. Figure 6 displays the distributions.




Figure 6. Absolute daily changes in the risk-neutral expected policy change (Panel A, basis points) and in the “No change” state price (Panel B, percentage points) on non-release days and on employment-report, CPI and FOMC days. Horizontal bars are group means.

6.4. Lead–lag relations with conventional markets

Table 6 reports the Granger tests of Hypothesis 3 on 110 trading days from 12 December 2025 to 21 May 2026. With the lag order chosen by information criteria, none of the eight directional tests rejects the null at the 10 percent level: the p -values range from 0.19 to 0.90 for the conventional F -test and from 0.20 to 0.88 for the HAC Wald test. The results are stable at two lags. At five lags two cells fall below 10 percent—the expected policy change to the VIX ($p = 0.059$, HAC $p = 0.033$) and to Bitcoin ($p = 0.006$, HAC $p = 0.317$)—but these appear in a block of 24 additional tests, are not supported by the information criteria, and in the Bitcoin case do not survive HAC inference. We read Table 6 as consistent with the null of no daily-horizon predictability in either direction.

Table 6. Granger causality between the Fed state-price panel and U.S. asset returns

Direction	VAR(p^*), $p^* = 1$ (AIC)			p -value at fixed lag		
	F	p	p (HAC)	$p = 1$	$p = 2$	$p = 5$
$\Delta \mathbb{E}^Q[\Delta r] \rightarrow 10\text{-year yield}$	0.19	0.666	0.534	0.666	0.252	0.370
$\Delta \mathbb{E}^Q[\Delta r] \rightarrow \text{SPY}$	0.90	0.345	0.198	0.345	0.364	0.253
$\Delta \mathbb{E}^Q[\Delta r] \rightarrow \text{VIX}$	0.39	0.533	0.505	0.533	0.373	0.059
$\Delta \mathbb{E}^Q[\Delta r] \rightarrow \text{BTC}$	1.75	0.188	0.264	0.188	0.629	0.006
10-year yield $\rightarrow \Delta \mathbb{E}^Q[\Delta r]$	0.13	0.718	0.635	0.718	0.277	0.706
SPY $\rightarrow \Delta \mathbb{E}^Q[\Delta r]$	0.35	0.557	0.228	0.557	0.816	0.908
VIX $\rightarrow \Delta \mathbb{E}^Q[\Delta r]$	0.29	0.589	0.463	0.589	0.369	0.835
BTC $\rightarrow \Delta \mathbb{E}^Q[\Delta r]$	0.02	0.897	0.884	0.897	0.983	0.747

Notes: Bivariate VAR Granger tests on $T = 110$ U.S. trading days (12 Dec 2025 – 21 May 2026). The Fed variable is the daily change in the risk-neutral expected June 2026 policy change (basis points), measured at the Polymarket close of the same calendar day (00:00 UTC, about four hours after the U.S. cash close); changes over weekends are assigned to Monday, as for the financial series. TNX is the daily change in the CBOE 10-year Treasury yield index (bps); SPY, VIX and BTC are daily log changes in percent. The lag order is selected by the AIC on the five-variable system (BIC selects 0, HQIC 0). F and p are the conventional Granger F -test; p (HAC) is the Wald test with Newey–West (5 lags) covariance. ***, **, *: 1%, 5%, 10%.

Table 7 reports the five-variable VAR(1). No lagged variable is significant in the equation for the expected policy change, whose R^2 is 0.018, and the lagged expected policy change is insignificant in the four other equations. The contemporaneous correlations, not shown in the table, are informative about the direction of the common shocks: the daily change in the expected policy change has a correlation of 0.26 with the change in the ten-year yield and of -0.15 with the SPY return, with near-zero correlations with the VIX and Bitcoin. The Polymarket contract on the June meeting and

the Treasury market thus respond to the same macroeconomic news within the day, in the direction one would expect—news that raises the expected policy rate raises long yields—but neither market’s daily change predicts the other’s. Figure 7 shows the complete matrix of p -values.

Table 7. VAR(1) estimates on the trading-day grid

Regressor	Dependent variable (day t)				
	$\Delta \mathbb{E}^Q[\Delta r]$	10-year yield	SPY	VIX	BTC
Constant	0.137 (1.27)	0.325 (0.82)	0.094 (1.18)	-0.123 (-0.20)	-0.087 (-0.30)
$\Delta \mathbb{E}^Q[\Delta r]_{t-1}$	-0.123 (-1.40)	-0.117 (-0.39)	-0.039 (-0.75)	0.157 (0.26)	-0.260 (-0.80)
10-year yield $_{t-1}$	0.002 (0.10)	-0.034 (-0.39)	-0.035 (-1.29)	0.364 (1.52)	-0.113* (-1.80)
SPY $_{t-1}$	-0.085 (-0.43)	0.781 (1.01)	0.099 (0.69)	-0.012 (-0.01)	0.161 (0.31)
VIX $_{t-1}$	0.003 (0.10)	-0.017 (-0.21)	0.025 (1.60)	-0.202 (-1.58)	0.058 (1.07)
BTC $_{t-1}$	0.020 (0.60)	-0.063 (-0.55)	0.010 (0.37)	-0.277 (-1.00)	-0.053 (-0.48)
R^2	0.018	0.036	0.054	0.057	0.052
Observations	109				

Notes: Equation-by-equation OLS of each variable on one lag of all five variables (the AIC-selected order; see Table 6). Newey–West (5 lags) t -statistics in parentheses. Variables are defined in Table 6.

***, **, *: 1%, 5%, 10%.

The earlier version of this paper reported that the Fed contract Granger-caused the VIX at the 1 percent level and the SPY at the 10 percent level, and that the VIX Granger-caused the Fed contract at the 5 percent level. Appendix C shows that these results were produced jointly by the one-day misalignment of the Polymarket calendar and by the selection, for each pair, of the lag with the smallest p -value among five; with the alignment corrected, the same selection procedure yields p -values of 0.06, 0.37 and 0.26 for the three relations, and with the lag fixed by information criteria the p -values are 0.53, 0.59 and 0.35.

6.5. Spreads and depth

Figure 8 completes the microstructure evidence. Panel A plots the Roll estimate of the effective spread for the two main claims over rolling twenty-day windows. For the “No change” claim the estimate averages 1.3 cents, falling from 1.6 cents over the first month to 1.1 cents over the last; in 32 percent of windows the serial covariance of price changes is positive and the estimator is set to zero, a familiar limitation of the Roll measure at daily frequency when price changes




Figure 7. Granger-causality p -values (bivariate VAR(1), 110 trading days) from the column variable to the row variable. “Fed” denotes the daily change in the Polymarket-implied expected policy change.

are dominated by information rather than by bounce. Panel B shows the cumulative depth of the order book on 22 May 2026 as a function of the distance from the midpoint. The modal claim is deep on both sides—about \$0.75 million bid and \$0.60 million offered within one cent of the best quotes—whereas the four low-priced claims are quoted at one-tick spreads with substantial depth on the bid side but only a few thousand dollars offered within one cent; a buyer of a complete set at market would therefore have moved the low-priced legs by several ticks even for a modest size. The relative spread ranges from 0.1 percent for the modal claim to 40 percent for the “50+ bps increase” claim quoted at 0.002–0.003: in relative terms the tails of the state space are expensive to trade, which is one more reason to treat the recorded prices of those claims, and the overround they generate, with caution.

7. The 2024 Election Market and Further Checks

7.1. The complete set of seventeen candidate claims

The “Presidential Election Winner 2024” event is the most heavily traded market in the platform’s history and provides a state space more than three times larger than the Fed group. The second column of Table 3 reports the pricing statistics over 304 daily closes from 5 January to 5 November




Figure 8. Panel A: Roll (1984) estimate of the effective spread (cents) for the “No change” and “25 bps decrease” claims, rolling twenty-day windows. Panel B: cumulative resting depth of the five order books on 22 May 2026 (USD thousand, log scale) as a function of the distance from the midpoint (cents, log scale); ask side solid, bid side dashed.

2024. The sum of the seventeen prices averages 1.0146 with a standard deviation of 0.0313 and a median of 1.0105; the mean overround of 1.46 percentage points has a HAC standard error of 0.20. The overround exceeds one on 92 percent of days, is within two percentage points of one on 74 percent of days and within five on 99 percent. As in the Fed group it shrinks towards resolution—1.022 over the first thirty days, 1.0077 (standard deviation 0.0011) over the last thirty—and the slope on time to resolution is positive (0.020 per year, HAC $t = 3.73$). Deviations are far less persistent than in the Fed group: the first-order autocorrelation is 0.06 and the Dickey–Fuller statistic -16.4 , so that a deviation on one day is essentially gone the next. The maximum of 1.515 is a single day, 9 April 2024, on which the claim on one minor candidate printed a trade at 0.50 against surrounding closes of 0.0015 and 0.0020; it illustrates once more that a recorded daily price on an inactive claim is not an executable quote, and the median-based statistics in the table are unaffected by it.

The lower bound is further from one in this market than in the Fed group because the horizon is longer: $e^{-r_t \tau_t}$ averages 0.978 and is 0.957 in January 2024 when the effective rate was 5.33 percent and the election ten months away. The complete set therefore traded on average 3.7 percentage points above its discounted value, and the entire overround is again on the upper side of the band. With seventeen claims, fifteen of which traded at less than two cents for most of the sample, the sum of half-spreads alone accounts for about one percentage point at a tick of 0.001 and for far more at the one-cent tick that applies to most of the price range on the platform. The evidence from the two state spaces is therefore consistent: the complete set is priced at or marginally above one,

the time value of money is not reflected in the prices, and the measured overround is of the order of the spread costs of the low-priced legs.

Figure 9 plots the five most traded candidate claims and the sum of all seventeen. The price of the Trump claim ranged from 0.405 to 0.705, the maximum of 0.705 being reached on 15 July 2024 after the attempted assassination of 13 July (0.595 on 12 July, 0.695 on 14 July). President Biden’s withdrawal on 21 July moved his claim from 0.0695 to 0.0005 and Vice-President Harris’s from 0.18 to 0.32 within a day, and the debate of 10 September moved the Trump claim from 0.5225 to 0.496 over two days. The Trump claim peaked at 0.6655 on 29 October and closed at 0.58 on the evening of 5 November, with the Harris claim at 0.4215 and the fifteen remaining claims summing to 0.0075. The earlier version of this paper described this final value as a post-election “convergence” to 0.58 and inferred a settlement delay; the description was wrong, since the sample ends before the result was known, and the 0.58 is simply the market’s last pre-result price. The market resolved to the Trump claim after the election, outside our sample.

7.2. Sensitivity of the pricing tests

The pricing statistics of Table 3 are insensitive to the treatment of the listing day except for the maximum and the standard deviation, as reported in the table note. Computing the lower bound with the upper bound of the target range (3.75 percent) instead of the effective rate shifts δ_t by at most 0.0006 and changes no conclusion. Replacing the sample mean by the median overround (0.93 percentage points for the Fed group, 1.05 for the election group) does not change the ordering of the two markets or the sign of the term-structure slope. Using a cost allowance of one cent on the complete set instead of half a cent raises the share of Fed days inside the executable band from 29 to 49 percent; using the contemporaneous Roll estimate of the effective spread of the modal claim as the allowance gives 42 percent. We have not attempted to reconstruct the executable bid-side sum over the whole sample, which would require historical order-book data that the public API does not provide; this is the most useful extension of the pricing analysis.

7.3. Sensitivity of the announcement and lead–lag results

The announcement-day comparisons use absolute changes and rank-based or median-centred tests, so that no single observation drives them. Dropping the two largest daily moves in the sample—18 March 2026 (FOMC) and 6 March 2026 (employment report)—leaves the pooled release-day variance ratio for the “No change” state price at 4.7 (Brown–Forsythe $p = 0.005$, Mann–Whitney $p = 0.023$) but reduces that for the expected policy change to 1.2 ($p = 0.38$; Mann–Whitney $p = 0.051$): the release-day effect on the modal state price is robust, whereas the effect on the first moment of the distribution rests on the two days on which the distribution shifted most. Valuing the




Figure 9. Upper panel: daily closing prices of the five most traded claims in the “Presidential Election Winner 2024” event, 4 January to 5 November 2024; dotted lines mark the attempted assassination (13 July), President Biden’s withdrawal (21 July), the Harris–Trump debate (10 September) and election day (5 November). Lower panel: price of the complete set of seventeen claims with the discounted lower bound (dashed) and one (solid).

open-ended states at ± 75 rather than ± 50 basis points rescales the expected-policy-change series by about 11 percent (correlation 0.99 between the two versions) and leaves the pooled tests unchanged (variance ratio 4.9, Brown–Forsythe $p = 0.004$). The lead–lag tests are reported at three lag orders in Table 6. Replacing the expected policy change by the change in the “No change” state price as the Polymarket variable produces one cell below the 10 percent level among the sixteen tests at one and two lags—the state price to SPY, with $p = 0.07$ at one lag and 0.04 at two—which is about what the null predicts for sixteen correlated tests; the seven other directions have p -values above 0.10 at both lag orders. Splitting the 110 trading days into two halves of 55 produces no rejection at one lag in either half in any direction.

7.4. Checks not performed

Several analyses that would strengthen the paper are not feasible with the data at hand and are recorded here as recommendations rather than results. Intraday price histories (the API offers minute and hourly fidelity for active markets) would allow the response to each release to be timed within the day and compared with the response of Treasury futures, which is the natural way to assess price discovery across the two venues; the daily data used here cannot speak to leads of minutes or hours. Federal funds futures or overnight index swap quotes for the June 2026 meeting would allow a direct comparison of the Polymarket-implied expected change with the futures-implied one and an estimate of the difference in risk premia, in the spirit of [Piazzesi and Swanson \(2008\)](#); we do not have those quotes. Historical order books would allow the executable bounds to be tested day by day. Finally, a cross-section of negative-risk groups—other FOMC meetings, other elections—would allow the term structure of the overround and its relation to depth to be estimated rather than described. None of these extensions requires new methods, only data we did not collect.

8. Discussion and Conclusion

8.1. What the state prices show

Negative-risk events on Polymarket provide complete sets of state claims on well-defined macroeconomic and political outcomes, and this paper has used two of them to ask how the Arrow–Debreu pricing restrictions fare in a blockchain market. The answer is in three parts. The complete set of claims is priced at one dollar within the bounds that the platform’s own mechanics make executable: the recorded sums of 1.017 (five Fed claims) and 1.015 (seventeen election claims) are statistically above one but economically within the half-spreads of the low-priced legs, they decline steadily towards resolution, they revert within days, and at the one order-book snapshot available the bid-side and ask-side sums bracket one. The normalised prices define a risk-neutral distribution over the

June 2026 FOMC decision whose mean moved from -15.5 basis points to zero over six months, and that mean moves two to five times more, in variance terms, on FOMC and employment-report days than on other days, while weekend days carry about a third of weekday variance. And daily changes in that mean neither lead nor lag the Treasury, equity, volatility or cryptocurrency markets, although they are contemporaneously correlated with Treasury yields.

Two interpretive points follow. First, the Arrow–Debreu metaphor is accurate but less demanding than it sounds. Completeness is a property of the contract design, not a finding; what the data confirm is the law of one price for the riskless complete set and the absence of exploitable deviations at the quotes. The one genuine departure from the frictionless benchmark—the complete set trades at one rather than at its discounted value—reflects the asymmetry of the executable arbitrage (the sell-side round trip settles immediately through the adapter, the buy-side requires holding to resolution) and the fact that collateral earns nothing on the platform. A trader who buys a complete set is lending at zero; the market prices this correctly by refusing to pay less than one for it. Second, the normalised prices are risk-neutral probabilities and nothing in the data identifies the physical ones. The literature on futures-implied policy expectations has shown that the risk premium in such measures is large and time-varying (Piazzesi and Swanson, 2008); there is no reason to believe the Polymarket measure is exempt, and our hypothesis tests are designed not to depend on the distinction.

8.2. Manipulation, oracles and other threats

Three institutional risks deserve explicit mention because they bear on the external validity of any inference drawn from these prices. The first is concentrated position-taking. During the 2024 election a single trader operating several accounts was reported to have taken positions of some \$30 million on the Trump claim, and the platform’s own prices moved with those purchases; theory and experiments suggest that such trades are absorbed when other participants anticipate them (Hanson et al., 2006), but the daily data used here cannot separate information from flow, and our pricing results are in any case silent on whether the *level* of any claim was right. The second is resolution risk. Polymarket outcomes are determined by UMA’s optimistic oracle, in which disputed resolutions are settled by token-holder vote; in March 2025 a market on a U.S.–Ukraine minerals agreement was resolved against the apparent facts after a concentrated vote, and the platform declined to compensate holders. For the two markets studied here the resolution criteria—an FOMC statement and an election result—are as unambiguous as such criteria can be, but the possibility of an oracle dispute is a basis risk that is priced, if at all, in the small residual mass on improbable states. The third is volume. Reported volumes on unregulated blockchain venues can be inflated by wash trading (Cong et al., 2023); we have used volume only descriptively, and the depth and spread evidence of Section 6 is a better guide to the economic size of these markets than the headline

figures.

8.3. Limitations

The study rests on two negative-risk groups observed at daily frequency. The Fed sample has 146 daily closes, thirteen scheduled releases and 110 trading days in common with U.S. markets; the announcement-day and lead-lag tests are correspondingly limited in power, and the absence of a detectable CPI effect or of a lead-lag relation is a statement about what daily data from one market can show, not a demonstration of informational equivalence. The pricing tests use recorded daily prices rather than executable quotes, and the single order-book snapshot cannot establish that the bounds held on every day. The alignment of the Polymarket close with the U.S. close leaves a residual window of three to four hours; a release in that window—there are none among the scheduled releases studied, but there could be unscheduled news—would appear in the Polymarket close of day t and the U.S. close of day $t + 1$. The discount factor is computed from the effective federal funds rate rather than from the opportunity cost actually faced by platform participants, which may be higher. And the volatility models of Appendix A, which an earlier version of this paper presented as a main result, turn out to be fragile: the ARCH-LM test does not reject homoskedasticity, the GARCH estimates sit on the boundary of the parameter space, and the asymmetry parameter is identified by two announcement days.

8.4. Conclusion

Blockchain prediction markets make complete state spaces tradable and observable. For the two state spaces studied here, the prices behave as Arrow–Debreu theory says state prices should—they sum to the price of a riskless claim within executable bounds, they define a well-behaved risk-neutral distribution, and that distribution responds to scheduled information. They do not behave as some of the enthusiasm around prediction markets would suggest: they do not lead the Treasury or equity markets at the daily horizon, they are not physical probabilities, and their recorded sums should not be read as arbitrage opportunities without the order book in hand. The most promising next steps are intraday, with the futures market as the benchmark, and cross-sectional, across the many negative-risk groups the platform now lists; the framework and the calendar conventions set out here are directly applicable to both.

References

- Akaike, Hirotugu**, “A New Look at the Statistical Model Identification,” *IEEE Transactions on Automatic Control*, 1974, 19 (6), 716–723.
- Amihud, Yakov**, “Illiquidity and Stock Returns: Cross-Section and Time-Series Effects,” *Journal of Financial Markets*, 2002, 5 (1), 31–56.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega**, “Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange,” *American Economic Review*, 2003, 93 (1), 38–62.
- Arrow, Kenneth J.**, “The Role of Securities in the Optimal Allocation of Risk-Bearing,” *Review of Economic Studies*, 1964, 31 (2), 91–96.
- Baillie, Richard T., G. Geoffrey Booth, Yiuman Tse, and Tatyana Zabotina**, “Price Discovery and Common Factor Models,” *Journal of Financial Markets*, 2002, 5 (3), 309–321.
- Berg, Joyce E., Forrest D. Nelson, and Thomas A. Rietz**, “Prediction Market Accuracy in the Long Run,” *International Journal of Forecasting*, 2008, 24 (2), 285–300.
- Bernanke, Ben S. and Kenneth N. Kuttner**, “What Explains the Stock Market’s Reaction to Federal Reserve Policy?,” *Journal of Finance*, 2005, 60 (3), 1221–1257.
- Bollerslev, Tim**, “Generalized Autoregressive Conditional Heteroskedasticity,” *Journal of Econometrics*, 1986, 31 (3), 307–327.
- Chernov, Mikhail, Vadim Elenev, and Dongho Song**, “The Comovement of Voter Preferences: Insights from U.S. Presidential Election Prediction Markets Beyond Polls,” NBER Working Paper 33339, National Bureau of Economic Research 2025.
- Cong, Lin William, Xi Li, Ke Tang, and Yang Yang**, “Crypto Wash Trading,” *Management Science*, 2023, 69 (11), 6427–6454.
- Corsi, Fulvio**, “A Simple Approximate Long-Memory Model of Realized Volatility,” *Journal of Financial Econometrics*, 2009, 7 (2), 174–196.
- Debreu, Gerard**, *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*, New Haven: Yale University Press, 1959.
- Dickey, David A. and Wayne A. Fuller**, “Distribution of the Estimators for Autoregressive Time Series with a Unit Root,” *Journal of the American Statistical Association*, 1979, 74 (366), 427–431.

- Ederington, Louis H. and Jae Ha Lee**, “How Markets Process Information: News Releases and Volatility,” *Journal of Finance*, 1993, 48 (4), 1161–1191.
- Engle, Robert F.**, “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation,” *Econometrica*, 1982, 50 (4), 987–1007.
- French, Kenneth R. and Richard Roll**, “Stock Return Variances: The Arrival of Information and the Reaction of Traders,” *Journal of Financial Economics*, 1986, 17 (1), 5–26.
- Glosten, Lawrence R. and Paul R. Milgrom**, “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders,” *Journal of Financial Economics*, 1985, 14 (1), 71–100.
- , **Ravi Jagannathan, and David E. Runkle**, “On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks,” *Journal of Finance*, 1993, 48 (5), 1779–1801.
- Gonzalo, Jesús and Clive W. J. Granger**, “Estimation of Common Long-Memory Components in Cointegrated Systems,” *Journal of Business & Economic Statistics*, 1995, 13 (1), 27–35.
- Granger, Clive W. J.**, “Investigating Causal Relations by Econometric Models and Cross-Spectral Methods,” *Econometrica*, 1969, 37 (3), 424–438.
- Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson**, “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 2005, 1 (1), 55–93.
- Hanson, Robin**, “Combinatorial Information Market Design,” *Information Systems Frontiers*, 2003, 5 (1), 107–119.
- , **Ryan Oprea, and David Porter**, “Information Aggregation and Manipulation in an Experimental Market,” *Journal of Economic Behavior & Organization*, 2006, 60 (4), 449–459.
- Harvey, Campbell R., Ashwin Ramachandran, and Joey Santoro**, *DeFi and the Future of Finance*, Hoboken: John Wiley & Sons, 2021.
- Hasbrouck, Joel**, “One Security, Many Markets: Determining the Contributions to Price Discovery,” *Journal of Finance*, 1995, 50 (4), 1175–1199.
- , “Intraday Price Formation in U.S. Equity Index Markets,” *Journal of Finance*, 2003, 58 (6), 2375–2400.

- Kuttner, Kenneth N.**, “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market,” *Journal of Monetary Economics*, 2001, 47 (3), 523–544.
- Kyle, Albert S.**, “Continuous Auctions and Insider Trading,” *Econometrica*, 1985, 53 (6), 1315–1335.
- Lütkepohl, Helmut**, *New Introduction to Multiple Time Series Analysis*, Berlin: Springer, 2005.
- MacKinlay, A. Craig**, “Event Studies in Economics and Finance,” *Journal of Economic Literature*, 1997, 35 (1), 13–39.
- Madhavan, Ananth**, “Market Microstructure: A Survey,” *Journal of Financial Markets*, 2000, 3 (3), 205–258.
- Makarov, Igor and Antoinette Schoar**, “Trading and Arbitrage in Cryptocurrency Markets,” *Journal of Financial Economics*, 2020, 135 (2), 293–319.
- Manski, Charles F.**, “Interpreting the Predictions of Prediction Markets,” *Economics Letters*, 2006, 91 (3), 425–429.
- Nakamura, Emi and Jón Steinsson**, “High-Frequency Identification of Monetary Non-Neutrality: The Information Effect,” *Quarterly Journal of Economics*, 2018, 133 (3), 1283–1330.
- Newey, Whitney K. and Kenneth D. West**, “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 1987, 55 (3), 703–708.
- Ng, Hunter, Lin Peng, Yubo Tao, and Dexin Zhou**, “Price Discovery and Trading in Modern Prediction Markets,” Working Paper 5331995, SSRN 2025.
- O’Hara, Maureen**, *Market Microstructure Theory*, Cambridge, MA: Blackwell, 1995.
- Ottaviani, Marco and Peter Norman Sørensen**, “Price Reaction to Information with Heterogeneous Beliefs and Wealth Effects: Underreaction, Momentum, and Reversal,” *American Economic Review*, 2015, 105 (1), 1–34.
- Page, Lionel and Robert T. Clemen**, “Do Prediction Markets Produce Well-Calibrated Probability Forecasts?,” *Economic Journal*, 2013, 123 (568), 491–513.
- Piazzesi, Monika and Eric T. Swanson**, “Futures Prices as Risk-Adjusted Forecasts of Monetary Policy,” *Journal of Monetary Economics*, 2008, 55 (4), 677–691.
- Reichenbach, Felix and Martin Walther**, “Exploring Decentralized Prediction Markets: Accuracy, Skill, and Bias on Polymarket,” Working Paper 5910522, SSRN 2025.

- Rhode, Paul W. and Koleman S. Strumpf**, “Historical Presidential Betting Markets,” *Journal of Economic Perspectives*, 2004, 18 (2), 127–142.
- Rigobon, Roberto and Brian Sack**, “The Impact of Monetary Policy on Asset Prices,” *Journal of Monetary Economics*, 2004, 51 (8), 1553–1575.
- Roll, Richard**, “A Simple Implicit Measure of the Effective Bid–Ask Spread in an Efficient Market,” *Journal of Finance*, 1984, 39 (4), 1127–1139.
- Rothschild, David**, “Forecasting Elections: Comparing Prediction Markets, Polls, and Their Biases,” *Public Opinion Quarterly*, 2009, 73 (5), 895–916.
- Schwarz, Gideon**, “Estimating the Dimension of a Model,” *Annals of Statistics*, 1978, 6 (2), 461–464.
- Sims, Christopher A.**, “Macroeconomics and Reality,” *Econometrica*, 1980, 48 (1), 1–48.
- Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz**, “Prediction Markets for Economic Forecasting,” in Graham Elliott and Allan Timmermann, eds., *Handbook of Economic Forecasting*, Vol. 2A, Amsterdam: Elsevier, 2013, pp. 657–687.
- Wolfers, Justin and Eric Zitzewitz**, “Prediction Markets,” *Journal of Economic Perspectives*, 2004, 18 (2), 107–126.
- **and** —, “Interpreting Prediction Market Prices as Probabilities,” NBER Working Paper 12200, National Bureau of Economic Research 2006.
- **and** —, “Prediction Markets in Theory and Practice,” NBER Working Paper 12083, National Bureau of Economic Research 2006.
- Yan, Bingcheng and Eric Zivot**, “A Structural Analysis of Price Discovery Measures,” *Journal of Financial Markets*, 2010, 13 (1), 1–19.

Appendices

A. Conditional Volatility of the Modal State Price

This appendix reports the volatility models that an earlier version of the paper presented as a principal result. The series is the daily log return on the “No change” claim over the main sample (145 calendar-day returns, listing day excluded), in percent; it has a mean of 0.69, a standard deviation of 6.67, a minimum of -31.6 (28 January 2026) and a maximum of $+30.9$ (18 March 2026), both FOMC days.

Before estimating conditional-variance models we test for their premise. Engle’s ARCH-LM statistic with five lags is 4.92 ($p = 0.43$) and the Ljung–Box statistic on squared returns at ten lags is 10.7 ($p = 0.38$): neither rejects conditional homoskedasticity. The Ljung–Box statistic on the returns themselves is 20.7 ($p = 0.023$), so there is some evidence of serial correlation in the mean, consistent with the underreaction predicted by [Ottaviani and Sørensen \(2015\)](#) or with stale prices on inactive days, but not of volatility clustering of the kind that GARCH models are designed to capture.

Table 8 reports GARCH(1,1) and GJR-GARCH(1,1) estimates with Gaussian and Student- t innovations and Bollerslev–Wooldridge robust standard errors. Three features stand out. The persistence $\alpha_1 + \beta_1$ (plus $\gamma_1/2$ for the asymmetric models) equals one in all four specifications, and the constant ω is zero to three decimal places in three of them: the estimates lie on the boundary of the parameter space, where the usual asymptotic inference does not apply. The asymmetry parameter γ_1 of the GJR model is about 0.30 in both distributions, but its robust standard error is 0.72 under Gaussian innovations and the symmetric ARCH coefficient is estimated at exactly zero; the Student- t degrees of freedom are below three, so that the fitted innovation distribution has infinite kurtosis and barely finite variance. The large likelihood gains of the asymmetric and Student- t specifications therefore reflect the fitting of two or three announcement-day observations rather than a pervasive leverage effect, and we do not interpret γ_1 as evidence of one. The earlier version reported a GJR coefficient of 0.30 as a “novel leverage effect” with an AIC improvement of 46 over the symmetric model; the coefficient is reproduced here, but the inference is not.

Table 9 reports the heterogeneous autoregressive regression of [Corsi \(2009\)](#) with the daily absolute return as the volatility proxy and, alternatively, the squared return, using the standard definitions of the weekly and monthly components as averages of the daily proxy over five and twenty-two days. With absolute returns the weekly component is positive and significant ($\beta_w = 0.63$, $t = 2.24$) and the regression explains 15 percent of the variation in next-day absolute returns; with squared returns, which are dominated by the two FOMC days, the fit collapses ($R^2 = 0.02$, joint $p = 0.49$). Persistence in the absolute size of daily moves at the weekly horizon is thus detectable,

Table 8. GARCH-family estimates for the “No change” contract (daily log returns, percent)

Parameter	GARCH normal	GJR-GARCH normal	GARCH Student- <i>t</i>	GJR-GARCH Student- <i>t</i>
μ	0.433 (0.273)	0.089 (0.104)	0.011 (0.033)	-0.003 (0.033)
ω	0.528 (0.862)	0.000 (0.175)	0.000 (0.019)	0.000 (0.018)
α_1	0.183** (0.083)	0.000 (0.890)	0.291** (0.141)	0.000 (0.055)
γ_1	—	0.304 (0.721)	—	0.334*** (0.084)
β_1	0.817*** (0.036)	0.848 (0.521)	0.709*** (0.147)	0.833*** (0.047)
ν	—	—	2.583*** (0.176)	2.820*** (0.237)
Log-likelihood	-461.2	-425.3	-382.4	-373.6
AIC	930.4	860.6	774.8	759.1
BIC	942.3	875.5	789.6	777.0
Persistence $\alpha_1 + \beta_1 + \gamma_1/2$	1.000	1.000	1.000	1.000
Ljung–Box p, z_t^2 (10 lags)	0.987	0.455	0.987	0.384
Observations	145			

Notes: Constant-mean models estimated by (quasi-)maximum likelihood on calendar-day log returns of the “No change” contract, listing day excluded; Bollerslev–Wooldridge robust standard errors in parentheses. The GJR variance equation is $\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \gamma_1 \varepsilon_{t-1}^2 \mathbf{1}\{\varepsilon_{t-1} < 0\} + \beta_1 \sigma_{t-1}^2$. The likelihood-ratio statistic for $\gamma_1 = 0$ (normal innovations) is 71.76 ($p = 0.000$). Engle’s ARCH-LM test (5 lags) on the returns: 4.92 ($p = 0.426$). ***, **, *: 1%, 5%, 10%.

and is plausibly linked to the clustering of macroeconomic releases within the month, but the evidence is modest. Figure 10 plots the returns and the fitted conditional standard deviations.

Table 9. HAR regressions for the “No change” contract

Regressor	$RV_t = r_t $		$RV_t = r_t^2$	
	Coef.	(SE)	Coef.	(SE)
Constant	0.994	(0.725)	32.296*	(19.078)
$RV_t^{(d)}$	-0.058	(0.110)	-0.075	(0.059)
$RV_t^{(w)}$	0.628**	(0.281)	0.325	(0.294)
$RV_t^{(m)}$	0.132	(0.262)	0.074	(0.310)
R^2	0.152		0.021	
Adjusted R^2	0.131		-0.004	
Joint F -test p -value	0.000		0.488	
Observations	123		123	

Notes: $RV_{t+1}^{(d)} = c + \beta_d RV_t^{(d)} + \beta_w RV_t^{(w)} + \beta_m RV_t^{(m)} + \varepsilon_{t+1}$ with $RV_t^{(w)}$ and $RV_t^{(m)}$ the averages of the daily measure over the past 5 and 22 calendar days (Corsi, 2009). Daily log returns in percent; listing day excluded. Newey–West (5 lags) standard errors. ***, **, *: 1%, 5%, 10%.

B. Variable Definitions

C. Reconciliation with the Earlier Version

The first version of this paper (May 2026) reported several results that the present version does not reproduce. This appendix records the differences and their causes so that readers of both versions can reconcile them; the full audit trail is in the repository’s revision log.

The earlier completeness statistics (mean 1.0205, standard deviation 0.0527, maximum absolute deviation 0.570, 146 observations) are reproduced exactly when the listing day is included: the seeded quotes of 10 December 2025 summed to 1.57 before any trade. Excluding that day gives the figures of Table 3. The earlier text attributed the 2 percent overround to transaction costs “and the time value of money”; as shown in Section 3, the time value of money lowers the price of the complete set and cannot contribute to a positive overround. The earlier text also stated that the effective federal funds rate was 4.25–4.50 percent over the sample; the FRED series shows 3.62–3.64 percent.

The earlier event study reported cumulative abnormal returns of -0.615 and -0.999 for individual contracts with p -values of 0.107 and 0.095 from two and three events respectively; the contract




Figure 10. Panel A: daily log returns (percent) on the “No change” claim, 12 December 2025 to 22 May 2026. Panel B: conditional standard deviation from the Gaussian GARCH(1,1) and the Student- t GJR-GARCH(1,1) of Table 8.

Symbol	Definition
$p_{j,t}$	Daily closing price of the YES token of contract j , in USDC per \$1 payoff, labelled by the UTC day it closes (the CLOB daily point stamped 00:00 UTC of day $t + 1$).
Σ_t	Price of the complete set, $\sum_j p_{j,t}$, computed on days when all N contracts have a close.
δ_t	Discounted value of \$1 at resolution, $\exp(-r_t \tau_t)$, with r_t the effective federal funds rate (FRED DFF) and τ_t the time to resolution in years.
$q_{j,t}$	Risk-neutral probability of state j , $p_{j,t}/\Sigma_t$.
$\mathbb{E}_t^{\mathbb{Q}}[\Delta r]$	Risk-neutral expected change in the upper bound of the target range at the June 2026 meeting, $\sum_j q_{j,t} x_j$ with $x = (-50, -25, 0, 25, 50)$ basis points.
$\Delta p_{\text{No change},t}$	Daily change in the “No change” state price, in percentage points.
r_t (returns)	Daily log return on the “No change” claim, in percent, used only in Appendix A.
TNX	Daily change in the CBOE ten-year Treasury yield index, in basis points, on the U.S. trading-day grid.
SPY, VIX, BTC	Daily log changes (percent) in the SPDR S&P 500 ETF, the CBOE Volatility Index and Bitcoin (USD), on the U.S. trading-day grid.
$\hat{S}_t$	Roll (1984) spread estimate, $2\sqrt{-\widehat{\text{Cov}}(\Delta p_t, \Delta p_{t-1})}$ over a rolling twenty-day window, zero when the covariance is positive.
$RV_t^{(d)}, RV_t^{(w)}, RV_t^{(m)}$	Daily volatility proxy ($ r_t $ or r_t^2) and its averages over the past 5 and 22 calendar days.

Table 10. Variable definitions.

labels in the text did not match those in the table. The design is replaced by the announcement-day analysis of Section 6, which uses all thirteen releases in the sample and a variance-based identification.

The earlier Granger tests reported $F = 4.96$ ($p = 0.001$) for the Fed contract to the VIX, $F = 2.73$ ($p = 0.049$) for the VIX to the Fed contract and $F = 2.42$ ($p = 0.072$) for the Fed contract to the SPY, described as using AIC-selected lags. The code selected, for each pair, the lag in $\{1, \dots, 5\}$ with the smallest p -value, and the Polymarket series was labelled one day late, so that a Polymarket “return” on day t was in fact the change to the evening of day $t - 1$ and the regression of VIX changes on lagged Polymarket changes effectively used information from the same evening. With the calendar corrected but the minimum- p selection retained, the three p -values become 0.064, 0.370 and 0.258; with the lag fixed by information criteria they are 0.533, 0.589 and 0.345 (Table 6). The earlier VAR table contained only the intercepts because the information criterion had selected zero lags, which the text did not report.

The earlier abstract and summary referred to Hasbrouck information shares, to Treasury futures, to a lead of “15–45 minutes” and to Amihud illiquidity ratios. None of these was computed: the data are daily, contain no Treasury futures and no daily volume, and the series are not cointegrated, so that information shares are undefined. The earlier description of the election market as converging to 0.58 after the election and of a settlement delay was incorrect; the sample ends on election evening (Section 7). The earlier robustness section described market-model, Patell, Student- t , EGARCH, subsample, cross-market, Johansen and GMM analyses whose results were not produced by the code; they have been removed, and the analyses that were feasible with the data (Student- t innovations, alternative lag orders, alternative valuation of open-ended states, listing-day treatment) are reported in Sections 7 and Appendix A.