Why So Serious?

Decomposing the Belief Volatility Smile in Prediction Markets

Andrea Bortolini

andrea.bortolini.2026@anderson.ucla.edu

March 18, 2026

Abstract

Prediction markets have moved from curiosity to infrastructure. Billions of dollars now trade on platforms such as Kalshi and Polymarket, regulators on both sides of the Atlantic are actively debating their role in financial markets, and their prices increasingly appear alongside traditional indicators in policy discussions and financial media. Yet the tools used to analyze them remain borrowed wholesale from equity and options markets — frameworks designed for unbounded, log-normally distributed assets, applied without adjustment to contracts whose prices are, by construction, bounded probabilities. This paper argues that the mismatch matters.

We develop a *belief-space framework* that takes the bounded nature of prediction-market prices seriously. The logit transform maps bounded probabilities to the real line, yielding a natural state space in which belief revisions are scale-free, unbounded, and directly interpretable as log-odds updates. Within this framework, we define *Implied Belief Volatility* (IBV) as the rate of daily belief revision, construct a cross-strike Shannon entropy measure from FOMC rate-target ladders that tracks the market's collective uncertainty about monetary policy outcomes, and characterize the geometry of the belief surface across moneyness and time-to-resolution.

Our main findings challenge conventional intuitions. The apparent volatility smile — IBV is higher for out-of-the-money contracts — turns out to be largely a geometric artifact of the logit transform rather than a genuine informational phenomenon: in probability space, out-of-the-money contracts are actually *less* volatile. Market entropy declines in a smooth, disciplined funnel as FOMC dates approach, with sharp drops around genuine monetary policy surprises. And the workhorse Gaussian assumption is decisively rejected: belief updating is episodic rather than continuous, with a large share of days information-inert and active days exhibiting dramatically fat-tailed dynamics. These departures from Gaussianity produce economically material pricing errors for anyone valuing prediction-market contracts under standard distributional assumptions.

Contents

1 Introduction

Prediction markets aggregate beliefs about future events into time-varying prices (Arrow et al., 2008; Snowberg et al., 2013; Manski, 2006; Wolfers and Zitzewitz, 2006). Unlike stock prices, prediction-market prices are bounded probabilities $p_t \in (0, 1)$, and treating them as if they live on the real line obscures fundamental geometric features of belief dynamics. The core question of this paper is: *What is the correct state space and stochastic law for modeling belief revisions in prediction markets?*

We argue that the answer is the logit line. If p_t is the market probability at time t, then $x_t = \log[p_t/(1 - p_t)]$ is the corresponding *logit belief state*. Revisions $\Delta x_t = x_t - x_{t-1}$ are unbounded, scale-free, and directly interpretable as log-odds updates. The absolute revision $|\Delta x_t|$ — which we call Implied Belief Volatility (IBV) — measures the rate of belief revision in natural units.

Our empirical laboratory is the Kalshi FOMC rate-target ladder: for each Federal Open Market Committee meeting, Kalshi lists binary contracts on each possible target rate outcome. These contracts form a complete probability distribution over the rate space, enabling us to compute both per-contract IBV and a cross-strike Shannon entropy. We use daily data (close-to-close logit returns) as our baseline, which is more transparent and defensible than intraday bars and avoids microstructure complications.

The paper has four main contributions:

1. **Logit-space framework.** We formalize the belief-space representation, define IBV and cross-strike entropy, and derive the Gaussian logit-updating benchmark as a natural baseline.

2. **Geometry vs. economics of the smile.** IBV is higher for OTM contracts (large $|x_t|$), producing an inverse smile in logit space. We decompose this into a *mechanical* component (the logit transform amplifies OTM moves even when probability-space volatility is flat or declining) and an *informational* residual. The decomposition reveals that the smile is largely mechanical: in probability space, OTM contracts have significantly *lower* volatility.

3. **Entropy funnel.** Cross-strike entropy declines monotonically as FOMC meetings approach, collapsing from $H \approx 1.0$ more than 60 days out to $H \approx 0.3$ in the final week. This funnel quantifies how the market's collective uncertainty about the rate outcome is resolved over time, with striking heterogeneity across meetings captured by the Information Value of a Meeting (IVM).

4. **Non-Gaussian dynamics and pricing correction.** Daily logit returns are decisively non-Gaussian: they exhibit a zero-inflated structure (33% of days have near-zero updates) and heavy-tailed active updates consistent with a Zero-Inflated Student-t (ZIT) model. This departure from Gaussianity generates material pricing errors: a Gaussian benchmark systematically overprices most contracts, with the largest error concentrated at short horizons and near-ATM positions (16.5 cents at $\tau = 1$ day), attenuating sharply with horizon (maximum 4.1 cents at $\tau = 30$ days).

Appendix B tests whether prediction-market entropy contains incremental information about FOMC-day equity volatility beyond the VIX, finding economically plausible but statistically inconclusive evidence with 28 meetings — illustrating both the promise and current data limitations of this approach.

2 Belief-Space Framework

2.1 Logit Belief State and IBV

Let $p_t \in (0,1)$ be the closing mid-price of a binary prediction-market contract on day t, interpreted as the market's probability that the underlying event resolves "Yes." We interpret p_t as a market probability following Manski (2006) and Wolfers and Zitzewitz (2006). Define the *logit belief state*:

$$x_t = \text{logit}(p_t) = \log\left(\frac{p_t}{1-p_t}\right). \tag{1}$$

The map $p_t \mapsto x_t$ is the canonical link from probability space to the real line (McCullagh and Nelder, 1989). It satisfies $p_t = \sigma(x_t)$ where $\sigma(\cdot)$ is the logistic function. We define *logit moneyness* as $|x_t|$, the absolute distance from ATM ($x_t = 0$) in logit space: contracts with $|x_t|$ large are deep out-of- or in-the-money.

Definition 1 (Implied Belief Volatility)**.** *The daily Implied Belief Volatility of contract k on day t is*

$$IBV_{k,t} = |\Delta x_{k,t}| = |x_{k,t} - x_{k,t-1}|. \tag{2}$$

IBV is the absolute log-odds revision per day. It is dimensionless and invariant to monotone re-parameterizations of the probability scale. IBV is the prediction-market analog of the realized volatility estimator of Andersen et al. (2003), applied to logit returns rather than log-price returns. The Jacobian of the logit transform,

$$\frac{\partial x}{\partial p} = \frac{1}{p(1-p)}, \tag{3}$$

is minimized at $p = \frac{1}{2}$ (ATM) and increases as probabilities approach the boundaries, diverging as $p \to 0$ or $p \to 1$. This amplification has a direct consequence: even if probability-space volatility $|\Delta p_t|$ is flat across strikes, logit-space IBV will appear higher at OTM/ITM, producing what we call the *mechanical smile*.

2.2 Probability-Space Decomposition

For each day, the IBV can be decomposed into a *mechanical* and an *informational* component. Let $\bar{p}_t = (p_t + p_{t-1})/2$ be the average probability. By a first-order Taylor expansion of the logit function around $\bar{p}_t$ (accurate to within 2% for $|\Delta p| < 0.2$):

$$\text{IBV}_{k,t} \approx \underbrace{\frac{|\Delta p_{k,t}|}{\bar{p}_t(1-\bar{p}_t)}}_{\text{mechanical (IBV\_mech)}} + \underbrace{\varepsilon_{k,t}}_{\text{informational (IBV\_info)}}. \tag{4}$$

The mechanical component rescales the probability-space revision by the Jacobian and is entirely determined by the current probability level. Any residual — IBV_info — reflects belief updating that cannot be explained by the geometric transformation alone. The key insight is that OTM contracts (large $|x|$) have *lower* probability-space volatility $|\Delta p|$ on average — they are information-inert near the boundaries — while their logit-space volatility appears large due to the Jacobian amplification. The negative regression slope on $|x|$ in probability space (Section 4.2) captures exactly this: the mechanical amplification accounts for the entire logit-space smile, with no residual informational component required.

2.3 Cross-Strike PMF and Shannon Entropy

For each FOMC meeting, Kalshi lists contracts $k = 1, \ldots, K$ corresponding to target-rate outcomes $r_1 < r_2 < \cdots < r_{K+1}$. The cross-strike implied PMF assigns probability π_k to outcome $r_k \leq R_T < r_{k+1}$, recovered from the no-arbitrage ladder prices. The *Shannon entropy* of this PMF is:

$$H(t) = -\sum_{k=1}^{K} \pi_k(t) \log \pi_k(t), \tag{5}$$

(Shannon, 1948). $H(t)$ measures the collective uncertainty about the rate outcome at time t; it is maximized at the uniform distribution ($H = \log K$) and zero at a degenerate one.

Definition 2 (Information Value of a Meeting)**.** *For meeting i, let $H_{pre}(i)$ be the mean cross-strike entropy over days with $DTE \in [2,7]$ and $H_{ann}(i)$ be the entropy on the announcement day ($DTE \leq 1$). The* Information Value of Meeting i *is:*

$$IVM_i := H_{pre}(i) - H_{ann}(i). \tag{6}$$

$IVM > 0$ indicates that final-week information resolved uncertainty before announcement; $IVM \approx 0$ indicates a fully anticipated outcome; $IVM < 0$ indicates that pre-announcement uncertainty increased *on the day of the decision, consistent with a genuine surprise.*

2.4 Gaussian Benchmark

A natural baseline model assumes that daily logit returns are i.i.d. Gaussian (Dalen, 2025):

$$\Delta x_t \overset{\text{iid}}{\sim} \mathcal{N}(\mu, \sigma^2). \tag{7}$$

Under this model, the cumulative τ-day logit return $X_\tau = \sum_{d=1}^{\tau} \Delta x_d$ is $\mathcal{N}(\mu\tau, \sigma^2\tau)$, and the price of a digital contract (paying 1 if $x_0 + X_\tau > 0$) is (Hull, 2018):

$$V_{\text{Gaussian}}(x_0, \tau) = \Phi\left(\frac{x_0 + \mu\tau}{\sigma\sqrt{\tau}}\right), \tag{8}$$

where Φ is the standard normal CDF. The Gaussian benchmark is elegant and tractable, but Section 6 documents that it is empirically inadequate.

3 Data and FOMC Laboratory

3.1 Sample

Our primary dataset consists of all Kalshi FOMC rate-target contracts for meetings from March 2022 through May 2025. For each meeting, Kalshi lists multiple binary contracts covering a range of possible target rates (typically 3–19 contracts per meeting depending on market conditions). We download the complete trade history via the Kalshi API, aggregate to daily bars, and compute close-to-close logit returns.

After filtering for contracts with at least 3 active trading days and removing bars with implausible price moves ($|\Delta x| > 8$), our pooled sample contains **13,532 daily logit returns** from **319 contracts** across **31 meetings** (2022–2025).

3.2 Inclusion Criteria

A daily observation is included if: (i) there is a valid closing price on both day t and day $t-1$; (ii) the bar is not flagged as a gap (no trading activity); (iii) $|\Delta x_t| \leq 8$ (removes metadata discontinuities). This is a simple, transparent filter that avoids arbitrary hourly-usability thresholds.

3.3 Summary Statistics

Table 1. Sample Summary: Daily FOMC Logit Returns. Active returns have $|\Delta x| > 0.001$; near-zero (information-inert) returns have $|\Delta x| \leq 0.001$. The 33.4% near-zero mass reflects days with no economically meaningful belief revision.

	All observations	Active ($|\Delta x| > 0.001$)	Near-zero ($|\Delta x| \leq 0.001$)
N (market-days)	13,532	9,015	4,517
Fraction	100%	66.6%	33.4%
Mean Δx	0.035	0.053	
Std Δx	0.470	0.576	
Excess kurtosis	17.90	10.93	
Skewness	0.641	0.433	

Three features of Table 1 are noteworthy. First, 33.4% of market-days have near-zero logit revisions ($|\Delta x| \leq 0.001$) — an information-inert mass that is structurally inconsistent with continuous Gaussian updating and already visible in the 50th percentile of zero. Second, the full-sample excess kurtosis of 17.90 implies dramatically fat-tailed active revisions: extreme belief updates occur far more often than Gaussianity predicts. Third, the skewness of 0.641 reflects episodic sharp upward logit revisions, particularly around surprise rate decisions. These three moments — zero-inflation, fat tails, and skewness — jointly motivate the Zero-Inflated Student-t model developed in Section 6.

4 Geometry vs. Economics of the Smile

4.1 IBV Surface in Logit Space

IBV organized by logit moneyness $|x_t|$ and days-to-settlement (DTE) reveals two patterns. First, OTM contracts (large $|x_t|$) exhibit significantly higher IBV than near-ATM contracts — an *inverse smile*. The OTM bucket $|x| > 3$ has mean IBV $\approx$ 0.51–0.63 depending on DTE, while near-ATM ($|x| < 0.5$) has mean IBV $\approx$ 0.25–0.39. Second, IBV is relatively stable across DTE buckets for a given moneyness level, suggesting that belief revision intensity is more strongly determined by position in the logit space than by time remaining.

A regression of log(IBV) on $|x|$ and $\log(1 + \text{DTE})$ with meeting fixed effects yields (estimated on active days with $|\Delta x| > 0.001$, for which log(IBV) is well-defined; the probability-space regression in eq. (10) uses $\log(1+\text{IBV\_prob})$ and covers the full sample):

$$\log(\text{IBV}) = \underbrace{0.289}_{(0.015)***} |x| \underbrace{+\ 0.028}_{(0.013)*} \log(1 + \text{DTE}) + \text{meeting FE} + \varepsilon, \tag{9}$$

confirming a strong, statistically significant positive relationship between logit moneyness and IBV. The small and only marginally significant DTE coefficient confirms the IBV surface is largely stable across the contract lifetime for a given moneyness level.

4.2 Probability-Space Decomposition

Is the logit-space smile a genuine economic phenomenon (higher OTM uncertainty) or an artifact of the logit transform (the Jacobian $1/[p(1-p)]$ diverges at extreme probabilities)? We answer this by computing IBV in both spaces.

Figure 1 shows the decomposition. The key finding: *in probability space, IBV_prob is decreasing in* $|x|$, with a slope of -0.017 ($p < 0.001$) from the regression:

$$\log(1 + \text{IBV\_prob}) = \underbrace{-0.017}_{(0.001)***} |x| + \log(1 + \text{DTE}) + \text{series FE} + \varepsilon. \tag{10}$$

This confirms that the IBV inverse smile is **largely mechanical**. Deep OTM contracts have low probability-space volatility precisely because they are information-inert near the $[0, 1]$ boundary — they rarely update. The logit transform then amplifies these small moves via the Jacobian, producing a large apparent IBV. The genuine informational component (IBV_info) is near zero across all moneyness bins.

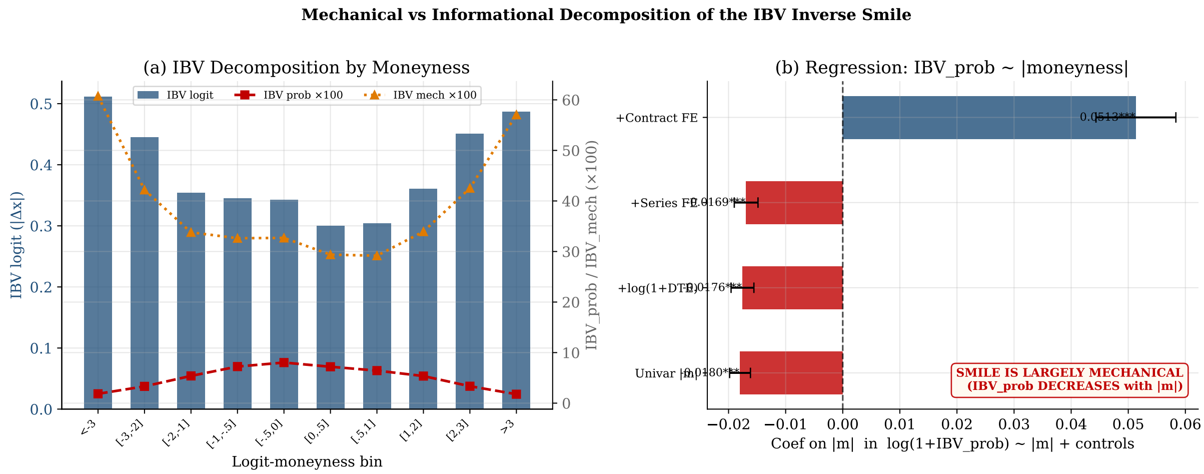

Figure 1. Mechanical vs. informational decomposition of the IBV inverse smile. Panel (a): IBV_logit (bars, left axis) is high for OTM contracts, while IBV_prob (circles, right axis) is flat or declining, confirming bounded-support inertia at the extremes. Panel (b): sequential regression coefficients of $\log(1+\text{IBV\_prob})$ on logit moneyness $|x|$, with controls for $\log(1+\text{DTE})$ and series fixed effects — all coefficients are negative and highly significant, confirming the smile is largely mechanical.

Interpretation. The mechanical decomposition does not diminish the contribution of the logit framework. Rather, it clarifies it: the logit transform is the *correct* state space for modeling belief dynamics. It is correct not because it produces a dramatic smile, but because it yields natural, unbounded, scale-free revisions. This boundary amplification is a feature, not a bug — it correctly captures that a move from $p = 0.50$ to $p = 0.51$ (one percentage point) corresponds to a logit revision of approximately 0.04 log-odds units, whereas a move from $p = 0.99$ to $p = 0.98$ (also one percentage point) corresponds to a revision of approximately 0.70 log-odds units. The same nominal probability change carries fundamentally different informational content depending on the starting belief.

4.3 Bounded Random-Walk Benchmark

A natural question is whether the IBV surface can be reproduced by *bounded support alone*, without any non-Gaussian updating. We test this with a reflected Gaussian random walk (RGRW) in probability space:

$$p^*_{t+1} = p_t + \sigma\varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1), \tag{11}$$

with $p_{t+1} = \mathcal{R}(p^*_{t+1})$ where $\mathcal{R}$ is the unit-interval reflection map ($p \mapsto -p$ if $p < 0$; $p \mapsto 2-p$ if $p > 1$). We set $\mu = 0$ and calibrate $\hat{\sigma}$ to the ATM empirical standard deviation of Δp ($\hat{\sigma} = 0.121$). For each observed (p_t, x_t), we simulate 2,000 one-step RGRW draws and compute model-implied IBV_logit and IBV_prob.

Result. Figure 2 compares the empirical smile to the RGRW model. The RGRW reproduces the U-shape of the IBV smile in logit space (correlation of cell-level means with empirical surface = 0.63, $R^2 = 40\%$). However, it *over-predicts* logit IBV at the boundaries: model-implied IBV_logit at $|x| > 3$ is 1.23 vs. empirical 0.51 (Figure 2, panel (a)). This reversal — the data is *less* volatile near the boundaries than the reflected model implies — indicates that deep OTM contracts are information-inert: they rarely update, consistent with the ZIT's 33% zero-inflation mass. The RGRW also fails to match probability-space IBV (RMSE = 0.042). Bounded geometry alone is insufficient.

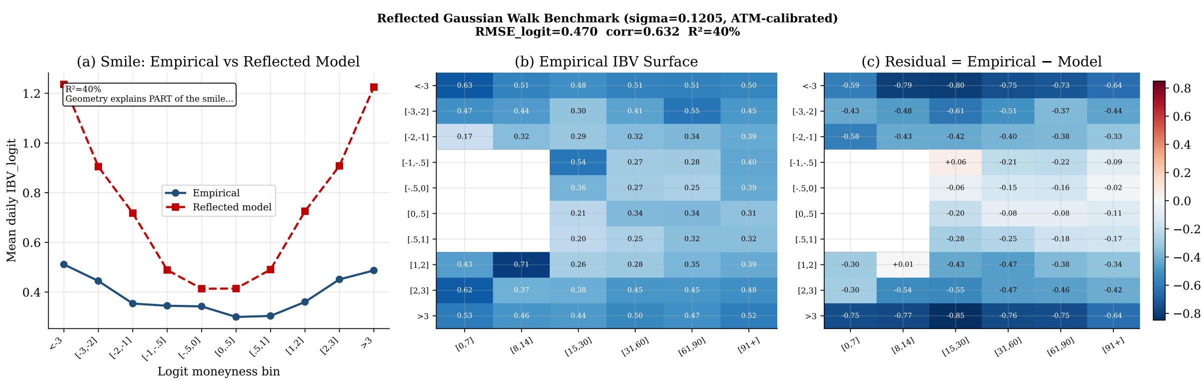

Figure 2. Reflected Gaussian random walk benchmark ($\hat{\sigma} = 0.121$, ATM-calibrated, 2,000 paths/observation). Panel (a): empirical vs. model smile by moneyness. Panel (b): empirical IBV surface. Panel (c): residual (empirical − model) — negative residuals at OTM indicate the data is *less* volatile near the boundaries than pure reflection predicts.

5 Entropy Around Scheduled Resolution

Figure 3 shows the cross-strike Shannon entropy $H(t)$ from daily FOMC ladder prices, organized by days-to-settlement. The pattern is striking: entropy declines monotonically as the meeting date approaches, collapsing from $H \approx 1.0$ for DTE > 60 to $H \approx 0.3$ in the final week. Of course, part of this decline is mechanical — the meeting date is prescheduled and known to all participants. The more informative object is the *heterogeneity*

Table 2. IBV Surface Fit: Model Comparison

Model	$\text{RMSE}_{\text{logit}}$	$\text{RMSE}_{\text{prob}}$	r_{logit}	R^2_{logit}	Role
Gaussian logit benchmark	0.118	0.015	-0.107	0.01	tractable baseline
Reflected Gaussian walk	0.470	0.042	$+0.632$	0.40	bounded geometry
ZIT / Student-t	0.185	0.018	-0.137	0.02	non-Gaussian updating
NIG	0.184	0.019	-0.037	0.00	best unconditional fit

RMSE and r computed over 52 surface cells ($n \geq 5$). All logit-space models (Gaussian, ZIT, NIG) predict IBV surfaces inconsistent with the empirical U-shape (near-zero or negative r); low RMSE reflects a correct *level* prediction but wrong *shape*. The RGRW is the only model that reproduces the U-shape ($r = 0.63$, $R^2 = 40\%$) via bounded-geometry mechanics, but over-predicts OTM magnitude (RMSE $= 0.47$): the data is *less* volatile near the boundaries than pure reflection predicts, consistent with the ZIT zero-inflation mass attenuating deep-OTM updates.

across meetings captured by IVM, which reveals whether a given meeting's uncertainty was resolved gradually in the final week or arrived as a shock on announcement day.

Table 3. Mean Cross-Strike Entropy by Days-to-Settlement (DTE)

DTE window	Mean H	N
> 60 days	1.008	1,935
30–60 days	0.681	613
7–30 days	0.453	456
≤ 7 days	0.308	168

Information Value by Meeting. Table 4 reports IVM_i for all 28 meetings with sufficient final-week ladder coverage. The cross-meeting distribution is striking: most meetings have IVM $\in [0, 0.1]$ (moderate resolution in the final week), but a handful of outliers tell informative stories. The June 2023 meeting (IVM $= 0.31$) saw the sharpest pre-announcement convergence: entropy collapsed from 0.59 to 0.28, consistent with a market that rapidly priced in a hold. The May and December 2022 meetings also show high IVM, when markets were rapidly reassessing the pace of the tightening cycle. By contrast, the June 2022 meeting has IVM $= -0.28$: entropy on announcement morning was *higher* than during the preceding week, consistent with the surprise 75 basis-point hike (the first since 1994). The September 2024 meeting (surprise 50bp cut) similarly shows IVM $= -0.08$. Negative IVM meetings are not failures of the measure — they are exactly the meetings economic history records as surprises.

This entropy funnel is consistent with a model where the market continuously incorporates new information — macroeconomic data releases, Fed communications, finan-

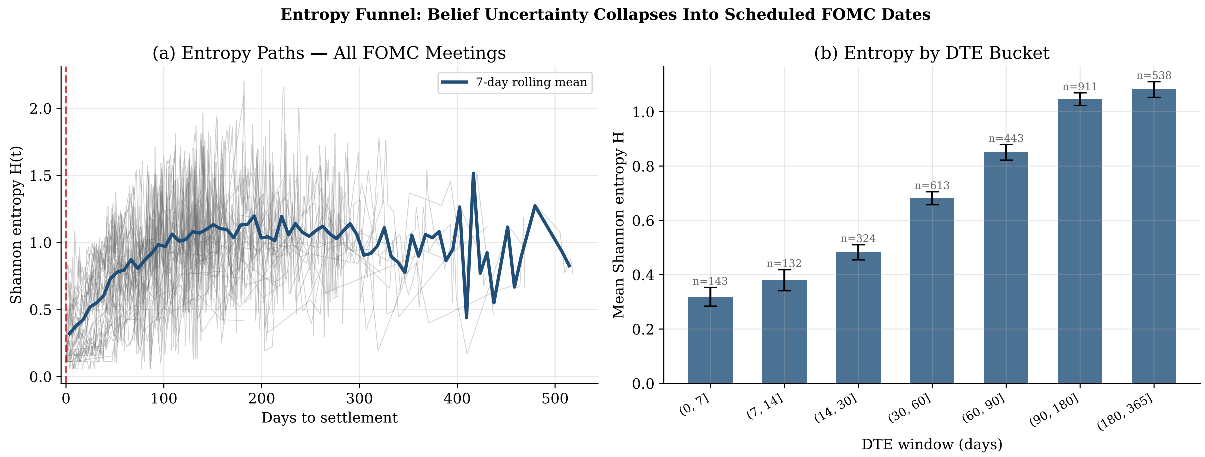

Figure 3. Entropy funnel: Shannon entropy $H(t)$ declines as FOMC meetings approach. Panel (a): individual meeting entropy paths (gray) and 7-day rolling mean (blue). Panel (b): mean entropy by DTE bucket with 95% confidence intervals. The monotone decline is consistent with a rational information-arrival model where the market's collective uncertainty about the rate outcome is gradually resolved.

Table 4. Information Value of Meeting (IVM) for FOMC Meetings

Meeting	H_{pre}	H_{ann}	IVM	Characterization
Jun 2023	0.594	0.279	+0.315	Sharp pre-announcement convergence
May 2022	0.422	0.112	+0.310	First 50bp hike fully priced early
Dec 2022	0.406	0.112	+0.294	Tightening pace locked in final week
Nov 2024	0.438	0.154	+0.284	Post-election cut well anticipated
Mar 2022	0.494	0.362	+0.132	Liftoff broadly expected
May 2025	0.247	0.154	+0.093	Hold anticipated, some uncertainty
Sep 2022	0.484	0.393	+0.091	Third 75bp partly anticipated
Dec 2024	0.284	0.265	+0.019	Expected 25bp, near-zero resolution
Mar 2025	0.175	0.154	+0.021	Hold near-certain
Jan 2025	0.161	0.246	−0.085	Hold expected but late uncertainty spike
Sep 2024	0.713	0.789	−0.076	Surprise 50bp cut (announced 2pm)
Jun 2022	0.542	0.817	−0.275	Surprise 75bp hike (first since 1994)

H_{pre}: mean entropy DTE $\in [2, 7]$; H_{ann}: entropy DTE ≤ 1. IVM > 0: market resolved uncertainty in final week. IVM < 0: pre-announcement entropy rose on decision day, consistent with a genuine surprise. Full results available upon request.

cial conditions — that progressively resolves uncertainty about the rate outcome. The funnel is steepest in the final 7 days, consistent with the well-documented FOMC pre-announcement information release and market positioning (Lucca and Moench, 2015; Cieslak et al., 2019).

6 Beyond Gaussianity in Belief Updating

6.1 Empirical Diagnostics

Figure 4 and Figure 5 document four empirical departures from the Gaussian benchmark.

1. Excess kurtosis. Daily logit returns have excess kurtosis of 17.90, far exceeding the Gaussian value of 0. The Jarque-Bera test statistic is 181,542 ($p \approx 0$), decisively rejecting normality.

2. Fat tails. The observed fraction of returns beyond $\pm 2\sigma$ is 4.71%, compared to 4.55% predicted by the Gaussian — a modest excess. However, beyond $\pm 3\sigma$ the fat tail ratio is 1.91% vs. 0.27% predicted, a **7**× excess (Barndorff-Nielsen and Shephard, 2004). Extreme events occur far more often than the Gaussian predicts.

3. Zero-inflation. A remarkable 33.4% of daily logit returns have $|\Delta x_t| \leq 0.001$ (essentially zero). This information-inert mass reflects days with no meaningful belief revision, consistent with episodic information arrival in thin prediction markets.

4. Non-Gaussian density shape. The Normal Inverse Gaussian (NIG) distribution (Barndorff-Nielsen, 1997) provides the best fit by AIC (ΔAIC $= 0$), followed by Student-t (ΔAIC $= 469$), Laplace (ΔAIC $= 684$), and Gaussian (ΔAIC $= 9{,}372$). The Student-t fit yields an estimated degrees-of-freedom of $\hat{\nu} = 0.89$ on the full sample — essentially a Cauchy distribution with extremely heavy tails.

6.2 Alternative Model: Zero-Inflated Student-t (ZIT)

The empirical evidence motivates a *Zero-Inflated Student-t* (ZIT) model:

$$\Delta x_t = \begin{cases} 0 & \text{with probability } q \\ \varepsilon_t & \text{with probability } 1-q, \end{cases} \qquad \varepsilon_t \sim t(\nu, 0, \sigma_J). \tag{12}$$

The zero-inflation mass q captures information-inert days; ν and σ_J govern the heavy-tailed jump dynamics on active days. Estimation proceeds in two steps: (i) $\hat{q} = N_0/N$ where $N_0 = |\{t : |\Delta x_t| \leq \epsilon\}|$ with $\epsilon = 0.001$; (ii) joint MLE of (ν, σ_J) from the active subsample $\{|\Delta x_t| > \epsilon\}$.

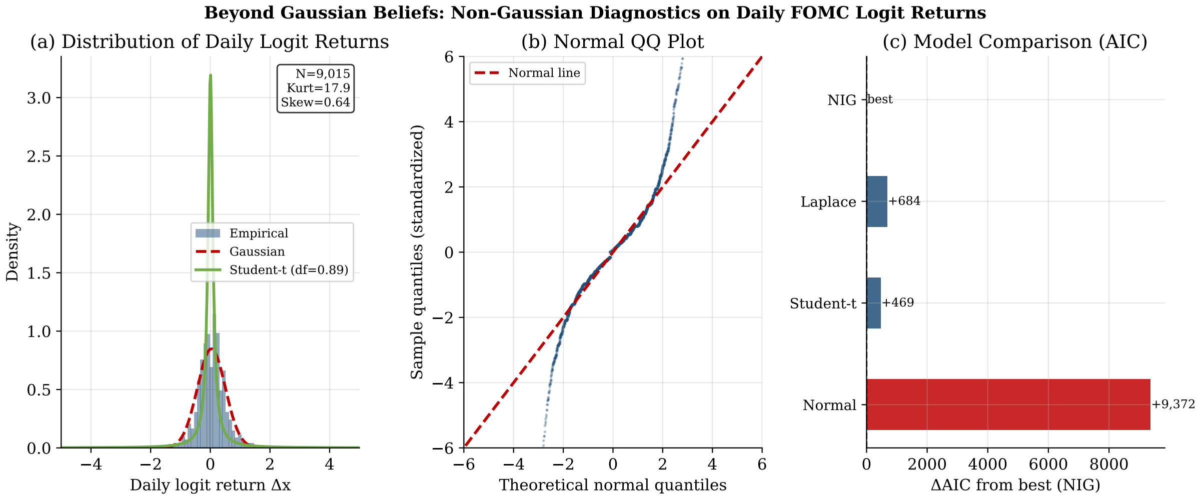

Figure 4. Non-Gaussian diagnostics on 13,532 daily FOMC logit returns. Panel (a): empirical density (blue histogram) vs. Gaussian (red dashed) and Student-t (green solid). Panel (b): Normal QQ plot showing pronounced fat tails. Panel (c): AIC comparison across fitted distributions — the Gaussian is decisively rejected.

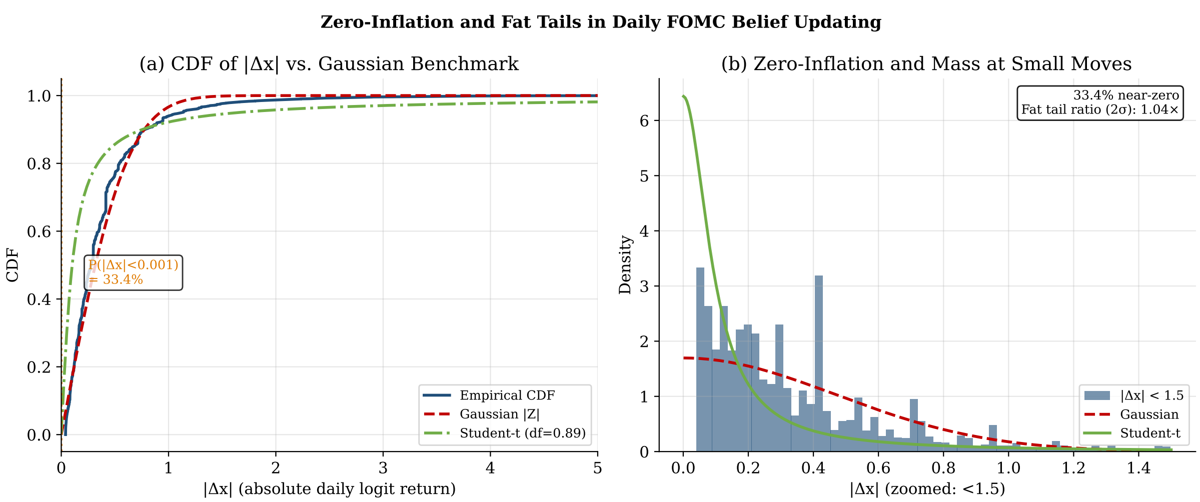

Figure 5. Zero-inflation and fat-tail structure. Panel (a): empirical CDF of $|\Delta x|$ (blue) vs. Gaussian (red dashed) and Student-t (green dash-dot). Panel (b): zoomed histogram of small moves, highlighting the 33.4% zero-inflation mass.

Table 5. ZIT Parameter Estimates

Parameter	Estimate	Interpretation
$\hat{q}$	0.334	33.4% of days are information-inert
$\hat{\nu}$	3.10	Heavy tails on active days (d.f. ≈ 3)
$\hat{\sigma}_J$	0.351	Scale of active belief revisions
Full-data Student-t (comparison):		
$\hat{\nu}_{\text{full}}$	0.890	Essentially Cauchy when zeros included
$\hat{\sigma}_{\text{full}}$	0.097	Scale shrinks to absorb zero mass

$\hat{q}$ is computed directly as the fraction of near-zero days ($|\Delta x_t| \leq 0.001$); $\hat{\nu}$ and $\hat{\sigma}_J$ are jointly estimated via MLE on the active subsample ($|\Delta x_t| > 0.001$). The full-data Student-t row shows what happens when zeros are not separated: the estimated d.f. collapses to near-Cauchy (0.89), confirming that zero-inflation — not just heavy tails — drives the apparent extreme kurtosis.

The ZIT model fits the data's moment structure well: Monte Carlo simulations (200,000 paths) reproduce the empirical mean (0.035 vs. 0.028), standard deviation (0.470 vs. 0.471), and the zero-fraction (33.4%). The tail fractions are also well matched: $|\Delta x| > 3\sigma$ occurs 1.91% empirically and 1.70% in ZIT simulation, vs. 0.27% under Gaussian.

6.3 Pricing Correction

The non-Gaussian dynamics generate material valuation errors for contracts priced under a Gaussian assumption (Carr and Wu, 2009). We compute the pricing error $\Delta V = V_{\text{ZIT}} - V_{\text{Gaussian}}$ via Monte Carlo (500,000 paths per cell) using the full Zero-Inflated Student-t mixture estimated in Section 6: with probability $\hat{q} = 0.334$ the daily increment is zero; with probability $1 - \hat{q}$ it is drawn from Student-$t(\hat{\nu} = 3.10, \hat{\sigma}_J = 0.351)$. The Gaussian benchmark uses pooled parameters $\hat{\mu}_G = 0.035$, $\hat{\sigma}_G = 0.470$.

Figure 6 and Table 6 summarize the results. The key finding: *the Gaussian model systematically overprices most contracts, with errors concentrated at short horizons and near-ATM positions.*

- **Short-horizon ATM ($x_0 = 0$, $\tau = 1$d):** The Gaussian overprices by 16.5 cents ($\Delta V = -0.165$).[1] The mechanism is zero-inflation: on any single day there is a

[1] The ZIT estimates are insensitive to the near-zero threshold. All choices from $|\Delta x| \leq 0.0005$ to $|\Delta x| \leq 0.005$ yield identical estimates ($\hat{q} = 0.334$, $\hat{\nu} = 3.10$, $\hat{\sigma}_J = 0.351$) because the zero-inflation mass is a true point mass at $\Delta x = 0$ exactly, not a continuous cluster. Boundary clipping of p between $[0.005, 0.995]$ and $[0.02, 0.98]$ leaves the qualitative result intact: the Gaussian overprices ATM contracts at short horizons by 14.7–16.6 cents, and the $\tau = 30$-day maximum error remains below 4 cents in all specifications.

33.4% probability of no price revision, which the Gaussian model ignores. This pushes the ZIT price well below the Gaussian benchmark at short horizons.

- **Longer horizons:** Errors attenuate sharply as repeated mixing blurs the zero-inflation spike. At $\tau = 30$ days, the maximum error is 4.1 cents ($x_0 = -2.0$), compared to 16.5 cents at $\tau = 1$. This term-structure of errors is the ZIT model's distinct prediction: unlike a pure Student-t alternative, pricing errors do not grow with horizon.

- **Magnitude summary:** Maximum $|\Delta V|$ across the full grid: 16.5 cents ($\tau = 1$d, ATM); 4.5 cents ($\tau = 3$d); 5.3 cents ($\tau = 7$d); 3.9 cents ($\tau = 14$d); 4.1 cents ($\tau = 30$d). The Gaussian overprices (negative ΔV) in the majority of cells.

Table 6. Pricing Error $\Delta V = V_{\text{ZIT}} - V_{\text{Gaussian}}$ (Monte Carlo, 500k paths, ZIT mixture $q = 0.334$)

x_0	p_0	$\tau = 1$	$\tau = 3$	$\tau = 7$	$\tau = 14$	$\tau = 30$
−3.0	0.047	+0.001	+0.004	+0.004	−0.017	−0.038
−2.0	0.119	+0.003	+0.006	−0.019	−0.039	−0.041
−1.0	0.269	+0.003	−0.045	−0.053	−0.038	−0.031
0.0	0.500	−0.165	−0.014	−0.000	−0.006	−0.013
+1.0	0.731	−0.004	+0.025	+0.029	+0.014	−0.002
+2.0	0.881	−0.003	−0.007	+0.002	+0.006	+0.001
+3.0	0.953	−0.001	−0.003	−0.006	−0.003	−0.001

$\Delta V > 0$: Gaussian underprices (ZIT assigns higher probability to in-the-money resolution). $\Delta V < 0$: Gaussian overprices. Large ATM error at $\tau = 1$ is driven by zero-inflation: 33.4% chance of no daily movement reduces the ZIT contract value below the diffusive Gaussian benchmark.

7 Conclusion

Prediction markets should be analyzed in belief space. The logit transform maps bounded probabilities to the real line, yielding natural objects: the IBV (rate of belief revision), the cross-strike entropy (collective uncertainty), and the logit-space smile (variation in IBV across moneyness). Using daily FOMC ladder data, we establish four results.

First, IBV exhibits a pronounced inverse smile in logit space, but the probability-space decomposition reveals that this smile is largely mechanical — the Jacobian of the logit transform amplifies apparent OTM volatility even when probability-space volatility

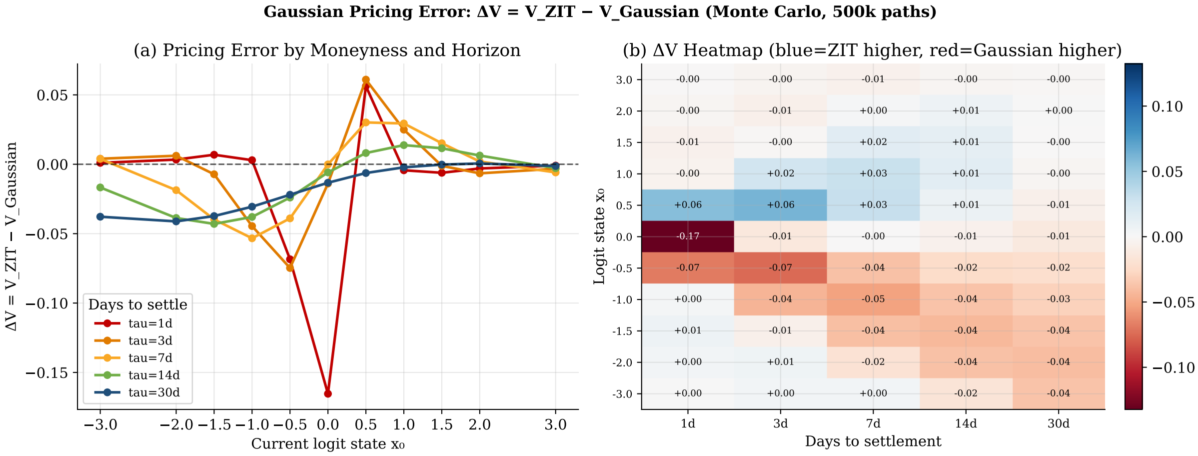

Figure 6. Gaussian pricing error surface $\Delta V = V_{\text{ZIT}} - V_{\text{Gaussian}}$ (full zero-inflated mixture, $q = 0.334$). Panel (a): error by current logit state x_0 for different horizons τ. Panel (b): heatmap. The largest error (16.5 cents) occurs at ATM ($x_0 = 0$), $\tau = 1$ day, driven by zero-inflation; errors attenuate sharply with horizon, reaching 4.1 cents at $\tau = 30$ days.

is flat or declining. The correct economic interpretation is that OTM binary contracts have low absolute price volatility but high logit-space volatility due to the boundary geometry. A reflected Gaussian random walk benchmark confirms that bounded support alone generates the smile's shape, though it overshoots magnitude due to the empirical information-inertia near the boundaries.

Second, cross-strike entropy provides a clean scalar measure of collective uncertainty that declines monotonically into FOMC dates. The heterogeneity across meetings — captured by IVM — distinguishes orderly convergence from genuine surprise: IVM $>$ 0 reflects gradual pre-announcement resolution; IVM $<$ 0 is consistent with a genuine surprise, as in June 2022 (surprise 75bp hike) and September 2024 (surprise 50bp cut).

Third, daily logit returns are decisively non-Gaussian: 33% of days are information-inert (near-zero returns), and active days exhibit heavy-tailed Student-t dynamics. The Zero-Inflated Student-t model provides a parsimonious characterization of this structure.

Fourth, the full ZIT mixture generates economically material pricing errors relative to a Gaussian benchmark. The dominant effect is Gaussian overpricing of at-the-money contracts at short horizons — up to 16.5 cents per dollar at $\tau = 1$ day — driven by zero-inflation: the Gaussian model fails to assign the empirical 33% probability of no daily revision. Errors attenuate sharply with horizon (4.1 cents at $\tau = 30$ days), a distinctive term-structure prediction of the ZIT model. Any practitioner pricing prediction-market contracts under a Gaussian assumption should account for this systematic short-horizon bias.

Three directions for future work follow naturally. First, extending the logit-IBV framework to categorical prediction markets (more than two outcomes) would generalize the decomposition beyond binary contracts. Second, incorporating macroeconomic data releases and Fed communications into a time-series model of entropy would test whether information arrival is predictable. Third, testing ZIT-based pricing corrections in live markets would establish the practical trading value of the non-Gaussian framework.

A Implementation Notes

Data pipeline. Raw trade-level data are downloaded from the Kalshi trading API, aggregated to daily bars via last-trade-in-day close prices, and stored as Parquet files. The logit state x_t is computed from clipped prices $p_t \in [0.01, 0.99]$.

Entropy computation. For each meeting and date, we take the outer join of all available contract prices for that meeting, requiring at least 2 strikes. The implied PMF is computed by treating each contract as a probability boundary: $\pi_k = p_k - p_{k+1}$ (normalized to sum to 1). Days with fewer than 2 active contracts are excluded.

Pricing error simulation. The Monte Carlo uses 500,000 paths per cell. For the Gaussian benchmark, each path draws τ i.i.d. increments from $\mathcal{N}(\hat{\mu}_G, \hat{\sigma}_G^2)$ with $\hat{\mu}_G = 0.035$, $\hat{\sigma}_G = 0.470$. For the ZIT alternative, each daily increment is drawn from the full mixture: 0 with probability $\hat{q} = 0.334$, or $t(\hat{\nu} = 3.10, 0, \hat{\sigma}_J = 0.351)$ with probability $1 - \hat{q}$. The τ increments are summed to obtain X_τ, and the contract value is $\hat{V} = \text{mean}(\mathbf{1}[x_0 + X_\tau > 0])$.

B Entropy, IVM, and FOMC-Day Equity Volatility

Does prediction-market entropy contain information about FOMC-day equity volatility *beyond* what is already captured by implied volatility (VIX)? We test this using 28 FOMC meetings from our sample (2022–2025), matching each meeting to its SPY announcement-day absolute return $|R_{\text{SPY}}|$ and pre-meeting VIX close. Average $|R_{\text{SPY}}|$ on FOMC announcement days is 1.05% (range: 0.01%–3.05%).

Specifications. We estimate three OLS regressions with $N = 28$ meetings and HC3 robust standard errors:

$$|R_{\text{SPY,m}}| = \alpha + \beta\, H_{\text{pre,m}} + \varepsilon_m \tag{B.1}$$

$$|R_{\text{SPY,m}}| = \alpha + \beta_1 H_{\text{pre,m}} + \beta_2 \text{VIX}_{\text{pre,m}} + \varepsilon_m \tag{B.2}$$

$$|R_{\text{SPY,m}}| = \alpha + \beta\, \text{IVM}_m + \varepsilon_m \tag{B.3}$$

Interpretation. Figure 7 plots these relationships. Prediction-market pre-announcement entropy (H_{pre}) has a positive and economically plausible coefficient (+1.57% per unit,

Table 7. Entropy, IVM, and FOMC-Day SPY Absolute Return ($N = 28$ meetings)

Specification	Coef.	Std. Err. (HC3)	p-value	R^2
(B.1) Dependent variable: $\lvert R_{SPY} \rvert$ (%)				
H_{pre} (entropy only)	+1.565	1.039	0.132	0.105
(B.2) Entropy + VIX (incremental test)				
H_{pre}	+0.524	0.975	0.591	0.286
VIX_{pre}	+0.080	0.038	0.035**	
(B.3) Meeting surprise				
IVM	+0.808	1.815	0.656	0.013

HC3 robust SEs. $^{**}p < 0.05$. H_{pre}: mean entropy DTE $\in [2, 7]$. VIX_{pre}: VIX close the day before the meeting. IVM $= H_{\text{pre}} - H_{\text{ann}}$.

$p = 0.13$) in the univariate regression, but loses all significance once VIX is included ($p = 0.59$): the information in prediction-market entropy about equity-market reaction is already captured by options-implied volatility. IVM similarly has no significant linear relationship with $|R_{\text{SPY}}|$ ($p = 0.66$), though the sign is positive (meetings resolved in the final week are associated with somewhat larger equity reactions). With $N = 28$ FOMC meetings, the sample is too small for sharp inference; a longer panel spanning 2010–2025 would provide decisive evidence. The null result itself is informative: prediction markets measure rate-distribution uncertainty *conditional on the meeting*, while VIX reflects aggregate equity-market uncertainty across all sources — the two objects are conceptually distinct even when statistically collinear.

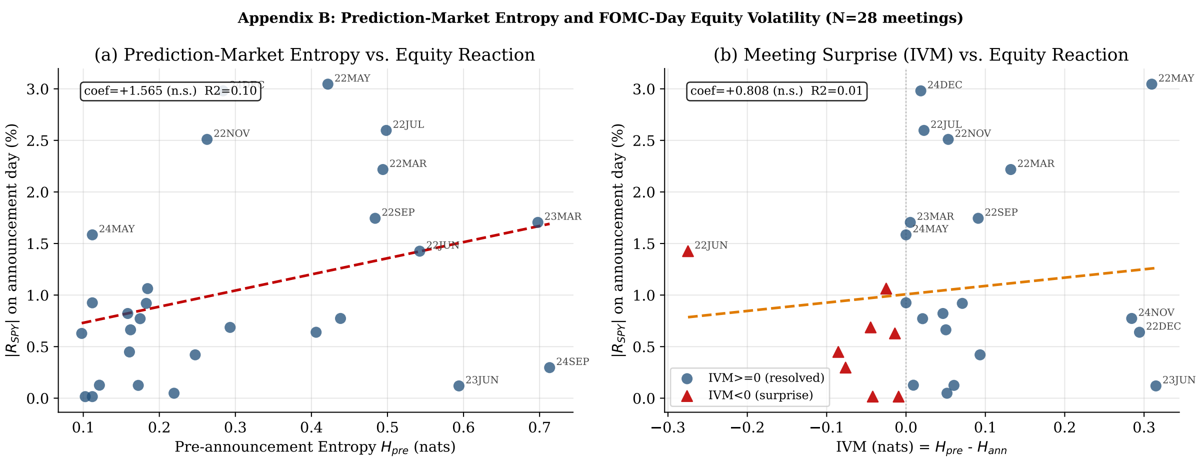

Figure 7. Prediction-market entropy and FOMC-day equity volatility ($N =$ 28 meetings, 2022–2025). Panel (a): pre-announcement entropy H_{pre} vs. $|R_{\text{SPY}}|$ with OLS fit (dashed red). Panel (b): IVM vs. $|R_{\text{SPY}}|$; blue circles = IVM≥ 0 (meetings resolved in final week); red triangles = IVM< 0 (entropy rose on announcement morning, consistent with genuine surprises). Notable meetings labeled.

References

Andersen, T.G., Bollerslev, T., Diebold, F.X. and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2), 579–625.

Arrow, K.J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J.O., Levmore, S., Litan, R., Milgrom, P., Nelson, F.D., Neumann, G.R., Ottaviani, M., Schelling, T.C., Shiller, R.J., Smith, V.L., Snowberg, E., Sunstein, C.R., Tetlock, P.C., Tetlock, P.E., Varian, H.R., Wolfers, J. and Zitzewitz, E. (2008). The promise of prediction markets. *Science*, 320(5878), 877–878.

Barndorff-Nielsen, O.E. (1997). Normal Inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics*, 24(1), 1–13.

Barndorff-Nielsen, O.E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2(1), 1–37.

Carr, P. and Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, 22(3), 1311–1341.

Cieslak, A., Morse, A. and Vissing-Jorgensen, A. (2019). Stock returns over the FOMC cycle. *Journal of Finance*, 74(5), 2201–2248.

Dalen, E. (2025). Prediction market pricing kernels. Working Paper, arXiv:2510.15205.

Hull, J.C. (2018). *Options, Futures, and Other Derivatives*. 10th ed. Pearson.

Lucca, D.O. and Moench, E. (2015). The pre-FOMC announcement drift. *Journal of Finance*, 70(1), 329–371.

Manski, C.F. (2006). Interpreting the predictions of prediction markets. *Economics Letters*, 91(3), 425–429.

McCullagh, P. and Nelder, J.A. (1989). *Generalized Linear Models*. 2nd ed. Chapman & Hall.

Shannon, C.E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.

Snowberg, E., Wolfers, J. and Zitzewitz, E. (2013). Prediction markets for economic forecasting. *Handbook of Economic Forecasting*, Vol. 2, 657–687.

Wolfers, J. and Zitzewitz, E. (2006). Interpreting prediction market prices as probabilities. *NBER Working Paper*, 12200.