Market Efficiency in Prediction Markets: A Comparison with Derivatives*

M. Fabi,† R. Marfè,‡ V. Ruffo,§ L. Schönleber¶

29th July 2026

Abstract

We study pricing efficiency in decentralized prediction markets by comparing Polymarket prices with option-implied prices derived from derivatives. We analyze nearly 5,000 Polymarket contracts written on Bitcoin (BTC) and Ethereum (ETH), two assets for which prediction-market contracts and liquid derivatives markets coexist. We find that Polymarket prices closely track option-implied prices, but systematic deviations remain: they are concentrated in tail and barrier contracts, vary with sentiment, volatility, demand, and market frictions, and imply more probability mass on extreme price movements. To guide our empirical investigation, we introduce a parsimonious limits-to-arbitrage model that formalizes equilibrium behavior in decentralized prediction markets.

Keywords: Prediction Markets, Market Efficiency, Risk-Neutral Pricing, Crypto Derivatives, Behavioral Biases, Decentralized Finance

JEL Classification Codes: G14, D84, G12, G13

*We received helpful comments and suggestions from Roman Kozhan, Dmitry Makarov (discussant), Da-Hea Kim (discussant), Ciamac Mollaemi, and Sebastian Mueller. We also thank the participants of the 14th International Moscow Finance Conference, the Frankfurt School PhD Brown Bag, the DeFi Workshop IP Grenoble & Paris, the 2025 Global AI Finance Research Conference, the DeFi & Crypto Workshop Scuola Normale Superiore Pisa, the 5th Workshop on Decentralized Finance (DeFi) in Association with Financial Cryptography 2026, the XXVII Workshop on Quantitative Finance, the International Conference – Tech 4 Finance 3: AI and Blockchain, the Designing DeFi Conference (D^2 2026), the Frontiers in DeFi 2026, the 3rd Structured Retail Products and Derivatives Conference, the Future Finance Fest (3f), and the 8th Future of Financial Information Conference. This project benefited from the invaluable assistance of Matteo Repici (École Polytechnique).

†Telecom Paris, 19 Pl. Marguerite Perey, 91120 Palaiseau, France. *michele.fabi@ensae.fr.*

‡Collegio Carlo Alberto and University of Turin, Piazza Vincenzo Arbarello 8, 10122 Turin, Italy. *roberto.marfe@carloalberto.org*

§Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt am Main, Germany. *v.ruffo@fs.de.*

¶Collegio Carlo Alberto and University of Turin, Piazza Vincenzo Arbarello 8, 10122 Turin, Italy. *lorenzo.schoenleber@carloalberto.org.*

1 Introduction

Prediction markets increasingly operate as financial markets. This development is reflected in the rapid growth of the two largest platforms, Polymarket and Kalshi, whose combined monthly trading volume reached approximately $25 billion in April 2026.[1] These platforms allow participants to trade fully collateralized contingent claims whose prices are commonly interpreted as market-implied probabilities. Many of these contracts have direct counterparts in derivatives markets: a claim paying one dollar if an asset exceeds a threshold at maturity resembles a digital option, while a claim paying one dollar if the asset reaches a threshold before maturity resembles a barrier-style derivative. This correspondence creates a natural finance question: do decentralized prediction markets price event risk consistently with the state prices implied by established derivatives markets?

We study this question by comparing cryptocurrency contracts traded on Polymarket with option-implied prices derived from European-style call and put options traded on Deribit, a leading centralized exchange for cryptocurrency options. Bitcoin (BTC) and Ethereum (ETH) provide a particularly useful setting because the same underlying risks are traded simultaneously through prediction-market contracts and liquid options. This parallel-market environment allows us to compare economically similar claims across two venues that differ substantially in market design, settlement, infrastructure, and trader composition. The option-implied benchmark is a risk-neutral valuation rather than an estimate of the event's objective probability. Accordingly, our primary objective is cross-market risk-neutral pricing consistency.

Our analysis covers 4,860 Polymarket contracts and approximately 4.4 million trades. We identify four contract types. *Above* contracts pay if the underlying exceeds a threshold at maturity, while *Range* contracts pay if it finishes within a specified interval. These terminal-value claims correspond to European digital options. *Reach* and *Dip* contracts instead pay if the underlying crosses an upper or lower barrier before expiration and therefore resemble first-passage or barrier-style digital claims.

We construct option-implied prices (OIPs) for these contracts using Deribit option surfaces and the double-exponential jump-diffusion model of Kou (2002). The model provides a common pricing framework for both terminal and path-dependent payoffs while allowing for asymmetric jumps and fat-tailed returns. Above and Range contracts are priced as terminal digital claims,

[1] https://www.theblock.co/data/decentralized-finance/prediction-markets

whereas Reach and Dip contracts are priced using barrier-crossing probabilities. Because barrier prices depend more strongly on assumptions about the underlying price process, inference for Reach and Dip contracts is more model-dependent than for Above and Range contracts. We validate the terminal-price implications of the model against risk-neutral densities recovered using the model-free approach of Breeden and Litzenberger (1978).

Our first finding is that Polymarket prices are strongly aligned with OIPs. Contract-level time-series correlations between CLOB and OIP prices average 0.88 for Above contracts, 0.77 for Range contracts, 0.88 for Reach contracts, and 0.86 for Dip contracts. Thus, prediction-market prices contain substantial information about the same underlying risks priced in options markets. The central result is not that Polymarket is disconnected from derivatives, but that deviations from derivative-implied prices are systematic and depend on the contract's payoff structure and market conditions.

Despite this broad alignment, Polymarket prices are generally higher than their option-implied counterparts. In the pooled BTC–ETH sample, the equal-contract mean CLOB–OIP wedge is 0.716 percentage points, and 72.6% of contract-level mean wedges are positive. The largest signed wedges occur for Reach and Dip contracts, indicating that the strongest disagreements arise for salient path-dependent events and tail outcomes.

The relationship between moneyness and the pricing wedge is payoff-specific. For Above and Range contracts, wedges are larger for more out-of-the-money (OTM) claims. For Reach and Dip contracts, by contrast, wedges increase as the underlying approaches the barrier. Additional evidence confirms this pattern: trading activity and wedges increase for Reach and Dip contracts, while wedges decline for Above and Range contracts as the underlying approaches the relevant terminal payoff boundary. These results are consistent with traders placing disproportionate attention on salient tail events.

The evolution of wedges over the contract life is also consistent with gradual learning and the resolution of disagreement. Wedges are generally largest near contract inception and decline toward maturity. The pattern is strongest for Reach contracts, whose average wedge falls from 0.0380 at the beginning of the contract life to approximately zero near expiration. Above and Range contracts display the same broad decline. Dip contracts also exhibit a substantial reduction in average wedges over the contract life, although their maturity coefficient is less stable in the panel regressions.

Market conditions further shape the degree of cross-market disagreement. Uncertainty, meas-

ured by the crypto volatility index, is associated with larger wedges for Above, Reach, and Dip contracts. Differences in the underlying reference prices across Deribit and Polymarket are particularly relevant for Range, Reach, and Dip contracts, while blockchain-related operational risk is most clearly associated with Above-contract wedges. Fear sentiment is associated with larger wedges for Above, Reach, and especially Dip contracts, while net buying pressure widens wedges for Above and Dip contracts. Announcements and weekend trading also matter, but their effects differ across payoff structures. These patterns suggest that pricing deviations reflect a combination of directional demand, the salience of events, and limits to cross-market integration, with the relative importance of each channel depending on the contract's payoff structure.

To organize these findings, we develop a parsimonious limits-to-arbitrage framework. A representative arbitrageur provides liquidity while facing inventory and execution costs. Demand shocks move Polymarket prices away from OIPs, while arbitrageurs only partially absorb these shocks because trading and warehousing positions are costly. We relate the latent demand shock to behavioral probability weighting, disagreement and learning, sentiment, and informed trading. The framework yields testable predictions for the effects of moneyness, maturity, directional demand, volatility, and market frictions. The empirical results broadly support these channels, while also showing that their importance and direction depend on contract design.

We next examine who trades in these markets. Trading activity and realized outcomes are highly concentrated across wallets. In the pooled resolved sample, the top 1% of wallets account for 73.9% of trades and 89.2% of notional. Similar to Akey et al. (2026), among profitable wallets, the top 1% capture 84.2% of total gains. Our option-implied benchmark allows us to go beyond realized outcomes and measure the pricing wedges associated with each wallet's trades. Larger wallets trade more in contracts with larger absolute CLOB–OIP wedges, linking differences across market participants directly to cross-market pricing deviations. These findings connect the contract-level pricing evidence to the distribution of activity across market participants.

Finally, we ask whether contract-level differences aggregate into distinct market-implied distributions of risk. We calibrate the Kou (2002) model directly to Polymarket contracts and construct daily Polymarket risk-neutral distributions, which we compare with corresponding distributions extracted from Deribit. The two markets agree on the broad shape of risk: both imply negative expected returns, negative skewness, and fat tails. Polymarket, however, assigns more probability outside Deribit's central return interval, particularly during market events. At the same time, its skewness and kurtosis are often less extreme than those implied by Deribit.

The distributional disagreement is therefore not a simple uniform increase in tail risk but reflects differences in how the two markets price particular regions of the return distribution.

Our paper makes four contributions. First, we introduce a cross-market approach to evaluating decentralized prediction markets. Existing work primarily studies whether prediction markets aggregate information or forecast outcomes accurately (Axén and Cortis, 2020; McCullough, 2022; Rothschild and Sethi, 2016; Wolfers and Zitzewitz, 2006). We instead treat prediction contracts as financial claims and compare their prices with risk-neutral valuations implied by a liquid derivatives market. This shifts the focus from realized forecasting accuracy to cross-market pricing consistency.

Second, we provide a unified mapping between prediction contracts and derivative payoffs. Above and Range contracts correspond to terminal digital claims, while Reach and Dip contracts correspond to path-dependent barrier claims. This mapping allows us to study how payoff design affects cross-market alignment and shows that deviations are especially pronounced where outcomes are tail-oriented, thresholds are salient, and valuation is path-dependent.

Third, we connect pricing wedges to market mechanisms and participant behavior. Building on limits-to-arbitrage and demand-based approaches (Bollen and Whaley, 2004; Brunnermeier and Pedersen, 2009; Shleifer and Vishny, 1997), we show that wedges vary systematically with moneyness, maturity, volatility, sentiment, directional demand, announcements, and market-friction proxies. We also document how trading activity, P&L, and exposure to larger pricing wedges are distributed across wallets.

Fourth, we extend the comparison from individual contracts to the full implied distribution. By estimating Polymarket-implied risk-neutral distributions and comparing them with Deribit distributions, we show that localized contract-level wedges have aggregate implications for the pricing of crypto tail risk.

Our interpretation is deliberately narrower than a blanket claim that one market is efficient and the other is not. Deribit OIPs are risk-neutral valuations rather than objective probabilities, and differences between the two markets may reflect risk premia, demand, liquidity, collateral, settlement conventions, or benchmark uncertainty. Moreover, the pricing of Reach and Dip contracts is more model-dependent than that of terminal claims. Subject to these qualifications, the evidence shows that decentralized prediction markets are highly informative and closely connected to derivatives markets, but that the two venues are not fully integrated: systematic differences remain in particular payoff structures and market states.

Literature Review

This paper contributes to the market microstructure literature on prediction markets by comparing the prices of Polymarket contracts and cryptocurrency options, an approach not taken in existing work to our knowledge. Previous studies typically evaluate prediction markets by asking whether prices forecast outcomes (Berg et al., 2008; Forsythe et al., 1992; Wolfers and Zitzewitz, 2004), are well calibrated—that is, whether events assigned a given market-implied probability occur with approximately that frequency (Wolfers and Zitzewitz, 2004, 2006), or outperform polls, surveys, and professional forecasts (Diercks et al., 2026; Forsythe et al., 1992; Rabetti et al., 2026). We instead examine whether contingent claims are priced consistently across prediction and derivatives markets. Beyond individual contracts, we compare the risk-neutral distributions implied by Polymarket with those extracted from cryptocurrency options.

The traditional motivation for prediction markets is that prices can aggregate information dispersed across heterogeneous participants. Organized election betting already played an important forecasting role in U.S. presidential races before modern polling (Rhode and Strumpf, 2004). The Iowa Electronic Markets later provided a modern real-money setting in which to study both forecasting and price formation (Berg et al., 2008; Forsythe et al., 1992), while laboratory experiments showed that suitably designed markets can aggregate private information into informative prices (Plott and Sunder, 1982, 1988). Research therefore shifted from asking whether prediction markets forecast accurately to examining how prices respond to information and which trading strategies contribute to price discovery (Gauriot and Page, 2025; Rothschild and Sethi, 2016). One influential explanation is that a relatively small group of sophisticated or marginal traders disciplines prices by correcting the errors of less informed participants (Forsythe et al., 1999; Gómez-Cram et al., 2026b; Oliven and Rietz, 2004).

The expansion of Kalshi and Polymarket has revived this literature by generating high-frequency data across a wide range of events. Earlier work on online prediction markets was pioneered partly in computer science, with an emphasis on artificial markets and automated market makers (Chen and Pennock, 2007; Hanson, 2003; Pennock et al., 2001). Diercks et al. (2026) show that Kalshi contracts provide continuously updated macroeconomic forecasts that are competitive with surveys and conventional market-based measures, while Angelini (2026) shows that they also reveal changes in uncertainty and tail risk. Related work uses blockchain prediction markets to study the incorporation of macroeconomic and central-bank information

(Eichengreen et al., 2025). In a corporate setting, Gómez-Cram et al. (2026a) and Rabetti et al. (2026) show that prediction-market prices provide informative earnings expectations and can outperform professional analysts. Akey et al. (2026) show that a small group of skilled traders drives much of Polymarket's accuracy. This evidence supports the usefulness of prediction markets but leaves open whether valuations are consistent across markets trading related risks.

Our approach connects this literature to state-contingent pricing and limits to arbitrage by mapping prediction contracts to digital and barrier-style claims valued using option prices (Aït-Sahalia and Lo, 1998; Aït-Sahalia et al., 2001; Axén and Cortis, 2020; Bakshi and Madan, 2000; Breeden and Litzenberger, 1978). In frictionless markets satisfying no-arbitrage, claims with identical state-contingent payoffs and settlement conditions must have the same price. In practice, however, transaction costs, capital constraints, and execution risk may prevent arbitrageurs from eliminating price differences (Gromb and Vayanos, 2010; Lamont and Thaler, 2003; Mitchell et al., 2002; Shleifer and Vishny, 1997). We therefore test whether pricing wedges between Polymarket and cryptocurrency options are systematic and whether they widen when arbitrage is more costly or less elastic.

A separate strand of research, rooted in behavioral finance, documents systematic distortions in the pricing of low-probability outcomes. In this context, the favorite–longshot bias refers to the tendency for unlikely events to be overpriced and highly likely events to be underpriced (Ali, 1977; Thaler and Ziemba, 1988). Subsequent work relates this pattern to nonlinear probability weighting (Kahneman and Tversky, 1979; Prelec, 1998; Snowberg and Wolfers, 2010; Tversky and Kahneman, 1992). Recent evidence from Polymarket and Kalshi suggests that these distortions vary with contract design, price level, and time to resolution (Deleep et al., 2026). Tail pricing may also reflect liquidity providers' protection against informed trading rather than purely behavioral biases (Shin, 1991, 1992, 1993). This literature is especially relevant for path dependent contracts, where the largest discrepancies arise for extreme path-dependent outcomes.

2 Institutional Setting and Contract Mapping

2.1 Polymarket

Polymarket operates a hybrid-decentralized CLOB for trading binary-event claims. Unlike centralized betting platforms or automated market makers (AMMs), it supports peer-to-peer trading through an off-chain matching engine and on-chain non-custodial settlement on Polygon. Par-

ticipants submit limit orders specifying prices and quantities, which are matched off-chain and settled on-chain.

Each market consists of complementary 'Yes" and 'No" tokens. The pair is fully collateralized, so their combined price equals 1.00 USDC. At resolution, the winning token redeems for 1.00 USDC and the losing token becomes worthless. Prices lie between 0 and 1 and are commonly interpreted as market-implied probabilities; for example, a "Yes" price of 0.42 indicates a 42% implied probability. Before resolution, participants may sell by accepting the prevailing bid or posting a limit order.

2.2 Contract Types and Derivative Mapping

Polymarket is a decentralized prediction market where users trade USDC-denominated claims on real-world events. For cryptocurrency markets, we identify four main contract types:

1. **Above Bets**: These are binary options structured around whether BTC will exceed a specified threshold by a fixed maturity date (e.g., "Will BTC be above $60,000 on July 31?"). These contracts resolve to "Yes" if the spot price at the market's expiry exceeds the stated level, and "No" otherwise.

2. **Range Bets**: These contracts specify whether BTC will lie within a predefined price interval at maturity (e.g., "Will BTC be between $50,000 and $60,000 on July 31?"). They effectively represent range digital options, with resolution contingent on the terminal price remaining strictly within the bounds.

3. **(Knock-In) Reach Bets**: These are forward-looking contracts on whether BTC will hit or exceed a particular price at any point before the stated expiry (e.g., "Will BTC reach $65,000 before August 1?"). They resemble American-style barrier options that activate and resolve immediately upon breach.

4. **(Knock-In) Dip Bets**: Similar in structure to the Reach contracts but applied to downward price movements, these markets ask whether BTC will dip below a certain price level before expiry (e.g., "Will BTC fall below $45,000 before August 1?"). These are useful for capturing market expectations of short-term downside risk.

Table 2.1 summarizes the derivative mapping. Above and Range contracts correspond to terminal digital claims, while Reach and Dip contracts resemble barrier digitals and are therefore

more model-dependent.

Contract	Payoff	Derivative Analog	Model Dependence	Considerations
Above	Pays 1 if $S_T > K$	European digital call	Low	Terminal claim; depends mainly on strike interpolation.
Range	Pays 1 if $K_L < S_T < K_U$	Range digital	Low	Terminal claim based on two strike probabilities.
Reach	Pays 1 if $S_t \geqslant K$ at any time before T	Upper-barrier digital	High	Path-dependent; depends on dynamics and barrier monitoring.
Dip	Pays 1 if $S_t \leqslant K$ at any time before T	Lower-barrier digital	High	Path-dependent; depends on dynamics and barrier monitoring.

Table 2.1: Contract Types and Derivative Mapping. This table maps the four Polymarket contract types to economically comparable derivative claims. Above and Range contracts depend only on the terminal price, whereas Reach and Dip contracts are path-dependent and therefore more model-dependent.

Polymarket markets are resolved through an on-chain oracle using externally verifiable data. Crypto contracts rely primarily on timestamped Binance prices: Above and Range contracts use the price at maturity, while Reach and Dip contracts use 1-minute candles. A typical Reach specification reads: "This market will immediately resolve to *Yes* if any Binance 1-minute candle for Bitcoin (BTC/USDT) between July 1, 2025, 00:00 and July 31, 2025, 23:59 ET has a final high price of $200,000 or higher. Otherwise, this market will resolve to *No*".

2.3 What Kind of Efficiency Do We Test?

Market efficiency can refer to forecasting accuracy, risk-neutral pricing consistency, or no-arbitrage efficiency. The latter asks whether price differences remain after accounting for trading costs, settlement frictions, and model uncertainty.

Our analysis primarily tests risk-neutral pricing consistency. We compare Polymarket prices with option-implied prices derived from Deribit, rather than evaluating whether Polymarket prices equal the physical probabilities of the corresponding events. The CLOB–OIP wedge therefore measures cross-market pricing disagreement and should not, by itself, be interpreted as an executable arbitrage opportunity.

The comparison is cleanest for Above and Range contracts, which map to terminal digital claims. Reach and Dip contracts are path-dependent, so their benchmarks depend more strongly on assumptions about the underlying price dynamics.

3 Data

This section describes the data sources and variable construction. Our analysis combines option-implied information from Deribit, obtained via Amberdata, with high-frequency prediction-market data from Polymarket, and is supplemented by measures of blockchain risk, macroeconomic announcements, sentiment, and cross-exchange price dislocations. We use Deribit surfaces to calibrate the Kou (2002) model and recover daily risk-neutral densities (RNDs) for BTC and ETH, which, in turn, provide benchmark prices for the contracts studied in the paper. These measures are then linked to hourly Polymarket outcomes and trading activity.

Sample timeline. The analysis uses three samples. Deribit surface data span January 1, 2023, to December 16, 2025, with Kou estimates beginning on January 8, 2023, because of the seven-day rolling window. The contract-level CLOB–OIP sample runs from January 23, 2024, to November 30, 2025, with BTC entering on March 17, 2024, and ETH on January 23, 2024. The daily RND comparison covers October 4, 2024, to November 27, 2025, for BTC and May 1, 2025, to November 27, 2025, for ETH.

3.1 Surface Data for Option-Implied Price (OIP) Construction

We obtain option data from Amberdata[2], which provides hourly fixed-delta implied volatility (IV) surfaces for BTC and ETH options at standardized maturities (e.g., 7 or 30 days). The data are collected from Deribit, the largest exchange for trading cryptocurrency options. On Deribit, options are cash-settled and quoted directly in the relevant cryptocurrency, not in USD. The surface consist of nine delta levels for both call and put options, ranging from (in absolute terms) 0.05 for far-OTM contracts to 0.5 for at-the-money (ATM) contracts. These surfaces are used to construct the option-implied benchmark prices described in Section 3.2.

3.2 Option-Implied Prices (OIP)

This section describes the calibration of the Kou (2002) double-exponential jump-diffusion model used to obtain parametric RNDs for BTC and ETH. The model flexibly captures asymmetric tail behavior through separate parameters for upward and downward jumps, while remaining sufficiently parsimonious to permit stable daily estimation. Because the risk-neutral drift is pinned down by the martingale condition, the free parameter vector reduces to $\vartheta = (\sigma, \lambda, p, \eta_1, \eta_2)$,

[2]https://www.amberdata.io/

where σ is diffusive volatility, λ is jump intensity, p is the probability of an upward jump, and η_1 and η_2 govern the magnitudes of positive and negative jumps, respectively.

To construct OIPs, we estimate the model daily using the Deribit surface archive. For comparability with the Polymarket calibration in Section 8.1, we report summary statistics for the overlapping sample.

Calibration is performed daily using 7-day constant maturity surfaces (Section 3.1). For each moneyness, we recover the corresponding strike using the spot price, the risk-free rate, the time to maturity, and the option type (call or put). European call and put option prices are then calculated under the risk-neutral Kou (2002) specification using the pricing routines described in Appendix A.3. To reduce noise, each daily estimate uses a 7-day rolling window of daily 7-day constant-maturity surfaces. With nine call-and-put delta levels per surface, this yields about 126 option quotes per day.

On each day, the parameters $\vartheta = (\sigma, \lambda, p, \eta_1, \eta_2)$ are chosen to minimize the weighted sum of squared pricing errors:

$$\min_{\vartheta} \sum_{j=1}^{N} \omega_j \left(\frac{V_j^{\text{model}}(\vartheta) - V_j^{\text{mkt}}}{S_{0,j}} \right)^2, \qquad \omega_j = \frac{1}{1 + |K_j/S_{0,j} - 1|^2}, \tag{1}$$

where N represents the total number of options in the calibration sample, $V_j^{\text{model}}(\vartheta)$ and V_j^{mkt} denote the model-implied price and the observed market price for the j-th option, respectively. Furthermore, $S_{0,j}$ is the spot price of the underlying asset, and K_j represents the option's strike price. The weights ω_j are specified so that pricing errors are normalized by the spot price and near-the-money contracts receive more weight. The admissible parameter space is bounded by $\sigma \in [0.01, 2]$, $\lambda \in [0, 10]$, $p \in [0.01, 0.99]$, $\eta_1 \in [1.01, 100]$, and $\eta_2 \in [0.01, 100]$. These bounds enforce economically meaningful estimates and, in particular, $\eta_1 > 1$ ensures that the exponential moment required by the martingale condition is finite.[3]

Table 3.1 summarizes the distribution of the estimated parameters for BTC and ETH on the same asset-specific dates used by the Polymarket RND tables below. BTC exhibits lower diffusive volatility than ETH, with an interquartile range for σ of 0.317–0.431, compared to 0.533–0.591 for ETH. BTC also displays lower jump intensity, with a mean λ equal to 6.912, compared to 8.145 for ETH. At the same time, BTC features substantially larger values of η_1

[3]We perform the optimization using L-BFGS-B under box constraints. Because the calibration is repeated across adjacent dates, we use the previous successful estimate as a warm start and apply small perturbations when the optimizer returns the same solution.

and η_2, indicating thinner jump tails on both sides than in ETH. ETH, by contrast, has a lower average probability of upward jumps, with a mean p equal to 0.318 compared to 0.411 for BTC, suggesting a more downside-oriented jump component. Overall, the calibration delivers economically interpretable and comparatively stable parameter estimates.[4]

Table 3.2 reports the average first four moments of the RND for BTC and ETH. Both assets exhibit negative annualized means, negative skewness, and kurtosis well above the Gaussian benchmark, indicating downside asymmetry and fat-tailed risk-neutral distributions. ETH has a more negative mean and higher volatility than BTC, while BTC displays slightly more negative skewness and higher kurtosis. Overall, the results suggest elevated risk-neutral uncertainty for both assets, with ETH showing greater dispersion and BTC somewhat stronger tail risk.[5]

Figure C.3 illustrates substantial time variation in the 7-day RNDs for both BTC and ETH. In both markets, the density is centered near zero returns, but its dispersion and skewness change markedly over time. BTC displays more pronounced spikes and stronger shifts in tail mass, while ETH appears somewhat flatter and more diffuse.

Once calibrated, we calculate benchmark prices for the four contract types. For Above and Range contracts, we rely on Section A.2, while for Reach and Dip contracts, we follow Section A.4. Pricing is done hourly using the calibrated parameters from the previous day.

As a validation, Section 9.2 compares the seven-day Kou density with a model-free Breeden and Litzenberger (1978) density recovered from the same Deribit surface. The two benchmarks are closely aligned near the center of the distribution, while Kou assigns more probability to both tails and consequently higher prices to OTM Above contracts. This supports the use of Kou as a flexible pricing benchmark, particularly for the path-dependent Reach and Dip contracts.

3.3 Polymarket

We collect prediction market data from Polymarket, which operates on the Polygon blockchain and offers a variety of markets, including those related to cryptocurrency price movements. We retrieve the data via the Polymarket API: First, we query the "markets" endpoint to retrieve a list of all active and closed markets, along with metadata (titles, timestamps, tags, resolution status, etc.). From this list, we filter for cryptocurrency-specific markets—specifically those referencing

[4] Figures C.1 and C.2 plot the parameters over time.

[5] Figures III.1 and III.2 report the time-series behavior of the quantile bands, annualized standard deviation, skewness, and kurtosis of the Deribit RND for BTC and ETH, respectively. Figure C.3 complements these plots with the full three-dimensional time evolution of the RND shape.

(a) BTC – Summary Statistics

Parameter	Q25	Q75	Min	Max	Mean	Std
σ	0.3169	0.4310	0.2432	0.6360	0.3734	0.0760
λ	6.1025	7.7561	4.1373	10.0000	6.9116	1.2635
p	0.3151	0.5033	0.1033	0.6930	0.4113	0.1218
η_1	10.5576	13.8865	6.9566	14.5731	12.3917	1.7717
η_2	9.4161	11.2364	6.7436	12.0133	10.3057	1.2202
BTC	90,606.34	110,475.60	60,274.69	124,722.25	99,462.21	14,870.34

(b) ETH – Summary Statistics

Parameter	Q25	Q75	Min	Max	Mean	Std
σ	0.5330	0.5907	0.4726	0.6837	0.5641	0.0422
λ	7.5356	8.8751	6.4306	9.3923	8.1452	0.7470
p	0.2584	0.3799	0.1352	0.4993	0.3183	0.0795
η_1	7.7514	8.5975	6.5899	9.3597	8.2155	0.6270
η_2	7.0795	8.7580	6.2323	10.0701	7.9719	0.9969
ETH	2,582.00	4,224.10	1,793.85	4,829.95	3,450.27	835.44

Table 3.1: Kou Parameters – Benchmark Calibration. This table reports summary statistics for the BTC and ETH parameters (σ, λ, p, η_1, η_2) estimated from the calibration using Deribit surfaces over the last seven days, together with the spot price. The Deribit sample is restricted to asset-specific periods for which the corresponding Polymarket RND table is available. Both samples terminate on 2025-11-27; the BTC sample starts on 2024-10-04, whereas the ETH sample starts on 2025-05-01. For each asset, the table reports the 25th percentile, the 75th percentile, the minimum, the maximum, and the standard deviation.

(a) BTC – RND Summary Statistics

Mean	Volatility	Skewness	Kurtosis
-0.1244	0.5032	-0.5720	9.6401
(0.0555)	(0.1009)	(0.4350)	(1.4054)

(b) ETH – RND Summary Statistics

Mean	Volatility	Skewness	Kurtosis
-0.2816	0.7548	-0.5545	8.7631
(0.0507)	(0.0712)	(0.2812)	(1.0854)

Table 3.2: Kou RND Moments – Benchmark Calibration. This table reports the average moments of the RND for a 7-day maturity, estimated from Deribit surfaces and restricted to the same asset-specific dates used in the corresponding Polymarket RND panels. Both samples terminate on 2025-11-27; the BTC sample starts on 2024-10-04, whereas the ETH sample starts on 2025-05-01. For each panel, we report the annualized mean and volatility, together with skewness and kurtosis; standard deviations are shown in parentheses. Panels A and B correspond to BTC and ETH, respectively.

BTC or ETH—using keyword matching and manual verification of market descriptions. For each relevant market, we used its CLOB-specific token IDs to query the "clob-token-timeseries" endpoint, which returns historical prices at hourly frequency.

For each market, we also source trade data at the tick level, recorded at the transaction level,

where multiple rows may share the same transaction hash when a taker order matches against several makers. Each row contains the market identifier (condition ID), the outcome tokens, the side (buy or sell), the size, the price, and basic trader metadata (wallet address, pseudonym).

For BTC and ETH, the dataset contains 4,386,994 trades across 4,860 unique contracts. Trading activity differs markedly across contract types. Above contracts are the most actively traded, with 1,682,365 trades, followed by Reach contracts (1,272,270), Range contracts (814,199), and Dip contracts (618,160). The table also reports total trading volume in notional USD, measured as the total number of contracts traded. Since each contract has a unit payoff of \$1, this measure corresponds directly to the notional amount exchanged. Above contracts account for the largest share of volume (USD 1,155.13 million), followed by Reach contracts (USD 791.14 million), Dip contracts (USD 400.79 million), and Range contracts (USD 348.68 million). Overall, total BTC and ETH trading volume equals USD 2,695.74 million.[6]

Bet Type	Trades	Volume	Unique Bets
Above	1,682,365	1,155.13	2,274
Range	814,199	348.68	1,808
Reach	1,272,270	791.14	407
Dip	618,160	400.79	371
All	4,386,994	2,695.74	4,860

Table 3.3: Bets – Summary Statistics. The table reports, for each contract type, the total number of trades, total trading volume (in millions of USD), and the number of unique contracts.

Following Bollen and Whaley (2004), we construct *net buying pressure (NBP)*, which is different from volume when buy and sell orders offset. NBP is calculated as the difference between the number of contracts bought and sold, divided by the total number of contracts. This calculation is performed for each contract and on an hourly frequency:

$$\text{NBP}_t = \frac{(\text{Buy Volume}_t - \text{Sell Volume}_t)}{\text{Total Volume}_t}.$$

We calculate the moneyness and maturity of each contract over time: For Above contracts, the moneyness is the ratio of the spot price to the strike price (S/K), treating them like call options. For Range contracts, the lower strike is a call with moneyness S/K_1, and the upper strike is a put with moneyness K_2/S; if the spot price lies between the two strikes (in the money), the average of the two is taken, while otherwise the side that is OTM is used. For Reach contracts, they are treated like calls, so moneyness is again S/K. Finally, for Dip contracts, which behave

[6] Figure B.2 plots trading volume over time by contract type.

like puts, moneyness is computed as K/S. Hence, for all contracts, if moneyness moves toward 1, it indicates that the market price is approaching the strike and the contract is becoming more ATM. Maturity is measured as the number of hours remaining until resolution.

3.4 Blockchain Related Risk Measures – L2 Risk

To proxy L2 operational risk, we use Polygon block-level data to construct a measure of inter-block intervals. We compute timestamp differences between consecutive blocks using metadata indexed through The Graph. The intervals are aggregated hourly, retaining the longest block gap in each hour. To capture persistent disruptions rather than isolated spikes, we take a 24-hour rolling average of these maximum gaps. We refer to this measure as *L2 Risk*.

3.5 Macroeconomic Risks

We construct an *Announcement* indicator capturing policy events, specifically tariff announcements made by President Donald Trump between November 2024 and April 2025, as well as Federal Open Market Committee (FOMC) meetings. We account for FOMC meeting days and tariff-related events documented by the American Institute for Economic Research.[7]

3.6 Sentiment

The Fear & Greed Index measures investor sentiment in the cryptocurrency market by aggregating several market indicators into a single scale ranging from 0 to 100. Lower values indicate extreme fear, reflecting heightened investor concern and potential selling pressure, while higher values indicate extreme greed, often associated with strong optimism and speculative behavior. We normalize the Fear & Greed Index by subtracting 50 and dividing by 100, centering the measure around zero so that negative values indicate fear and positive values indicate greed.

3.7 Other Risks

Polymarket uses a Binance data feed to provide verifiable off-chain information for smart contract resolution. Deribit relies on a proprietary feed drawing from multiple centralized exchanges. To capture cross-exchange risk for the same underlying, we focus on BTC price divergence between Binance and Deribit. We define hourly relative deviation as the Binance-Deribit price

[7]We source the FOMC meetings from https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm. The data for the announcements can be found here: https://aier.org/article/how-equity-markets-reacted-to-trumps-tariff-announcements-the-data

difference, normalized by the Binance price, and measure short-term volatility using its 24-hour rolling standard deviation. This *Diff Underlying* measure captures frictions in price discovery, informational latency, and short-term market segmentation.

We calculate the returns and volatilities of BTC, ETH, and USDC. To ensure neutrality, we rely on price data from a source independent of Polymarket, Binance, and Deribit. Specifically, we obtain hourly BTC-USD and USDC-USD prices from CryptoCompare, from which we compute the first and second moments: The variable Δ BTC (%) denotes the percentage change in BTC (its return), while BTC Vol_{30} captures the 30-day rolling volatility of BTC. Similarly, Δ USDC (%) measures the percentage change in USDC, and USDC Vol_{30} represents the 30-day rolling volatility of USDC. We repeat the analysis for ETH.

Internet Appendix II.3 analyzes Polygon gas costs using transaction receipts. We compute realized fees from gas use and effective gas prices, convert them into USD at the contemporaneous MATIC–USD rate, and find that they average only \$0.0163 per trade. Because Polymarket uses a gasless model, these costs are borne by the platform rather than traders.

Internet Appendix II.4 describes the construction of bid–ask spread measures for the Polymarket CLOB. BTC markets generally have tight median spreads for both YES and NO outcomes, with effective spreads concentrated around 1 cent and relative spreads near zero. Both distributions are right-skewed, especially for YES, indicating occasional episodes of much wider spreads and weaker liquidity.

4 Theoretical Framework

To guide the empirical analysis, we present a parsimonious model of equilibrium behavior in decentralized prediction markets. The model is not intended as a general theory of limit-order-book arbitrage, a subject already extensively treated in the literature (see, e.g., Grossman and Miller, 1988). Instead, it provides a simple framework for understanding how demand imbalances can translate into persistent deviations from a frictionless benchmark when arbitrageurs face inventory risk, execution frictions, and a finite trading horizon.

Our framework builds on the limits-to-arbitrage tradition (Bollen and Whaley, 2004; Brunnermeier and Pedersen, 2009; Shleifer and Vishny, 1997). A representative arbitrageur supplies liquidity against non-arbitrage demand but cannot instantaneously eliminate all pricing discrepancies. The reason is twofold. First, holding inventory exposes the arbitrageur to fluctuations

in the benchmark value of the claim. Second, adjusting inventory is costly because trading is subject to time-varying execution frictions. We first derive the equilibrium pricing wedge as a function of aggregate demand shocks and arbitrage frictions. We then use the resulting relation to organize the empirical hypotheses.

4.1 The Environment

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{Q})$ be a filtered probability space. Let S_t denote the underlying asset price, such as BTC or ETH. A prediction-market contract pays one unit if a specified event occurs and zero otherwise.

More generally, let $X_T \in \{0,1\}$ be an $\mathcal{F}_T$-measurable event payoff associated with a contract that settles at the finite deadline T. For an Above contract, $X_T = \mathbf{1}_{\{S_T>K\}}$, whereas Range contracts depend on whether S_T lies in a specified interval, and Reach or Dip contracts may depend on whether the underlying crosses a barrier before T. The analysis below therefore applies to both terminal and path-dependent binary event claims.

We take as frictionless benchmark the risk-neutral value of the event claim,

$$q_t = \mathbb{E}^{\mathbb{Q}}[X_T \mid \mathcal{F}_t].$$

Thus $q_t \in [0,1]$ is the benchmark probability of the event under the information and pricing measure reflected in the reference derivatives market.

For a terminal threshold event, q_t coincides with the value of the corresponding digital claim, while for path-dependent events it is the option-implied value of the associated barrier-style payoff.[8]

By the law of iterated expectations, q_t is a bounded $\mathbb{Q}$-martingale. At settlement, all uncertainty is resolved: $q_T = X_T \in \{0,1\}$.

We deliberately keep the benchmark process abstract. In particular, q_t may be a general square-integrable càdlàg martingale with both continuous fluctuations and jumps. We assume that its predictable quadratic variation is absolutely continuous:

$$d\langle q\rangle_t = \nu_t\, dt,$$

[8]For a terminal Above claim and a sufficiently regular call-price surface $C_t(K)$, the Breeden and Litzenberger (1978) relation gives $q_t = -e^{r(T-t)}\partial_K C_t(K)$. For path-dependent claims, the benchmark is instead obtained from the Kou (2002) model used to value the corresponding barrier payoff. The theoretical framework only requires the existence of the benchmark process q_t, not a particular parametric representation.

where $\nu_t \geqslant 0$ denotes the instantaneous variance rate of the benchmark claim. For tractability, the arbitrageur treats ν_t as an exogenous, bounded, deterministic function of calendar time. This preserves the finite-dimensional state space while allowing inventory risk to vary over the life of the contract.[9]

The instantaneous variance rate need not vanish pointwise as maturity approaches. What vanishes is the remaining conditional uncertainty about the terminal event. Since $q_T = X_T$ and X_T is binary,

$$\mathrm{Var}_t^{\mathbb{Q}}(X_T) = q_t(1-q_t) = \mathbb{E}_t^{\mathbb{Q}}\left[\langle q\rangle_T - \langle q\rangle_t\right], \tag{2}$$

which converges to zero at settlement.

The prediction market operates as a central limit order book. Let p_t denote the prediction-market mid-price, and define the cross-market pricing wedge as

$$\Delta_t \equiv p_t - q_t.$$

The market is populated by a representative arbitrageur and by non-arbitrage traders. Non-arbitrage traders generate aggregate net demand D_t, driven by a latent demand shock Z_t and partially elastic with respect to the wedge:

$$D_t = Z_t - \eta\Delta_t, \qquad \eta > 0. \tag{3}$$

Here, η captures the price elasticity of non-arbitrage demand. A positive Z_t represents excess demand for the prediction-market claim relative to the benchmark, while a larger positive wedge dampens that demand.

4.2 The Arbitrageur's Problem

The arbitrageur is competitive and takes the cross-market pricing wedge Δ_t as given when choosing her trading rate. The wedge is therefore not a control variable in the arbitrageur's problem. Instead, it is pinned down in equilibrium by combining the arbitrageur's optimal liquidity supply with market clearing.

Let y_t denote the representative arbitrageur's inventory in the prediction-market claim, net

[9] For example, if $dq_t = \sigma_t^q\, dW_t + \int_{\mathbb{R}} u\, \widetilde{N}(dt, du)$, where $\widetilde{N}$ is a compensated jump measure with compensator $\lambda_t(du)\, dt$, then $\nu_t = (\sigma_t^q)^2 + \int_{\mathbb{R}} u^2 \lambda_t(du)$. Hence ν_t captures both continuous mark-to-market risk and jump-driven repricing.

of the corresponding benchmark hedge. The arbitrageur chooses a trading rate v_t, so that $dy_t = v_t\,dt$. She chooses this rate to maximize expected discounted trading profits net of inventory and execution costs over the finite campaign duration $[t, T]$.

The frictions take two forms. First, the arbitrageur incurs the quadratic inventory-risk cost $\frac{\gamma_t}{2} y_t^2,\ \gamma_t \equiv \gamma\nu_t,\ \gamma > 0$. The parameter γ measures the arbitrageur's aversion to inventory risk, while the exogenous coefficient ν_t measures the instantaneous mark-to-market variance of the benchmark claim. A larger ν_t therefore raises the cost of warehousing a given inventory position. Equation (2), rather than pointwise convergence of ν_t, captures the resolution of uncertainty at the terminal date.

Second, the arbitrageur faces the time-varying quadratic execution cost $\frac{\Lambda_t}{2} v_t^2,\ \Lambda_t > 0$, which captures instantaneous liquidity and execution costs in the decentralized trading infrastructure. As with ν_t, we treat Λ_t as an exogenous, bounded, deterministic function of time.

Given the pricing wedge Δ_s, the arbitrageur chooses a trading policy $\{v_s\}_{s\in[t,T]}$ to solve

$$V(y_t, Z_t, t) = \max_{\{v_s\}_{s\in[t,T]}} \mathbb{E}_t \left[\int_t^T e^{-r(s-t)} \left(-\Delta_s v_s - \frac{\Lambda_s}{2} v_s^2 - \frac{\gamma\nu_s}{2} y_s^2 \right) ds \right]$$

subject to $dy_s = v_s\,ds$.

The terminal condition is $V(y_T, Z_T, T) = 0$. This condition does not require forced liquidation. Because the prediction-market claim and its benchmark hedge deliver the same realized event payoff at settlement, residual hedged inventory has zero terminal mark-to-market value. Moreover,

$$p_T = q_T = X_T, \qquad \Delta_T = 0. \tag{4}$$

The market-clearing relation derived below applies on $[0, T)$; at T, the contract settles mechanically at X_T.

Because the problem has a finite horizon and time-varying coefficients, the value function depends explicitly on calendar time. Its Hamilton–Jacobi–Bellman equation is

$$V_t - rV + \max_v \left\{ -\Delta_t v - \frac{\Lambda_t}{2} v^2 - \frac{\gamma\nu_t}{2} y_t^2 + V_y v \right\} + \mathcal{L}_Z V = 0,$$

where $\mathcal{L}_Z$ is the infinitesimal generator of the latent demand shock. The first-order condition with respect to v gives

$$v_t^* = \frac{V_y - \Delta_t}{\Lambda_t}. \tag{5}$$

When $V_y > \Delta_t$, the arbitrageur increases her long position; when $V_y < \Delta_t$, she reduces it. The execution-cost coefficient Λ_t governs the speed at which inventory can adjust.

4.3 Equilibrium

Equilibrium closes the model by requiring the arbitrageur's inventory to absorb the net demand generated by non-arbitrage traders:

$$y_t + D_t = 0, \qquad t < T. \tag{6}$$

Combining (6) with (3) gives

$$\Delta_t = \frac{Z_t + y_t}{\eta}. \tag{7}$$

Substituting (7) into (5) yields $v_t^* = \frac{V_y - (Z_t + y_t)/\eta}{\Lambda_t}$. Because the objective is linear-quadratic, the state dynamics are linear, and the horizon is finite, we conjecture a value function that is quadratic in the state variables with time-varying coefficients. Substituting this conjecture and the market-clearing condition into the Hamilton–Jacobi–Bellman equation yields a system of Riccati ordinary differential equations evaluated backward from the terminal conditions at T. Solving this non-stationary system, as detailed in Appendix I, produces an equilibrium wedge that is linear in the aggregate demand shock:

$$\Delta_t = \Phi_t\big((\nu_s, \Lambda_s)_{s\in[t,T]}; \gamma, \eta\big)\, Z_t, \qquad t < T,$$

where $\Phi_t > 0$ is the time-dependent inverse elasticity of arbitrage capital over the remaining life of the contract.

Equation (4) guarantees that the wedge is zero at settlement. Before settlement, however, its time profile depends jointly on the evolution of demand, benchmark variance, and execution costs. The finite terminal date alone does not imply that the wedge must decline monotonically before T.

Under the standard monotonicity conditions of the linear-quadratic problem, a local increase in either inventory risk or execution costs makes arbitrage supply less elastic: $\frac{\partial \Phi_t}{\partial \nu_t} > 0, \qquad \frac{\partial \Phi_t}{\partial \Lambda_t} > 0$. Here each derivative is shorthand for a local increase in the corresponding coefficient, holding the remainder of its path fixed. Higher benchmark variance raises the cost of warehousing inventory, while higher execution costs slow inventory adjustment. A given demand imbalance

therefore produces a larger pricing wedge.

4.4 Microfoundations of Demand Shocks

To bridge the theoretical benchmark with the empirical analysis, we open the "black box" of the aggregate demand shock Z_t. Rather than treating this imbalance as purely exogenous noise, we allow its functional form to vary with the economic mechanism generating demand pressure. In this sense, Z_t is a latent demand imbalance that can be mapped to several behavioral and rational channels.

Behavioral Probability Weighting. One channel is that retail traders distort probabilities, generating directionally biased demand. Consistent with prospect theory and the weighting function developed by Prelec (1998), participants may overweight small probabilities and underweight sufficiently large probabilities. Let $\omega(q) = \exp\left[-\beta(-\ln q)^{\alpha}\right], \qquad \alpha \in (0,1), \quad \beta > 0.$ We capture the associated demand component as $Z_t^{\text{bias}} = \omega(q_t) - q_t$. For parameter values that generate the usual inverse-S weighting pattern, $Z_t^{\text{bias}} > 0$ for sufficiently low-probability events. This mechanism can therefore generate excess demand for extreme Reach and Dip contracts. For high-probability events, the sign is generally reversed.

Bayesian Learning. A second interpretation focuses on the gradual resolution of disagreement through the accumulation of information. At contract inception, participants may hold heterogeneous priors, creating a speculative imbalance proportional to initial disagreement $\mathcal{D}_0$. As public signals arrive, beliefs converge through Bayesian updating. We represent this component as $Z_t^{\text{learn}} = \mathcal{D}_0 \left(\frac{\tau_0}{\tau_0 + t\tau_s}\right)$, where τ_0 is the precision of initial priors and τ_s is the precision of the incoming signal flow. This demand component declines with elapsed time: $\frac{\partial Z_t^{\text{learn}}}{\partial t} = -\frac{\mathcal{D}_0 \tau_0 \tau_s}{(\tau_0 + t\tau_s)^2} < 0$ when $\mathcal{D}_0 > 0$.

Exogenous Sentiment Shocks. Alternatively, demand imbalances may arise from broader sentiment cycles or salient but non-informative news. To avoid confusing the sentiment state with the underlying asset price S_t, let M_t denote latent market sentiment: $Z_t^{\text{sent}} = \chi M_t,\ dM_t = -\kappa_M M_t\, dt + \sigma_M\, dW_t^M$. Periods of euphoria or panic can therefore push the prediction-market price away from the benchmark, while κ_M determines the persistence of the imbalance.

Informed Trading and Adverse Selection. Finally, some traders may possess private information not yet reflected in the derivatives-market benchmark. Following Kyle (1985), informed order flow introduces the component $Z_t^{\text{info}} = \psi\left(\mathbb{E}[X_T \mid \mathcal{F}_t^{\text{private}}] - q_t\right)$. In this setting, the arbitrageur interprets observed demand as a noisy signal of future benchmark revisions. To avoid being picked off, she supplies liquidity less aggressively. The observed wedge may then reflect both an arbitrage friction and the price-discovery process through which private information enters prediction-market prices.

4.5 Testable Implications

The equilibrium relationship $\Delta_t = \Phi_t\big((\nu_s, \Lambda_s)_{s\in[t,T]}; \gamma, \eta\big)\, Z_t$ provides a structural foundation for the empirical analysis. Holding the path of arbitrage frictions fixed, the sensitivity of the wedge to a demand driver x_k is $\frac{\partial \Delta_t}{\partial x_k}\Big|_{\nu,\Lambda} = \Phi_t \frac{\partial Z_t^k}{\partial x_k}$, where Z_t^k is the demand component associated with channel k. Since $\Phi_t > 0$, the direction of the demand-induced wedge is determined by the sign of the corresponding demand response, while its magnitude is amplified by limits to arbitrage.

Pricing of Tail Events. Under behavioral probability weighting, $Z_t^{\text{bias}} = \omega(q_t) - q_t$. For parameter values that imply overweighting of small probabilities, Z_t^{bias} is positive for q_t near zero, so the model predicts a positive wedge for sufficiently unlikely events. This mechanism is especially relevant for deep out-of-the-money Reach and Dip contracts.

Learning and the Resolution of Disagreement. The learning mechanism gives the following ceteris-paribus maturity effect:

$$\frac{\partial \Delta_t}{\partial t}\bigg|_{\Phi_t} = \Phi_t \frac{\partial Z_t^{\text{learn}}}{\partial t} = -\Phi_t \frac{\mathcal{D}_0 \tau_0 \tau_s}{(\tau_0 + t\tau_s)^2} < 0. \tag{8}$$

Equation (8) isolates the learning channel while holding arbitrage conditions fixed. The total time derivative is instead $\frac{d\Delta_t}{dt} = \dot{\Phi}_t Z_t + \Phi_t \dot{Z}_t$, so the observed maturity profile also depends on the evolution of benchmark variance and execution frictions through Φ_t.

Sentiment, Adverse Selection, and Aggregate Demand. For the sentiment and informed-trading specifications, $\frac{\partial \Delta_t}{\partial Z_t} = \Phi_t > 0$. An increase in directional demand raises the wedge that the arbitrageur must absorb, whether the imbalance originates in sentiment, salient news, or private information.

Inventory Risk and Execution Frictions. Holding $Z_t > 0$ fixed, the local comparative statics are

$$\frac{\partial \Delta_t}{\partial \nu_t} = Z_t \frac{\partial \Phi_t}{\partial \nu_t} > 0, \qquad \frac{\partial \Delta_t}{\partial \Lambda_t} = Z_t \frac{\partial \Phi_t}{\partial \Lambda_t} > 0.$$

Higher benchmark variance raises inventory risk, while higher execution costs make rebalancing more expensive. In either case, arbitrageurs supply liquidity less elastically, and a given demand shock generates a larger wedge.

Empirical Specification and Hypotheses Summary

To test these theoretical sensitivities, we use the pooled ordinary-least-squares panel specification estimated in Section 5. For contract i, written on asset a, at time t, the pricing wedge $\Delta_{i,a,t}$, defined as the central-limit-order-book price minus the option-implied probability, is modeled as

$$\begin{aligned}\Delta_{i,a,t} = \beta_0 &+ \beta_{\text{Moneyness}} \,\text{Moneyness}_{i,a,t} + \beta_{\text{Maturity}} \,\text{Maturity}_{i,a,t} + \beta_{\text{NBP}} \,\text{NBP}_{i,a,t} \\ &+ \beta_{\text{Sentiment}} \,\text{Sentiment}_t + \beta_{\text{Volatility}} \,\text{Volatility}_{a,t} + \beta_{\text{Frictions}} \,\text{Frictions}_t \\ &+ \beta_{\text{Controls}} \,\text{Controls}_{i,a,t} + \varepsilon_{i,a,t}.\end{aligned}$$

Mapping the theoretical channels to these empirical proxies yields four main hypotheses:

- **H1 (Behavioral Probability Weighting).** The pricing wedge varies systematically with moneyness, with larger positive wedges for sufficiently low-probability events. The expected sign of $\beta_{\text{Moneyness}}$ depends on the construction of the moneyness variable and the payoff type.

- **H2 (Bayesian Learning).** Holding arbitrage conditions fixed, the wedge is wider earlier in the life of the contract and declines as disagreement is resolved. Since elapsed time is inversely related to time to expiration, a maturity variable measuring time remaining implies $\beta_{\text{Maturity}} > 0$. This is a reduced-form prediction: the total maturity effect may also incorporate changes in ν_t and Λ_t.

- **H3 (Information and Sentiment).** The wedge responds to aggregate directional demand and salient news. The baseline prediction is $\beta_{\text{NBP}} > 0$, while the sign of $\beta_{\text{Sentiment}}$ depends on whether the contract payoff is bullish or bearish and on the direction of the sentiment shock.

- **H4 (Limits to Arbitrage).** Increases in benchmark variance or execution frictions steepen the arbitrageur's supply curve and widen the pricing wedge, conditional on a positive demand imbalance. The baseline prediction is $\beta_{\text{Volatility}} > 0$ when the empirical volatility measure is a valid proxy for ν_t. Direct friction measures such as L2 Risk and Diff Underlying are also expected to enter positively, whereas the Weekend coefficient may vary with payoff structure and with the particular market-segmentation channel captured by the proxy.

In Section 5, we test these hypotheses across the four betting structures—Above, Range, Reach, and Dip—allowing the empirical relationship between demand, arbitrage frictions, and the cross-market wedge to vary with the payoff profile.

5 Cross-Market Price Alignment

This section establishes the degree of alignment between Polymarket prices and OIPs for BTC and ETH contracts. We summarize the pricing sample and wedge distribution, compare representative price paths, and examine contract-level correlations across Above, Range, Reach, and Dip contracts. Overall, the two price series move closely together, although alignment is more heterogeneous for path-dependent contracts. Section 6 then examines the magnitude and determinants of systematic pricing wedges.

5.1 Alignment of Probabilities from Prediction and Option Markets

Table 5.1 summarizes the distribution and timing across the four betting types: Above, Range, Reach, and Dip. The pricing sample comprises 2,550 contracts on BTC and 2,310 on ETH, totaling 4,860 contracts. Exact active dates run from January 23, 2024, to November 30, 2025, with BTC entering on March 17, 2024; we summarize this contract-level sample as January 2024 to November 2025. The above contracts are the most prevalent, followed by Range, Reach, and Dip. Maturities vary notably across types: Above and Range contracts have short average maturities (measured from contract inception until resolution) of roughly a week. In contrast, Reach and Dip contracts are longer-dated, with average maturities of approximately eleven days to two weeks, reflecting positions targeting more gradual or less likely price movements.

Table 5.2 reports the distribution of the CLOB–OIP wedge. The average wedge is positive for both BTC and ETH across all contract types. In the pooled BTC–ETH sample, the equal-

(a) BTC – Summary Statistics

Bet Type	Number of Bets	Avg. Maturity	Start Date	End Date
Above	1209	6.50	2024-03-17	2025-11-30
Range	932	6.59	2025-02-07	2025-11-30
Reach	221	14.95	2024-04-02	2025-11-24
Dip	188	14.29	2024-04-25	2025-11-21
All	2550	7.92	2024-03-17	2025-11-30

(b) ETH – Summary Statistics

Bet Type	Number of Bets	Avg. Maturity	Start Date	End Date
Above	1065	6.65	2024-02-21	2025-11-30
Range	876	6.58	2025-04-11	2025-11-30
Reach	186	10.97	2024-08-06	2025-11-24
Dip	183	11.89	2024-01-23	2025-11-24
All	2310	7.39	2024-01-23	2025-11-30

Table 5.1: Bets – Summary Statistics. This table reports the number, average time to maturity in days, and active date ranges of contracts across the four prediction market contract types: Above, Range, Reach, and Dip, plus an aggregate All row. The start and end dates reflect the earliest and latest contract appearances in the data. Panel (a) shows BTC contracts, while Panel (b) shows ETH contracts. The different types of contracts are discussed in Section 2.2.

contract mean wedge is 0.716 percentage points, the median is 0.630 percentage points, and 72.6% of contract-level mean wedges are positive. The largest signed wedges appear in Reach and Dip contracts, consistent with prediction-market prices that embed a premium for salient barrier events relative to the option-implied benchmark.

Asset	Type	Bets	Obs.	Mean	Median	P25	P75	Mean Abs.	Pos. (%)	p
BTC	Above	1,209	181,638	0.396	0.416	-0.558	1.535	2.739	62.696	0.0013
BTC	Range	932	158,618	0.588	0.565	-0.033	1.048	1.752	74.034	< 0.0001
BTC	Reach	221	82,919	1.415	1.267	0.384	2.561	2.750	85.520	< 0.0001
BTC	Dip	188	67,023	1.340	1.014	0.363	2.354	2.730	81.383	< 0.0001
BTC	All	2,550	490,198	0.624	0.587	-0.202	1.408	2.378	70.196	< 0.0001
ETH	Above	1,065	181,642	0.580	0.579	-0.280	1.429	2.013	68.732	< 0.0001
ETH	Range	876	148,955	0.955	0.745	0.247	1.456	1.924	82.877	< 0.0001
ETH	Reach	186	51,817	1.052	1.137	0.280	2.239	2.693	79.570	< 0.0001
ETH	Dip	183	54,838	1.304	0.721	-0.114	2.338	3.034	71.038	< 0.0001
ETH	All	2,310	437,252	0.817	0.700	0.006	1.540	2.115	75.152	< 0.0001
BTC+ETH	All	4,860	927,450	0.716	0.630	-0.110	1.463	2.253	72.551	< 0.0001

Table 5.2: CLOB–OIP Wedge – Summary Statistics. This table reports summary statistics for the CLOB–OIP wedge in percentage points. The wedge is computed hourly as $100 \times (\text{CLOB} - \text{OIP})$ and first averaged within each contract; the reported moments (Mean, Median, P25, P75), positive-share (Pos. (%)), and one-sided p-values are then computed from contract-level means, so each contract receives equal weight. Obs. denotes the number of hourly observations.

Figure 5.1 shows illustrative examples of price dynamics over time for the four BTC-related contract types (Section 2.2) across the two pricing sources, CLOB and OIP. Panel (a) shows

the "BTC above 97,000" contract, whose price declines through early July before rising sharply and converging to one as the event becomes increasingly likely and ultimately occurs. Panel (b) tracks prices for the "BTC between 97,000 and 99,000" Range contract. Panel (c) shows the "BTC reach 130,000" contract, whose probability spikes in mid-January before declining steadily toward zero. Panel (d) shows the "BTC dip 90,000" contract, whose probability falls sharply in early May and then gradually converges toward zero. Across all panels, the CLOB and OIP pricing series closely track each other, though minor discrepancies are evident.

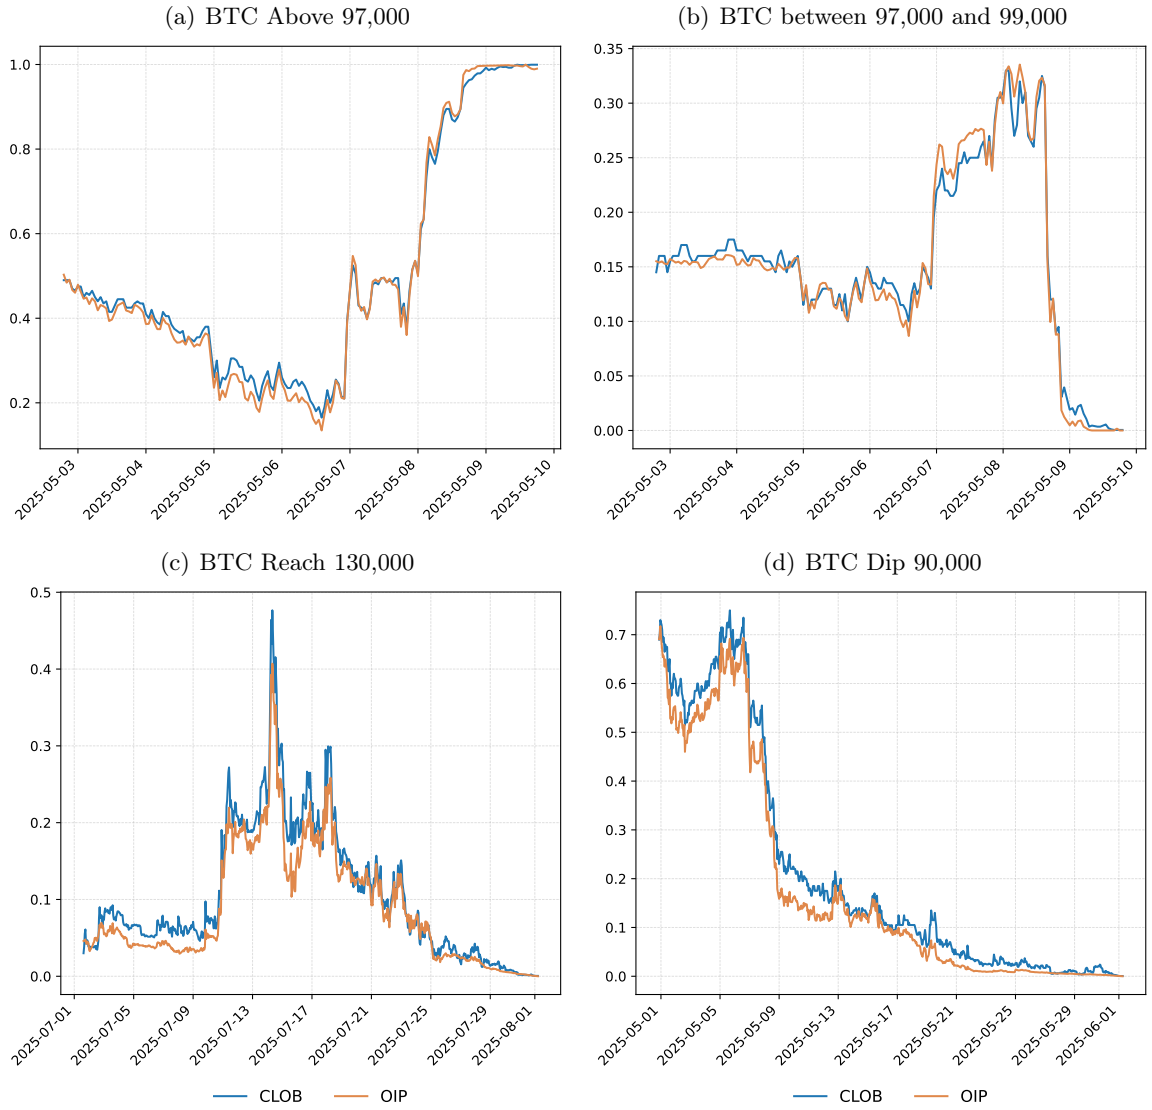


Figure 5.1: Example – Different Betting Types – CLOB vs. OIP. The figure displays time series of CLOB and OIP prices for one representative event from each of the four betting types (Above, Range, Reach, and Dip) shown in Panels (a) through (d), respectively. The different types of contracts are discussed in Section 2.2.

Figure 5.2 shows the distribution of contract-level time-series correlations between CLOB

and OIP prices. Above contracts display the strongest and most concentrated alignment, with a mean correlation of 0.88. Range contracts are more dispersed, with a mean of 0.77. Reach and Dip contracts have similarly high means of 0.88 and 0.86, but substantially wider distributions, including some low or negative correlations. Overall, CLOB and OIP prices are closely aligned on average, although agreement is more heterogeneous for path-dependent contracts.

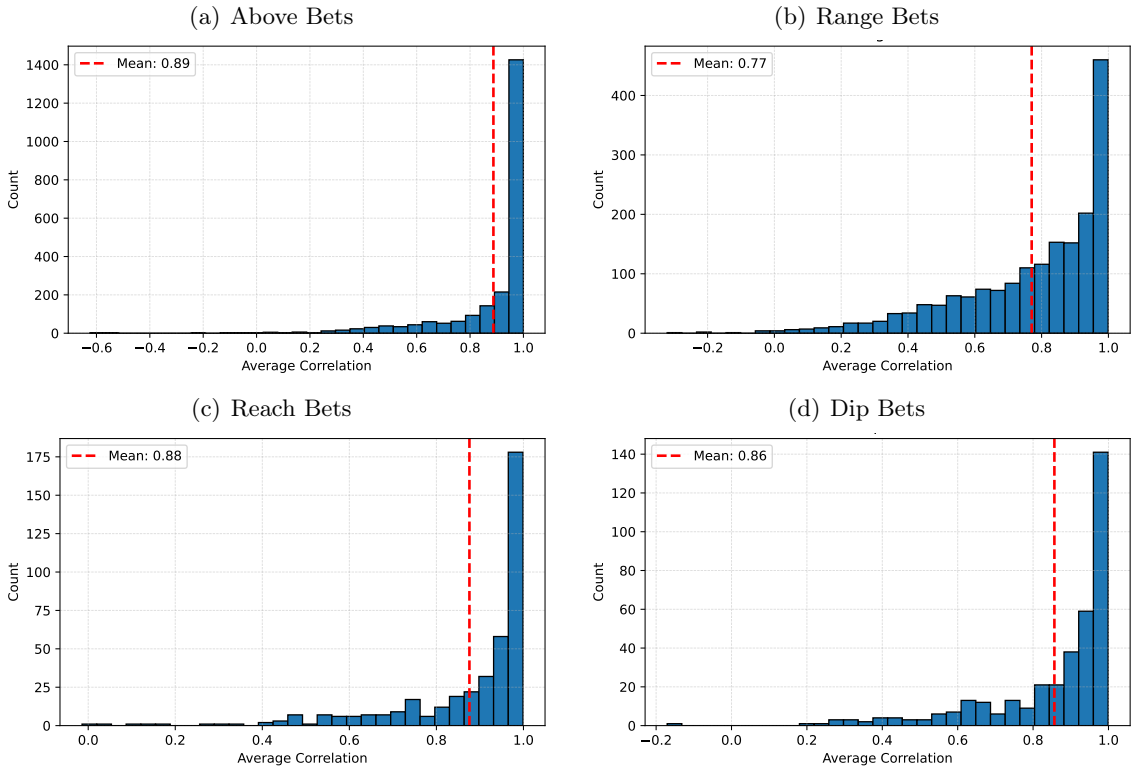


Figure 5.2: Histogram – Average Correlations per Betting Type. The figure displays the histogram of the average correlation of the CLOB and OIP for each of the four betting types (Above, Range, Reach, and Dip) shown in Panels (a) through (d), respectively. The different types of contracts are discussed in Section 2.2.

The evidence so far establishes strong average alignment between CLOB and OIP prices. High correlations, however, do not rule out systematic differences in price levels. We therefore turn to the magnitude and determinants of the CLOB–OIP pricing wedge and assess how these deviations vary across payoff structures.

6 Pricing Deviations and Market Efficiency

This section examines the magnitude, direction, and determinants of the CLOB–OIP pricing wedge. We first test the hypotheses developed in Section 4 using panel regressions across the four contract types. We then provide additional evidence on tail-event pricing, salient moneyness

states, and the evolution of wedges over the contract life. Together, these analyses assess when and why Polymarket prices depart systematically from option-implied benchmarks.

6.1 Regression Analysis

Section 4 delivers four testable predictions for the wedge between Polymarket prices and option-implied benchmarks. First, behavioral probability weighting implies stronger mispricing in tail states, so moneyness should matter, with the sign depending on contract design. Second, Bayesian learning predicts that the wedge is larger early in the contract's life and declines as uncertainty resolves, implying a positive role for maturity. Third, information and sentiment shocks should affect the wedge through directional demand, sentiment, and announcement effects. Fourth, limits to arbitrage suggest that higher volatility and trading frictions should widen the wedge. This section tests these predictions across Above, Range, Reach, and Dip contracts, asking whether the strength and sign of these effects vary by payoff structure.

We estimate a pooled OLS regression. Asset-specific differences are controlled for by including an ETH dummy variable that equals one for ETH contracts and zero otherwise. The goal is to explain the CLOB–OIP price wedge using the following explanatory variables: returns and volatilities of BTC, ETH, and USDC; Deribit's BTC and ETH volatility indices, which we denote VIX; blockchain operational risk (L2 Risk); Binance–Deribit underlying-price dislocations; NBP; sentiment (Fear & Greed); and the Announcement indicator.[10]

Table 6.1 reports correlations among the explanatory variables. The 30-day volatility measure is positively correlated with the VIX ($\rho = 0.487$) and USDC volatility ($\rho = 0.411$). The VIX is also related to L2 Risk ($\rho = 0.251$), while L2 Risk correlates with Diff Underlying ($\rho = 0.465$) and Fear & Greed ($\rho = 0.506$). USDC volatility is positively associated with the VIX ($\rho = 0.380$) and Diff Underlying ($\rho = 0.323$). Fear & Greed is positively correlated with the VIX ($\rho = 0.191$) and Diff Underlying ($\rho = 0.276$), but negatively with Vol_{30} ($\rho = -0.181$). NBP and Announcement are weakly correlated with most regressors, although NBP correlates with Fear & Greed ($\rho = -0.115$) and Vol_{30} ($\rho = -0.113$). Returns are also weakly correlated with the risk variables, with absolute coefficients generally around 0.05 or less.

[10] Figure B.1 presents the main explanatory variables used in the panel regressions. Panels (a) and (b) plot BTC's and ETH's hourly percentage change and 30-day rolling volatility, showing pronounced return fluctuations alongside periods of elevated volatility. Panel (c) shows the corresponding measures for USDC. As expected for a stablecoin, the return series is tightly centered on zero, but short bursts of volatility are visible during peg stress episodes. Panel (d) shows the L2 Risk, constructed from Polygon block-level data. The measure varies substantially over time, with clear spikes that coincide with periods of congestion and higher on-chain activity.

	USDC Vol_{30}	L2 Risk	NBP	Fear & Greed	Announcement	Δ USDC (%)	Vol_{30}	VIX	Diff Underlying	Δ (%)
USDC Vol_{30}	1.000	0.274	-0.049	0.219	-0.046	0.005	0.411	0.380	0.323	0.006
L2 Risk	0.274	1.000	-0.079	0.506	-0.047	-0.000	0.045	0.251	0.465	0.007
NBP	-0.049	-0.079	1.000	-0.115	-0.030	0.007	-0.113	-0.047	-0.081	-0.054
Fear & Greed	0.219	0.506	-0.115	1.000	-0.010	0.000	-0.181	0.191	0.276	0.009
Announcement	-0.046	-0.047	-0.030	-0.010	1.000	0.000	0.043	-0.043	-0.036	0.003
Δ USDC (%)	0.005	-0.000	0.007	0.000	0.000	1.000	0.001	-0.001	0.000	0.051
Vol_{30}	0.411	0.045	-0.113	-0.181	0.043	0.001	1.000	0.487	0.177	-0.004
VIX	0.380	0.251	-0.047	0.191	-0.043	-0.001	0.487	1.000	0.121	0.009
Diff Underlying	0.323	0.465	-0.081	0.276	-0.036	0.000	0.177	0.121	1.000	0.010
Δ (%)	0.006	0.007	-0.054	0.009	0.003	0.051	-0.004	0.009	0.010	1.000

Table 6.1: Regressors – Correlation Matrix. This table reports the unconditional correlations among the explanatory variables: returns, USDC returns, volatility, USDC volatility, and additional market-structure variables: L2 Risk, Diff Underlying, NBP, and Fear & Greed. The construction of the variables is explained in Section 3. The data range from January 2024 to November 2025.

Table 6.2 reports pooled OLS results for Above contracts with monthly fixed effects. Moneyness is negative and highly significant throughout, implying larger wedges for more OTM contracts, while maturity is positive and highly significant, indicating that wedges narrow toward expiration. The VIX and L2 Risk enter positively and significantly, consistent with option-market uncertainty and blockchain-related frictions widening the wedge. By contrast, Vol_{30} and Diff Underlying enter negatively, while USDC volatility is positive. NBP is positive and significant in its individual specification, and Announcement is positive and highly significant whenever included. Fear & Greed is negative throughout but statistically significant only in the smaller-sample specification.

The information and event variables further reinforce the interpretation that Above-contract wedges respond to demand and salient market conditions. NBP is positive and significant in its individual specification, consistent with the theoretical prediction that directional buying pressure widens the CLOB–OIP wedge for upside contracts. Announcement is positive and highly significant throughout, indicating that salient macroeconomic and policy events increase cross-market discrepancies. The Fear & Greed Index enters negatively, suggesting that the wedge is larger outside periods of broad crypto-market optimism, when uncertainty or caution may be more relevant for pricing. Overall, the results indicate that Above-contract wedges are systematically explained by contract characteristics, option-market uncertainty, L2 execution risk, directional buying pressure, and announcement effects.

Table 6.3 reports pooled OLS results for Range contracts with monthly fixed effects. Moneyness is negative and highly significant in the main specifications, indicating larger wedges for more extreme range contracts, while maturity is positive and highly significant, consistent with wider discrepancies for longer-dated contracts. These signs reverse only in the smaller-

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		-0.058***							
		(0.006)							
USDC Vol_{30}		0.027***							
		(0.003)							
VIX			0.031***				0.036***	0.041***	0.081***
			(0.008)				(0.012)	(0.012)	(0.012)
L2 Risk				0.017***			0.027***		0.045***
				(0.003)			(0.004)		(0.006)
Diff Underlying					-0.036***			-0.043***	-0.039***
					(0.004)			(0.004)	(0.007)
NBP									0.008**
									(0.003)
Fear & Greed						-0.006	-0.009	-0.005	-0.030***
						(0.004)	(0.006)	(0.006)	(0.008)
Announcement						0.064***	0.069***	0.067***	0.089***
						(0.008)	(0.008)	(0.008)	(0.013)
Δ (%)	-0.001	-0.001	-0.002	-0.001	-0.001	-0.001	-0.002	-0.001	-0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Δ USDC (%)	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Weekend						0.008	0.010	0.012*	0.013
						(0.007)	(0.007)	(0.007)	(0.009)
Moneyness	-0.121***	-0.120***	-0.119***	-0.121***	-0.122***	-0.119***	-0.118***	-0.119***	-0.128***
	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.004)
Maturity	0.056***	0.057***	0.056***	0.056***	0.056***	0.056***	0.055***	0.056***	0.061***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)	(0.004)
ETH	0.037***	0.134***	-0.019	0.037***	0.059***	0.037***	-0.027	-0.011	-0.112***
	(0.004)	(0.010)	(0.016)	(0.004)	(0.005)	(0.004)	(0.022)	(0.022)	(0.023)
R^2 Adj	0.028	0.029	0.028	0.028	0.029	0.029	0.029	0.029	0.070
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	363,280	363,280	363,243	363,280	363,280	363,280	363,243	363,243	100,241

Table 6.2: Panel Regression – Above Bets – CLOB–OIP. This table reports the results from a pooled OLS panel regression of the CLOB–OIP price wedge for the Above contracts on Moneyness, Maturity, volatility, USDC volatility, and additional market structure variables: L2 Risk, Diff Underlying, NBP, Fear & Greed, and Announcement. We include returns, USDC returns, and weekend and ETH dummy variables as controls. The construction of the variables is explained in Section 3. The data range from January 2024 to November 2025. Standard errors are clustered by time and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All continuous variables are standardized to have a mean of 0 and a variance of 1; indicator variables remain in levels.

sample specification. These findings support the role of contract characteristics in explaining the CLOB–OIP wedge. The risk and friction variables also contain informative patterns: Vol_{30} and Diff Underlying enter positively, suggesting that realized volatility and cross-exchange price fragmentation widen discrepancies for Range contracts. At the same time, USDC Vol_{30}, the VIX, and L2 Risk enter negatively, indicating that central-interval contracts respond differently to broad uncertainty than directional or tail-oriented contracts.

The additional controls further highlight the distinct pricing behavior of Range contracts. The Weekend dummy is uniformly positive and highly significant, pointing to systematically wider discrepancies on weekends. NBP and the Fear & Greed Index are not statistically significant, suggesting that Range-contract wedges are driven less by directional order flow or broad sentiment and more by contract structure, trading conditions, and fragmentation. Overall, Range contracts provide selective but useful support for the theoretical framework: the wedge is

systematic, but its drivers reflect the specific nature of central-interval payoffs rather than the tail-risk channels that dominate Reach and Dip contracts.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		0.066***							
		(0.016)							
USDC Vol_{30}		-0.035***							
		(0.011)							
VIX			-0.060***				-0.081***	-0.072***	-0.095***
			(0.017)				(0.018)	(0.018)	(0.014)
L2 Risk				-0.017**			-0.024***		0.003
				(0.009)			(0.009)		(0.006)
Diff Underlying					0.027***			0.028***	0.013*
					(0.010)			(0.010)	(0.007)
NBP									-0.004
									(0.004)
Fear & Greed						0.008	-0.003	-0.003	-0.008
						(0.008)	(0.009)	(0.009)	(0.008)
Announcement						0.019	0.016	0.016	-0.069***
						(0.024)	(0.024)	(0.025)	(0.012)
Δ (%)	0.002	0.002	0.002	0.002	0.002	0.001	0.002	0.001	-0.003
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)	(0.003)
Δ USDC (%)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.003)
Weekend						0.054***	0.051***	0.050***	0.059***
						(0.017)	(0.017)	(0.016)	(0.008)
Moneyness	-0.027***	-0.029***	-0.028***	-0.027***	-0.027***	-0.026***	-0.028***	-0.028***	0.023***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.005)
Maturity	0.066***	0.066***	0.066***	0.066***	0.066***	0.064***	0.065***	0.064***	-0.048***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)	(0.005)
ETH	0.056***	-0.058**	0.167***	0.056***	0.037***	0.056***	0.206***	0.171***	0.210***
	(0.008)	(0.025)	(0.033)	(0.008)	(0.008)	(0.008)	(0.035)	(0.036)	(0.030)
R^2 Adj	0.009	0.010	0.009	0.009	0.009	0.010	0.010	0.010	0.009
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	307,573	307,573	307,536	307,573	307,573	307,573	307,536	307,536	63,864

Table 6.3: Panel Regression – Range Bets – CLOB–OIP. This table reports the results from a pooled OLS panel regression of the CLOB–OIP price wedge for the Range contracts on Moneyness, Maturity, volatility, USDC volatility, and additional market structure variables: L2 Risk, Diff Underlying, NBP, Fear & Greed, and Announcement. We include returns, USDC returns, and weekend and ETH dummy variables as controls. The construction of the variables is explained in Section 3. The data range from January 2024 to November 2025. Standard errors are clustered by time and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All continuous variables are standardized to have a mean of 0 and a variance of 1; indicator variables remain in levels.

Table 6.4 reports pooled OLS results for Reach contracts with monthly fixed effects. Moneyness and maturity are both positive and highly significant throughout, indicating larger wedges for contracts closer to the relevant barrier and for longer-dated contracts. These results provide strong support for the behavioral probability-weighting and Bayesian-learning predictions. The risk and friction variables also point to important limits-to-arbitrage effects. The VIX and Diff Underlying enter positively and significantly, whereas Vol_{30}, USDC volatility, and L2 Risk enter negatively. Thus, broad option-market uncertainty and cross-exchange fragmentation widen Reach-contract wedges, while the other risk and friction proxies have the opposite association.

The additional controls further highlight the payoff-specific nature of Reach-contract pricing. The Weekend dummy is uniformly negative and highly significant, indicating smaller wedges on

weekends, while the Fear & Greed Index is also negative and significant, suggesting that the wedge is larger outside periods of broad market optimism. Announcement effects are limited, but positive in one of the larger specifications. Overall, Reach contracts provide clear evidence that contract characteristics, maturity, market-wide uncertainty, and cross-exchange fragmentation are central drivers of the wedge, supporting the view that upside tail contracts are especially sensitive to probability weighting, learning, and limits to arbitrage.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		-0.089***							
		(0.013)							
USDC Vol_{30}		-0.052***							
		(0.009)							
VIX			0.093***				0.061***	0.057***	0.100***
			(0.012)				(0.011)	(0.011)	(0.012)
L2 Risk				-0.073***			-0.062***		-0.048***
				(0.008)			(0.008)		(0.008)
Diff Underlying					0.143***			0.142***	0.234***
					(0.012)			(0.012)	(0.016)
NBP									-0.004
									(0.003)
Fear & Greed						-0.063***	-0.040***	-0.047***	-0.063***
						(0.010)	(0.010)	(0.010)	(0.009)
Announcement						0.024	0.017	0.040**	-0.004
						(0.017)	(0.017)	(0.017)	(0.014)
Δ (%)	0.005	0.006	0.005	0.005	0.005	0.005	0.005	0.004	-0.001
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)
Δ USDC (%)	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.003
	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Weekend						-0.146***	-0.145***	-0.144***	-0.132***
						(0.010)	(0.010)	(0.010)	(0.011)
Moneyness	0.155***	0.154***	0.155***	0.158***	0.159***	0.159***	0.159***	0.162***	0.193***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)
Maturity	0.399***	0.394***	0.397***	0.404***	0.397***	0.400***	0.402***	0.396***	0.416***
	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
ETH	0.123***	0.283***	-0.065**	0.123***	0.081***	0.124***	-0.000	-0.034	-0.100***
	(0.006)	(0.025)	(0.026)	(0.006)	(0.006)	(0.006)	(0.024)	(0.024)	(0.027)
R^2 Adj	0.282	0.285	0.283	0.284	0.286	0.288	0.289	0.292	0.341
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	134,736	134,736	134,711	134,736	134,736	134,736	134,711	134,711	89,419

Table 6.4: Panel Regression – Reach Bets – CLOB–OIP. This table reports the results from a pooled OLS panel regression of the CLOB–OIP price wedge for the Reach contracts on Moneyness, Maturity, volatility, USDC volatility, and additional market structure variables: L2 Risk, Diff Underlying, NBP, Fear & Greed, and Announcement. We include returns, USDC returns, and weekend and ETH dummy variables as controls. The construction of the variables is explained in Section 3. The data range from January 2024 to November 2025. Standard errors are clustered by time and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All continuous variables are standardized to have a mean of 0 and a variance of 1; indicator variables remain in levels.

Table 6.5 reports pooled OLS results for Dip contracts with monthly fixed effects. Moneyness is positive and highly significant throughout, indicating larger wedges for contracts closer to the relevant downside barrier. Maturity is small and unstable across specifications, becoming slightly negative in the richer specifications and providing little evidence of a robust maturity effect. The risk and fragmentation variables provide strong evidence that downside-tail wedges are related to market uncertainty. The VIX enters positively and highly significantly in all specifications in

which it is included, and Diff Underlying is positive and highly significant throughout, consistent with the view that market-wide uncertainty and cross-exchange fragmentation widen the CLOB–OIP wedge for downside barrier contracts.

The information and sentiment variables also support a downside-protection interpretation. NBP is positive and significant, indicating that buying pressure for Dip contracts is associated with larger wedges. The Fear & Greed Index is negative and highly significant, implying that wedges are larger during periods of fear rather than greed, as expected for bearish tail contracts. The Announcement indicator is positive and highly significant throughout, indicating that salient macroeconomic and policy events increase discrepancies for Dip contracts. The Weekend dummy is uniformly negative and highly significant, indicating smaller wedges on weekends. Overall, the results suggest that Dip contracts embed a downside tail-risk premium: wedges are largest when contracts move closer to the downside barrier, when market-wide uncertainty and cross-exchange fragmentation are elevated, and when sentiment reflects fear-driven demand for protection.

Taken together, Tables 6.2–6.5 show that the framework from Section 4 is informative, but strongly payoff-specific. Cross-market price differences between Polymarket and Deribit are systematically related to contract characteristics, market uncertainty, sentiment, announcements, and market frictions, but the sign and strength of these effects vary across contract types. The moneyness prediction is the most robust result across specifications: wedges are larger for more OTM Above and Range contracts, and for Reach and Dip contracts that move closer to their relevant barriers. Maturity is positive for Above, Range, and Reach contracts, consistent with larger discrepancies earlier in the contract life, while Dip contracts display a small and unstable maturity effect, suggesting a more short-horizon downside-protection component.

The risk and friction channels also depend on the payoff structure. The VIX is positive and significant for Above, Reach, and Dip contracts, indicating that option-market uncertainty widens wedges, especially for directional and tail-oriented contracts. Cross-exchange fragmentation, measured by Diff Underlying, is particularly important for Range, Reach, and Dip contracts, while L2 Risk matters most clearly for Above contracts. Sentiment and event effects are similarly asymmetric. Fear & Greed is negative for Above, Reach, and Dip contracts, suggesting that wedges are larger outside periods of broad crypto-market optimism, while announcements increase discrepancies, especially for Above and Dip contracts. Overall, the results suggest that cross-market misalignment is systematic, but its drivers depend crucially on payoff structure and on whether traders are pricing central outcomes, upside tails, or downside tails.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		-0.026*							
		(0.015)							
USDC Vol_{30}		-0.009							
		(0.006)							
VIX			0.180***				0.096***	0.081***	0.084***
			(0.015)				(0.015)	(0.016)	(0.027)
L2 Risk				-0.067***			-0.025***		0.010
				(0.008)			(0.008)		(0.007)
Diff Underlying					0.192***			0.161***	0.189***
					(0.010)			(0.010)	(0.015)
NBP									0.011***
									(0.003)
Fear & Greed						-0.165***	-0.142***	-0.135***	-0.245***
						(0.011)	(0.012)	(0.011)	(0.015)
Announcement						0.075***	0.069***	0.080***	0.137***
						(0.019)	(0.019)	(0.018)	(0.020)
Δ (%)	0.003	0.003	0.002	0.003	0.002	0.002	0.001	0.001	0.007
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)
Δ USDC (%)	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)
Weekend						-0.199***	-0.197***	-0.198***	-0.189***
						(0.010)	(0.010)	(0.010)	(0.013)
Moneyness	0.093***	0.092***	0.099***	0.094***	0.090***	0.076***	0.082***	0.078***	0.072***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.007)
Maturity	0.009**	0.010**	0.002	0.011***	0.004	0.002	0.000	-0.004	-0.006
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)
ETH	-0.017***	0.027	-0.347***	-0.017**	-0.063***	-0.029***	-0.202***	-0.214***	-0.241***
	(0.007)	(0.026)	(0.029)	(0.007)	(0.007)	(0.007)	(0.029)	(0.029)	(0.048)
R^2 Adj	0.060	0.060	0.064	0.062	0.067	0.079	0.081	0.085	0.129
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	121,861	121,861	121,826	121,861	121,861	121,861	121,826	121,826	71,470

Table 6.5: Panel Regression – Dip Bets – CLOB–OIP. This table reports the results from a pooled OLS panel regression of the CLOB–OIP price wedge for the Dip contracts on Moneyness, Maturity, volatility, USDC volatility, and additional market structure variables: L2 Risk, Diff Underlying, NBP, Fear & Greed, and Announcement. We include returns, USDC returns, and weekend and ETH dummy variables as controls. The construction of the variables is explained in Section 3. The data range from January 2024 to November 2025. Standard errors are clustered by time and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All continuous variables are standardized to have a mean of 0 and a variance of 1; indicator variables remain in levels.

The panel regressions identify systematic, payoff-specific drivers of the CLOB–OIP wedge. We next complement these average relationships with more direct evidence on the mechanisms underlying the observed deviations.

6.2 Additional Results

This subsection examines the four channels from Section 4 in greater detail. We first study extreme-event pricing and trading around salient moneyness states, and then examine how wedges evolve over the contract life and across market conditions.

6.2.1 Behavioral Probability Weighting

This subsection provides additional evidence that pricing wedges are largest for salient tail events. We first examine representative Reach and Dip contracts with very low option-implied

probabilities and then study how trading activity and wedges change when contracts approach economically salient payoff boundaries.

We begin by examining whether Polymarket assigns relatively high prices to extreme outcomes. Figure 6.1 compares CLOB and OIP prices for four representative tail-event contracts. Panels (a) and (b) show the BTC Reach 200,000 and Dip 40,000 contracts, while Panels (c) and (d) report the corresponding ETH Reach 8,000 and Dip 1,000 contracts. For the BTC Reach contract, CLOB prices remain above OIP throughout the sample, starting near 3% and gradually declining toward zero. Short-lived spikes in the CLOB price suggest episodes of heightened optimism or speculative demand, while OIP remains lower and close to zero near resolution. The BTC Dip contract displays a similar pattern, with CLOB prices peaking above 2% before declining, whereas OIP movements remain smaller. The ETH contracts show comparable dynamics: CLOB prices are generally higher and more volatile than their option-implied counterparts and exhibit occasional spikes before converging toward zero. Overall, these examples are consistent with probability weighting or narrative-driven speculation increasing the prices assigned to highly unlikely outcomes.

The representative contracts suggest that extreme outcomes may receive disproportionate weight. We next examine whether this pattern extends beyond individual examples by studying trading activity and pricing wedges when contracts approach their relevant payoff boundaries.

We focus on the subset of contracts that approach their ATM payoff boundary. We define the event as the first active hour in which the spot price is within two percent of that boundary: for Reach, $S_t/K \geqslant 0.98$; for Dip, $K/S_t \geqslant 0.98$; for Above, $|S_t/K - 1| \leqslant 0.02$; and for Range, $\min\{|S_t/K_L - 1|, |S_t/K_U - 1|\} \leqslant 0.02$. Figure 6.2 then uses the available contract-hours in the 24 hours before and after this event, pools Reach/Dip and Above/Range separately, and plots transaction counts and CLOB–OIP wedges in percentage points, after subtracting each series' own pre-window mean at the contract level.

The left panel shows the activity spike behind the event-time exercise: transaction counts for Reach/Dip contracts rise sharply as the barrier enters the two-percent band, whereas the Above/Range activity response is much smaller. The right panel separates this activity timing from directional disagreement. Around the same event, the Reach/Dip wedge rises relative to its own pre-event baseline, while the Above/Range series falls. The joint pattern is consistent with path-dependent contracts drawing attention near salient moneyness states and exhibiting a different wedge response from terminal payoff-boundary contracts.

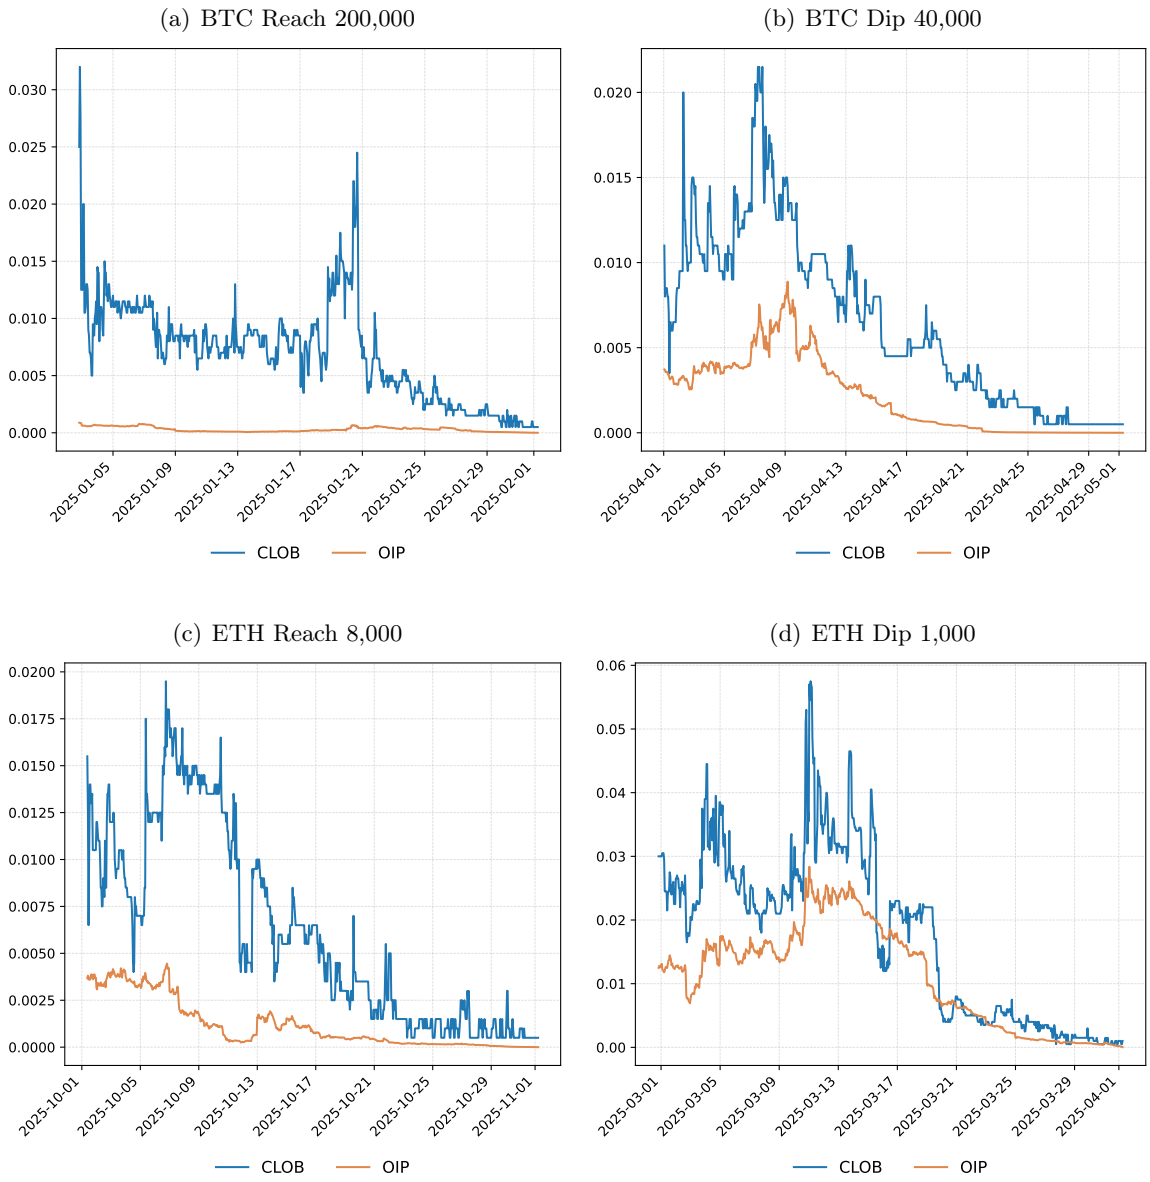


Figure 6.1: Behavioral Biases for Extreme Events. The figure displays the CLOB and OIP over time for four representative contracts. Panel (a) corresponds to the contract "BTC reaches 200,000 before February 2025," while Panel (b) corresponds to "BTC dips below 40,000 before May 2025." Panel (c) shows the contract "Will Ethereum reach 8,000 in October?", and Panel (d) shows "Will Ethereum dip to 1,000 by March 31?". The different types of contracts are discussed in Section 2.2.

Taken together, the event-time evidence mirrors the regression results in Tables 6.2–6.5: as contracts approach their payoff boundaries, wedges rise for Reach and Dip contracts, consistent with their positive moneyness coefficients, but fall for Above and Range contracts, consistent with their negative coefficients. This contrast reinforces the payoff-specific role of moneyness. We next examine whether wedges decline as uncertainty is resolved over the contract life.

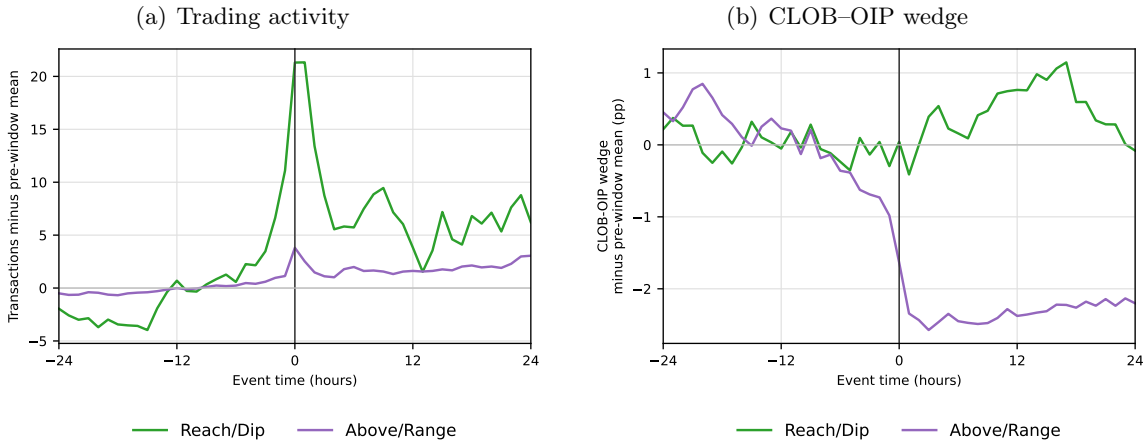


Figure 6.2: Trading Activity and CLOB–OIP Around Events. Panel (a) plots event-time means of transaction counts for Reach/Dip contracts and Above/Range contracts. Transaction-count activity is the number of observed Polymarket trades in the contract-hour. Panel (b) plots the corresponding CLOB–OIP wedge in percentage points. Event time zero is the first hour in which the contract reaches the two-percent moneyness event. The window spans 24 hours before and 24 hours after the event, and each series subtracts its own pre-window mean.

6.2.2 Bayesian Learning

The Bayesian-learning hypothesis predicts that pricing wedges are largest early in the contract life and decline as information accumulates. Table 6.6 compares average CLOB–OIP wedges across the beginning, middle, and end of each contract. Most differences are statistically significant at the 1% level.

Wedges are largest at inception and generally decline toward maturity. The pattern is strongest for *Reach* contracts, where the wedge falls from 0.0380 to essentially zero, and for *Dip* contracts, where it declines from 0.0317 to 0.0087. *Above* and *Range* contracts exhibit the same broad pattern, although their end-of-period wedges remain positive. Overall, the evidence is consistent with pricing discrepancies dissipating as uncertainty is resolved.

As a complementary timing dimension, Table 6.7 compares weekday and weekend wedges. *Above* and *Range* contracts exhibit slightly larger wedges on weekends, whereas *Reach* and *Dip* contracts have larger wedges on weekdays. In the pooled sample, the average wedge is marginally higher on weekdays, 0.0083 versus 0.0075 on weekends. Thus, weekend effects are payoff-specific rather than uniform.

Taken together, these results complete the contract-level analysis by showing that pricing wedges vary systematically across payoff structures, contract stages, and trading conditions. We next turn to the participants behind this activity and examine how trading and realized outcomes are distributed across wallets.

(a) Above Bets

Segment	Mean	p-value	Obs.
Beginning	0.0074	0.0000	52957
Middle	0.0053	0.0000	255046
End	0.0007	0.0055	55359

(b) Range Bets

Segment	Mean	p-value	Obs.
Beginning	0.0236	0.0000	44089
Middle	0.0043	0.0000	219576
End	0.0085	0.0000	43956

(c) Reach Bets

Segment	Mean	p-value	Obs.
Beginning	0.0380	0.0000	9677
Middle	0.0180	0.0000	115183
End	0.0001	0.7400	9913

(d) Dip Bets

Segment	Mean	p-value	Obs.
Beginning	0.0317	0.0000	8689
Middle	0.0050	0.0000	104171
End	0.0087	0.0000	9041

Table 6.6: CLOB–OIP Wedges Across Contract Stages. Panels (a)–(d) report the average CLOB–OIP wedge, p-value, and number of observations for the beginning, middle, and end of Above, Range, Reach, and Dip contracts. The sample runs from January 2024 to November 2025.

	Mean			
Bet Type	Weekend	Weekday	Diff	p-value
Above	0.0055	0.0047	0.0008	0.0000
Range	0.0099	0.0067	0.0032	0.0000
Reach	0.0135	0.0200	-0.0065	0.0000
Dip	0.0008	0.0096	-0.0089	0.0000
All	0.0075	0.0083	-0.0007	0.0000

Table 6.7: CLOB–OIP Wedges on Weekdays and Weekends. This table reports average CLOB–OIP wedges by contract type and trading period. Differences are computed as Weekend minus Weekday.

7 Wallet-Level Trading

This section examines intraday patterns in trading activity and pricing wedges before turning to market participants. We then study the concentration of trading activity and P&L across wallets and whether larger wallets are more exposed to contracts with larger CLOB–OIP wedges.

7.1 Intraday Trading Activity and Pricing Wedges

Figures 7.2 and 7.1 provide a descriptive UTC-hour profile of the CLOB–OIP wedge and trading intensity. Mean transaction count peaks at 14:00 UTC for both BTC and ETH, at 10.10 and 5.48 transactions, respectively, inside the fixed Europe–United States band. The mean CLOB–OIP wedge peaks later, at 17:00 UTC, at 1.19 percentage points for BTC and 1.27 percentage points for ETH. Thus, the timing of maximum average wedge does not coincide with trading activity. The colored windows are fixed descriptive UTC reference periods rather than holiday- or daylight-saving-adjusted exchange calendars.

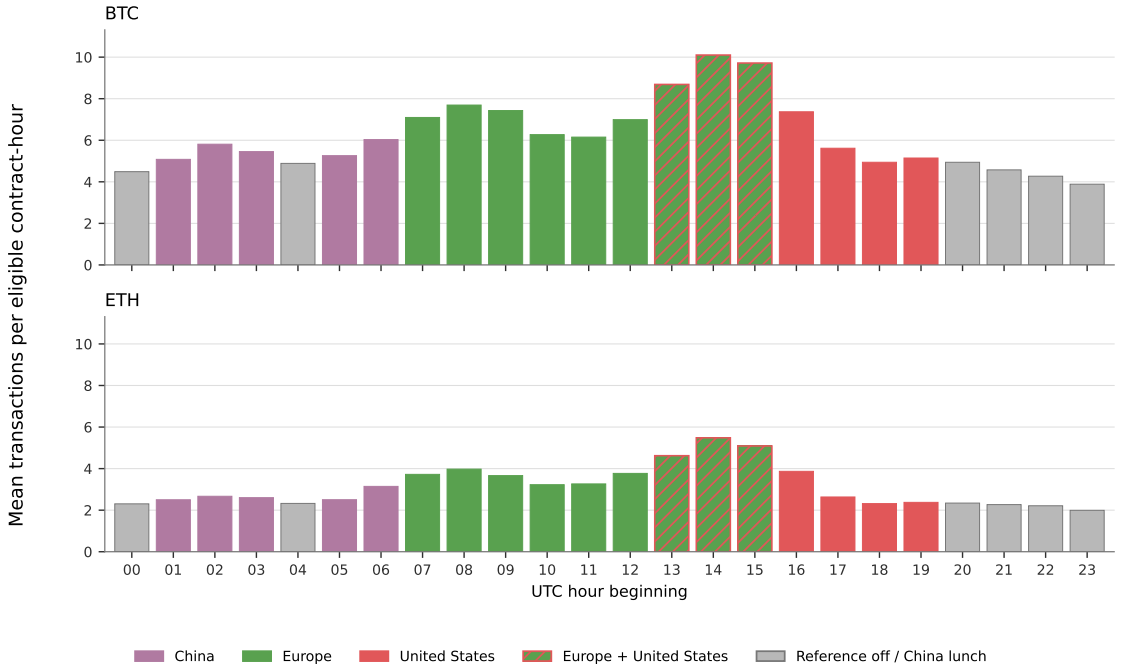


Figure 7.1: Polymarket Trading Activity by UTC Hour. The panels report the mean number of transactions per BTC and ETH contract-hour on all calendar days. Transactions are aggregated to contract-hours and linked to the same sample used in Figure 7.2; hours without a transaction are assigned zero. Bar colors mark the same fixed descriptive UTC reference windows.

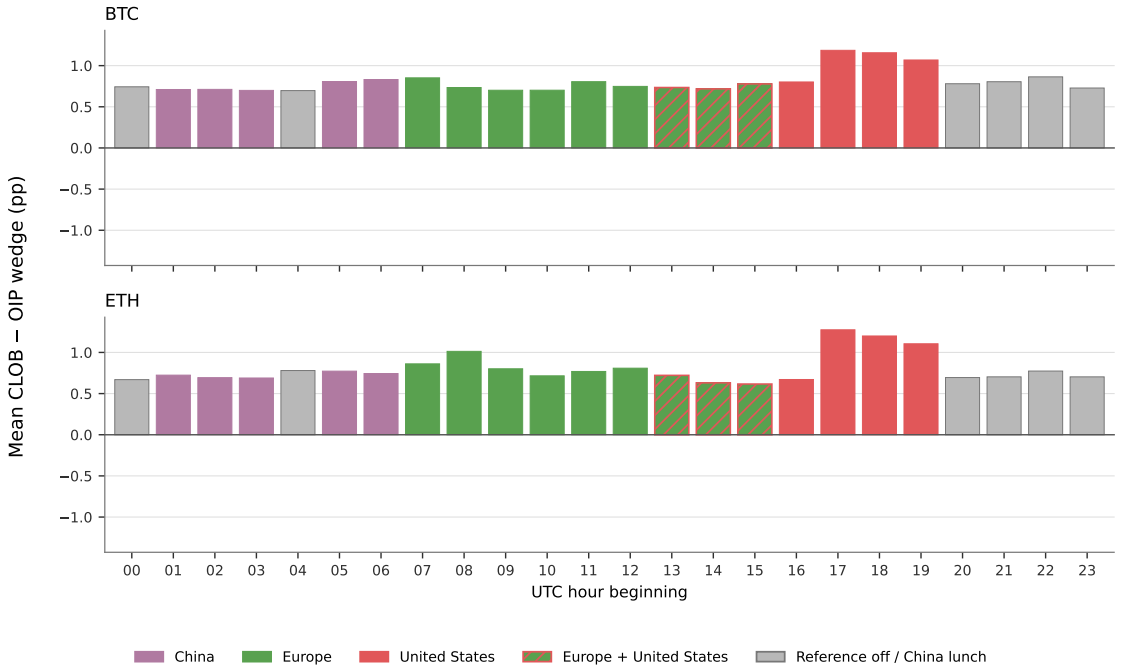


Figure 7.2: CLOB–OIP Wedge by UTC Hour. The panels report the mean CLOB–OIP wedge in percentage points for BTC and ETH on all calendar days. Each contract-hour receives equal weight. Bar colors mark fixed descriptive UTC reference windows; hatched bars mark the Europe–United States overlap. The windows are not adjusted for holidays or daylight-saving changes.

7.2 Concentration of Trading Activity and P&L

We next ask whether the trading gains associated with these markets are broadly distributed across participants or concentrated among a small set of wallets. This provides a direct distributional counterpart to the contract-level evidence and connects our sample to the broader Polymarket evidence in Akey et al. (2026), who document that, across the full platform, the top 0.1% of profitable users capture 51.2% of gains and the top 1% capture 76.5%.

We combine BTC and ETH trade histories with contract-level resolution data and restrict the P&L calculation to markets with observed final outcomes. For each wallet, P&L equals the cash paid or received when trading plus the payoff of its resulting position at resolution. We then aggregate realized P&L, wallet-leg notional, and trade counts at the wallet-address level for BTC, ETH, and the pooled BTC–ETH sample.

Table 7.1 shows that gains are highly concentrated in BTC and ETH markets. Winner Share and Loser Share report the fraction of wallets with positive and negative realized P&L, respectively. The top-gain shares are computed within profitable wallets: for example, Top 1% Gains is the fraction of total positive realized P&L captured by the highest-P&L 1% of profitable wallets. The loss share is defined analogously within losing wallets, ranking losses by absolute magnitude. In the pooled sample, the top 0.1% of profitable wallets capture 58.8% of gains, while the top 1% capture 84.2%. Losses are also concentrated, though less sharply: the top 1% of losing wallets account for 68.8% of pooled losses. These magnitudes exceed the concentration documented by Akey et al. (2026), suggesting that distributional skew is at least as pronounced in crypto contracts as in the broader Polymarket universe.

Symbol	Wallets	Winner Share	Loser Share	Top 0.1% Gains	Top 1% Gains	Top 5% Gains	Top 1% Losses
BTC	565,100	60.6%	38.8%	58.3%	84.9%	95.9%	69.6%
ETH	335,994	66.3%	33.0%	57.6%	84.3%	95.6%	67.4%
Pooled	641,979	58.7%	40.7%	58.8%	84.2%	95.5%	68.8%

Table 7.1: Wallet-Level P&L Concentration. This table reports resolved-market wallet P&L concentration for BTC, ETH, and the pooled BTC–ETH sample. Winner and loser shares report the fraction of wallets with positive and negative realized P&L; top shares report the fraction of total gains or losses accounted for by the corresponding tail of profitable or losing wallets.

Figure 7.3 plots concentration of gains and losses for the pooled sample. The top 5% of profitable wallets account for 95.5% of total gains, while the top 5% of losing wallets account for 88.9% of total losses. The figure shows that concentration is not confined to the very top rank cutoff; most realized gains and losses are generated by a small upper tail of wallets on both sides of the P&L distribution.

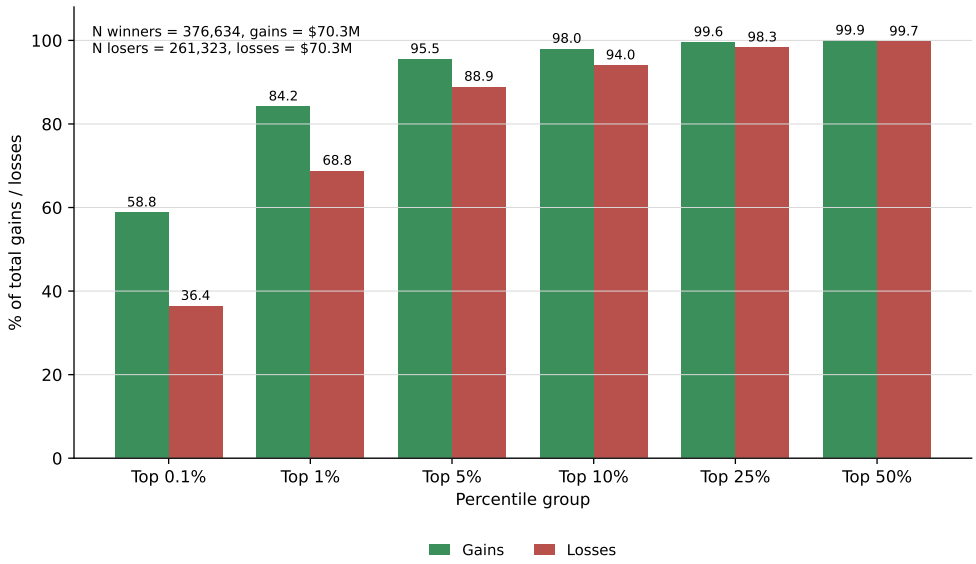


Figure 7.3: Wallet-Level Concentration of Gains and Losses. This figure reports the cumulative share of total realized gains and losses accounted for by the top 0.1%, 1%, 5%, 10%, 25%, and 50% of profitable and losing wallets in the pooled resolved BTC–ETH sample. P&L is computed using the role-expanded taker-counterparty cash-flow convention described in the text.

Trading activity is similarly concentrated. In the pooled resolved sample, the top 1% of wallets account for 73.9% of trades and 89.2% of notional. Less active wallets have negative aggregate P&L and execute over 70% of trades at prices below 0.10 or above 0.90, whereas wallets with at least 100 trades earn positive aggregate P&L and trade less at extreme prices. This pattern is consistent with the longshot-trading channel in Akey et al. (2026). Since the data identify wallets rather than investors and do not separately classify liquidity provision, we interpret the results as descriptive evidence on trading-outcome concentration rather than persistent skill or maker-taker profitability.

The concentration of activity and outcomes raises a natural question: do larger wallets also trade disproportionately in contracts with larger deviations from OIPs?

7.3 Wallet Size and CLOB–OIP Exposure

To study this, we compare wallet size with the CLOB–OIP wedges of the contracts they trade. We restrict the sample to contract-hours with maturities between 3 and 45 days. For each bet, we compute its mean absolute CLOB–OIP wedge, in percentage points, over this window and rank bets into pooled BTC–ETH wedge deciles. Wallets are ranked into pooled BTC–ETH size deciles based on their in-sample gross notional.

Table 7.2 reports the pooled BTC–ETH exposure pattern by wallet-size decile. Moving from

D1 to D10, the notional-weighted absolute wedge traded rises from 0.69 to 4.70 percentage points. The notional share allocated to wedge-decile-10 bets rises from 5.2% to 44.2%, and the share allocated to wedge deciles 8–10 rises from 15.9% to 73.1%. Thus, the concentration documented above is not only a measure of total activity: within each wallet-size group, larger wallets allocate a substantially larger fraction of their own notional to bets with larger price-benchmark wedges.

Wallet-size decile	Wallets	Wallet-leg notional	Unique bets traded	Notional-weighted absolute wedge (pp)	Share in wedge decile 10	Share in wedge deciles 8–10
D1	32,339	1,036.7	401	0.69	5.2%	15.9%
D2	43,035	9,301.1	552	0.88	8.7%	18.5%
D3	37,497	33,903.5	760	1.62	15.5%	29.1%
D4	33,057	76,611.8	822	2.15	22.6%	35.8%
D5	31,696	152,538.9	984	2.39	24.2%	40.7%
D6	29,383	277,777.9	1,271	3.30	34.3%	53.5%
D7	27,221	602,144.8	1,593	3.72	41.6%	63.0%
D8	29,766	1,746,622.7	2,195	3.95	48.6%	70.2%
D9	26,277	4,735,249.8	2,628	4.35	50.6%	74.3%
D10	22,491	435,811,163.2	4,313	4.70	44.2%	73.1%
All wallets	312,762	443,446,350.5	4,313	4.69	44.3%	73.0%

Table 7.2: Wallet Size and Wedge Exposure. This table reports wallet-leg activity in the pooled BTC–ETH sample for bets 3–45 days from maturity. Wallet-size deciles are pooled BTC–ETH deciles of gross notional, with D10 denoting the largest-wallet decile. Wedge deciles are pooled BTC–ETH bet-level deciles of the mean absolute CLOB–OIP wedge in percentage points over the same maturity window. Wallet-leg notional is YES-payoff notional and is not one-sided exchange volume. Notional-weighted absolute wedge is the wallet-leg-notional-weighted absolute CLOB–OIP wedge in percentage points.

Figure 7.4 gives the allocation view. Each bar sums to 100% of a wallet-size decile's notional and partitions that notional into wedge deciles 1–7, 8–9, and 10. The lower wallet deciles are concentrated in contracts with smaller wedges, while the upper wallet deciles shift toward contracts with larger wedges. This pattern shows that the wallets responsible for most activity are also more exposed to the contracts where Polymarket prices differ most from the OIP.

Taken together, these results show that larger wallets trade more in contracts with larger CLOB–OIP wedges. Internet Appendix Table II.4 repeats the exercise for all maturities; the upper wallet-size deciles remain more exposed to larger-wedge bets, although wedge deciles are re-ranked within that broader sample.

The wallet-level analysis shows that exposure to pricing wedges is concentrated among larger wallets. We now move from individual contracts and traders to the aggregate distribution of implied risk, asking whether these cross-market differences also appear in the full RND.

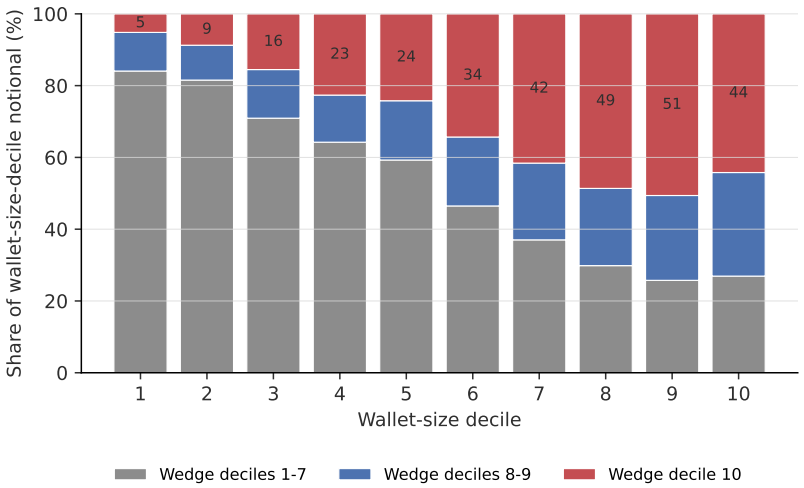


Figure 7.4: Wallet-Size Allocation Across Wedge Tiers. Each bar represents 100% of a wallet-size decile's role-expanded wallet-leg notional in the pooled BTC–ETH 3–45 day sample. The stacked segments allocate that notional across bets in wedge deciles 1–7, 8–9, and 10, where wedge deciles are bet-level deciles of mean absolute CLOB–OIP wedges.

8 Comparing RNDs from Deribit and Polymarket

This section asks whether the cross-market differences documented at the contract level also appear in the full implied distribution. We calibrate the Kou (2002) model directly to Polymarket contracts to construct daily 7-day RNDs and compare them with the corresponding distributions extracted from Deribit option surfaces.

The main finding is that Polymarket and Deribit broadly agree on the direction of risk: both imply negative expected returns, negative skewness, and fat tails. However, the Polymarket RNDs are more dispersed across calibration dates and place more probability outside Deribit's central return band, while their moment-based skewness and kurtosis generally imply milder downside asymmetry than the Deribit benchmark. Regression evidence shows that aggregate distributional disagreement is primarily driven by risk conditions, especially VIX and Vol_{30}. At the moment level, Polymarket reacts more strongly to VIX and USDC Vol_{30} in the mean and volatility regressions, indicating heightened sensitivity to tail-risk states and stablecoin stress. Deribit, by contrast, responds more strongly to Diff Underlying and, for skewness, to Fear & Greed and Daily NBP. Overall, the results suggest that Polymarket prices encode similar broad distributional risks as Deribit options, but the two markets differ in how strongly they respond to stablecoin frictions, sentiment, and underlying-market dislocations.

8.1 Kou Calibration for Polymarket

In parallel with the Deribit calibration in Section 3.2, we also calibrate the Kou (2002) model directly to Polymarket contracts. This yields a Polymarket RND comparable to the Deribit benchmark, but derived from prediction-market prices rather than volatility surfaces.

As before, the free parameter vector is $\vartheta = (\sigma, \lambda, p, \eta_1, \eta_2)$. Calibration is performed daily on the four contract types jointly, using a 7-day rolling window of Polymarket observations. Above and Range contracts are priced as European digitals following Section A.2, while Reach and Dip contracts are priced as single-touch barrier digitals following Section A.4. Each observed Polymarket price is mapped to its model counterpart using the contract characteristics, spot price, risk-free rate, and time to maturity. To reduce the influence of stale or illiquid quotes, we filter the hourly panel by removing contracts with unchanged prices for at least 24 hours, retaining only end-of-day observations, restricting prices to $(0.10, 0.90)$, requiring maturities of 3–60 days, and trimming extreme moneyness observations. Each rolling window typically contains a few hundred contract-day observations across the four contract types.

On each calibration date, we estimate $\vartheta = (\sigma, \lambda, p, \eta_1, \eta_2)$ by minimizing the weighted least-squares pricing-error criterion in (1), with near-the-money contracts receiving greater weight. For Range contracts, moneyness is computed using the mid-strike $(K_{1,j} + K_{2,j})/2$. Since Polymarket prices are in $[0, 1]$, pricing errors are not normalized by the spot price. The optimization is performed using L-BFGS-B under box constraints, warm-started on each date by the Deribit Kou (2002) estimates of the previous calibration date.[11] The warm start anchors the Polymarket estimates near the benchmark parameters and substantially improves the stability of the daily estimates relative to a cold start.

Table 8.1 summarizes the distribution of the Polymarket-calibrated parameters for BTC and ETH and compares them to the Deribit benchmark. Overall, the estimates remain within the same economic range, but the Polymarket calibration is more dispersed, consistent with the noisier prediction-market panel.[12] Relative to the matched Deribit benchmark, Polymarket implies slightly higher diffusive volatility and lower jump intensity. The probability of upward jumps is lower for BTC but nearly identical for ETH, while the larger and more volatile estimates of η_2 suggest differences in the pricing of downside jump risk across the two markets.

Table 8.2 reports the average moments of the Polymarket 7-day RND for BTC and ETH.

[11] Upon failure, the optimizer is restarted up to two times from perturbed starting values, and, whenever the returned solution is boundary-degenerate, a small grid of interior starting points is explored to escape the bounds.

[12] Figures C.4 and C.5 plot the Polymarket-calibrated parameters over time for BTC and ETH, respectively.

Consistent with the Deribit benchmark, both assets display negative annualized means, negative skewness, and excess kurtosis, indicating asymmetric downside risk and fat tails.[13] Quantitatively, however, the Polymarket RNDs are less negatively skewed and less leptokurtic than the Deribit benchmark, especially for ETH. This suggests that Polymarket prices reflect the same broad pattern of downside risk and fat tails, but imply less extreme moment-based skewness and kurtosis than the Deribit benchmark.

(a) BTC – Summary Statistics

Parameter	Q25	Q75	Min	Max	Mean	Std
σ	0.3266	0.4634	0.2244	0.7970	0.4079	0.1007
λ	5.1509	7.9971	0.6151	9.5843	6.2706	2.2312
p	0.2534	0.4653	0.0325	0.9033	0.3611	0.1564
η_1	10.0234	13.8887	5.2601	19.3909	12.0751	2.7143
η_2	9.9919	13.9682	6.5352	45.3977	14.2896	8.4174
BTC	90,967.09	110,547.40	60,198.78	125,097.86	99,461.94	14,718.05

(b) ETH – Summary Statistics

Parameter	Q25	Q75	Min	Max	Mean	Std
σ	0.5716	0.6368	0.4666	0.7443	0.6011	0.0589
λ	7.5479	8.3915	0.4134	9.3990	7.4853	1.8637
p	0.2474	0.3376	0.0342	0.9416	0.3146	0.1301
η_1	8.4104	10.0083	3.8288	18.7979	9.6183	2.2970
η_2	7.7482	10.0247	2.6960	45.7213	11.8693	9.1353
ETH	2,592.38	4,223.66	1,775.80	4,851.55	3,451.13	829.71

Table 8.1: Kou Parameters – Polymarket Calibration. This table reports summary statistics for the BTC and ETH parameters (σ, λ, p, η_1, η_2) estimated from the calibration to Polymarket digital contracts over a 7-day rolling window, together with the spot price. For each asset, the table reports the 25th percentile, the 75th percentile, the minimum, the maximum, and the standard deviation.

8.2 Tail-Probability Comparison

We next move from moment-based comparisons to tail probabilities. To compare tail regions in probability units, we fix the Deribit 7-day log-return cutoffs and evaluate the Polymarket RND at those same cutoffs. For asset a and date t, let $Q^D_{a,t}(0.05)$ and $Q^D_{a,t}(0.95)$ denote the Deribit 5th and 95th percentiles. We compute for both tails (L and R)

$$\Delta P^L_{a,t} = F^{PM}_{a,t}\big(Q^D_{a,t}(0.05)\big) - 0.05, \qquad \Delta P^R_{a,t} = 1 - F^{PM}_{a,t}\big(Q^D_{a,t}(0.95)\big) - 0.05.$$

[13] Figures III.3 and III.4 show the time-series behavior of the quantile bands, annualized standard deviation, skewness, and kurtosis of the Polymarket RND for BTC and ETH, respectively. Figure C.6 complements these plots with the full three-dimensional time evolution of the RND shape.

(a) BTC – RND Summary Statistics

Mean	Volatility	Skewness	Kurtosis
-0.1359	0.5167	-0.4554	8.2036
(0.0598)	(0.1080)	(0.4607)	(3.0635)

(b) ETH – RND Summary Statistics

Mean	Volatility	Skewness	Kurtosis
-0.2738	0.7433	-0.4499	6.9578
(0.0615)	(0.0877)	(0.3516)	(2.0084)

Table 8.2: Kou RND Moments – Polymarket Calibration. This table reports the average moments of the Polymarket RND for a 7-day maturity, estimated from the Kou calibration to Polymarket digital contracts. For each panel, we report the annualized mean and volatility, together with skewness and kurtosis; standard deviations are shown in parentheses. Panels A and B correspond to BTC and ETH, respectively.

These are probability differences, measured in percentage points. Positive values indicate that Polymarket assigns a probability greater than 5% to returns below the Deribit 5th-percentile cutoff or above the Deribit 95th-percentile cutoff, respectively.

As an illustration, if the Deribit 5th-percentile cutoff is -8%, then Deribit assigns 5 percent probability to returns below -8%. If Polymarket assigns 7 percent probability to returns below the same cutoff, then $\Delta P^L_{a,t} = 0.07 - 0.05 = 0.02$, meaning that Polymarket places 2 percentage points more probability mass below the left-tail cutoff. The right-tail measure is interpreted analogously: if Polymarket assigns 6 percent probability to returns above the Deribit 95th-percentile cutoff, then $\Delta P^R_{a,t} = 0.06 - 0.05 = 0.01$, meaning that Polymarket places 1 percentage point more mass above the right-tail cutoff.

Figure 8.1 reports these daily probability gaps for BTC and ETH. The dashed horizontal lines show sample averages. The updated average gaps are positive in all four cases: BTC has mean gaps of 0.85 percentage points in the left tail and 0.82 percentage points in the right tail, while ETH has corresponding gaps of 0.21 and 0.48 percentage points. Using Newey–West standard errors with seven daily lags, all four averages are statistically different from zero. Economically, Polymarket places less mass inside Deribit's 5–95 percent return band and more mass in the tails outside it. This indicates a greater Polymarket probability mass outside the Deribit moderate-tail band.

To make the comparison economically tangible, Figure 8.2 translates the BTC 5th and 95th percentile log-return cutoffs into BTC terminal prices. In the average BTC distribution, normalizing the spot price to \$100,000, the Deribit 5–95 percent interval is approximately \$90.1k–

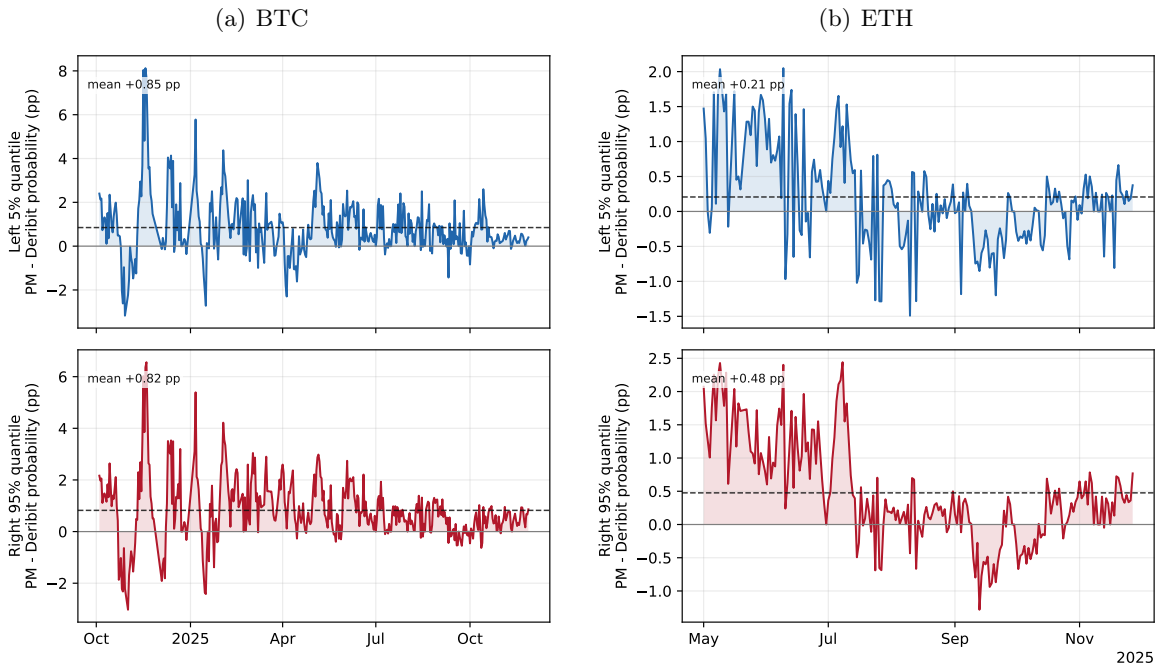


Figure 8.1: RND Tail-Probability Differences. Panels (a) and (b) report BTC and ETH, respectively. Within each panel, the upper plot reports $F^{PM}_{a,t}(Q^{D}_{a,t}(0.05)) - 0.05$, and the lower plot reports $1 - F^{PM}_{a,t}(Q^{D}_{a,t}(0.95)) - 0.05$. These are daily Polymarket-minus-Deribit tail-probability differences at Deribit's seven-day RND log-return cutoffs, reported in percentage points. Dashed horizontal lines show the sample average within each plot.

$110.1k, while the Polymarket interval is approximately $89.5k–$110.7k. The average difference is therefore small in dollar terms, consistent with the modest average tail-probability gaps above.

The selected event rows in Figure 8.2 focus on dates where this comparison is especially visible. The first two rows, November 18–19, 2024, cover the post-election BTC rally and the launch window for Nasdaq-listed IBIT options, when attention to BTC upside and derivatives access was unusually high; they should therefore be read as one short episode rather than independent shocks. The January 6, 2025 row captures BTC's move back above $100,000, while the February 2, 2025 row captures the tariff-shock weekend, when crypto repriced while traditional markets were closed. These dates are descriptive events rather than causal tests. For example, on November 18, 2024, when BTC spot was about $90.5k, Deribit implied a 5–95 percent interval of approximately $79.1k–$103.0k, whereas Polymarket implied $74.5k–$109.1k. During these episodes, Polymarket implies wider moderate tails than Deribit, extending further on both the downside and upside.

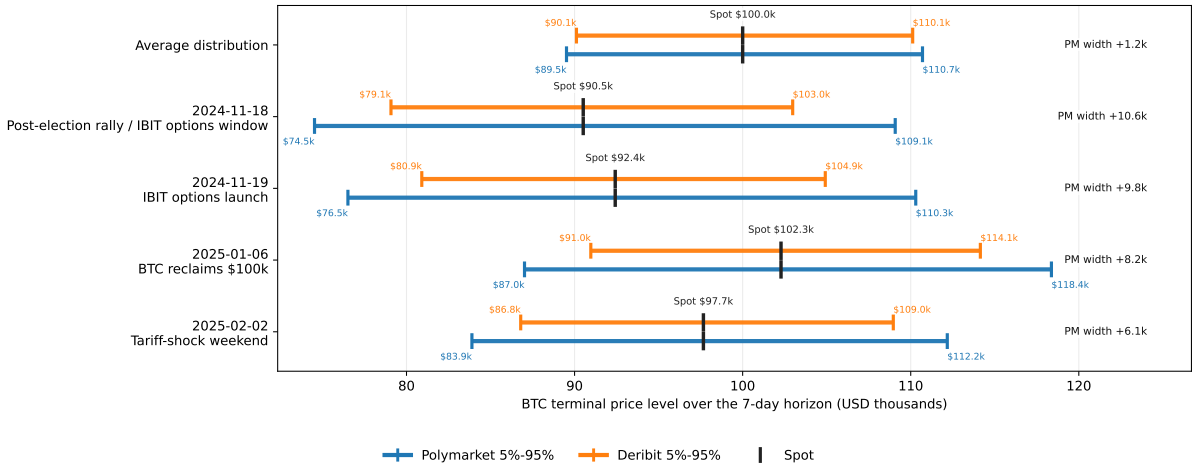


Figure 8.2: BTC RND Tail Price Intervals. This figure translates the 5th and 95th percentiles of the 7-day BTC RNDs into terminal BTC price levels. The average row normalizes spot to $100,000; event rows use the daily BTC spot price. The event rows are November 18–19, 2024, covering the post-election rally and IBIT options launch; January 6, 2025, when BTC moved back above $100,000; and February 2, 2025, the tariff-shock weekend. Horizontal intervals report the implied 5–95 percent price range for Polymarket and Deribit, and vertical tick marks denote the interval endpoints and spot price.

9 Robustness

We conduct several robustness tests.

9.1 Regression Analysis with Contract Fixed Effects

Appendix E re-estimates the nine baseline specifications with contract and calendar-month fixed effects. The contract fixed effects absorb time-invariant differences across individual bets, so the coefficients are identified from within-contract variation. Tables E.1–E.4 show that several central results remain robust.

For Above contracts, moneyness remains negative and maturity positive across all specifications. The positive VIX and Announcement effects, as well as the negative effects of Vol_{30}, Diff Underlying, and Fear & Greed, also remain. For Range contracts, moneyness remains negative and maturity positive in the main sample, while the positive Weekend and negative VIX effects are preserved. However, several other risk measures, including Vol_{30}, USDC volatility, and Diff Underlying, change sign relative to the pooled estimates.

For Reach contracts, maturity remains strongly positive, and the signs of the VIX, Vol_{30}, USDC volatility, L2 Risk, Diff Underlying, and Weekend coefficients are broadly unchanged. The moneyness coefficient, however, becomes negative in the richer specifications, and the Fear & Greed coefficient changes sign. For Dip contracts, the main results are particularly stable: moneyness, VIX, Diff Underlying, NBP, and Announcement remain positive, Weekend remains

negative, and maturity remains small and unstable. Overall, the fixed-effect estimates confirm that the main relationships are not driven solely by persistent differences across contracts, although some risk and sentiment coefficients remain sensitive to the specification.

9.2 Benchmark Validation

We validate Kou (2002) model against a simpler Breeden and Litzenberger (1978) benchmark. This validation is particularly relevant because Reach and Dip contracts require a model: their payoffs depend on whether the underlying crosses a barrier before expiration, which cannot be identified from the terminal risk-neutral distribution alone. For each asset and date, we compare the seven-day Kou risk-neutral density with a model-free risk-neutral density (Breeden and Litzenberger (1978)) recovered from the same Deribit volatility surface over the same finite support of 50 strikes. Across 1,074 matched dates per asset, the average Kolmogorov–Smirnov distance is 0.036 for BTC and 0.032 for ETH, while the corresponding total-variation distances are 0.064 and 0.059. Relative to the model-free benchmark, Kou assigns approximately one percentage point less probability to an Above payoff at the forward, but between 1.30 and 1.72 percentage points more probability to outcomes below $0.9F$ and above $1.1F$. The two densities are therefore broadly aligned around the center of the distribution, while Kou places more probability mass in both tails.

Figure 9.1 translates this difference into seven-day terminal Above probabilities. For every date, we identify the available Above thresholds, retain strictly out-of-the-money observations, price each threshold under both densities, and aggregate by $K/S_0 - 1$. Kou produces a higher mean price in every displayed bin, with peak mean differences of 3.03 percentage points for BTC and 2.65 percentage points for ETH. The comparison supports Kou as a richer tail-sensitive pricing benchmark. Appendix D reports the density comparison, summary statistics, and two illustrative examples.

10 Conclusion

This paper evaluates pricing efficiency in decentralized prediction markets by comparing Polymarket prices with option-implied benchmarks from Deribit. Focusing on Bitcoin- and Ethereum-related contracts across four payoff structures, *Above*, *Range*, *Reach*, and *Dip*, we document substantial alignment between prediction-market prices and derivative-implied valuations.

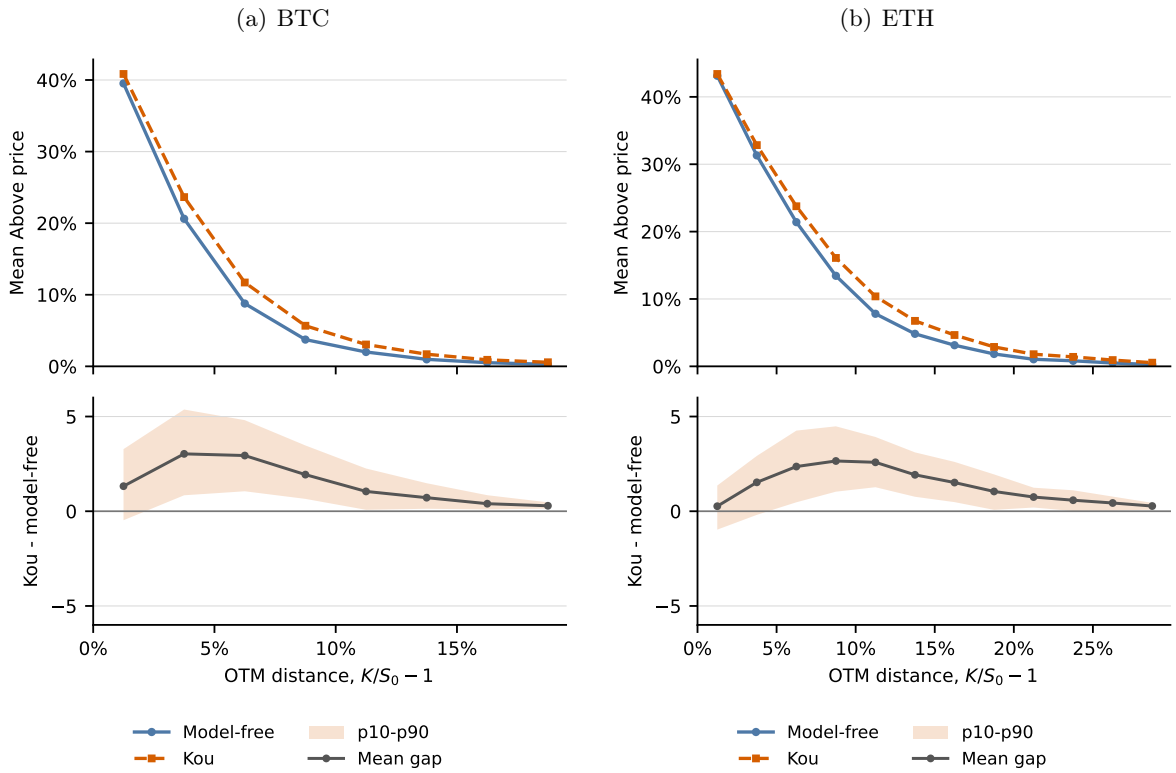


Figure 9.1: OTM Above Contracts – Model-Free vs. Kou. Panels (a) and (b) report the BTC and ETH results, respectively. Within each panel, the upper plot shows mean terminal seven-day Above probabilities from the model-free Breeden and Litzenberger (1978) and Kou (2002) densities by daily OTM distance. The lower plot shows the mean Kou-minus-model-free difference; the shaded area is the 10th–90th percentile range.

At the same time, CLOB–OIP wedges are systematic rather than random. They are largest early in the contract life, especially pronounced for tail and barrier contracts, and vary with volatility, sentiment, directional demand, announcements, and market-friction proxies. Weekday–weekend effects are not uniform; their sign and magnitude depend on the payoff structure.

A distribution-level comparison further shows that Polymarket and Deribit encode similar broad crypto risks, including negative expected returns, negative skewness, and fat tails. However, the two markets differ in the amount and shape of tail risk, as well as in their sensitivity to stablecoin stress, volatility, sentiment, and underlying-market dislocations.

Overall, decentralized prediction markets display substantial informational alignment with derivatives markets, while their payoff design, user base, and infrastructure generate persistent deviations from derivative-implied values. These insights contribute to a deeper understanding of market integration, belief formation, and the boundaries of informational efficiency in emerging decentralized financial ecosystems.

References

Joseph Abate and Peter P. Valkó. Multi-precision laplace transform inversion. *International Journal for Numerical Methods in Engineering*, 60(5):979–993, 2004.

Yacine Aït-Sahalia and Andrew W. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53(2):499–547, 1998. doi: 10.1111/0022-1082.215228.

Yacine Aït-Sahalia, Yubo Wang, and Francis Yared. Do option markets correctly price the probabilities of movement of the underlying asset? *Journal of Econometrics*, 102(1):67–110, 2001. doi: 10.1016/S0304-4076(00)00091-9.

Pat Akey, Vincent Grégoire, Nicolas Harvie, and Charles Martineau. Who wins and who loses in prediction markets? evidence from polymarket. Unpublished working paper, Working Paper, April 2026. URL `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6443103`. Available at SSRN 6443103.

Mukhtar M. Ali. Probability and utility estimates for racetrack bettors. *Journal of Political Economy*, 85(4):803–815, 1977. doi: 10.1086/260601.

Giovanni Angelini. The shape of macroeconomic beliefs. Working Paper 2606.30040, arXiv, 2026.

David Axén and Dominic Cortis. Hedging on betting markets. *Risks*, 8(3):88, 2020.

Gurdip Bakshi and Dilip Madan. Spanning and derivative-security valuation. *Journal of Financial Economics*, 55(2):205–238, 2000. ISSN 0304-405X.

Joyce E. Berg, Forrest D. Nelson, and Thomas A. Rietz. Prediction market accuracy in the long run. *International Journal of Forecasting*, 24(2):285–300, 2008. doi: 10.1016/j.ijforecast.2008.03.007.

N. P. Bollen and R. E. Whaley. Does net buying pressure affect the shape of implied volatility functions? *The Journal of Finance*, 59(2):711–753, 2004.

Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4):621–651, 1978. ISSN 0021-9398.

M. K. Brunnermeier and L. H. Pedersen. Market liquidity and funding liquidity. *The Review of Financial Studies*, 22(6):2201–2238, 2009.

Yiling Chen and David M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pages 49–56, 2007. doi: 10.5555/3020488.3020495.

Rama Cont and Peter Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, 2004.

Philip J. Davis and Philip Rabinowitz. *Methods of Numerical Integration*. Dover, 2 edition, 2007.

Avaneesh Deleep, J. Bai, J. Lee, D. Suresh, and H. Dhawan. How wise is the crowd? bias and edge in prediction markets. Working Paper 6322678, SSRN, 2026. Working-paper metadata should be updated if a revised version is released.

Anthony M. Diercks, Jared Dean Katz, and Jonathan H. Wright. Kalshi and the rise of macro markets. NBER Working Paper 34702, National Bureau of Economic Research, January 2026. URL `http://www.nber.org`.

Barry Eichengreen, Ganesh Viswanath-Natraj, Junxuan Wang, and Zijie Wang. Under pressure? central bank independence meets blockchain prediction markets. July 2025. 65 pages. Posted: 30 Jul 2025.

Robert Forsythe, Forrest Nelson, George R. Neumann, and Jack Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–1161, 1992.

Robert Forsythe, Thomas A. Rietz, and Thomas W. Ross. Wishes, expectations and actions: A survey on price formation in election stock markets. *Journal of Economic Behavior & Organization*, 39(1):83–110, 1999. doi: 10.1016/S0167-2681(99)00027-X.

Romain Gauriot and Lionel Page. How market prices react to information: Evidence from binary options markets. *The Economic Journal*, 2025.

Helyette Geman, Nicole El Karoui, and Jean-Charles Rochet. Changes of numeraire, changes of probability measure and option pricing. *Journal of Applied Probability*, 32(2):443–458, 1995. doi: https://doi.org/10.2307/3215299.

J. Gil-Pelaez. Note on the inversion theorem. *Biometrika*, 38(3-4):481–482, 1951. doi: https://doi.org/10.1093/biomet/38.3-4.481.

Roberto Gómez-Cram, Yunhan Guo, Theis Ingerslev Jensen, and Howard Kung. Financial prediction markets: A new measure of earnings expectations. Unpublished working paper, Working Paper, 2026a. Available at SSRN.

Roberto Gómez-Cram, Yunhan Guo, Theis Ingerslev Jensen, and Howard Kung. Prediction market accuracy: Crowd wisdom or informed minority? Unpublished working paper, Working Paper, 2026b. Available at SSRN.

Denis Gromb and Dimitri Vayanos. Limits of arbitrage: The state of the theory. In *Annual Review of Financial Economics*, volume 2, pages 251–275. 2010. doi: 10.1146/annurev-financial-073009-104107.

S. J. Grossman and M. H. Miller. Liquidity and market structure. *The Journal of Finance*, 43 (3):617–633, 1988.

Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1): 107–119, 2003. doi: 10.1023/A:1022058209073.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979. doi: 10.2307/1914185.

Steven G. Kou. A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101, 2002. doi: https://doi.org/10.1287/mnsc.48.8.1086.166.

Steven G. Kou and Hui Wang. First passage times of a jump diffusion process. *Advances in Applied Probability*, 35(2):504–531, 2003. doi: https://doi.org/10.1239/aap/1051201658.

Steven G. Kou and Hui Wang. Option pricing under a double exponential jump diffusion model, 2004. Working paper; widely circulated manuscript, available on author pages.

A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

Andreas E. Kyprianou. *Fluctuations of Lévy Processes with Applications*. Springer, 2 edition, 2014.

Owen A. Lamont and Richard H. Thaler. Can the market add and subtract? mispricing in tech stock carve-outs. *Journal of Political Economy*, 111(2):227–268, 2003. doi: 10.1086/367683.

Amanda McCullough. Polymarket historical accuracy and bias, 2022. Retrieved from `https://finance.yahoo.com/news/polymarket-achieves-94-accuracy-predicting-010322785.html`.

Mark Mitchell, Todd Pulvino, and Erik Stafford. Limited arbitrage in equity markets. *Journal of Finance*, 57(2):551–584, 2002. doi: 10.1111/1540-6261.00434.

Kenneth Oliven and Thomas A. Rietz. Suckers are born but markets are made: Individual rationality, arbitrage, and market efficiency on an electronic futures market. *Management Science*, 50(3):336–351, 2004. doi: 10.1287/mnsc.1040.0191.

David M. Pennock, Steve Lawrence, C. Lee Giles, and Finn Årup Nielsen. The real power of artificial markets. *Science*, 291(5506):987–988, 2001. doi: 10.1126/science.291.5506.987.

Charles R. Plott and Shyam Sunder. Efficiency of experimental security markets with insider information: An application of rational-expectations models. *Journal of Political Economy*, 90(4):663–698, 1982. doi: 10.1086/261084.

Charles R. Plott and Shyam Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56(5):1085–1118, 1988. doi: 10.2307/1911360.

D. Prelec. The probability weighting function. *Econometrica*, 66(3):497–527, 1998.

William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 3 edition, 2007.

Daniel Rabetti, Jiaqi Shao, and Che Zhang. Beating the earnings game: Why do prediction markets outperform professional analysts? Working Paper 6649938, SSRN, 2026.

Paul W. Rhode and Koleman S. Strumpf. Historical presidential betting markets. *Journal of Economic Perspectives*, 18(2):127–141, 2004. doi: 10.1257/0895330041371277.

David M. Rothschild and Rajiv Sethi. Trading strategies and market microstructure: Evidence from a prediction market. *Journal of Prediction Markets*, 10(1):1–25, 2016.

Hyun Song Shin. Optimal betting odds against insider traders. *Economic Journal*, 101(408):1179–1185, 1991. doi: 10.2307/2234434.

Hyun Song Shin. Prices of state contingent claims with insider traders, and the favourite–longshot bias. *Economic Journal*, 102(411):426–435, 1992. doi: 10.2307/2234523.

Hyun Song Shin. Measuring the incidence of insider trading in a market for state-contingent claims. *Economic Journal*, 103(420):1141–1153, 1993. doi: 10.2307/2234240.

A. Shleifer and R. W. Vishny. The limits of arbitrage. *The Journal of Finance*, 52(1):35–55, 1997.

Erik Snowberg and Justin Wolfers. Explaining the favorite–longshot bias: Is it risk-love or misperceptions? *Journal of Political Economy*, 118(4):723–746, 2010. doi: 10.1086/655844.

Richard H. Thaler and William T. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988. doi: 10.1257/jep.2.2.161.

A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.

Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18 (2):107–126, 2004. doi: 10.1257/0895330041371321.

Justin Wolfers and Eric Zitzewitz. Interpreting prediction market prices as probabilities. *Economics Letters*, 91(2):225–229, 2006.

A Kou Model

A.1 Risk-neutral Kou model

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geqslant 0}, \mathbb{Q})$ support a Brownian motion W and an independent Poisson process N with intensity $\lambda > 0$. Under $\mathbb{Q}$, the Kou double-exponential jump–diffusion is

$$S_t = S_0 e^{X_t}, \qquad X_t = \mu t + \sigma W_t + \sum_{j=1}^{N_t} Y_j,$$

where $\sigma > 0$, $\mu \in \mathbb{R}$, and the i.i.d. jumps $(Y_j)_{j\geqslant 1}$ are independent of (W, N). The Kou (2002) jump law is

$$\mathbb{Q}(Y \geqslant 0) = p, \quad Y \mid (Y \geqslant 0) \sim \text{Exp}(\eta_1), \qquad \mathbb{Q}(Y < 0) = 1 - p, \quad -Y \mid (Y < 0) \sim \text{Exp}(\eta_2), \tag{A.1}$$

with $p \in (0, 1)$ and $\eta_1, \eta_2 > 0$. For $\theta \in (-\eta_2, \eta_1)$, its MGF is

$$M_Y(\theta) = \mathbb{E}[e^{\theta Y}] = \frac{p\eta_1}{\eta_1 - \theta} + \frac{(1-p)\eta_2}{\eta_2 + \theta}.$$

A.1.1 Martingale condition and drift

Let r be the constant short rate and q the continuous dividend yield. Risk-neutrality requires

$$e^{-(r-q)t} S_t \text{ to be a } \mathbb{Q}\text{-martingale,}$$

equivalently $\mathbb{E}^{\mathbb{Q}}[S_t] = S_0 e^{(r-q)t}$, or $\mathbb{E}^{\mathbb{Q}}[e^{X_t}] = e^{(r-q)t}$. Independence and the Poisson PGF imply

$$\mathbb{E}^{\mathbb{Q}}[e^{X_t}] = \exp\big(t\left(\mu + \tfrac{1}{2}\sigma^2 + \lambda(M_Y(1) - 1)\right)\big),$$

so

$$\mu = (r - q) - \tfrac{1}{2}\sigma^2 - \lambda\big(M_Y(1) - 1\big),$$

the drift used throughout the pricing routines.

A.1.2 Characteristic function

For $u \in \mathbb{R}$, let $\varphi_T(u) = \mathbb{E}^{\mathbb{Q}}[e^{\mathrm{i}uX_T}]$. The Lévy–Khintchine formula gives

$$\varphi_T(u) = \exp(T\psi(u)), \qquad \psi(u) = \mathrm{i}u\mu - \tfrac{1}{2}\sigma^2u^2 + \lambda\big(M_Y(\mathrm{i}u) - 1\big). \tag{A.2}$$

The tractability of M_Y is the model's key analytical feature.

A.2 European (non path-dependent) cash-or-nothing digitals

We now derive the procedures for pricing path-dependent digitals based on whether the process hits a barrier during $[0, T]$.

A.2.1 Payoffs and probability representations

Fix $T > 0$, define $X_T = \log(S_T/S_0)$, and for $K > 0$ set $k = \log(K/S_0)$. We consider

$$\begin{aligned}
\text{Up digital:} \quad & \mathbf{1}_{\{S_T \geqslant K\}} = \mathbf{1}_{\{X_T \geqslant k\}}, \\
\text{Down digital:} \quad & \mathbf{1}_{\{S_T \leqslant K\}} = \mathbf{1}_{\{X_T \leqslant k\}}, \\
\text{Between digital:} \quad & \mathbf{1}_{\{K_1 \leqslant S_T \leqslant K_2\}} = \mathbf{1}_{\{k_1 \leqslant X_T \leqslant k_2\}}, \qquad k_i = \log(K_i/S_0).
\end{aligned}$$

Their time-0 prices are

$$\begin{aligned}
D^{\uparrow}(K,T) = e^{-rT}\mathbb{Q}(X_T \geqslant k), \qquad D^{\downarrow}(K,T) = e^{-rT}\mathbb{Q}(X_T \leqslant k), \\
D^{\mathrm{btw}}(K_1,K_2,T) = e^{-rT}\mathbb{Q}(k_1 \leqslant X_T \leqslant k_2).
\end{aligned}$$

Let $F_Q(x) = \mathbb{Q}(X_T \leqslant x)$ be the $\mathbb{Q}$-CDF of X_T. Then

$$\begin{aligned}
D^{\downarrow}(K,T) = e^{-rT}F_Q(k), \qquad D^{\uparrow}(K,T) = e^{-rT}\big(1 - F_Q(k)\big), \\
D^{\mathrm{btw}}(K_1,K_2,T) = e^{-rT}\big(F_Q(k_2) - F_Q(k_1)\big).
\end{aligned}$$

Thus, pricing European digitals reduces to computing F_Q.

A.2.2 Gil–Pelaez inversion for the CDF

Let X have characteristic function $\varphi(u) = \mathbb{E}[e^{\mathrm{i}uX}]$. If $F(x) = \mathbb{P}(X \leqslant x)$ is continuous at x, the Gil–Pelaez inversion formula (Gil-Pelaez, 1951) is

$$F(x) = \frac{1}{2} - \frac{1}{\pi} \int_0^\infty \Im\left(\frac{e^{-\mathrm{i}ux}\varphi(u)}{u} \right) \mathrm{d}u. \tag{A.3}$$

For $\sigma > 0$, Gaussian convolution gives X_T a continuous density, so (A.3) applies with φ_T from (A.2).

A.2.3 Numerical evaluation of Gil–Pelaez

Equation (A.3) is an oscillatory integral on $[0, \infty)$. Our implementation follows standard numerical inversion practice (see, e.g., Abate and Valkó, 2004; Press et al., 2007):

- **Truncation:** replace ∞ with $U_{\max}$; the error is controlled by the decay of $\varphi_T(u)$. In (A.2), $e^{-\frac{1}{2}\sigma^2 u^2 T}$ ensures rapid decay.
- **Removable singularity at $u = 0$:** since $e^{-\mathrm{i}ux}\varphi(u) = 1 + \mathcal{O}(u)$ as $u \downarrow 0$, its imaginary part is $\mathcal{O}(u)$, so the singularity in (A.3) is removable; we set the integrand to 0 at $u = 0$.
- **Quadrature:** discretize $[0, U_{\max}]$ uniformly and apply composite Simpson's rule, a high-order Newton–Cotes formula (see, e.g., Davis and Rabinowitz, 2007).
- **Windowing:** multiply the integrand by a smooth window, such as Hann, that vanishes at $U_{\max}$ to reduce truncation ringing, a Fourier-inversion stabilization (Press et al., 2007).

A.3 European calls and puts via change of numéraire

The same building blocks yield the European call and put prices used in the routines.

A.3.1 Probability decomposition

Let $k = \log(K/S_0)$. The call payoff is

$$(S_T - K)^+ = (S_0 e^{X_T} - K)\mathbf{1}_{\{X_T > k\}} = S_0 e^{X_T}\mathbf{1}_{\{X_T > k\}} - K\mathbf{1}_{\{X_T > k\}}.$$

Hence

$$C = e^{-rT}\mathbb{E}^{\mathbb{Q}}[(S_T - K)^+] = e^{-rT}\left(S_0 \mathbb{E}^{\mathbb{Q}}[e^{X_T}\mathbf{1}_{\{X_T > k\}}] - K\mathbb{Q}(X_T > k) \right), \tag{A.4}$$

where the first term is an exponentially tilted tail expectation.

A.3.2 Share measure (stock numéraire)

Define the share measure $\mathbb{Q}^S$ by (Geman et al., 1995)

$$\frac{d\mathbb{Q}^S}{d\mathbb{Q}}\Big|_{\mathcal{F}_T} = \frac{e^{X_T}}{\mathbb{E}^{\mathbb{Q}}[e^{X_T}]} = \frac{S_T}{\mathbb{E}^{\mathbb{Q}}[S_T]}.$$

For any integrable $\mathcal{F}_T$-measurable H,

$$\mathbb{E}^{\mathbb{Q}}[e^{X_T}H] = \mathbb{E}^{\mathbb{Q}}[e^{X_T}]\mathbb{E}^{\mathbb{Q}^S}[H].$$

Since $\mathbb{E}^{\mathbb{Q}}[e^{X_T}] = e^{(r-q)T}$, setting $H = \mathbf{1}_{\{X_T>k\}}$ gives

$$\mathbb{E}^{\mathbb{Q}}[e^{X_T}\mathbf{1}_{\{X_T>k\}}] = e^{(r-q)T}\mathbb{Q}^S(X_T > k).$$

Substitution into (A.4) yields

$$C = S_0e^{-qT}\mathbb{Q}^S(X_T > k) - Ke^{-rT}\mathbb{Q}(X_T > k).$$

With $F_Q(x) = \mathbb{Q}(X_T \leqslant x)$ and $F_S(x) = \mathbb{Q}^S(X_T \leqslant x)$,

$$C = S_0e^{-qT}\big(1 - F_S(k)\big) - Ke^{-rT}\big(1 - F_Q(k)\big), \qquad P = Ke^{-rT}F_Q(k) - S_0e^{-qT}F_S(k). \quad \text{(A.5)}$$

A.3.3 Computing F_S by a complex shift

Under $\mathbb{Q}^S$, the characteristic function has the Esscher/numéraire shift

$$\varphi_T^S(u) \equiv \mathbb{E}^{\mathbb{Q}^S}[e^{\mathrm{i}uX_T}] = \frac{\mathbb{E}^{\mathbb{Q}}[e^{(1+\mathrm{i}u)X_T}]}{\mathbb{E}^{\mathbb{Q}}[e^{X_T}]} = \frac{\varphi_T(u-\mathrm{i})}{\varphi_T(-\mathrm{i})}.$$

This identity is standard for numéraire changes (Geman et al., 1995) and exponential tilting of Lévy processes (Cont and Tankov, 2004). F_S follows from (A.3) with $\varphi = \varphi_T^S$ and enters (A.5).

A.4 Barrier (path-dependent) cash-or-nothing digitals

We now derive the procedures for pricing path-dependent digitals based on whether the process hits a barrier during $[0, T]$.

A.4.1 Stopping times and payoffs

Let $H > S_0$ and $L < S_0$ be upper and lower barriers, with first hitting times τ_H and τ_L and first exit time $\tau_{L,H} = \tau_L \wedge \tau_H$:

$$\tau_H = \inf\{t \in [0,T] : S_t \geqslant H\},$$
$$\tau_L = \inf\{t \in [0,T] : S_t \leqslant L\},$$
$$\tau_{L,H} = \inf\{t \in [0,T] : S_t \notin (L,H)\}.$$

The corresponding cash-or-nothing digitals are

$$\begin{aligned} \text{Up-hit digital:} \quad & D^{\uparrow}_{\text{hit}}(H,T) = e^{-rT}\mathbb{Q}(\tau_H \leqslant T), \\ \text{Down-hit digital:} \quad & D^{\downarrow}_{\text{hit}}(L,T) = e^{-rT}\mathbb{Q}(\tau_L \leqslant T), \\ \text{Double no-touch (DNT):} \quad & D_{\text{dnt}}(L,H,T) = e^{-rT}\mathbb{Q}(\tau_{L,H} > T). \end{aligned} \tag{A.6}$$

In log space, $b = \log(H/S_0) > 0$ and $\ell = \log(L/S_0) < 0$, so these events concern the running supremum and infimum of X.

A.4.2 Why Laplace transforms in maturity are natural

Barrier pricing requires first-passage and two-sided survival probabilities of a Lévy process. A standard approach Laplace-transforms maturity and solves a boundary-value problem for the generator's resolvent (Cont and Tankov, 2004; Kyprianou, 2014). In the Kou model, the double-exponential jump kernel reduces the resulting integro-differential equations to finite-dimensional linear algebra and closed-form Laplace transforms; see Kou and Wang (2003) for first-passage times and Kou and Wang (2004) for option-pricing applications.

A.4.3 Laplace exponent and root structure

For real arguments, define the cumulant exponent

$$\kappa(\beta) \equiv \log \mathbb{E}^{\mathbb{Q}}[e^{\beta X_1}] = \mu\beta + \tfrac{1}{2}\sigma^2\beta^2 + \lambda\big(M_Y(\beta) - 1\big), \qquad \beta \in (-\eta_2, \eta_1).$$

For $s > 0$, consider the characteristic equation $\kappa(\beta) = s$. Because M_Y is rational in β and $\frac{1}{2}\sigma^2\beta^2$ is convex, $\kappa(\beta)$ has four real roots: $0 < \beta_2(s) < \beta_1(s)$ and $-\alpha_1(s) < -\alpha_2(s) < 0$, with

$\alpha_1(s) > \alpha_2(s) > 0$ (Kou and Wang, 2003, 2004). Exponential ansätze $e^{\beta x}$ solve the homogeneous Laplace-domain equations, with coefficients fixed by the boundary conditions.

A.4.4 Single-barrier survival transforms (structure used in the code)

The upper-barrier survival probability in log space is

$$q^{\uparrow}(T;b) \equiv \mathbb{Q}\Big(\sup_{0\leqslant t\leqslant T} X_t < b\Big) = \mathbb{Q}(\tau_H > T).$$

Its Laplace transform in T is

$$\widehat{q}^{\uparrow}(s;b) = \int_0^{\infty} e^{-sT} q^{\uparrow}(T;b)\,\mathrm{d}T, \qquad s > 0.$$

For the double-exponential jump diffusion, Kou and Wang (2003) obtain

$$\widehat{q}^{\uparrow}(s;b) = \frac{1}{s}\left(A_1(s)e^{-\beta_1(s)b} + A_2(s)e^{-\beta_2(s)b}\right), \tag{A.7}$$

where $\beta_1(s), \beta_2(s)$ are the positive roots of $\kappa(\beta)$, and $(A_1(s), A_2(s))$ solve the 2×2 system implied by the barrier boundary and jump-kernel matching conditions in (A.1). Its coefficients are explicit rational functions of the model parameters and roots; see (Kou and Wang, 2003, Theorem 1 and subsequent corollaries).

The hitting probability in (A.6) follows from

$$\int_0^{\infty} e^{-sT}\mathbb{Q}(\tau_H \leqslant T)\,\mathrm{d}T = \frac{1}{s} - \widehat{q}^{\uparrow}(s;b),$$

so (A.7) immediately gives the Laplace transform of the up-hit probability.

The lower-barrier survival probability $q^{\downarrow}(T;\ell) = \mathbb{Q}(\inf_{0\leqslant t\leqslant T} X_t > \ell)$ analogously satisfies

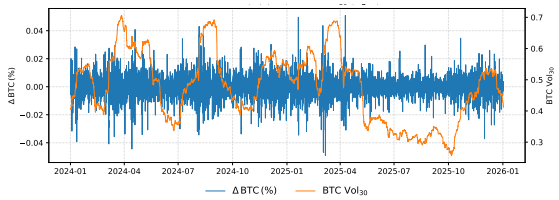
$$\widehat{q}^{\downarrow}(s;\ell) = \frac{1}{s}\left(\widetilde{A}_1(s)e^{\alpha_1(s)\ell} + \widetilde{A}_2(s)e^{\alpha_2(s)\ell}\right),$$

where $\alpha_1(s), \alpha_2(s)$ correspond to the negative roots of $\kappa(\beta)$; see Kou and Wang (2003).

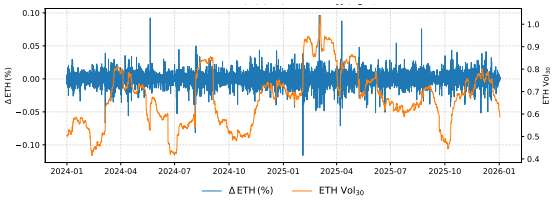
B Additional Tables and Figures

B.1 Returns, Volatility, L2 Risk, and Diff Underlying

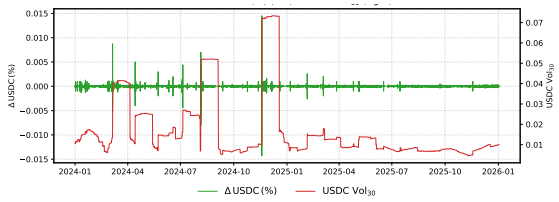
(a) Δ BTC (%) and BTC Vol_{30}



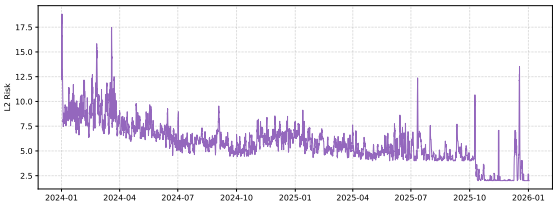
(b) Δ ETH (%) and ETH Vol_{30}



(c) Δ USDC (%) and USDC Vol_{30}



(d) L2 Risk



(e) Diff Underlying

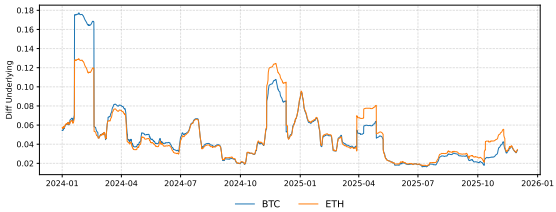
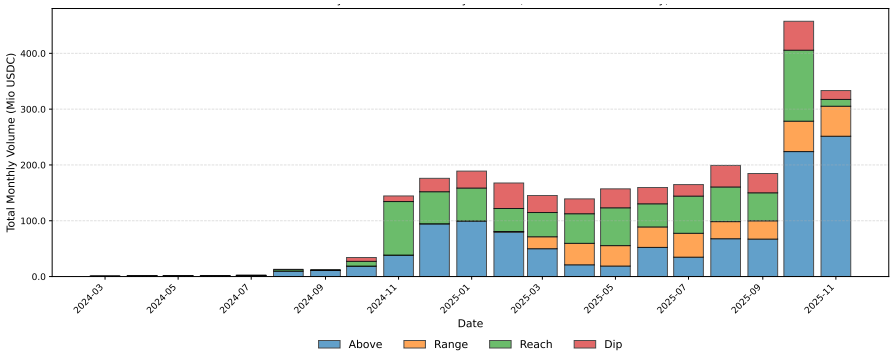


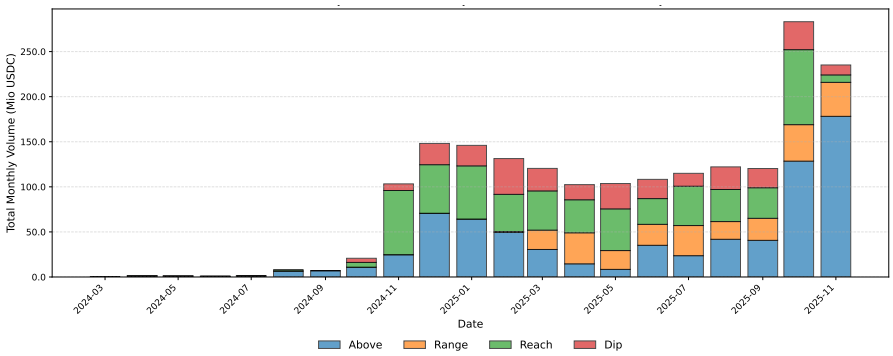
Figure B.1: Returns, Volatility, L2 Risk, and Diff Underlying. The figure displays the percentage change (left-axis) and the 30-day rolling volatility (right axis) of BTC (Panel (a)), ETH (Panel (b)) and USDC (Panel (c)). Panel (d) plots the L2 Risk as described in Section 3.4. Panel (e) plots Diff Underlying, the 24-hour rolling Binance–Deribit relative deviation, for BTC and ETH. The BTC and USDC data are obtained from CryptoCompare, while the Polygon block information is obtained from The Graph. Diff Underlying is constructed from Binance and Deribit underlying prices. The data is resampled to an hourly frequency and ranges from January 2024 to November 2025.

B.2 Polymarket Transaction Volume

(a) BTC and ETH



(b) BTC



(c) ETH

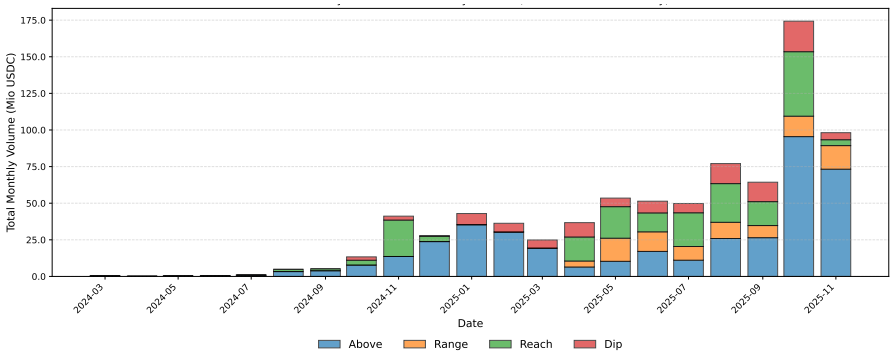


Figure B.2: Total Monthly Transaction Volume by Betting Type. The figure shows total monthly transaction volume (in million USDC), disaggregated by betting type. Trading volume is measured in notional USD terms as the total number of contracts traded; since each contract pays out 1 USD, this corresponds directly to the total notional amount exchanged. The sample covers the period from March 2024 to November 2025. Panel (a) shows total volume (BTC and ETH), while Panels (b) and (c) show BTC and ETH volume separately.

C Calibration Details

C.1 Deribit Calibration – Kou Parameters

C.1.1 BTC

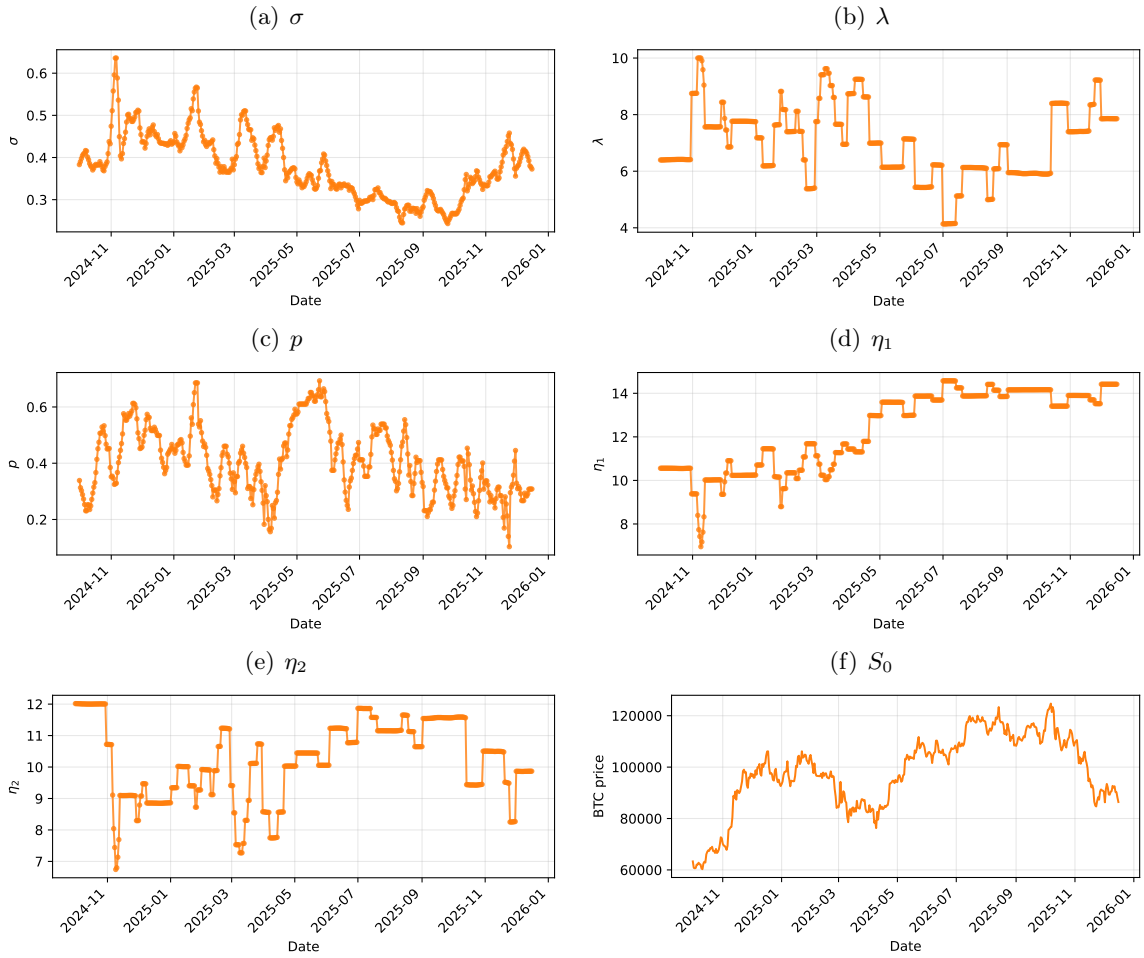


Figure C.1: BTC Kou Parameters – Deribit Calibration. The figure displays the time series of the estimated Kou (2002) parameters obtained from a calibration to the 7-day Deribit surfaces. Panels (a) through (e) report the estimates of σ, λ, p, η_1, and η_2, respectively. Panel (f) shows the BTC price.

C.1.2 ETH

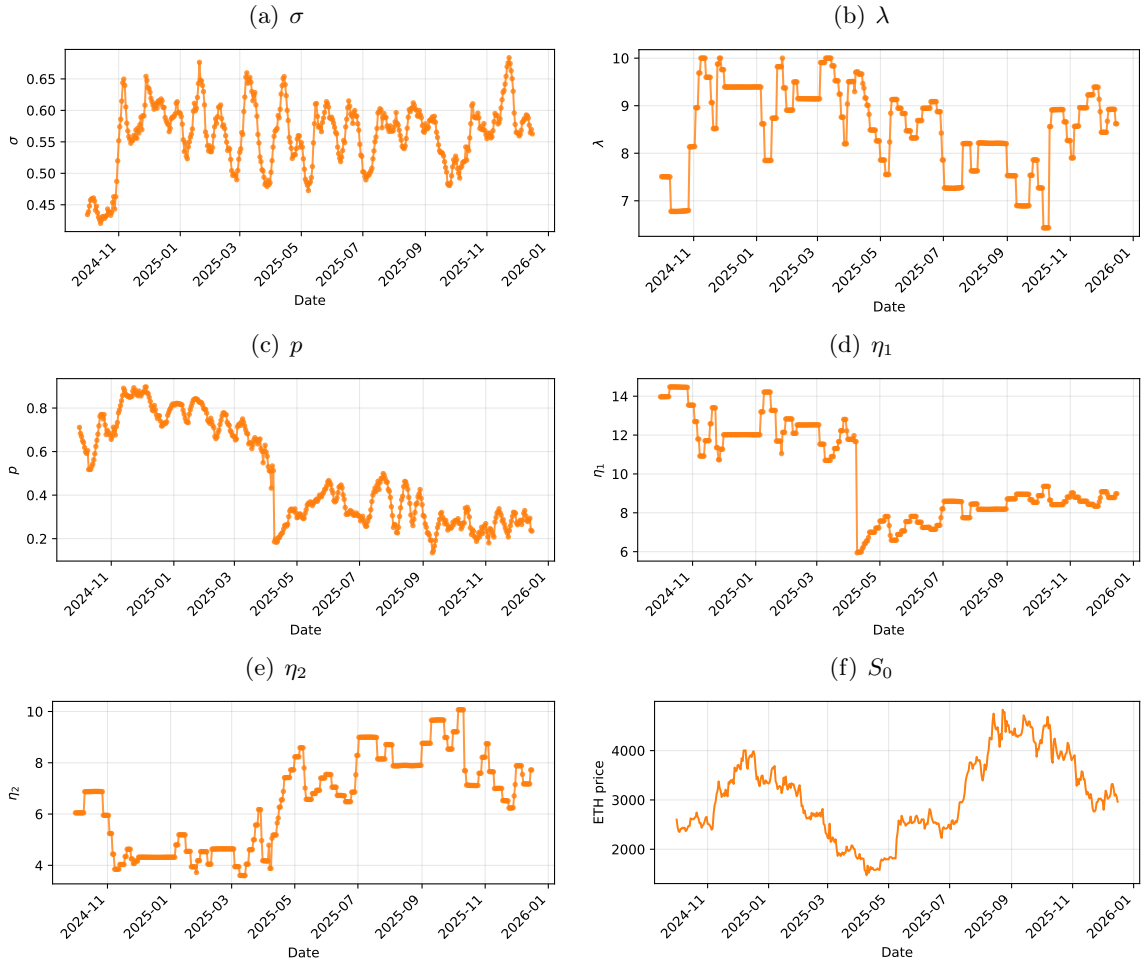
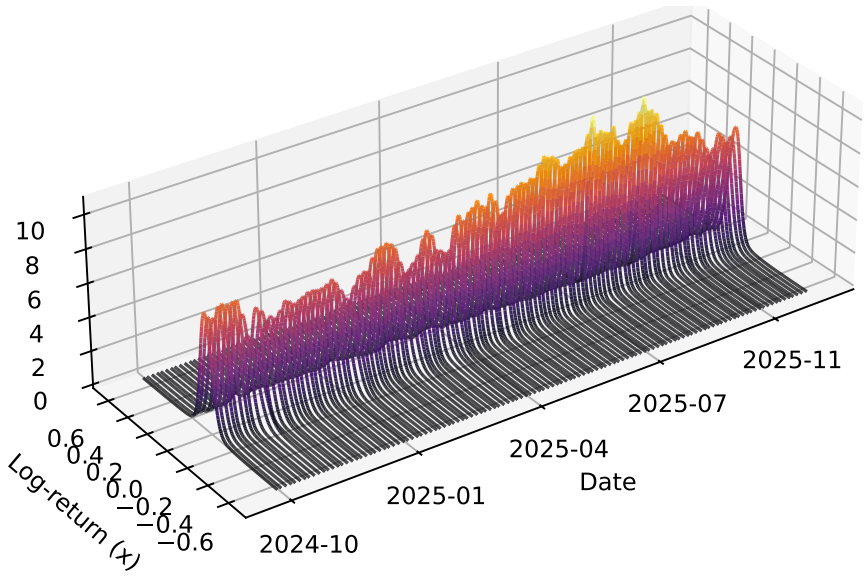


Figure C.2: ETH Kou Parameters – Deribit Calibration. The figure displays the time series of the estimated Kou (2002) parameters obtained from a calibration to the 7-day Deribit surfaces. Panels (a) through (e) report the estimates of σ, λ, p, η_1, and η_2, respectively. Panel (f) shows the ETH price.

C.1.3 RNDs Over Time

(a) BTC



(b) ETH

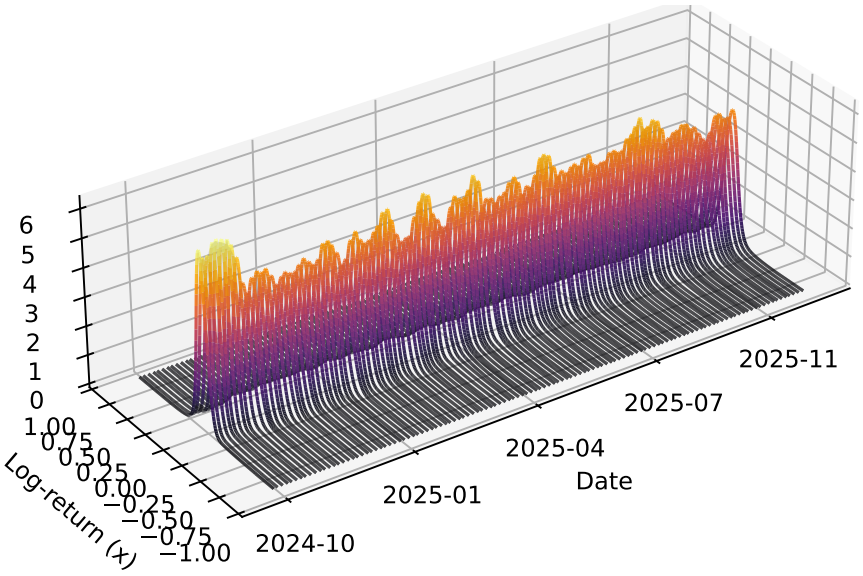


Figure C.3: Kou RND Over Time – Deribit Calibration. This figure presents the time-series evolution of the RND for the 7-day maturity, calibrated to Deribit surfaces. Panels (a) and (b) report the RND for BTC and ETH, respectively.

C.2 Polymarket Calibration – Kou Parameters

C.2.1 BTC

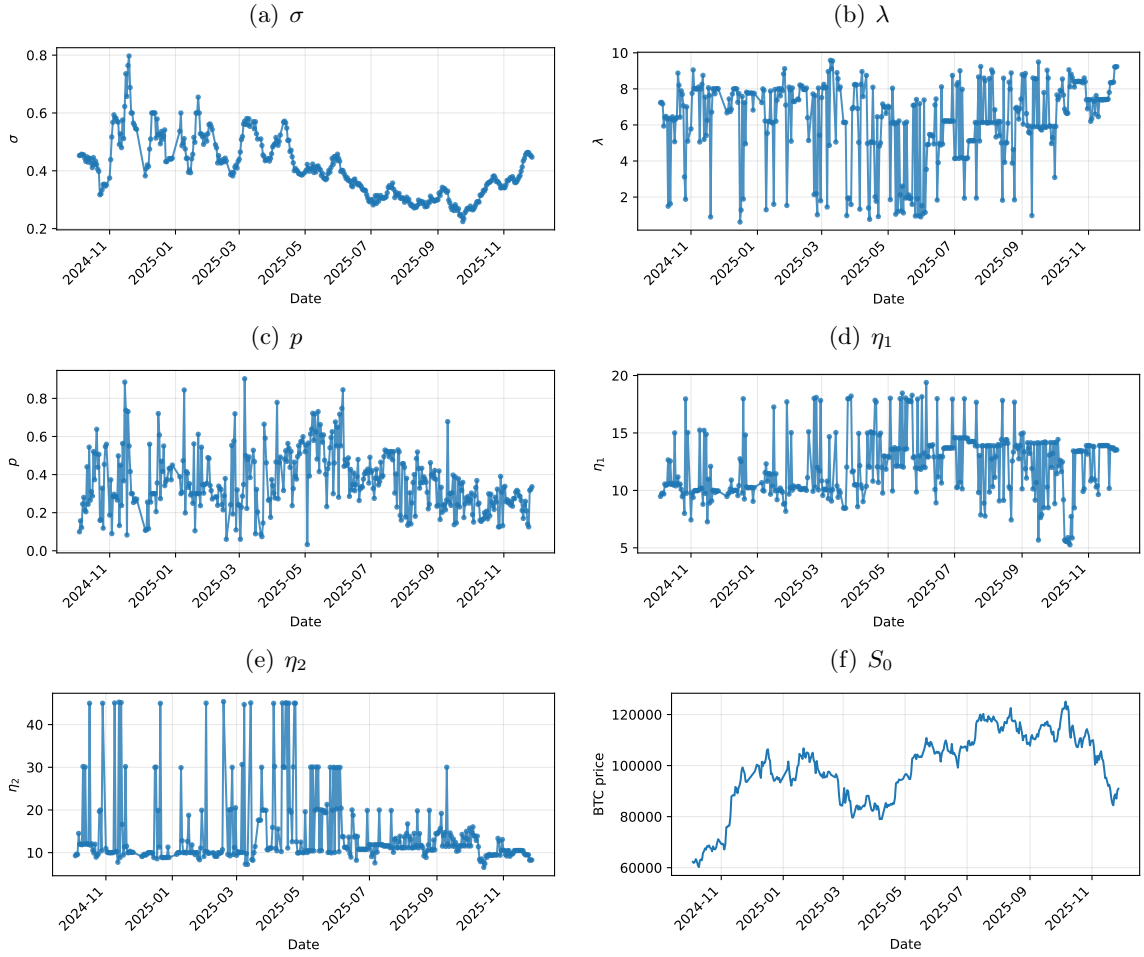


Figure C.4: BTC Kou Parameters – Polymarket Calibration. The figure displays the time series of the estimated Kou (2002) parameters obtained from a calibration to Polymarket digital contracts over a 7-day rolling window. Panels (a) through (e) report the estimates of σ, λ, p, η_1, and η_2, respectively. Panel (f) shows the BTC price.

C.2.2 ETH

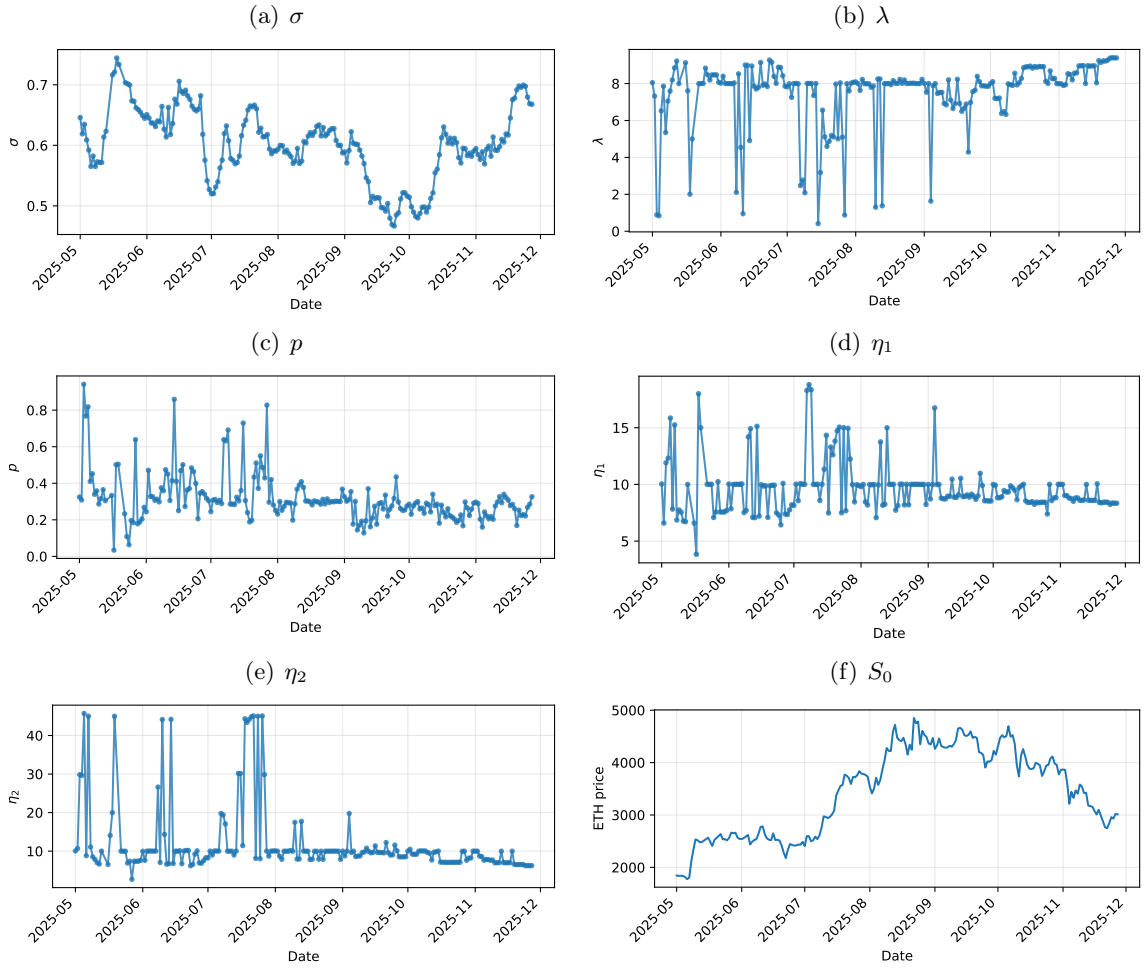
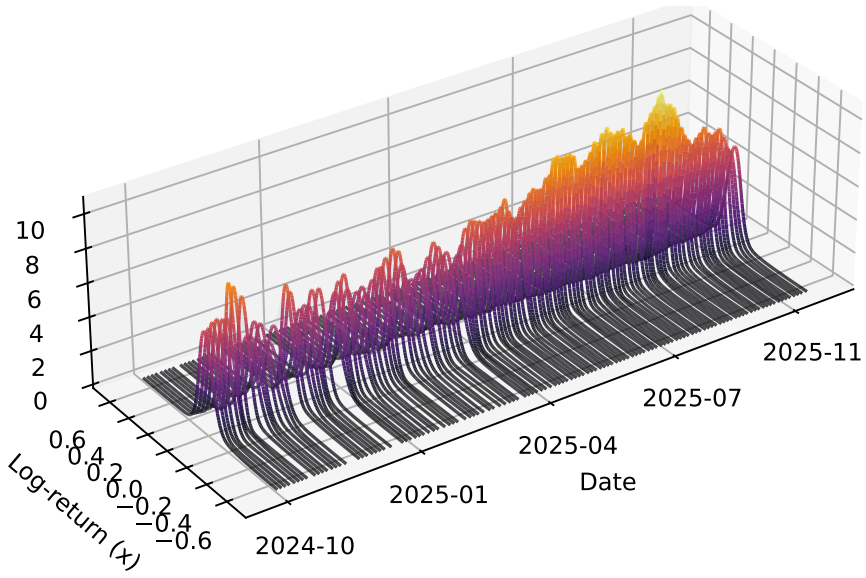


Figure C.5: ETH Kou Parameters – Polymarket Calibration. The figure displays the time series of the estimated Kou (2002) parameters obtained from a calibration to Polymarket digital contracts over a 7-day rolling window. Panels (a) through (e) report the estimates of σ, λ, p, η_1, and η_2, respectively. Panel (f) shows the ETH price.

C.2.3 RNDs Over Time

(a) BTC



(b) ETH

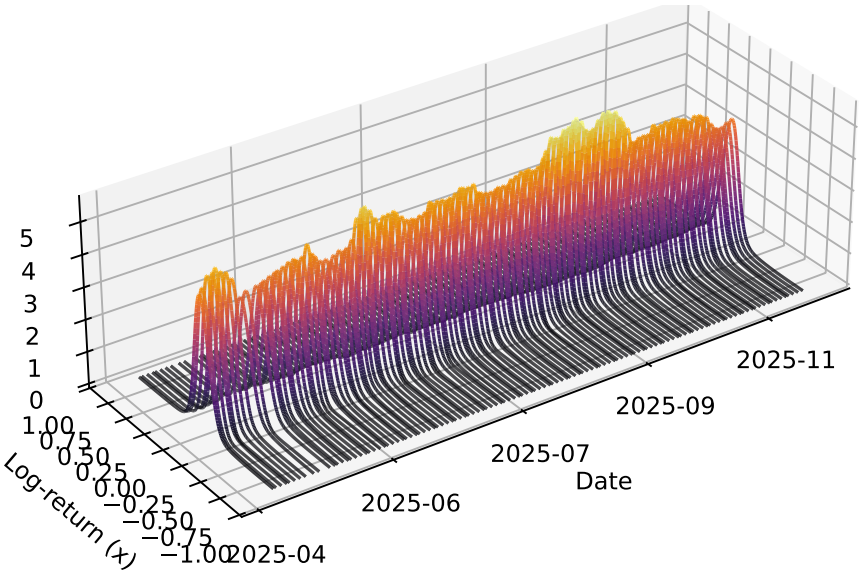


Figure C.6: Kou RND Over Time – Polymarket Calibration. This figure presents the time-series evolution of the Polymarket RND for the 7-day maturity, calibrated to Polymarket digital contracts. Panels (a) and (b) report the RND for BTC and ETH, respectively.

D Kou Benchmark Validation

In this section, we compare daily seven-day Kou densities with a simple model-free Breeden and Litzenberger (1978) benchmark recovered from the same Deribit implied volatility surfaces. The model-free construction recovers implied strikes, fits a quadratic volatility smile in strike space, differentiates the fitted call-price curve twice, and normalizes the resulting density over 50 observed-support strike points. The density is then compared with the Kou density on the exact same strike grid.

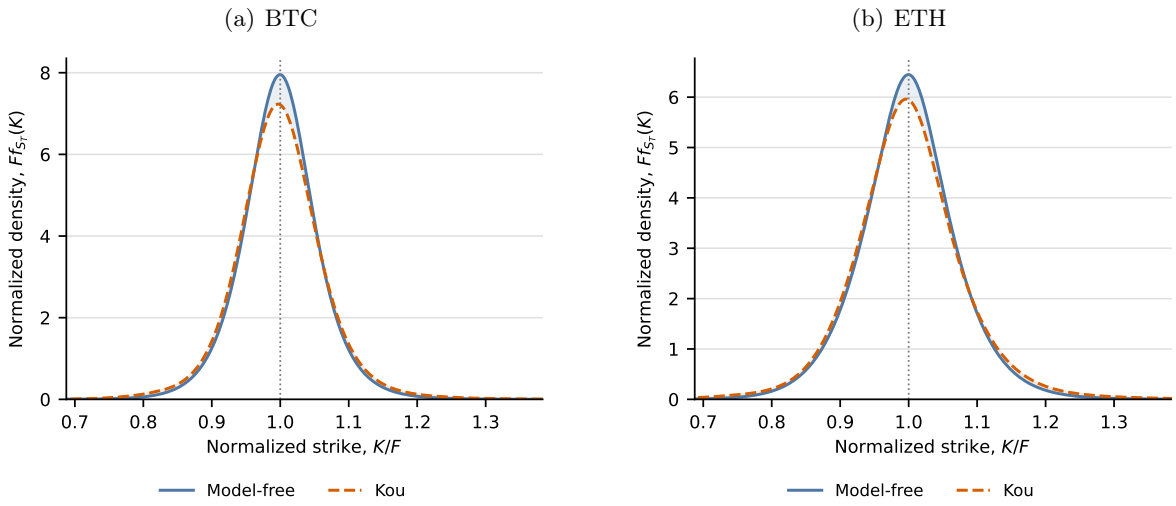


Figure D.1: Normalized Seven-Day Model-Free and Kou Risk-Neutral Densities. Panels (a) and (b) show averages for BTC and ETH, respectively. The horizontal axis is normalized strike, K/F, where F is the maturity-matched forward. The vertical axis is the corresponding normalized density, $Ff_{S_T}(K)$, so each curve integrates to one with respect to K/F. Both densities retain their finite observed strike support; no tail extrapolation is added.

Table D.1 summarizes the comparison across 1,074 matched dates for each asset. The average distributional distances are modest, but the location of the remaining mass differs economically. Kou prices an Above claim at the forward about one percentage point lower on average while assigning 1.30–1.37 percentage points more probability below $0.9F$ and 1.30–1.72 percentage points more probability above $1.1F$.

Figure D.2 illustrates the same tail difference for two fixed contracts: A BTC Above \$98,000 threshold and an ETH Above \$4,400 threshold through seven dates. At the highlighted far-OTM dates, the BTC probability is 2.79% under BL and 6.19% under Kou, while the ETH probability is 7.97% and 13.08%.

Table D.1 summarizes the comparison across 1,074 matched dates for each asset. The average distributional distances are modest, but the location of the remaining mass differs economically.

Asset	Dates	KS	TV	ATM gap	Left-tail gap	Right-tail gap
BTC	1,074	0.036	0.064	-1.08	+1.30	+1.30
ETH	1,074	0.032	0.059	-0.97	+1.37	+1.72

Table D.1: Seven-Day Kou and Model-Free RND Distributional and Tail Comparison. This table compares daily seven-day Kou and model-free densities on the same 50-point finite observed strike support. Dates is the number of matched daily BTC or ETH density pairs. KS is the Kolmogorov–Smirnov distance and TV is total variation. ATM gap is the Kou-minus-model-free difference for an Above claim struck at the maturity-matched forward, $K = F$. Left-tail and right-tail gaps are the corresponding differences in $\Pr(S_T < 0.9F)$ and $\Pr(S_T > 1.1F)$. Probability gaps are in percentage points and are averaged equally across dates.

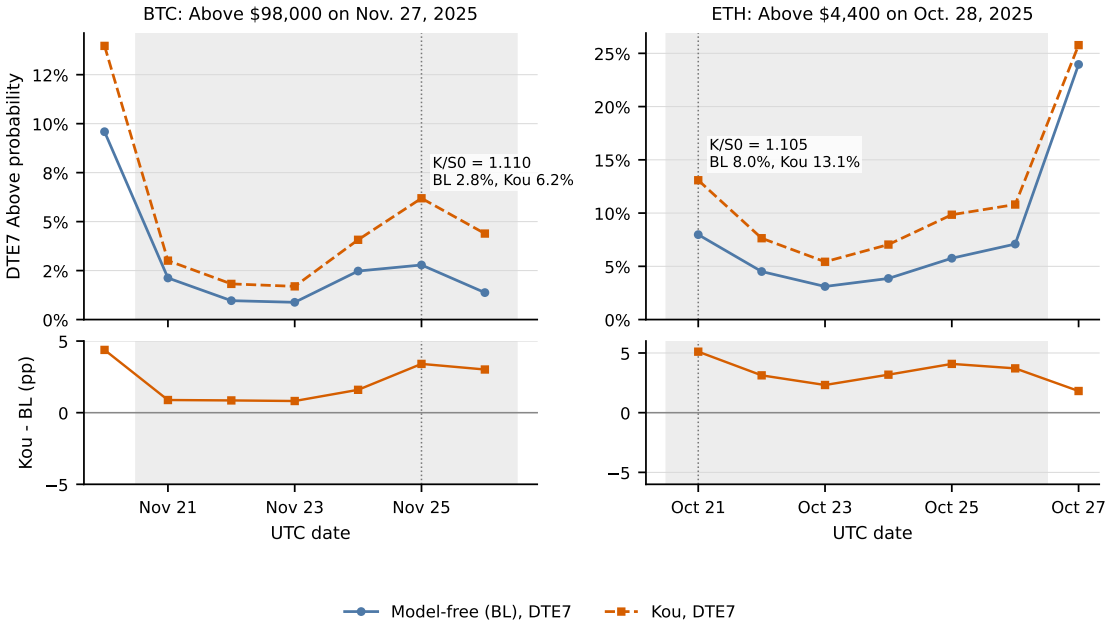


Figure D.2: Illustrative Fixed-Strike Above Prices from the Daily RNDs. Each panel prices one fixed contract definition from the daily seven-day model-free and Kou distributions. Gray shading marks dates on which $K/S_0 \geqslant 1.10$; dotted lines mark the largest positive Kou-minus-model-free gap within those dates. Resolution-date RNDs are excluded.

Kou prices an Above claim at the forward about one percentage point lower on average while assigning 1.30–1.37 percentage points more probability below $0.9F$ and 1.30–1.72 percentage points more probability above $1.1F$.

E Regressions with Contract Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		-0.054***							
		(0.019)							
USDC Vol_{30}		-0.028***							
		(0.005)							
VIX			0.203***				0.172***	0.183***	0.278***
			(0.024)				(0.029)	(0.029)	(0.034)
L2 Risk				0.000			0.006**		0.008
				(0.003)			(0.003)		(0.006)
Diff Underlying					-0.002			-0.028***	-0.142***
					(0.009)			(0.010)	(0.027)
NBP									0.005*
									(0.003)
Fear & Greed						-0.053***	-0.037***	-0.039***	-0.028**
						(0.008)	(0.010)	(0.009)	(0.014)
Announcement						0.052***	0.047***	0.045***	0.060***
						(0.009)	(0.009)	(0.009)	(0.013)
Δ (%)	0.001	0.001	0.000	0.001	0.001	0.001	0.000	0.000	-0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Δ USDC (%)	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Weekend						0.002	0.005	0.006	0.026***
						(0.006)	(0.006)	(0.006)	(0.009)
Moneyness	-0.125***	-0.129***	-0.095***	-0.125***	-0.125***	-0.097***	-0.080***	-0.079***	-0.221***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.011)
Maturity	0.049***	0.048***	0.054***	0.049***	0.049***	0.054***	0.056***	0.057***	0.056***
	(0.003)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.005)
ETH	–	–	–	–	–	–	–	–	–
Within R^2	0.006	0.007	0.007	0.006	0.006	0.007	0.008	0.008	0.025
Contract FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	363,280	363,280	363,243	363,280	363,280	363,280	363,243	363,243	100,241
Contracts	2,274	2,274	2,274	2,274	2,274	2,274	2,274	2,274	2,127
Time clusters	15,298	15,298	15,294	15,298	15,298	15,298	15,294	15,294	14,251

Table E.1: Panel Regression with Contract Fixed Effects – Above Bets – CLOB–OIP. This table reports contract-fixed-effect estimates for the same nine specifications and column-specific complete-case samples as the pooled Above regressions in Table 6.2. The dependent variable is the CLOB–OIP price wedge. The regressors include Moneyness, Maturity, volatility, USDC volatility, L2 Risk, Diff Underlying, NBP, Fear & Greed, Announcement, underlying and USDC returns, and the Weekend indicator as shown across columns. Contract and calendar-month fixed effects are included. The contract-invariant ETH indicator is absorbed and therefore not estimated. Standard errors are clustered by time and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All continuous variables are standardized within each specification to have a mean of 0 and a variance of 1; indicator variables remain in levels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		-0.199***							
		(0.067)							
USDC Vol_{30}		0.050**							
		(0.023)							
VIX			-0.151***				-0.188***	-0.151***	-0.260***
			(0.045)				(0.041)	(0.040)	(0.037)
L2 Risk				-0.007			-0.005		-0.000
				(0.010)			(0.010)		(0.007)
Diff Underlying					-0.105***			-0.116***	-0.064***
					(0.021)			(0.020)	(0.025)
NBP									-0.009**
									(0.004)
Fear & Greed						-0.007	-0.031*	-0.042***	-0.012
						(0.017)	(0.016)	(0.016)	(0.011)
Announcement						-0.003	0.004	-0.002	-0.057***
						(0.027)	(0.027)	(0.027)	(0.013)
Δ (%)	0.002	0.002	0.003	0.002	0.003	0.001	0.002	0.002	-0.003
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)	(0.003)
Δ USDC (%)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.003)
Weekend						0.050***	0.047***	0.053***	0.027***
						(0.017)	(0.017)	(0.017)	(0.008)
Moneyness	-0.028***	-0.031***	-0.037***	-0.028***	-0.033***	-0.026***	-0.035***	-0.037***	-0.040***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.009)
Maturity	0.069***	0.067***	0.069***	0.070***	0.071***	0.067***	0.068***	0.069***	-0.029***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.005)
ETH	–	–	–	–	–	–	–	–	–
Within R^2	0.005	0.006	0.006	0.005	0.006	0.006	0.006	0.007	0.008
Contract FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	307,573	307,573	307,536	307,573	307,573	307,573	307,536	307,536	63,864
Contracts	1,808	1,808	1,808	1,808	1,808	1,808	1,808	1,808	1,705
Time clusters	6,939	6,939	6,937	6,939	6,939	6,939	6,937	6,937	6,833

Table E.2: Panel Regression with Contract Fixed Effects – Range Bets – CLOB–OIP. This table reports contract-fixed-effect estimates for the same nine specifications and column-specific complete-case samples as the pooled Range regressions in Table 6.3. The dependent variable is the CLOB–OIP price wedge. The regressors include Moneyness, Maturity, volatility, USDC volatility, L2 Risk, Diff Underlying, NBP, Fear & Greed, Announcement, underlying and USDC returns, and the Weekend indicator as shown across columns. Contract and calendar-month fixed effects are included. The contract-invariant ETH indicator is absorbed and therefore not estimated. Standard errors are clustered by time and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All continuous variables are standardized within each specification to have a mean of 0 and a variance of 1; indicator variables remain in levels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		-0.104***							
		(0.012)							
USDC Vol_{30}		-0.088***							
		(0.009)							
VIX			0.342***				0.384***	0.378***	0.532***
			(0.017)				(0.018)	(0.018)	(0.018)
L2 Risk				-0.031***			-0.024***		-0.014**
				(0.006)			(0.006)		(0.006)
Diff Underlying					0.084***			0.074***	0.118***
					(0.012)			(0.012)	(0.016)
NBP									0.011***
									(0.002)
Fear & Greed						0.033***	0.097***	0.095***	0.124***
						(0.010)	(0.010)	(0.010)	(0.010)
Announcement						0.005	-0.008	-0.001	-0.019
						(0.015)	(0.015)	(0.015)	(0.013)
Δ (%)	0.006	0.007*	0.004	0.006	0.006	0.006	0.004	0.004	0.000
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)
Δ USDC (%)	0.000	0.001	0.001	0.000	0.001	0.001	0.001	0.001	0.002
	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)
Weekend						-0.143***	-0.141***	-0.143***	-0.110***
						(0.008)	(0.008)	(0.008)	(0.009)
Moneyness	0.022	0.038**	0.061***	0.031*	0.043**	-0.031	-0.062***	-0.049**	-0.074***
	(0.018)	(0.018)	(0.017)	(0.018)	(0.018)	(0.022)	(0.022)	(0.022)	(0.021)
Maturity	0.387***	0.398***	0.348***	0.395***	0.375***	0.379***	0.335***	0.318***	0.330***
	(0.009)	(0.010)	(0.010)	(0.010)	(0.009)	(0.009)	(0.010)	(0.010)	(0.010)
ETH	–	–	–	–	–	–	–	–	–
Within R^2	0.164	0.169	0.172	0.164	0.165	0.170	0.180	0.181	0.223
Contract FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	134,736	134,736	134,711	134,736	134,736	134,736	134,711	134,711	89,419
Contracts	407	407	407	407	407	407	407	407	406
Time clusters	13,454	13,454	13,451	13,454	13,454	13,454	13,451	13,451	12,071

Table E.3: Panel Regression with Contract Fixed Effects – Reach Bets – CLOB–OIP. This table reports contract-fixed-effect estimates for the same nine specifications and column-specific complete-case samples as the pooled Reach regressions in Table 6.4. The dependent variable is the CLOB–OIP price wedge. The regressors include Moneyness, Maturity, volatility, USDC volatility, L2 Risk, Diff Underlying, NBP, Fear & Greed, Announcement, underlying and USDC returns, and the Weekend indicator as shown across columns. Contract and calendar-month fixed effects are included. The contract-invariant ETH indicator is absorbed and therefore not estimated. Standard errors are clustered by time and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All continuous variables are standardized within each specification to have a mean of 0 and a variance of 1; indicator variables remain in levels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		-0.021							
		(0.014)							
USDC Vol_{30}		-0.004							
		(0.007)							
VIX			0.678***				0.666***	0.628***	0.855***
			(0.037)				(0.045)	(0.045)	(0.081)
L2 Risk				-0.020**			-0.012		0.025***
				(0.008)			(0.008)		(0.008)
Diff Underlying					0.209***			0.177***	0.241***
					(0.009)			(0.010)	(0.014)
NBP									0.006**
									(0.003)
Fear & Greed						-0.099***	0.011	0.017	-0.029
						(0.010)	(0.014)	(0.013)	(0.023)
Announcement						0.053***	0.041**	0.054***	0.113***
						(0.019)	(0.019)	(0.018)	(0.018)
Δ (%)	0.005	0.005	0.003	0.005	0.004	0.005	0.003	0.003	0.011**
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.004)	(0.005)
Δ USDC (%)	-0.005	-0.005	-0.005	-0.005	-0.005	-0.004	-0.005	-0.005	-0.003
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Weekend						-0.170***	-0.153***	-0.155***	-0.149***
						(0.009)	(0.009)	(0.009)	(0.011)
Moneyness	0.151***	0.149***	0.169***	0.145***	0.147***	0.042*	0.171***	0.176***	0.241***
	(0.021)	(0.021)	(0.021)	(0.021)	(0.020)	(0.023)	(0.028)	(0.028)	(0.043)
Maturity	0.002	0.003	-0.010	0.006	-0.007	0.018**	-0.009	-0.019**	-0.064***
	(0.007)	(0.008)	(0.007)	(0.007)	(0.008)	(0.007)	(0.008)	(0.008)	(0.010)
ETH	–	–	–	–	–	–	–	–	–
Within R^2	0.013	0.013	0.036	0.013	0.022	0.023	0.043	0.049	0.088
Contract FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	121,861	121,861	121,826	121,861	121,861	121,861	121,826	121,826	71,470
Contracts	371	371	371	371	371	371	371	371	366
Time clusters	11,603	11,603	11,600	11,603	11,603	11,603	11,600	11,600	10,321

Table E.4: Panel Regression with Contract Fixed Effects – Dip Bets – CLOB–OIP. This table reports contract-fixed-effect estimates for the same nine specifications and column-specific complete-case samples as the pooled Dip regressions in Table 6.5. The dependent variable is the CLOB–OIP price wedge. The regressors include Moneyness, Maturity, volatility, USDC volatility, L2 Risk, Diff Underlying, NBP, Fear & Greed, Announcement, underlying and USDC returns, and the Weekend indicator as shown across columns. Contract and calendar-month fixed effects are included. The contract-invariant ETH indicator is absorbed and therefore not estimated. Standard errors are clustered by time and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All continuous variables are standardized within each specification to have a mean of 0 and a variance of 1; indicator variables remain in levels.

Internet Appendix

I Model Details - Equilibrium Riccati System

We begin by restating the arbitrageur's value function and HJB equation. Given a pricing wedge Δ_s, the value function is

$$V(y_t, Z_t) = \max_{\{v_s\}_{s\geqslant t}} \mathbb{E}_t \left[\int_t^\infty e^{-r(s-t)} \left(-\Delta_s v_s - \frac{\Lambda_s}{2} v_s^2 - \frac{\gamma_s}{2} y_s^2 \right) ds \right], \qquad dy_s = v_s ds.$$

The corresponding HJB equation is

$$rV = \max_v \left\{ -\Delta v - \frac{\Lambda}{2} v^2 - \frac{\gamma}{2} y^2 + V_y v \right\} + \mathcal{L}_Z V. \tag{I.1}$$

Market clearing dictates that the arbitrageur absorbs the net demand, giving the pricing relation:

$$\Delta = \frac{y + Z}{\eta}. \tag{I.2}$$

We solve (I.1) by conjecturing a quadratic value function:

$$V(y, Z) = -\frac{1}{2} A y^2 - B y Z - \frac{1}{2} C Z^2 + V_0. \tag{I.3}$$

The relevant partial derivatives are $V_y = -Ay - BZ$ and $V_Z = -By - CZ$. The coefficient C does not enter V_y, and therefore does not directly affect the trading rule or the equilibrium pricing wedge.

For the main economic logic to hold, let

$$G_{yZ} \equiv [\mathcal{L}_Z V]_{yZ}$$

denote the yZ coefficient generated by the demand-shock generator when it is applied to the quadratic value function in (I.3). We keep the law of motion of Z_t fully generic and impose only the following restriction on this coefficient.

Assumption I.1 (Demand generator contribution)**.** *The generator coefficient satisfies*

$$G_{yZ} < \frac{r}{\eta}, \qquad \frac{dG_{yZ}}{d\theta} \geqslant 0 \quad \textit{for} \quad \theta \in \{\gamma, \Lambda\}.$$

Assumption I.1 is a restriction on the generator coefficient entering the Riccati system. The first inequality rules out generator effects so large that the cross term overturns the usual inventory-absorption logic. The second inequality is a non-offsetting condition: as frictions rise, the autonomous demand-shock dynamics do not move against the inventory-cost channel.

I.1 Derivation of the Pricing Wedge

Isolating the v-dependent terms in the HJB equation yields the optimal trading rate:

$$v^* = \frac{V_y - \Delta}{\Lambda}.$$

Substituting V_y and Δ, the optimal control is

$$v^* = -\frac{1}{\Lambda}\left[\left(A + \frac{1}{\eta}\right) y + \left(B + \frac{1}{\eta}\right) Z\right].$$

The stationary inventory target $y^*(Z)$ is defined as the point where the arbitrageur ceases trading ($v^* = 0$), which implies $V_y = \Delta$. Equating the marginal value of inventory to the pricing wedge yields

$$-Ay^*(Z) - BZ = \frac{1}{\eta}\left(y^*(Z) + Z\right).$$

Solving for $y^*(Z)$ gives

$$y^*(Z) = -\left(\frac{B + 1/\eta}{A + 1/\eta}\right) Z. \tag{I.4}$$

Substituting this equilibrium target inventory back into the market-clearing relation (I.2) yields the reduced-form equilibrium pricing wedge:

$$\Delta_t = \left[\frac{1}{\eta}\left(1 - \frac{B + 1/\eta}{A + 1/\eta}\right)\right] Z_t = \left[\frac{1}{\eta}\left(\frac{A - B}{A + 1/\eta}\right)\right] Z_t \equiv \Phi(\gamma, \Lambda) Z_t. \tag{I.5}$$

Equation (I.5) isolates the liquidity multiplier Φ, demonstrating how parameters affecting the model's frictions map into the equilibrium price impact of aggregate demand shocks.

I.2 Comparative Statics

The previous subsection showed that

$$y^*(Z) = -\left(\frac{B + 1/\eta}{A + 1/\eta}\right) Z, \qquad \Phi = \frac{1}{\eta}\left(\frac{A - B}{A + 1/\eta}\right).$$

Hence, the sign and comparative statics of the equilibrium wedge depend on the Riccati coefficients A and B. The coefficient C does not enter V_y, and therefore does not affect v^*, $y^*(Z)$, or Φ.

We first characterize A. Substituting v^* back into the HJB isolates the maximized trading term:

$$\frac{1}{2\Lambda}(V_y - \Delta)^2 = \frac{1}{2\Lambda}\left[\left(A + \frac{1}{\eta}\right) y + \left(B + \frac{1}{\eta}\right) Z\right]^2. \tag{I.6}$$

Matching the y^2 coefficients in (I.1) gives

$$-\frac{rA}{2} = \frac{1}{2\Lambda}\left(A + \frac{1}{\eta}\right)^2 - \frac{\gamma}{2}.$$

Equivalently, A solves $f(A, \gamma, \Lambda) = 0$, where

$$f(A, \gamma, \Lambda) \equiv A^2 + \left(r\Lambda + \frac{2}{\eta}\right) A + \left(\frac{1}{\eta^2} - \gamma\Lambda\right). \tag{I.7}$$

For the economically meaningful root, $f_A > 0$. The implicit-function theorem then gives

$$\frac{\partial A}{\partial \gamma} = \frac{\Lambda}{2A + r\Lambda + \frac{2}{\eta}} > 0. \tag{I.8}$$

Similarly, using (I.7) to rewrite the numerator,

$$\frac{\partial A}{\partial \Lambda} = \frac{\left(A + \frac{1}{\eta}\right)^2}{\Lambda\left(2A + r\Lambda + \frac{2}{\eta}\right)} > 0. \tag{I.9}$$

Thus, both inventory risk and execution costs increase the inventory-cost coefficient A.

We now characterize B. From (I.6), the maximized trading term contributes

$$\frac{1}{\Lambda}\left(A + \frac{1}{\eta}\right)\left(B + \frac{1}{\eta}\right)$$

to the yZ coefficient. The remaining contribution is the generator coefficient G_{yZ}. Matching the yZ coefficients in (I.1) gives

$$-rB = \frac{1}{\Lambda}\left(A + \frac{1}{\eta}\right)\left(B + \frac{1}{\eta}\right) + G_{yZ}. \tag{I.10}$$

This equation determines the cross coefficient B while keeping the law of motion of Z_t unres-

tricted. For compactness, define

$$\xi \equiv \frac{1}{\eta}, \qquad M_A \equiv A + \xi.$$

Then (I.10) can be written as

$$-rB = \frac{M_A}{\Lambda}(B + \xi) + G_{yZ}. \tag{I.11}$$

Solving (I.11) for $B + \xi$ yields

$$B + \xi = \frac{r\xi - G_{yZ}}{r + M_A/\Lambda}.$$

For the economically meaningful root $M_A > 0$, the first part of Assumption I.1 implies

$$B + \frac{1}{\eta} > 0. \tag{I.12}$$

Therefore, by (I.4), a positive demand imbalance requires the arbitrageur to take the opposite inventory position. Using the notation above, the pricing multiplier in (I.5) is

$$\Phi = \xi \frac{A - B}{M_A}. \tag{I.13}$$

Now, consider first inventory risk. Differentiating (I.11) with respect to γ gives

$$\left(r + \frac{M_A}{\Lambda}\right)\frac{dB}{d\gamma} = -\frac{B + \xi}{\Lambda}\frac{dA}{d\gamma} - \frac{dG_{yZ}}{d\gamma}.$$

By (I.8), $dA/d\gamma > 0$; by (I.12), $B + \xi > 0$; and by Assumption I.1, $dG_{yZ}/d\gamma \geqslant 0$. Hence $dB/d\gamma < 0$. Differentiating (I.13) gives

$$\frac{d\Phi}{d\gamma} = \xi \frac{B + \xi}{M_A^2}\frac{dA}{d\gamma} - \frac{\xi}{M_A}\frac{dB}{d\gamma} > 0. \tag{I.14}$$

Thus, higher inventory risk strictly increases the pricing multiplier and hence widens the equilibrium pricing wedge for a given demand imbalance.

Consider next execution costs. Differentiating (I.11) with respect to Λ gives

$$\left(r + \frac{M_A}{\Lambda}\right)\frac{dB}{d\Lambda} = -\frac{\Lambda\frac{dA}{d\Lambda} - M_A}{\Lambda^2}(B + \xi) - \frac{dG_{yZ}}{d\Lambda}.$$

Using (I.9), the terms involving $dA/d\Lambda$ cancel in the derivative of Φ. Therefore,

$$\frac{d\Phi}{d\Lambda} = \frac{\xi}{M_A\left(r + \frac{M_A}{\Lambda}\right)} \frac{dG_{yZ}}{d\Lambda} \geqslant 0. \tag{I.15}$$

The inequality follows from Assumption I.1; it is strict whenever $dG_{yZ}/d\Lambda > 0$. Hence, execution costs weakly increase the pricing multiplier, and strictly increase it when the generator coefficient responds strictly to execution frictions.

Combining (I.14) and (I.15), inventory risk strictly widens the equilibrium price gap, while execution costs weakly widen it under the regular generator-response condition.

II Additional Tables and Figures

II.1 Average Correlations per Betting Type

II.1.1 BTC

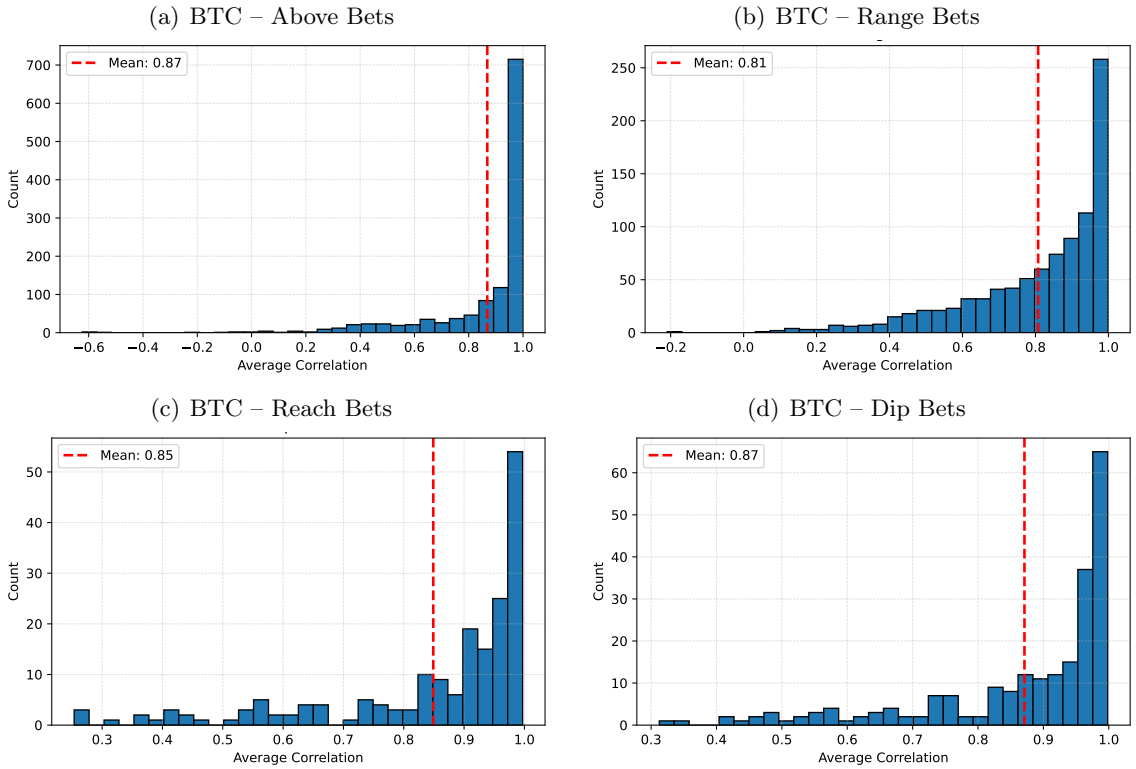


Figure II.1: Histogram – Average Correlations per Betting Type, BTC contracts. The figure displays the histogram of the average correlation of the CLOB and OIP for each of the four betting types (Above, Range, Reach, and Dip) shown in Panels (a) through (d), respectively. The different types of contracts are discussed in Section 2.2.

II.1.2 ETH

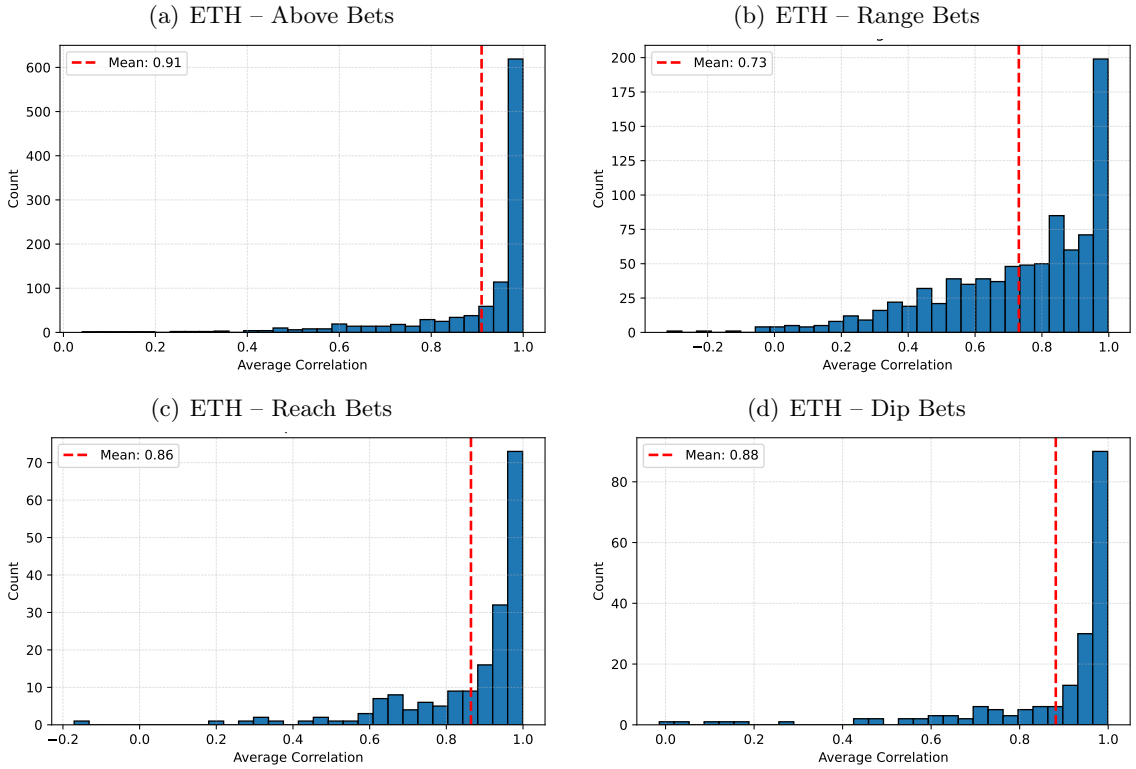


Figure II.2: Histogram – Average Correlations per Betting Type, ETH contracts. The figure displays the histogram of the average correlation of the CLOB and OIP for each of the four betting types (Above, Range, Reach, and Dip) shown in Panels (a) through (d), respectively. The different types of contracts are discussed in Section 2.2.

II.1.3 Partial Correlations per Betting Type

Figure II.3 reports the corresponding partial correlations between CLOB and OIP prices. For each contract, both price series are first residualized on moneyness and maturity, and the figure plots the correlation between the two residual series. The distributions remain concentrated in positive territory across the four betting types, indicating that the strong alignment between CLOB and OIP is not driven only by common variation in option-level state variables.

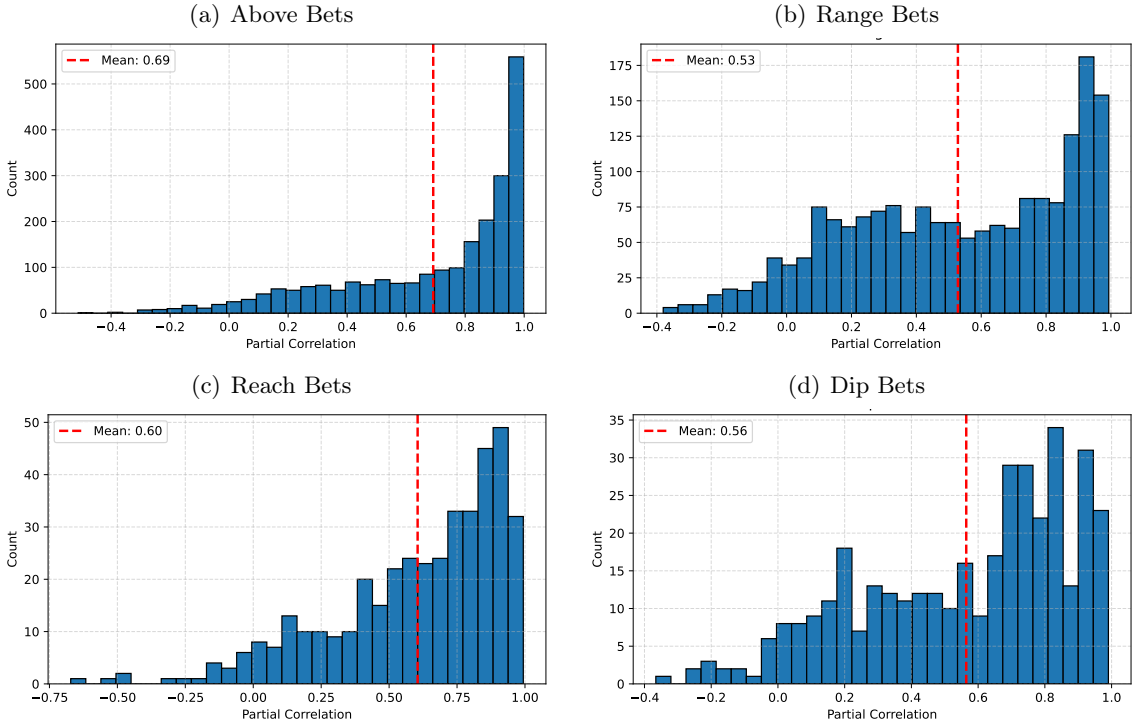


Figure II.3: Histogram – Partial Correlations per Betting Type. The figure displays the histogram of the partial correlation of the CLOB and OIP for each of the four betting types (Above, Range, Reach, and Dip) shown in Panels (a) through (d), respectively. For each contract, CLOB and OIP prices are residualized on moneyness and maturity before computing the correlation.

II.2 Difference per Betting Type and Time Segment

II.2.1 BTC

(a) BTC – Above Bets

Segment	Mean	p-value	Obs.
Beginning	0.0064	0.0000	27001
Middle	0.0047	0.0000	125240
End	0.0002	0.6597	29446

(b) BTC – Range Bets

Segment	Mean	p-value	Obs.
Beginning	0.0203	0.0000	22716
Middle	0.0036	0.0000	113273
End	0.0026	0.0000	22653

(c) BTC – Reach Bets

Segment	Mean	p-value	Obs.
Beginning	0.0411	0.0000	5252
Middle	0.0212	0.0000	72308
End	0.0008	0.0722	5379

(d) BTC – Dip Bets

Segment	Mean	p-value	Obs.
Beginning	0.0334	0.0000	4415
Middle	0.0059	0.0000	58051
End	0.0068	0.0000	4575

Table II.1: BTC Summary Statistics: Difference (CLOB - OIP) per Betting Type and Time Segment. Panels (a) to (d) display the average return (Mean), p-value, and number of observations (n obs) by time segment — beginning, middle, and end — for the four contract types: Above, Range, Reach, and Dip. See Section 2.2 for more details. The data range from March 2024 to November 2025.

II.2.2 ETH

(a) ETH – Above Bets

Segment	Mean	p-value	Obs.
Beginning	0.0085	0.0000	25956
Middle	0.0060	0.0000	129806
End	0.0013	0.0001	25913

(b) ETH – Range Bets

Segment	Mean	p-value	Obs.
Beginning	0.0272	0.0000	21373
Middle	0.0049	0.0000	106303
End	0.0146	0.0000	21303

(c) ETH – Reach Bets

Segment	Mean	p-value	Obs.
Beginning	0.0343	0.0000	4425
Middle	0.0128	0.0000	42875
End	-0.0007	0.2126	4534

(d) ETH – Dip Bets

Segment	Mean	p-value	Obs.
Beginning	0.0299	0.0000	4274
Middle	0.0039	0.0000	46120
End	0.0106	0.0000	4466

Table II.2: ETH Summary Statistics: Difference (CLOB - OIP) per Betting Type and Time Segment. Panels (a) to (d) display the average return (Mean), p-value, and number of observations (n obs) by time segment — beginning, middle, and end — for the four contract types: Above, Range, Reach, and Dip. See Section 2.2 for more details. The data range from January 2024 to November 2025.

II.3 Gas Fees

Bet Type	Transactions	Total Fees	Priority Fee (%)	Avg Fees
Above	1,682,365	30,862.20	28.6	0.0183
Range	814,199	12,852.14	32.6	0.0158
Reach	1,272,270	17,714.15	32.0	0.0139
Dip	618,160	10,244.38	26.8	0.0166
All	4,386,994	71,672.87	29.9	0.0163

Table II.3: Transaction Fees by Bet Type. The table reports transaction counts, total gas fees, priority fee shares, and average fees per transaction by contract type.

We analyze per-transaction gas costs for trades executed on Polygon, retrieved via Web3 from Infura endpoints, and the corresponding transaction receipts. From the receipt, we obtain the actual *gas used* and the *effective gas price*. The on-chain gas fees in USD for a transaction are calculated as

$$\text{gas fee} = \text{gas used} \times \text{effective gas price}/10^{18} \times P^{\text{MATIC/USD}},$$

where $P^{\text{MATIC/USD}}$ denotes the contemporaneous MATIC–USD exchange rate; MATIC is the native token of Polygon and is used to pay transaction fees.[14] Using the block-level base fee, we compute the priority fee as

$$\text{priority fee} = \text{gas used} \times (\text{effective gas price} - \text{base fee})/10^{18} \times P^{\text{MATIC/USD}},$$

which captures the portion of transaction costs paid directly to validators, while the base fee component is burned by the protocol. Average fees remain low across all contract types, with a mean fee per trade of \$0.0163; importantly, under Polymarket's gasless transaction model, these fees are borne by Polymarket rather than by traders.

[14] Under the EIP-1559 fee mechanism, the effective gas price consists of a base fee and a priority fee (validator tip).

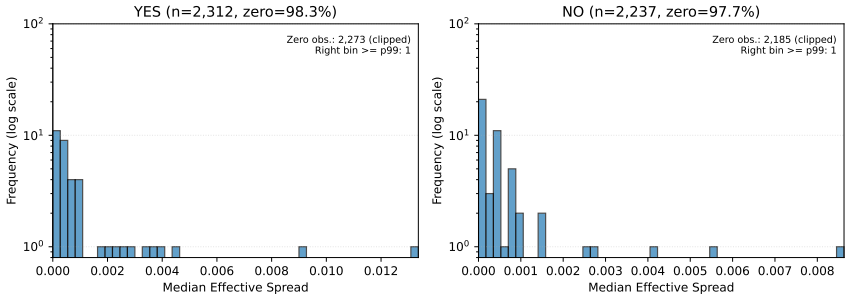
II.4 Spreads

For each condition ID, we load the taker-only trade data and construct bid–ask spreads at the 1-minute frequency separately for each outcome $o \in \{\text{YES}, \text{NO}\}$. Since the dataset contains taker executions only, a trade with side = SELL corresponds to a market sell order hitting the standing bid, while a trade with side = BUY corresponds to a market buy order hitting the standing ask. Therefore, within each 1-minute time interval t and outcome o, we proxy the bid quote as the maximum execution price among all sell trades, and we proxy the ask quote as the minimum execution price among all buy trades. We retain only those intervals for which both $\text{Bid}_{t,o}$ and $\text{Ask}_{t,o}$ are observed. Using these quantities, we compute

$$\begin{aligned} \text{Effective Spread}_{t,o} &= \text{Ask}_{t,o} - \text{Bid}_{t,o}, \\ \text{Mid}_{t,o} &= \frac{\text{Bid}_{t,o} + \text{Ask}_{t,o}}{2}, \\ \text{Relative Spread}_{t,o} &= \frac{\text{Effective Spread}_{t,o}}{\text{Mid}_{t,o}}. \end{aligned}$$

The results are displayed in Figure II.4 and Figure II.5, which show the empirical distributions of median bid–ask spreads for BTC and ETH markets, distinguishing between the YES and NO outcomes. Panels (a) report histograms of the median *effective spread* (in USDC), while Panels (b) report histograms of the median *relative spread* (in percentage). In Panels (a), the median effective spreads for both outcomes are largely clustered around one cent, suggesting that bid and ask prices are generally close. However, in Figure II.4 Panel (a), a pronounced right tail is visible, implying the presence of markets with higher median effective spreads. The YES distribution is more dispersed and reaches about 0.06, whereas the NO distribution is more concentrated and reaches about 0.04. Panel (b) shows a similar pattern for median relative spreads: the bulk of observations lies near zero, confirming that spreads are small relative to the mid-price for most markets. At the same time, the right tails remain substantial, particularly for the YES outcome, where some markets reach relative spreads close to one. Overall, markets exhibit tight spreads for most observations, but a small subset—especially on the YES side—shows much larger median spreads, indicating occasional liquidity deterioration.

(a) BTC – Median Effective Spreads



(b) BTC – Median Relative Spreads

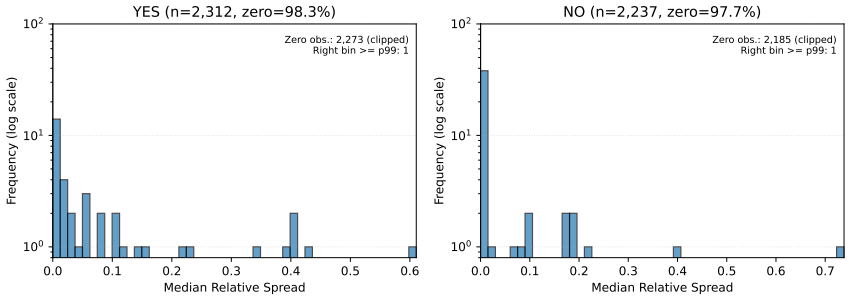
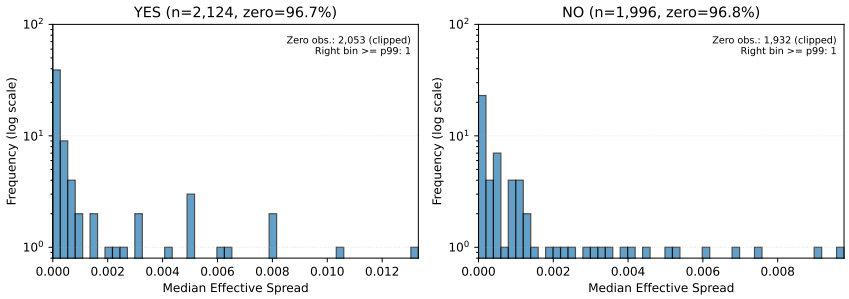


Figure II.4: Histogram – Median spreads per side. The figure displays the histograms of the median spread across BTC contracts, computed separately for YES and NO. Panel (a) reports effective spreads, and Panel (b) reports relative spreads.

(a) ETH – Median Effective Spreads



(b) ETH – Median Relative Spreads

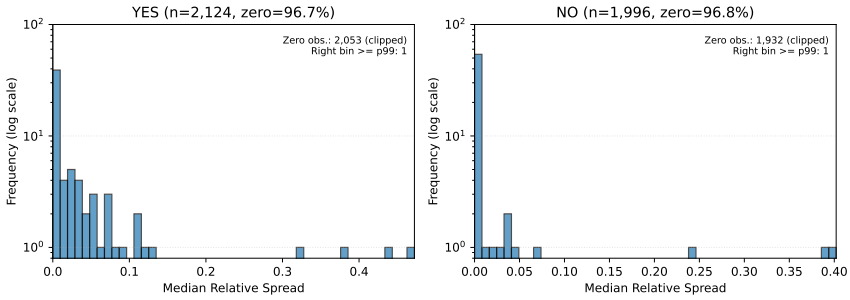


Figure II.5: Histogram – Median spreads per side. The figure displays the histograms of the median spread across ETH contracts, computed separately for YES and NO. Panel (a) reports effective spreads, and Panel (b) reports relative spreads.

II.5 Wallet-Wedge Exposure Robustness

Wallet-size decile	Wallets	Wallet-leg notional	Unique bets traded	Notional-weighted absolute wedge (pp)	Share in wedge decile 10	Share in wedge deciles 8–10
D1	49,353	1,535.3	1,681	0.67	16.0%	33.1%
D2	51,181	12,610.0	2,081	0.97	12.3%	30.9%
D3	46,124	48,468.3	2,506	1.97	13.9%	35.4%
D4	43,989	127,807.1	2,962	2.34	18.8%	42.5%
D5	41,072	241,431.1	3,311	2.89	23.6%	47.2%
D6	42,549	565,981.5	3,807	3.48	28.0%	56.6%
D7	43,866	1,385,277.7	4,153	3.52	29.9%	59.3%
D8	44,918	4,229,352.3	4,404	3.73	30.9%	63.9%
D9	43,036	14,951,023.8	4,617	3.53	26.5%	59.7%
D10	39,319	1,739,661,003.0	4,856	4.86	29.7%	59.4%
All wallets	445,407	1,761,224,490.0	4,856	4.85	29.7%	59.4%

Table II.4: Wallet Size and Wedge Exposure: Full Panel. This table repeats Table 7.2 using all contract-hours in the BTC–ETH panel rather than restricting to the 3–45 day maturity window. Wallet-size deciles are the same pooled BTC–ETH gross-notional deciles. Wedge deciles are re-computed as pooled BTC–ETH bet-level deciles of mean absolute CLOB–OIP wedges in the full-panel sample, so decile cutoffs need not match the 3–45 day table. Wallet-leg notional is role-expanded canonical YES-payoff notional. Notional-weighted absolute wedge is the wallet-leg-notional-weighted absolute CLOB–OIP wedge in percentage points.

II.6 Announcement Dates

Evaluation Day	Event
11/5/2024	Election Day
1/21/2025	Inauguration Day (Jan 20, 2025)
2/3/2025	On February 1, Trump issued EO announcing tariffs on Canada, Mexico, and China.
2/10/2025	Trump announced 25 percent import tariffs on steel and separate proclamation imposing 25 percent tariffs on aluminum as of March 12.
3/4/2025	EOs to raise the new tariffs on all imports from China from 10 percent to 20 percent, impose 10 percent tariffs on imports of Canadian oil and energy products and 25 percent tariffs on the remainder of imports from Canada.
3/25/2025	The White House issued secondary tariffs on third countries importing Venezuelan oil.
3/26/2025	The White House imposed 25 percent tariffs on automobiles and certain automobile parts.
4/3/2025	On April 2, the White House invoked IEEPA to impose baseline 10 percent tariff starting April 5 and then reciprocal tariffs starting April 9.
4/8/2025	The White House amended to impose additional 50 percent tariff on imports from China, increasing to 84 percent.
4/9/2025	The US imposed an additional country-specific tariff on China; then paused other reciprocal tariffs for 90 days, except for China. China will now face 125 percent of tariffs.
4/11/2025	The White House issued a list of products, including smartphones and semiconductors, to be excluded from the April 2 executive order.
4/29/2025	The White House issued a proclamation and an executive order to address concerns over stacking tariffs and avoiding the cumulative tariffs. The proclamation also amended previous tariffs under Section 232 regarding automobiles and automobile parts.
5/12/2025	After the U.S.-China Geneva meeting, the White House issued EO 14298 suspending the heightened China reciprocal tariffs for 90 days and instead applying a 10 percent additional duty effective May 14.
6/3/2025	The White House issued a proclamation raising Section 232 tariffs on steel and aluminum from 25 percent to 50 percent, effective June 4.
7/7/2025	The White House extended the July 9 expiry of the reciprocal-tariff suspension to August 1 and announced new reciprocal tariff rates for many countries.
7/30/2025	The White House suspended duty-free de minimis treatment for commercial low-value shipments globally and separately imposed a 50 percent tariff on semi-finished copper products and intensive copper derivative products.
7/31/2025	The White House further modified reciprocal tariff rates for certain countries; countries listed in Annex I received country-specific rates, while countries not listed remained subject to a 10 percent rate.
8/11/2025	The White House extended the suspension of the heightened China reciprocal tariffs until November 10 while maintaining a 10 percent additional rate during the suspension.
9/5/2025	The White House modified Annex II exclusions under the April 2 reciprocal-tariff order and created a framework for tariff relief under future trade deals.
9/29/2025	The White House issued a Section 232 proclamation on timber, lumber, and derivative wood products, including a 10 percent tariff on softwood lumber and 25 percent tariffs on certain furniture, cabinets, and vanities.
10/17/2025	The White House issued a Section 232 proclamation imposing a 25 percent tariff on medium- and heavy-duty trucks and key parts, and a 10 percent tariff on buses.
11/4/2025	The White House announced a U.S.-China economic and trade arrangement, extending the suspension of the heightened China reciprocal tariffs to November 10, 2026; in a separate order it reduced the fentanyl-related China duty from 20 percent to 10 percent effective November 10, 2025.
11/14/2025	The White House modified the scope of the reciprocal tariffs so that certain qualifying agricultural products were excluded.
12/31/2025	The White House delayed by one year the scheduled January 1, 2026 tariff increases on certain upholstered furniture, kitchen cabinets, and vanities, leaving the current 25 percent rate in place.

Table II.5: Trump's Tariff Announcements Dates. The table lists the key tariff-related policy communication dates identified in the AIER analysis ("How Equity Markets Reacted to Trump's Tariff Announcements," American Institute for Economic Research). Together with the FOMC announcement dates (Table II.6), these dates form the basis for constructing the "Announcement" dummy variable, which equals 1 on each identified event day and 0 otherwise, allowing us to isolate market reactions specifically tied to these policy signals.

Year	Month	Meeting Dates
2024	Jan	30, 31
2024	Mar	19, 20
2024	Apr	30
2024	May	01
2024	Jun	11, 12
2024	Jul	30, 31
2024	Sep	17, 18
2024	Nov	06, 07
2024	Dec	17, 18
2025	Jan	28, 29
2025	Mar	18, 19
2025	May	06, 07
2025	Jun	17, 18
2025	Jul	29, 30
2025	Aug	22
2025	Sep	16, 17
2025	Oct	28, 29
2025	Dec	09, 10

Table II.6: FOMC Announcement Dates. This table reports the Federal Open Market Committee's scheduled monetary policy announcement dates, sourced from the Federal Reserve's official meeting calendar. These dates, together with the tariff–related policy communication dates from the AIER analysis, are used to construct the "Announcement" dummy variable, which equals 1 on each identified event day and 0 otherwise, thereby isolating equity market reactions associated with major policy disclosures.

III Calibration Details – Kou RND Moments

III.1 Deribit Calibration

III.1.1 BTC

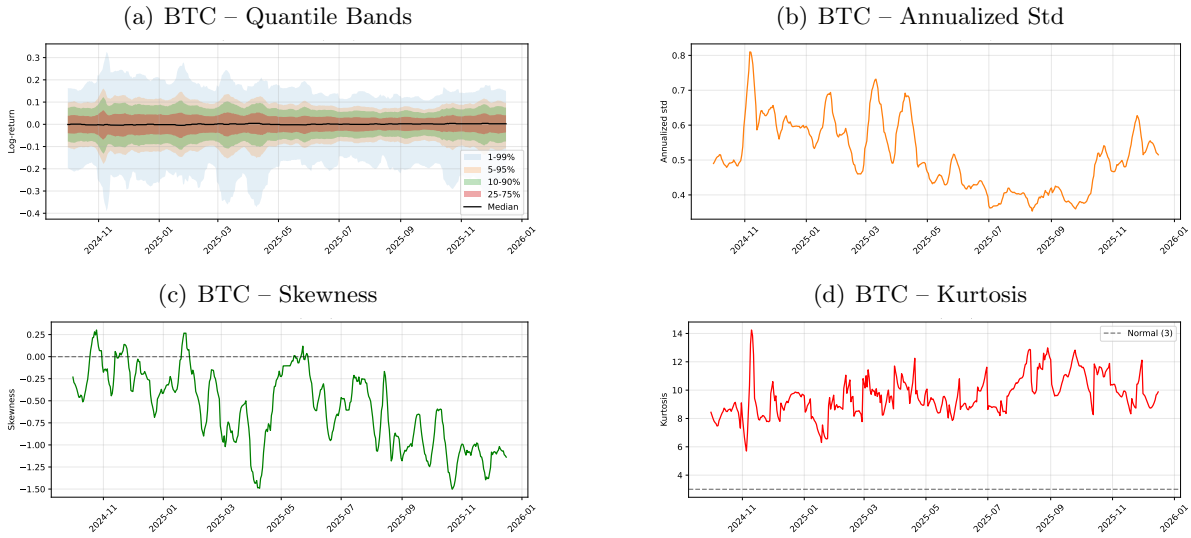


Figure III.1: BTC Kou RND Characteristics – Deribit Calibration. This figure presents selected characteristics of the RND for the 7-day maturity, calibrated to Deribit surfaces. Panels (a)–(d) report the time-series evolution of the quantile bands, annualized standard deviation, skewness, and kurtosis, respectively.

III.1.2 ETH

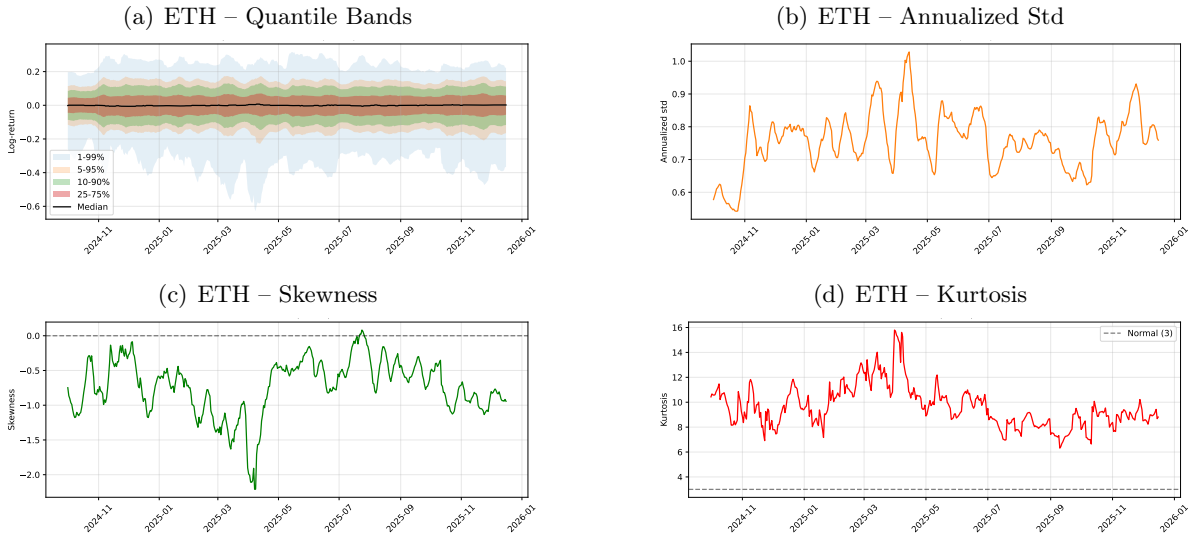
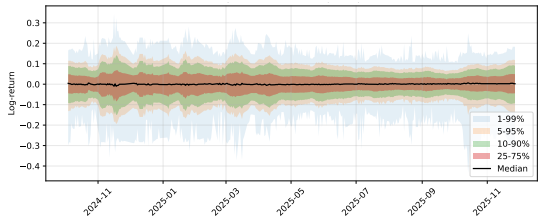


Figure III.2: ETH Kou RND Characteristics – Deribit Calibration. This figure presents selected characteristics of the RND for the 7-day maturity, calibrated to Deribit surfaces. Panels (a)–(d) report the time-series evolution of the quantile bands, annualized standard deviation, skewness, and kurtosis, respectively.

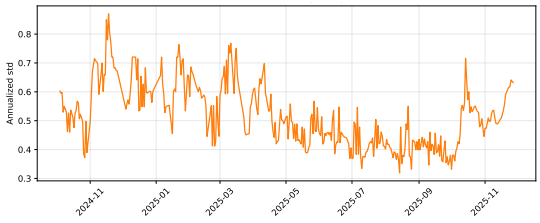
III.2 Polymarket Calibration

III.2.1 BTC

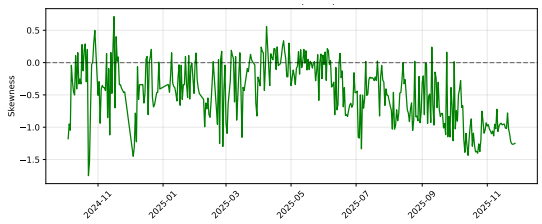
(a) BTC – Quantile Bands



(b) BTC – Annualized Std



(c) BTC – Skewness



(d) BTC – Kurtosis

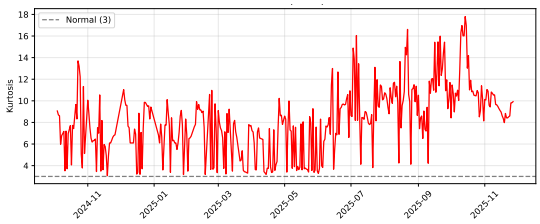
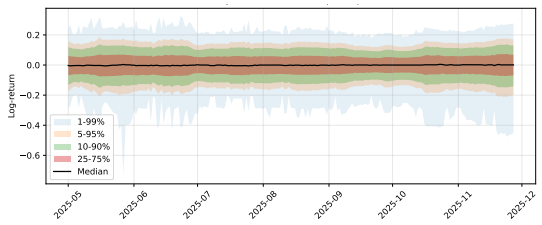


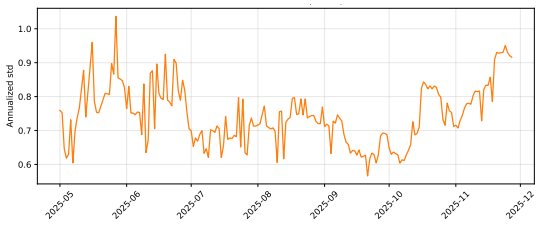
Figure III.3: BTC Kou RND Characteristics – Polymarket Calibration. This figure presents selected characteristics of the Polymarket RND for the 7-day maturity, calibrated to Polymarket digital contracts. Panels (a)–(d) report the time-series evolution of the quantile bands, annualized standard deviation, skewness, and kurtosis, respectively.

III.2.2 ETH

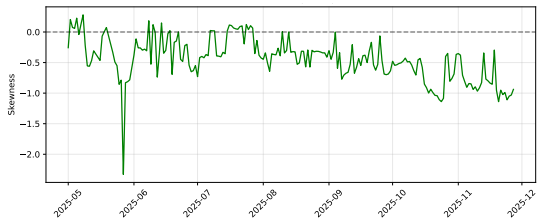
(a) ETH – Quantile Bands



(b) ETH – Annualized Std



(c) ETH – Skewness



(d) ETH – Kurtosis

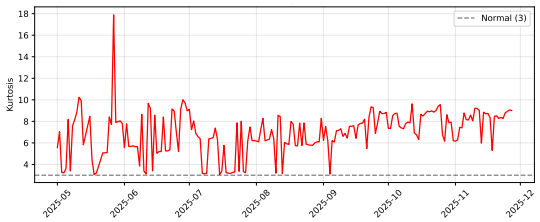


Figure III.4: ETH Kou RND Characteristics – Polymarket Calibration. This figure presents selected characteristics of the Polymarket RND for the 7-day maturity, calibrated to Polymarket digital contracts. Panels (a)–(d) report the time-series evolution of the quantile bands, annualized standard deviation, skewness, and kurtosis, respectively.

IV RNDs from Deribit and Polymarket - Regression Analysis

We next conduct a regression analysis analogous to that in Section 6.1, but shift the focus from individual contracts to the risk-neutral distributions themselves. Specifically, we examine whether disagreement between Polymarket and Deribit is visible directly at the distributional level. To this end, we compare daily 7-day RND estimates constructed from Polymarket contracts with those extracted from Deribit surfaces. The unit of observation is therefore a currency-day rather than a contract. This exercise abstracts from contract-level variation in moneyness and maturity, and instead asks whether aggregate distributional disagreement varies systematically with risk, market frictions, and attention.

Before estimating the regressions, we select a set of meaningful specifications guided by the controls' correlation structure. The correlation matrix in Table IV.1 shows that most regressors are only weakly correlated, suggesting limited multicollinearity concerns outside the risk block. The main exception is the strong correlation between Vol_{30} and VIX of 0.875, indicating that volatility measures capture closely related market conditions. The ETH dummy is also strongly correlated with Vol_{30} and VIX, reflecting systematic asset-level differences between BTC and ETH. Among the remaining variables, L2 Risk is positively correlated with Fear & Greed and negatively correlated with Daily NBP, while USDC volatility is moderately correlated with the difference in the underlying price. Overall, the matrix motivates estimating variables in both block and horse-race specifications, rather than relying exclusively on a fully saturated model.

	Δ (%)	Δ USDC (%)	ETH	Vol_{30}	USDC Vol_{30}	VIX	L2 Risk	Diff Underlying	Fear & Greed	Weekend	Announcement	Daily NBP
Δ (%)	1.000	-0.075	0.040	0.033	0.033	0.023	0.044	0.004	0.067	0.035	-0.033	-0.069
Δ USDC (%)	-0.075	1.000	-0.004	-0.007	-0.018	-0.003	0.000	-0.007	0.036	0.095	0.029	-0.006
ETH	0.040	-0.004	1.000	0.739	-0.182	0.804	-0.214	-0.246	-0.062	-0.011	0.006	0.297
Vol_{30}	0.033	-0.007	0.739	1.000	0.086	0.875	-0.071	0.216	-0.163	-0.012	0.029	0.065
USDC Vol_{30}	0.033	-0.018	-0.182	0.086	1.000	0.064	0.280	0.497	0.310	-0.014	-0.030	-0.267
VIX	0.023	-0.003	0.804	0.875	0.064	1.000	-0.064	0.171	-0.059	0.012	-0.013	0.158
L2 Risk	0.044	0.000	-0.214	-0.071	0.280	-0.064	1.000	0.175	0.579	0.018	0.008	-0.469
Diff Underlying	0.004	-0.007	-0.246	0.216	0.497	0.171	0.175	1.000	0.076	0.024	-0.017	-0.170
Fear & Greed	0.067	0.036	-0.062	-0.163	0.310	-0.059	0.579	0.076	1.000	-0.029	-0.007	-0.453
Weekend	0.035	0.095	-0.011	-0.012	-0.014	0.012	0.018	0.024	-0.029	1.000	-0.215	0.096
Announcement	-0.033	0.029	0.006	0.029	-0.030	-0.013	0.008	-0.017	-0.007	-0.215	1.000	-0.046
Daily NBP	-0.069	-0.006	0.297	0.065	-0.267	0.158	-0.469	-0.170	-0.453	0.096	-0.046	1.000

Table IV.1: RND Regressors – Daily Correlation Matrix. This table reports unconditional correlations among the daily regressors used in the RND regressions. The sample is the daily currency-level RND sample used in Tables IV.2–IV.6.

For each day, we first restrict both RND grids to the 7-day horizon and retain dates observed for both markets. The Polymarket and Deribit densities are then interpolated to a common log-return grid and normalized to integrate to one. Let $f_{a,t}^{PM}$ and $f_{a,t}^{D}$ denote the resulting densities, and let $F_{a,t}^{PM}$ and $F_{a,t}^{D}$ denote the corresponding CDFs. The aggregate distribution-distance

outcome is the first Wasserstein distance,

$$W_{1,a,t} = \int_{\mathbb{R}} \left| F^{PM}_{a,t}(x) - F^{D}_{a,t}(x) \right| \, dx,$$

computed numerically on the common log-return grid. In the moment regressions, the mean, volatility, skewness, and kurtosis are computed from the same normalized densities; moment differences are defined as Polymarket minus Deribit. Controls are constructed from the hourly panel and collapsed to one currency-day before merging with the RND outcomes. BTC/ETH spot and USDC returns are recomputed from the last available daily prices, yielding close-to-close returns. Other state variables—30-day realized volatility, Deribit VIX, L2 Risk, Binance–Deribit spot dislocation, and the Fear & Greed index—enter as their latest available value on each calendar date. Daily NBP is computed by first averaging contract-level NBP within each hour and then averaging across hours within the day. Continuous variables and the dependent variable are standardized; indicators remain at levels; month fixed effects are included; and standard errors are clustered by date.

Table IV.2 reports regression results relating the Wasserstein distance between Polymarket and Deribit RNDs to market conditions. The results show that this disagreement is strongly related to risk conditions. The VIX enters positively and significantly in both the standalone specification and across all horse-race specifications, indicating that the two implied distributions diverge more in high-tail-risk environments. Realized volatility and USDC volatility are likewise positive and significant when included in the volatility block, suggesting that broader market instability and stablecoin stress are associated with larger RND discrepancies.

In the next step, we move beyond a single aggregate distance measure and compare the time-series behavior of the first four moments of the two RNDs. This allows us to identify which distributional dimensions drive the overall disagreement documented in Table IV.2. Specifically, we examine whether Polymarket and Deribit differ systematically in their implied mean, volatility, skewness, and kurtosis, and whether these differences vary with the same risk, friction, and attention variables used above.

Table IV.3 reports the results for the mean of the 7-day RND, separately for Polymarket, Deribit, and their difference. Since continuous regressors and the dependent variable are standardized, coefficient magnitudes can be compared directly across continuous covariates.

Columns (1)–(4) show that the Polymarket mean declines in high-risk states: VIX and Vol_{30}

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Vol_{30}		0.268***							
		(0.069)							
USDC Vol_{30}		0.191**							
		(0.091)							
VIX			0.363***				0.388***	0.395***	0.402***
			(0.111)				(0.125)	(0.121)	(0.121)
L2 Risk				-0.019			-0.012		-0.010
				(0.045)			(0.045)		(0.047)
Diff Underlying					0.140			0.145	0.165
					(0.100)			(0.102)	(0.106)
Daily NBP									0.111*
									(0.060)
Fear & Greed						-0.030	0.038	0.042	0.065
						(0.056)	(0.061)	(0.061)	(0.061)
Announcement						-0.027	-0.024	-0.017	-0.015
						(0.092)	(0.090)	(0.085)	(0.086)
Δ (%)	0.016	0.010	0.022	0.016	0.016	0.016	0.025	0.025	0.028
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
Δ USDC (%)	-0.021	-0.017	-0.022	-0.022	-0.021	-0.017	-0.020	-0.019	-0.019
	(0.028)	(0.023)	(0.031)	(0.027)	(0.023)	(0.028)	(0.033)	(0.027)	(0.025)
Weekend						-0.068	-0.070	-0.071	-0.099
						(0.071)	(0.071)	(0.071)	(0.070)
ETH	0.408***	-0.152	-0.399	0.408***	0.378***	0.408***	-0.455	-0.502*	-0.558**
	(0.057)	(0.147)	(0.254)	(0.057)	(0.062)	(0.057)	(0.283)	(0.276)	(0.278)
R^2 Adj	0.439	0.467	0.447	0.438	0.443	0.438	0.445	0.450	0.452
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	597	597	597	597	597	597	597	597	597

Table IV.2: RND Divergence Regression – Wasserstein Distance. This table reports daily currency-level regressions where the dependent variable is the first Wasserstein distance between the Polymarket and Deribit RND for a 7-day horizon, computed as the integrated absolute distance between the two CDFs on a common log-return grid. Standard errors are clustered by date and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

are negative and highly significant, with VIX having the largest effect. USDC Vol_{30} is also negative and significant, indicating that stablecoin stress lowers the Polymarket mean, while Fear & Greed shifts it upward. Columns (5)–(8) show a similar pattern for Deribit, although the VIX effect is smaller in magnitude and USDC Vol_{30} is insignificant. Columns (9)–(12) confirm these differences: USDC Vol_{30} is negative and significant, implying a stronger Polymarket response to stablecoin stress, whereas Diff Underlying is negative and significant, indicating a stronger Deribit response to Binance–Deribit price wedges. Overall, the Polymarket mean is more sensitive to VIX and USDC Vol_{30}, whereas the Deribit mean reacts more strongly to Diff Underlying.

Table IV.4 reports the results for the volatility of the 7-day RND. Since all explanatory variables are standardized, coefficient magnitudes can be compared directly. Columns (1)–(4) show that Polymarket volatility increases in high-risk states: Vol_{30} and VIX are positive and highly significant, with VIX having the largest effect. USDC Vol_{30} is also positive and significant, indicating that stablecoin stress raises Polymarket volatility, while Fear & Greed lowers it. Columns (5)–(8) show a similar pattern for Deribit, although the Vol_{30} and VIX coefficients

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Outcome	Polymarket	Polymarket	Polymarket	Polymarket	Deribit	Deribit	Deribit	Deribit	Difference	Difference	Difference	Difference
Vol_{30}	-0.539***	-0.500***			-0.463***	-0.437***			-0.153	-0.123		
	(0.058)	(0.057)			(0.042)	(0.039)			(0.099)	(0.102)		
USDC Vol_{30}	-0.065*	-0.077**			0.017	0.006			-0.208**	-0.212**		
	(0.033)	(0.034)			(0.013)	(0.012)			(0.089)	(0.090)		
VIX			-0.748***	-0.702***			-0.675***	-0.636***			-0.128	-0.115
			(0.070)	(0.076)			(0.054)	(0.057)			(0.157)	(0.183)
L2 Risk	0.041		0.057**		0.060***		0.073***		-0.054		-0.048	
	(0.026)		(0.024)		(0.017)		(0.016)		(0.058)		(0.056)	
Diff Underlying	0.221***	0.218***	0.048	0.044	0.321***	0.315***	0.197***	0.190***	-0.280**	-0.276**	-0.396***	-0.387***
	(0.055)	(0.055)	(0.052)	(0.052)	(0.036)	(0.036)	(0.037)	(0.036)	(0.123)	(0.122)	(0.116)	(0.115)
Daily NBP	0.038	0.058*	0.017	0.033	0.018	0.039*	-0.010	0.007	0.048	0.044	0.068	0.064
	(0.030)	(0.031)	(0.029)	(0.030)	(0.022)	(0.021)	(0.020)	(0.019)	(0.064)	(0.064)	(0.064)	(0.064)
Fear & Greed		0.165***		0.089**		0.145***		0.083***		0.039		0.007
		(0.037)		(0.038)		(0.024)		(0.025)		(0.069)		(0.081)
Announcement	-0.033	-0.031	-0.056	-0.054	-0.004	-0.003	-0.032	-0.031	-0.073	-0.072	-0.057	-0.055
	(0.074)	(0.069)	(0.059)	(0.058)	(0.052)	(0.048)	(0.043)	(0.044)	(0.127)	(0.126)	(0.125)	(0.125)
Δ (%)	0.019	0.025	-0.001	0.002	0.012	0.017	-0.005	-0.003	0.018	0.021	0.011	0.013
	(0.019)	(0.018)	(0.018)	(0.018)	(0.019)	(0.018)	(0.018)	(0.018)	(0.042)	(0.042)	(0.042)	(0.042)
Δ USDC (%)	0.021	0.015	0.022	0.019	0.004	-0.001	0.005	0.001	0.041	0.040	0.044	0.044
	(0.027)	(0.029)	(0.016)	(0.016)	(0.017)	(0.018)	(0.009)	(0.009)	(0.034)	(0.034)	(0.034)	(0.034)
Weekend	-0.017	-0.011	0.041	0.041	-0.038	-0.034	0.008	0.009	0.057	0.061	0.082	0.083
	(0.039)	(0.038)	(0.037)	(0.037)	(0.029)	(0.029)	(0.028)	(0.028)	(0.087)	(0.087)	(0.088)	(0.088)
ETH	-0.879***	-0.967***	-0.297*	-0.404**	-1.235***	-1.295***	-0.665***	-0.757***	1.008***	0.946***	0.992***	0.961**
	(0.126)	(0.123)	(0.167)	(0.181)	(0.089)	(0.084)	(0.127)	(0.134)	(0.213)	(0.216)	(0.357)	(0.413)
R^2 Adj	0.790	0.796	0.800	0.801	0.906	0.910	0.919	0.918	0.137	0.136	0.117	0.116
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	597	597	597	597	597	597	597	597	597	597	597	597

Table IV.3: RND Regression – Mean. This table reports daily currency-level regressions where the dependent variables are the Polymarket RND mean, the Deribit RND mean, and their difference. Standard errors are clustered by date and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

are smaller and USDC Vol_{30} is weak or insignificant. Columns (9)–(12) confirm these differences: USDC Vol_{30} is positive and highly significant, implying a stronger Polymarket response to stablecoin stress, while Diff Underlying is positive and significant in the difference regressions, indicating that the two markets respond differently to Binance–Deribit price discrepancies. Overall, Polymarket is more sensitive to VIX, Vol_{30}, and USDC Vol_{30}, whereas Deribit reacts more strongly to Diff Underlying.

Table IV.5 reports the results for the 7-day RND skewness. Since lower skewness corresponds to greater downside risk, negative coefficients indicate that the distribution becomes more left-skewed, whereas positive coefficients indicate a reduction in downside asymmetry. Columns (1)–(4) show that Polymarket skewness becomes more negative in high-risk states: Vol_{30} is negative and significant, while VIX is only weakly significant and loses explanatory power once additional controls are included. Fear & Greed enters positively and significantly, indicating that more favorable sentiment reduces downside skewness. Columns (5)–(8) show a related pattern for Deribit, but with stronger sensitivity to sentiment and market activity. In particular, Fear & Greed is larger and highly significant, while Daily NBP is negative and significant. USDC Vol_{30} is positive and significant for Deribit, but not for Polymarket. Columns (9)–(12) compare the two markets directly. USDC Vol_{30} is negative and significant, implying that Polymarket skewness becomes more negative relative to Deribit during stablecoin stress. By contrast, the negat-

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Outcome	Polymarket	Polymarket	Polymarket	Polymarket	Deribit	Deribit	Deribit	Deribit	Difference	Difference	Difference	Difference
Vol_{30}	0.571***	0.526***			0.489***	0.452***			0.194*	0.177		
	(0.060)	(0.058)			(0.044)	(0.040)			(0.107)	(0.111)		
USDC Vol_{30}	0.055**	0.069**			-0.021*	-0.008			0.208***	0.209***		
	(0.028)	(0.029)			(0.011)	(0.011)			(0.076)	(0.077)		
VIX			0.785***	0.726***			0.731***	0.670***			0.109	0.115
			(0.065)	(0.071)			(0.053)	(0.053)			(0.148)	(0.174)
L2 Risk	-0.047*		-0.064***		-0.063***		-0.076***		0.047		0.037	
	(0.025)		(0.023)		(0.017)		(0.016)		(0.057)		(0.055)	
Diff Underlying	-0.211***	-0.207***	-0.032	-0.028	-0.294***	-0.288***	-0.165***	-0.159***	0.237**	0.233**	0.363***	0.358***
	(0.052)	(0.052)	(0.049)	(0.049)	(0.033)	(0.033)	(0.034)	(0.033)	(0.117)	(0.116)	(0.109)	(0.108)
Daily NBP	-0.017	-0.040	0.007	-0.013	-0.006	-0.031	0.024	0.002	-0.029	-0.024	-0.046	-0.040
	(0.030)	(0.030)	(0.028)	(0.029)	(0.022)	(0.021)	(0.019)	(0.019)	(0.066)	(0.065)	(0.065)	(0.065)
Fear & Greed		-0.188***		-0.112***		-0.182***		-0.117***		-0.006		0.019
		(0.038)		(0.039)		(0.026)		(0.027)		(0.075)		(0.085)
Announcement	0.016	0.013	0.041	0.039	0.011	0.009	0.042	0.039	0.011	0.010	-0.003	-0.003
	(0.069)	(0.063)	(0.054)	(0.053)	(0.053)	(0.048)	(0.043)	(0.043)	(0.118)	(0.118)	(0.115)	(0.115)
Δ (%)	-0.020	-0.027	0.001	-0.004	-0.006	-0.013	0.012	0.007	-0.037	-0.038	-0.030	-0.030
	(0.019)	(0.018)	(0.018)	(0.018)	(0.018)	(0.017)	(0.017)	(0.017)	(0.038)	(0.039)	(0.038)	(0.039)
Δ USDC (%)	-0.017	-0.010	-0.018	-0.014	-0.004	0.002	-0.005	-0.000	-0.034	-0.034	-0.037	-0.038
	(0.024)	(0.025)	(0.013)	(0.014)	(0.017)	(0.019)	(0.010)	(0.009)	(0.030)	(0.030)	(0.030)	(0.030)
Weekend	0.026	0.019	-0.034	-0.035	0.034	0.028	-0.015	-0.016	-0.022	-0.024	-0.051	-0.051
	(0.039)	(0.037)	(0.037)	(0.037)	(0.029)	(0.028)	(0.028)	(0.028)	(0.088)	(0.088)	(0.088)	(0.089)
ETH	0.827***	0.926***	0.227	0.365**	1.173***	1.258***	0.533***	0.673***	-0.990***	-0.955***	-0.849**	-0.862**
	(0.124)	(0.121)	(0.153)	(0.166)	(0.091)	(0.084)	(0.121)	(0.122)	(0.219)	(0.226)	(0.338)	(0.395)
R^2 Adj	0.816	0.824	0.827	0.829	0.914	0.921	0.930	0.930	0.139	0.139	0.118	0.117
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	597	597	597	597	597	597	597	597	597	597	597	597

Table IV.4: RND Regression – Volatility. This table reports daily currency-level regressions where the dependent variables are the Polymarket RND volatility, the Deribit RND volatility, and their difference. Standard errors are clustered by date and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

ive coefficient on Fear & Greed shows that Deribit skewness reacts more strongly to favorable sentiment. Announcement is positive and significant, suggesting that announcement days shift Polymarket skewness upward relative to Deribit. Overall, the skewness results suggest that Polymarket is more sensitive to USDC Vol_{30}, while Deribit responds more strongly to Fear & Greed and Daily NBP.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Outcome	Polymarket	Polymarket	Polymarket	Polymarket	Deribit	Deribit	Deribit	Deribit	Difference	Difference	Difference	Difference
Vol_{30}	-0.319***	-0.246**			-0.313***	-0.163**			-0.033	-0.096		
	(0.114)	(0.110)			(0.081)	(0.072)			(0.097)	(0.106)		
USDC Vol_{30}	-0.048	-0.068			0.149***	0.109***			-0.183***	-0.166***		
	(0.059)	(0.059)			(0.021)	(0.021)			(0.062)	(0.062)		
VIX			-0.205*	-0.032			-0.219**	0.143			-0.005	-0.161
			(0.116)	(0.122)			(0.111)	(0.089)			(0.120)	(0.128)
L2 Risk	0.039		0.056		0.054		0.072		-0.010		-0.009	
	(0.048)		(0.047)		(0.043)		(0.046)		(0.050)		(0.051)	
Diff Underlying	0.206**	0.202**	0.103	0.111	0.051	0.045	0.019	0.042	0.159*	0.160*	0.085	0.071
	(0.090)	(0.089)	(0.077)	(0.077)	(0.041)	(0.040)	(0.040)	(0.035)	(0.093)	(0.093)	(0.087)	(0.085)
Daily NBP	-0.047	-0.017	-0.061	-0.022	-0.147***	-0.092**	-0.189***	-0.113***	0.087	0.067	0.111	0.081
	(0.074)	(0.072)	(0.069)	(0.069)	(0.048)	(0.042)	(0.047)	(0.042)	(0.081)	(0.078)	(0.077)	(0.076)
Fear & Greed		0.273***		0.286***		0.530***		0.585***		-0.209***		-0.246***
		(0.063)		(0.068)		(0.052)		(0.052)		(0.071)		(0.073)
Announcement	0.139	0.143*	0.127	0.138*	-0.078	-0.070	-0.112	-0.089	0.209*	0.205*	0.228**	0.218**
	(0.086)	(0.085)	(0.081)	(0.081)	(0.094)	(0.081)	(0.088)	(0.079)	(0.109)	(0.105)	(0.106)	(0.102)
Δ (%)	0.022	0.034	0.014	0.030	-0.087***	-0.064***	-0.092***	-0.060***	0.101***	0.091***	0.098***	0.084**
	(0.034)	(0.033)	(0.033)	(0.033)	(0.024)	(0.022)	(0.025)	(0.023)	(0.034)	(0.035)	(0.035)	(0.035)
Δ USDC (%)	-0.011	-0.020	-0.010	-0.020	0.006	-0.012	0.004	-0.015	-0.016	-0.009	-0.014	-0.006
	(0.029)	(0.031)	(0.027)	(0.029)	(0.017)	(0.017)	(0.020)	(0.017)	(0.029)	(0.028)	(0.030)	(0.027)
Weekend	-0.079	-0.068	-0.047	-0.043	0.029	0.051	0.050	0.059	-0.104	-0.114	-0.092	-0.096
	(0.072)	(0.069)	(0.073)	(0.070)	(0.060)	(0.054)	(0.060)	(0.055)	(0.071)	(0.071)	(0.073)	(0.074)
ETH	0.979***	0.815***	0.795***	0.396	1.114***	0.781***	0.967***	0.132	-0.038	0.101	-0.087	0.273
	(0.219)	(0.211)	(0.271)	(0.285)	(0.161)	(0.144)	(0.266)	(0.215)	(0.189)	(0.207)	(0.287)	(0.303)
R^2 Adj	0.362	0.380	0.355	0.373	0.647	0.715	0.634	0.710	0.321	0.331	0.307	0.320
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	597	597	597	597	597	597	597	597	597	597	597	597

Table IV.5: RND Regression – Skewness. This table reports daily currency-level regressions where the dependent variables are the Polymarket RND skewness, the Deribit RND skewness, and their difference. Standard errors are clustered by date and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table IV.6 reports the results for the kurtosis of the 7-day RND. Higher kurtosis corresponds to fatter tails, while lower kurtosis indicates thinner-tailed distributions. Columns (1)–(4) show that Polymarket kurtosis is only weakly related to most regressors. VIX is negative and significant, suggesting lower Polymarket tail thickness in high tail-risk states, while Fear & Greed is negative and significant only in the richer VIX specification. The remaining friction variables are not significant. Columns (5)–(8) show stronger effects for Deribit: VIX is negative and highly significant, USDC Vol_{30} is negative and significant, and L2 Risk is also negative and significant when included. Fear & Greed is negative and significant, indicating lower kurtosis in more favorable sentiment states. Columns (9)–(12) show that the differences between Polymarket and Deribit are generally not statistically significant. Overall, the kurtosis regressions provide weaker evidence of systematic cross-market differences than the regressions on mean, volatility, and skewness.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Outcome	Polymarket	Polymarket	Polymarket	Polymarket	Deribit	Deribit	Deribit	Deribit	Difference	Difference	Difference	Difference
Vol_{30}	-0.113	-0.131			-0.080	-0.090			-0.079	-0.093		
	(0.098)	(0.100)			(0.104)	(0.107)			(0.094)	(0.099)		
USDC Vol_{30}	-0.058	-0.052			-0.184***	-0.175***			0.033	0.036		
	(0.045)	(0.045)			(0.048)	(0.049)			(0.053)	(0.054)		
VIX			-0.272**	-0.359***			-0.402***	-0.535***			-0.081	-0.104
			(0.118)	(0.130)			(0.135)	(0.142)			(0.132)	(0.148)
L2 Risk	-0.034		-0.035		-0.105**		-0.114**		0.018		0.022	
	(0.056)		(0.055)		(0.043)		(0.045)		(0.055)		(0.055)	
Diff Underlying	-0.003	0.000	-0.056	-0.059	0.176*	0.185**	0.085	0.090	-0.094	-0.096	-0.104	-0.110
	(0.078)	(0.077)	(0.071)	(0.070)	(0.095)	(0.094)	(0.083)	(0.082)	(0.097)	(0.097)	(0.091)	(0.090)
Daily NBP	-0.061	-0.074	-0.058	-0.078	0.016	-0.010	0.041	0.001	-0.073	-0.074	-0.083	-0.084
	(0.059)	(0.059)	(0.056)	(0.057)	(0.047)	(0.047)	(0.046)	(0.046)	(0.065)	(0.065)	(0.063)	(0.063)
Fear & Greed		-0.092		-0.146**		-0.139**		-0.242***		-0.026		-0.030
		(0.064)		(0.069)		(0.057)		(0.061)		(0.076)		(0.081)
Announcement	0.009	0.008	0.009	0.004	0.050	0.050	0.067	0.060	-0.016	-0.017	-0.025	-0.027
	(0.104)	(0.102)	(0.105)	(0.102)	(0.095)	(0.093)	(0.094)	(0.092)	(0.111)	(0.111)	(0.110)	(0.110)
Δ (%)	-0.003	-0.006	-0.009	-0.017	0.058	0.055	0.048	0.038	-0.033	-0.035	-0.035	-0.038
	(0.030)	(0.029)	(0.030)	(0.031)	(0.035)	(0.035)	(0.037)	(0.037)	(0.033)	(0.033)	(0.033)	(0.033)
Δ USDC (%)	-0.003	0.000	-0.002	0.003	-0.011	-0.006	-0.009	0.000	0.003	0.003	0.002	0.003
	(0.026)	(0.025)	(0.026)	(0.025)	(0.018)	(0.018)	(0.019)	(0.020)	(0.029)	(0.029)	(0.029)	(0.029)
Weekend	0.009	0.006	0.025	0.023	-0.000	-0.002	0.020	0.018	0.010	0.008	0.016	0.015
	(0.071)	(0.071)	(0.071)	(0.070)	(0.069)	(0.069)	(0.069)	(0.069)	(0.081)	(0.081)	(0.081)	(0.081)
ETH	-0.631***	-0.590***	-0.253	-0.052	-0.881***	-0.854***	-0.144	0.165	-0.215	-0.185	-0.194	-0.141
	(0.200)	(0.201)	(0.282)	(0.304)	(0.211)	(0.217)	(0.312)	(0.320)	(0.192)	(0.200)	(0.311)	(0.344)
R^2 Adj	0.311	0.313	0.315	0.319	0.338	0.338	0.334	0.342	0.243	0.243	0.244	0.244
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	597	597	597	597	597	597	597	597	597	597	597	597

Table IV.6: RND Regression – Kurtosis. This table reports daily currency-level regressions where the dependent variables are the Polymarket RND kurtosis, the Deribit RND kurtosis, and their difference. Standard errors are clustered by date and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.