

The Extremity Premium

Sentiment Regimes and Adverse Selection in Cryptocurrency Markets

Murad Farzulla^{1,2,*}

¹Disensus, London, UK ²King’s College London, London, UK

*Correspondence: murad@disensus.ai ORCID: 0009-0002-7164-8704

February 2026

Abstract

Using the Crypto Fear & Greed Index and Bitcoin daily data, we document that sentiment *extremity* predicts excess uncertainty beyond realized volatility. Extreme fear and extreme greed regimes exhibit significantly higher spreads than neutral periods—a phenomenon we term the “extremity premium.” Extended validation on the full Fear & Greed history (February 2018–January 2026, $N = 2,896$) confirms the finding: within-volatility-quintile comparisons show a premium ($p < 0.001$, pooled volatility-demeaned Cohen’s $d = 0.21$ —a *post-hoc, exploratory* test, as the pre-specified within-quintile endpoint does not survive multiple-testing correction; raw pooled extreme-vs-neutral $d = 0.40$), Granger causality runs from uncertainty to spreads (primary-sample $F = 12.79$; the extended-sample $F = 211$ is partly mechanical, sharing a high-low input with the spread measure), and placebo tests reject the null ($p < 0.0001$). The effect replicates on Ethereum and across 6 of 7 market cycles. However, the premium is sensitive to functional form: comprehensive regression controls absorb regime effects, while nonparametric stratification preserves them. We interpret this as evidence that sentiment extremity captures volatility-regime interactions not fully represented by parametric controls—consistent with, but not conclusively separable from, the F&G Index’s embedded volatility component. An agent-based model is included as an illustrative device that reproduces the pattern qualitatively; because its spread-uncertainty link is coded rather than emergent, it does no inferential work, and the inferential weight rests entirely on the empirical analysis (the formal moment-matching test we report validates a separate simplified model, not the full agent specification). The results suggest that intensity, not direction, drives uncertainty-linked liquidity withdrawal in cryptocurrency markets, though identification of “pure” sentiment effects from volatility remains an open challenge.

Keywords: extremity premium, sentiment regimes, adverse selection, market microstructure, cryptocurrency, agent-based modeling

JEL Classification: C63, G12, G14

1 Introduction

Cryptocurrency markets present a distinctive challenge for market microstructure analysis: sentiment signals exhibit substantial uncertainty arising from both model limitations and inherent market noise. Traditional market microstructure models (Glosten and Milgrom, 1985; Avellaneda and Stoikov, 2008) predict that market makers should widen spreads when uncertainty about asset value increases, but

existing frameworks do not decompose sentiment uncertainty into its epistemic (model-related) and aleatoric (inherent noise) components. This decomposition is critical for understanding how market makers respond to information quality in cryptocurrency markets, where sentiment signals are inherently noisy and model confidence varies substantially.

This paper investigates a central question in market microstructure: *Do market makers respond more to sentiment direction or to sentiment uncertainty?* The inferential weight of the paper rests on an empirical analysis of real cryptocurrency data, which documents that sentiment *extremity* predicts excess spread-setting uncertainty beyond realized volatility. We complement this empirical core with an agent-based model calibrated to 739 days of real Bitcoin market data. We are explicit upfront that the agent-based model is an *illustrative* device, not a source of inferential evidence: its spread-uncertainty link is coded rather than emergent (Section 7), so it can only confirm that the assumed mechanism behaves as designed—it cannot independently validate the mechanism. The formal goodness-of-fit test we report (Simulated Method of Moments, Section 5.13) is run on a separate, simplified reduced-form model rather than on the full agent-based specification, and we say so wherever that result is invoked. The empirical regime finding, not the simulation, carries the contribution.

1.1 Motivation

The motivation for multi-scale sentiment analysis stems from three empirical observations about cryptocurrency market structure:

First, information fragmentation. Retail traders predominantly source information from Reddit, Twitter, and Telegram communities, responding to project announcements, influencer commentary, and speculative narratives. Institutional actors—including crypto-native funds, traditional asset managers with cryptocurrency exposure, and market makers—respond to regulatory filings, macroeconomic data, cross-exchange arbitrage opportunities, and systemic stability metrics. Farzulla (2026a) demonstrates that whitepaper claims and project fundamentals explain only a fraction of cross-sectional return variation, suggesting that sentiment and narrative factors dominate cryptocurrency price formation. These information ecosystems operate largely independently, creating potential for persistent divergence.

Second, asymmetric response times. Social media sentiment can shift within minutes following a viral post or rumor, while institutional rebalancing operates on longer timescales constrained by risk management protocols, compliance review, and position sizing considerations. This temporal asymmetry suggests that the *relative* weight of retail versus institutional sentiment should itself be time-varying and regime-dependent.

Third, uncertainty heterogeneity. The quality of sentiment signals varies dramatically, and this variation affects market microstructure outcomes. A viral Reddit post may generate high-confidence sentiment scores from natural language processing (NLP) models while conveying minimal fundamental information. Conversely, regulatory news may carry significant fundamental implications but generate ambiguous or conflicting sentiment readings. Appropriately responding to these signals requires decomposing uncertainty into its epistemic (model-related) and aleatoric (inherent noise) components, following Kendall and Gal (2017).

1.2 Theoretical Motivation

The framework is motivated by a conjecture we term the *multi-scale divergence hypothesis*: when retail and institutional sentiment diverge significantly, subsequent market volatility may increase. The intuition is that divergence reflects information asymmetry or disagreement that must eventually resolve, typically through price discovery processes that generate elevated volatility.

Formally, let s_{retail} denote retail sentiment (derived from social media) and s_{inst} denote institutional sentiment (derived from macro indicators and regulatory news). A natural divergence measure is:

$$D_t = |s_{retail,t} - s_{inst,t}| \quad (1)$$

This hypothesis would predict that D_t correlates positively with forward realized volatility σ_{t+k} for some horizon k . While we implement infrastructure for divergence tracking, **we do not validate this hypothesis in the current work**—our simulation did not produce sufficient divergence events for meaningful statistical analysis.

1.3 Contributions

This paper makes six primary contributions to market microstructure theory and sentiment analysis:

1. **The extremity premium:** We document that extreme sentiment regimes (both greed and fear) exhibit elevated uncertainty relative to neutral, even after volatility control. Extreme greed adds +5.5 percentage points uncertainty above baseline; extreme fear adds +3.9 percentage points. All significant effects survive Bonferroni correction for multiple comparisons.
2. **Negative result as contribution:** On our heuristic decomposition, aleatoric uncertainty carries the larger share of the mean index (81.6% versus 18.4% epistemic; Table 4)—a split that is conditional on the decomposition weights and proxy scales rather than estimated (indeed the raw proxy means alone do not order the two components this way), so we read it as suggestive and not load-bearing. Taken tentatively, it is consistent with cryptocurrency sentiment being structurally noisy rather than information-asymmetric, which would redirect research effort from signal refinement toward regime detection—simple macro indices may suffice.
3. **Directional identification:** Granger causality tests show uncertainty predicts spreads (primary sample, 3-day lag: $F = 12.79$, $p < 0.001$; extended sample, 1-day lag: $F = 211$) but not vice versa ($F = 0.82$, $p = 0.49$). While instrumental variables proved weak, OLS and IV estimates are nearly identical, suggesting minimal endogeneity bias.
4. **SMM-based validation:** Rather than informal “stylized facts matching,” we validate a simplified representation of the model’s key mechanisms via Simulated Method of Moments. The J-test ($p = 0.36$) indicates the reduced-form model is not rejected—it replicates volatility clustering, kurtosis, and spread-volatility correlations without hard-coding them.
5. **Out-of-sample validation:** The extremity premium holds directionally in 2022 bear market data (93% fear regimes)—same direction as main sample, though not statistically significant due to regime imbalance. This is suggestive but not statistically confirmed; the consistency is directional rather than inferentially robust.
6. **Multiple spread estimators:** Both Corwin-Schultz (2012) and a close-to-close Abdi-Rinaldo (2017) variant show consistent uncertainty correlations, ruling out estimation-specific artifacts.

1.4 Primary Hypothesis and Endpoints

The primary hypothesis is that sentiment *extremity* (distance from neutral) predicts uncertainty more than sentiment *direction* (bullish vs. bearish). We operationalize this through two pre-specified primary endpoints:

1. **Extremity premium:** Mean uncertainty in extreme regimes (greed or fear) minus mean uncertainty in neutral regimes, tested via two-sided t -test with Bonferroni correction for multiple comparisons.

2. **Volatility-controlled effect:** The extremity premium conditional on realized volatility quintile, tested via within-quintile t -tests to rule out mechanical volatility confounding.

Secondary analyses include Granger causality tests, IV exploration, cross-asset replication, and out-of-sample validation. These are explicitly exploratory and do not affect the primary conclusions.

Power considerations. With $N = 715$ complete cases (170 extreme, 110 neutral), the primary two-sample comparison has approximately 98% power to detect a medium effect ($d = 0.5$) at $\alpha = 0.05$. The extended sample ($N = 2,896$) provides adequate power for stratified analyses. Within-quintile tests ($n \approx 30$ – 60 per cell) are powered for large effects ($d \geq 0.8$) but underpowered for small effects—we acknowledge this limitation.

1.5 Paper Organization

Section 2 reviews related work. Section 3 presents the methodological framework. Section 4 describes data sources and implementation. Section 5 presents empirical results. Section 6 interprets findings. Section 7 discusses limitations. Section 8 concludes.

2 Literature Review

2.1 Agent-Based Market Models

Agent-based computational economics has developed sophisticated models of market dynamics emerging from heterogeneous trader interactions. The Santa Fe Artificial Stock Market (Palmer et al., 1994) demonstrated that realistic market properties—including the stylized facts documented by Cont (2001): volatility clustering, fat-tailed returns, and long-range dependence—can emerge from simple agent learning rules. LeBaron (2006) provides a comprehensive survey, identifying key design choices including agent heterogeneity, learning mechanisms, and market clearing protocols.

Order book dynamics have received particular attention. Cont et al. (2010) develop a stochastic model matching empirical order flow patterns, while Paddrik et al. (2012) apply ABM to flash crash analysis. Recent work has developed quantitative agent-based models calibrated to real economies, demonstrating that ABMs can match and sometimes outperform mainstream macro models on leverage cycles, bubbles, and crisis dynamics (Farmer, 2025). Complexity-economics applications extend ABM and network models to labor-market transitions and automation shocks, showing how network structure shapes occupational mobility (del Rio-Chanona et al., 2021). Pangallo and del Rio-Chanona (2024) argue that ABMs require stronger empirical calibration to move beyond stylized-fact matching toward genuine forecasting and policy evaluation. Our calibration to CryptoCompare BTC/USDT data via Simulated Method of Moments (Section 5), with explicit moment-matching targets for volatility clustering, return autocorrelation, and spread dynamics, represents precisely the kind of data-driven grounding they advocate—while our uncertainty decomposition adds a mechanistic layer that purely statistical calibration cannot provide. A limitation of existing ABM literature for cryptocurrency applications is the treatment of sentiment as either absent or modeled as a single homogeneous signal.

2.2 Market Microstructure Theory

The seminal work of Glosten and Milgrom (1985) establishes that bid-ask spreads arise from adverse selection: market makers face informed traders with superior information and must widen spreads to compensate for expected losses. Kyle (1985) develops a model wherein an informed trader optimally conceals private information through strategic order submission.

Avellaneda and Stoikov (2008) extend this framework to high-frequency market making, developing optimal quote-setting strategies that balance inventory risk against adverse selection. Their model

shows that market makers should widen spreads when uncertainty about fair value increases—a prediction directly relevant to our uncertainty-driven spread adjustment mechanism. [Barucca and Lillo \(2017\)](#) examine agent reflexivity in price formation, showing how feedback loops between price dynamics and agent beliefs shape microstructure outcomes. Recent financial-computing work characterizes limit order book forecasting and persistence, providing empirical benchmarks for microstructure modeling ([Briola et al., 2025a,b](#)). A pilot exploration by the author suggests that cryptocurrency liquidity responds differentially to event types, with infrastructure failures (exchange collapses, protocol exploits) producing substantially larger spread increases than regulatory announcements—indicating that structural uncertainty may dominate sentiment-driven uncertainty in determining liquidity provision ([Farzulla, 2026b](#)).

2.3 Sentiment Analysis in Finance

[Tetlock \(2007\)](#) demonstrates that media pessimism predicts stock market returns and trading volume. [Antweiler and Frank \(2004\)](#) find that message board activity predicts volatility, though not returns, suggesting sentiment contains information about uncertainty rather than direction. [Loughran and McDonald \(2011\)](#) develop domain-specific word lists for financial texts. More recent work establishes that sentiment indicators—particularly happiness and fear indices—serve as robust nonlinear predictors of cryptocurrency returns, with predictive power concentrated at extreme market states ([Naeem et al., 2021](#)). [Dias et al. \(2022\)](#) confirm these findings using quantile regression, showing that investor emotions predict both returns and volatility with regime-dependent effect sizes.

[Gal and Ghahramani \(2016\)](#) show that Monte Carlo dropout enables approximate Bayesian inference in neural networks, producing prediction distributions rather than point estimates. [Kendall and Gal \(2017\)](#) distinguish epistemic uncertainty (reducible through more data) from aleatoric uncertainty (irreducible inherent noise). Our framework applies this decomposition to sentiment analysis.

2.4 Cryptocurrency Market Structure

[Makarov and Schoar \(2020\)](#) document significant and persistent price dislocations across exchanges, indicating fragmented liquidity and limited arbitrage. [Bouri et al. \(2017\)](#) examine Bitcoin’s hedging and safe-haven properties, finding that cryptocurrency price dynamics depend on broader market conditions. [Bourghelle et al. \(2022\)](#) demonstrate that collective emotions drive Bitcoin volatility through regime-switching dynamics, with sentiment effects varying in sign and magnitude across calm, bubble-formation, and bubble-collapse phases—establishing a sentiment-volatility channel that our analysis explicitly controls for when isolating the spread-uncertainty relationship. [Chen and Hafner \(2019\)](#) extend this analysis to bubble formation, using sentiment as a regime-switching variable and finding that volatility increases as sentiment deteriorates. [Kyriazis et al. \(2022\)](#) show that Twitter-based uncertainty measures influence cryptocurrency volatility nonlinearly, with effects most pronounced at extreme quantiles—supporting our focus on sentiment extremity rather than direction.

Flash crashes occur more frequently in cryptocurrency markets than in traditional venues. [Golub et al. \(2012\)](#) analyze mini flash crashes, attributing them to liquidity withdrawal cascades. [Farzulla \(2025a\)](#) documents asymmetric volatility responses to positive versus negative sentiment shocks in cryptocurrency markets, finding that negative shocks produce larger and more persistent volatility increases—a pattern consistent with leverage effects observed in traditional markets but amplified by cryptocurrency market structure. [Jia et al. \(2022\)](#) find that extreme sentiment regimes amplify herding behavior in cryptocurrency markets, with both euphoria and dysphoria increasing the magnitude of herd-driven price movements. [Chen and Nguyen \(2024\)](#) provide complementary evidence that news sentiment has divergent effects on herding versus anti-herding behavior, with optimism amplifying coordination failures.

Gurdgiev and O’Loughlin (2020) document that anchoring biases are especially pronounced during periods of high uncertainty, linking behavioral effects to the fear-driven regimes central to our analysis. Rognone et al. (2020) compare cryptocurrency and foreign exchange markets, finding that Bitcoin reacts positively to both positive and negative news during bubble periods—suggesting information asymmetry is exacerbated in crypto relative to traditional markets. Koutmos (2022) uses transaction-level data to show that rising sentiment is robustly associated with price dynamics and liquidity shifts. The 24/7 trading environment, absence of circuit breakers, and high retail participation create conditions conducive to extreme price movements.

2.5 Systemic Risk and Macro Signals

Adrian and Brunnermeier (2016) introduce CoVaR, measuring value-at-risk conditional on systemic distress. Gudgeon et al. (2020) analyze the “decentralized financial crisis” of March 2020, documenting how DeFi liquidation cascades amplified market stress. Macro-financial ABMs with explicit banking and interbank markets show how liquidity freezes, policy rules, and network structure shape instability and crisis propagation (Popoyan et al., 2020). Empirical studies of correlation structure and volatility co-movement offer complementary evidence on how shocks propagate across assets, and information filtering networks provide a backbone for extracting sparse dependence structures from high-dimensional markets (Aste, 2025; Samal et al., 2021). Farzulla and Maksakov (2026) develop an aggregated systemic risk index specifically for cryptocurrency markets, combining network topology, liquidity concentration, and cross-exchange contagion measures—providing the macro-level risk signals that complement our micro-level sentiment uncertainty decomposition (live dashboard: asri.dissensus.ai). Network-theoretic analysis of financial instability further demonstrates that heterogeneous portfolio allocations can trigger stability-instability phase transitions, with diversification having non-monotonic effects depending on network connectivity (Forer et al., 2025). The bidirectional nature of risk transmission between traditional and decentralized finance—termed ‘crosstagon’ by Aufiero et al. (2025)—further motivates micro-level modeling of how sentiment-driven uncertainty propagates across market structures.

2.6 Research Gap

Despite substantial progress, no existing work combines multi-scale sentiment analysis with uncertainty decomposition in an agent-based market microstructure model calibrated to real cryptocurrency data. Our framework addresses this gap.

3 Methodology

3.1 Multi-Scale Signal Architecture

The framework processes sentiment information through two parallel layers:

Macro Layer: Institutional-level signals derived from the Crypto Fear & Greed Index, which aggregates: volatility (25%), market momentum/volume (25%), social media engagement (15%), surveys (15%), Bitcoin dominance (10%), and Google Trends (10%).

Micro Layer: Retail-level signals derived from social media sentiment analysis using CryptoBERT with Monte Carlo dropout for uncertainty quantification.

3.1.1 Fear & Greed Index Processing

The Fear & Greed Index produces daily values from 0 (extreme fear) to 100 (extreme greed). We convert to a normalized sentiment score $s \in [-1, 1]$:

$$s_{macro} = \frac{\text{Fear\&Greed} - 50}{50} \quad (2)$$

We classify regimes based on index thresholds:

- Extreme fear: < 25
- Fear: 25–44
- Neutral: 45–55
- Greed: 56–75
- Extreme greed: > 75

Index Composition and Circularity Concerns. The Fear & Greed Index aggregates seven components: volatility (25%), market momentum (25%), social media (15%), surveys (15%), BTC dominance (10%), and Google Trends (10%). Volatility is computed from 30/90-day historical price ranges, distinct from the intraday Parkinson volatility we measure. A potential concern is circularity: if F&G embeds volatility, correlating F&G-based regimes with volatility-derived uncertainty may be mechanical. Section 5.10.12 addresses this directly: DVOL-based regime classification (pure implied volatility) does *not* replicate the extremity premium—the premium is specific to sentiment-based regimes. Furthermore, 57% concordance between F&G and DVOL extreme classifications indicates they measure related but distinct phenomena.

3.1.2 ASRI Integration for Macro Sentiment

The Aggregated Systemic Risk Index (Farzulla and Maksakov, 2026) provides institutional-level macro signals through four data channels, complementing the Fear & Greed Index:

- **DeFi Health (35%):** Total Value Locked trends, stablecoin peg deviations, and protocol-level stress indicators from DeFiLlama.
- **Regulatory Opacity (40%):** Sentiment analysis of regulatory news via Google News RSS, with keyword filtering for SEC, CFTC, and enforcement-related terminology.
- **TradFi Linkage (25%):** Traditional finance spillover via FRED macroeconomic indicators. Note: ASRI uses traditional finance proxies for systemic risk; the uncertainty decomposition (Section 3.2) uses Deribit DVOL as primary crypto-native volatility (35%), with VIX as secondary spillover proxy (15%).

ASRI generates alert levels (low, moderate, elevated, high, critical) used for regime detection in the blending weights (Table 1). During crisis regimes ($ASRI > 70$), macro signals receive 60% weight; during regulatory events (elevated regulatory score), 70% weight. This regime-adaptive weighting reflects the empirical observation that institutional signals dominate during market stress, while retail sentiment is more informative during calm periods.

Data visualization is available at asri.dissensus.ai, with source code at github.com/studiofarzulla/asri.

3.1.3 CryptoBERT with MC Dropout

The micro layer processes social media text through CryptoBERT, a RoBERTa-based model fine-tuned on 3.2 million cryptocurrency-related posts. Following Gal and Ghahramani (2016), we enable dropout at inference time and run $T = 50$ forward passes for each input text, producing:

1. **Mean sentiment:** $\bar{s} = \frac{1}{T} \sum_{t=1}^T s_t$
2. **Epistemic uncertainty:** $\sigma_{epi}^2 = \frac{1}{T} \sum_{t=1}^T (s_t - \bar{s})^2$

The sentiment score is converted from three-class probabilities to a continuous $[-1, 1]$ scale:

$$s_{micro} = p_{bullish} - p_{bearish} \quad (3)$$

EWMA Smoothing. Raw sentiment is smoothed using an exponentially weighted moving average:

$$s_t^{smooth} = \alpha \cdot s_t^{raw} + (1 - \alpha) \cdot s_{t-1}^{smooth} \quad (4)$$

with $\alpha = 0.1$, corresponding to approximately 5-minute half-life.

3.1.4 Signal Blending

The blended sentiment score combines macro and micro signals:

$$s_{blend} = w_{macro}(r) \cdot s_{macro} + w_{micro}(r) \cdot s_{micro} \quad (5)$$

where weights depend on the detected regime r (Table 1).

Table 1: Regime-adaptive blending weights. Crisis and regulatory regimes weigh institutional (macro) signals more heavily; normal market conditions favor retail (micro) sentiment.

Regime	w_{macro}	w_{micro}
Crisis	0.60	0.40
Regulatory	0.70	0.30
Bullish	0.25	0.75
Bearish	0.35	0.65
Neutral	0.30	0.70

Implementation Note: The micro text layer (CryptoBERT with MC dropout) is described for methodological completeness. The current empirical analysis uses only the macro Fear & Greed layer; micro-layer integration is reserved for future work. Consequently, the epistemic uncertainty components dependent on micro-layer signals (MC Variance σ_{mc} and text-derived Shannon Entropy $H(p)$) are not computed in the current analysis. The empirical uncertainty decomposition (Table 4) uses only macro-available proxies: aleatoric sources (Deribit DVOL, VIX spillover, stablecoin peg deviation) and epistemic sources derivable from price data (cross-exchange dispersion as regulatory opacity proxy).

3.2 Uncertainty Decomposition

Following Kendall and Gal (2017), we decompose total uncertainty into epistemic and aleatoric components.

3.2.1 Epistemic Uncertainty Sources

Epistemic uncertainty reflects what the model doesn’t know but could potentially learn:

Regulatory Opacity. We proxy regulatory opacity using cross-exchange price dispersion, computed as the daily standard deviation of BTC spot prices across major exchanges (Binance, Coinbase, Kraken, Bitstamp) at 00:00 UTC:

$$\sigma_{reg} = \text{normalize}(\text{arbitrage_index}) \quad (6)$$

Data Availability. We compute a data completeness score:

$$\sigma_{data} = 1 - \frac{\text{available_sources}}{\text{expected_sources}} \quad (7)$$

MC Variance. The variance across Monte Carlo dropout passes:

$$\sigma_{mc}^2 = \frac{1}{T} \sum_{t=1}^T (s_t - \bar{s})^2 \quad (8)$$

The total epistemic uncertainty is:

$$\sigma_{epi} = \gamma_1 \sigma_{reg} + \gamma_2 \sigma_{data} + \gamma_3 \sigma_{mc} \quad (9)$$

with weights $(\gamma_1, \gamma_2, \gamma_3) = (0.3, 0.2, 0.5)$. These weights are heuristic rather than derived from calibration or cross-validation.

3.2.2 Aleatoric Uncertainty Sources

Aleatoric uncertainty reflects inherent noise that cannot be reduced:

Deribit DVOL (Primary). The crypto-native implied volatility index serves as our primary aleatoric source (35% weight):

$$\sigma_{dvol} = \text{normalize}(\text{DVOL}) \quad (10)$$

Unlike VIX, DVOL is derived from Deribit BTC options and reflects crypto-specific implied volatility without cross-asset spillover confounds.

VIX Spillover (Secondary). Traditional finance contagion proxy (15% weight):

$$\sigma_{vix} = \text{normalize}(\text{VIX}) \quad (11)$$

Peg Deviation. Stablecoin peg deviations indicate DeFi instability:

$$\sigma_{peg} = |\text{stablecoin_price} - 1.0| \quad (12)$$

Shannon Entropy. Text ambiguity via the entropy of sentiment probability distribution:

$$H(p) = - \sum_i p_i \log p_i \quad (13)$$

The total aleatoric uncertainty is:

$$\sigma_{ale} = \delta_1 \sigma_{dvol} + \delta_2 \sigma_{vix} + \delta_3 \sigma_{peg} + \delta_4 H(p) \quad (14)$$

with weights $(\delta_1, \delta_2, \delta_3, \delta_4) = (0.35, 0.15, 0.25, 0.25)$. These weights are heuristic rather than calibrated; sensitivity analysis is a priority for future work.

Clarification on Shannon Entropy Source. In the current empirical implementation, $H(p)$ is computed from the Fear & Greed Index distribution across a rolling 30-day window—not from text-derived sentiment probabilities via the micro layer (which was not implemented). This macro-derived entropy captures regime stability: stable sentiment periods exhibit low entropy, while volatile regime-switching periods exhibit high entropy.

3.2.3 Total Uncertainty

Total uncertainty combines both components:

$$\sigma_{total} = \sigma_{epi} + \sigma_{ale} \quad (15)$$

Note on Units and Scaling. The σ terms are *normalized uncertainty indices* in $[0, 1]$ rather than true standard deviations, and the implementation combines them *linearly*—as (weighted) sums of normalized indices—throughout, in both the empirical pipeline and the simulation. A variance-addition (quadratic) combination $\sigma_{total}^2 = \sigma_{epi}^2 + \sigma_{ale}^2$ would be the probabilistically motivated rule for true standard deviations of independent components, but is not what the implementation computes. This heuristic linear aggregation lacks rigorous probabilistic grounding—we acknowledge this as a methodological limitation.

3.3 Agent Specifications

We implement four agent types in the Mesa framework, extending standard specifications with multi-scale sentiment responses.

3.3.1 Market Makers

Market makers provide liquidity with uncertainty-aware spread adjustment:

$$p_{bid} = p_{mid} - \frac{s_{base}}{2} - \gamma Q - \delta \sigma_{total} \quad (16)$$

$$p_{ask} = p_{mid} + \frac{s_{base}}{2} - \gamma Q + \delta \sigma_{total} \quad (17)$$

where s_{base} is the base spread, Q is inventory, γ is inventory aversion, and δ scales spread widening with uncertainty. This extends [Avellaneda and Stoikov \(2008\)](#) by making spread adjustment scale with decomposed uncertainty.

3.3.2 Informed Traders

Informed traders act on sentiment signals when confidence is high:

$$\text{action} = \begin{cases} \text{buy } V & \text{if } s_{blend} > \tau \text{ and } \sigma_{epi} < \bar{\sigma}_{epi} \\ \text{sell } V & \text{if } s_{blend} < -\tau \text{ and } \sigma_{epi} < \bar{\sigma}_{epi} \\ \text{hold} & \text{otherwise} \end{cases} \quad (18)$$

where $\tau = 0.3$ is the sentiment threshold and $\bar{\sigma}_{epi} = 0.5$ is maximum acceptable epistemic uncertainty.

3.3.3 Noise Traders

Noise traders arrive according to a Poisson process with weak sentiment influence:

$$\text{direction} \sim \text{Bernoulli}(0.5 + \beta s_{blend}) \quad (19)$$

where $\beta = 0.1$ provides weak sentiment influence.

3.3.4 Arbitrageurs

Arbitrageurs exploit price dislocations and are sentiment-agnostic:

$$\text{action} = \begin{cases} \text{buy} & \text{if } p < p_{fair} - \varepsilon \\ \text{sell} & \text{if } p > p_{fair} + \varepsilon \\ \text{hold} & \text{otherwise} \end{cases} \quad (20)$$

3.4 Order Book Dynamics

The simulation uses Mesa with a continuous double auction mechanism.

Order Submission. Agents submit limit orders with price and quantity. Market orders execute as aggressive limit orders.

Matching Engine. The order book matches orders using price-time priority.

Price Updates. The mid-price is updated after each trade as the average of best bid and ask.

4 Data and Implementation

4.1 Data Sources

We analyze 739 days of Bitcoin market data from January 1, 2024 to January 8, 2026, combining two primary sources:

Binance BTC/USDT: Daily OHLCV data obtained via the public Binance API. This provides: open, high, low, close prices; trading volume; and derived metrics including returns and volatility.

Fear & Greed Index: Daily sentiment readings from Alternative.me. This composite indicator aggregates multiple sentiment signals and is freely available without authentication, ensuring full reproducibility.

Scope Note. The empirical analysis uses only the macro sentiment layer (Fear & Greed Index). The micro layer (CryptoBERT with MC Dropout, Section 3.1.3) is presented as part of the theoretical framework for completeness, but social media data were not collected for this study. Future work will integrate real-time Twitter/X data to test the full multi-scale model.

4.2 Data Preprocessing and Quality Control

Exchange Selection. We use Binance BTC/USDT as the primary data source due to: (1) highest global trading volume for BTC pairs, reducing microstructure noise; (2) USDT denomination avoiding USD banking frictions; (3) continuous 24/7 trading without market closures; and (4) public API access ensuring reproducibility.

Missing Data Handling. Of 739 calendar days, we observe complete OHLCV data for all days. The Fear & Greed Index has no missing values in the sample period. The uncertainty measure requires a rolling warm-up window, so it cannot be computed for the first 24 days at the series boundary; analyses that use the uncertainty series (the regime comparisons of Section 5.4, and the Granger and regime-transition tests, which additionally lag it) therefore run on $N = 715$ complete cases. This is the same 24-observation exclusion referred to as “missing uncertainty data” elsewhere in the paper. All sample sizes are reported explicitly throughout.

Outlier Treatment. We do not winsorize or trim outliers. Extreme values are informative for our research question (extreme sentiment regimes). The Corwin-Schultz spread estimator produces 23 days (3.1%) with negative estimated spreads (set to zero per standard practice). Results are robust to excluding these observations.

Time Alignment. The Fear & Greed Index is published daily at 00:00 UTC. Binance OHLCV data uses UTC midnight-to-midnight candles. Both series are thus naturally aligned without interpolation or time-zone adjustments.

4.3 Spread Estimation from OHLCV Data

Market microstructure spreads are not directly observable in daily OHLCV data. We estimate spreads using two established high-low estimators:

Corwin-Schultz (2012) Estimator. Exploits the insight that daily high-low range reflects both volatility and bid-ask spread. Two-day high-low ratios separate these components:

$$\hat{S}_{CS} = \frac{2(e^\alpha - 1)}{1 + e^\alpha} \quad (21)$$

where α is derived from $\beta = \mathbb{E} \left[\sum_{j=0}^1 (\ln(H_{t+j}/L_{t+j}))^2 \right]$ and $\gamma = (\ln(H_{t,t+1}/L_{t,t+1}))^2$, with $H_{t,t+1}$ and $L_{t,t+1}$ being the two-day high and low respectively. The estimator exploits the fact that volatility scales with the square root of time while spread does not (Corwin and Schultz, 2012).

Roll (1984) Measure. The Roll (1984) estimator uses negative serial covariance in returns as a spread proxy:

$$\hat{S}_{Roll} = 2\sqrt{-\text{Cov}(r_t, r_{t-1})} \quad (22)$$

when the covariance is negative (set to zero otherwise). This assumes the spread induces bid-ask bounce in transaction prices.

For our analysis, we primarily use the Corwin-Schultz estimator due to its superior performance in high-frequency cryptocurrency markets, where bid-ask bounce effects are less pronounced.

4.4 Sample Period Characteristics

The sample period (January 2024 to January 2026) captures a significant bull market, with Bitcoin appreciating from approximately \$44,000 to \$91,000 (+106%). This creates potential selection bias, as contrarian patterns may differ in bear markets. The sample includes:

- Bitcoin ETF approval (January 2024)
- Multiple all-time high breaches
- Periods of elevated regulatory uncertainty
- Several sharp corrections (>10%) within the broader uptrend

4.5 Implementation

The framework is implemented in Python 3.11.4 using the following packages (version numbers for reproducibility):

- **Mesa 2.1.1:** Agent-based modeling framework
- **Transformers 4.35.0:** HuggingFace library for CryptoBERT
- **pandas 2.1.3 / numpy 1.26.2:** Data manipulation
- **scipy 1.11.4 / statsmodels 0.14.1:** Statistical analysis and HAC standard errors
- **arch 6.2.0:** GARCH modeling for volatility estimation

All statistical tests use statsmodels implementations with Newey-West HAC standard errors (5-lag truncation) unless otherwise noted. Bootstrap and permutation tests use `numpy.random.seed(42)` for reproducibility; Monte Carlo weight simulations use `seed(2024)`. Source code is available at github.com/studiofarzulla/sentiment-microstructure-abm.

5 Results

5.1 Empirical Validation: Real Spread-Uncertainty Correlation

Before examining simulation results, we validate the spread-uncertainty relationship in real market data. Using Corwin-Schultz spreads estimated from 739 days of Binance BTC/USDT OHLCV data (Section 4), we test correlation with observable uncertainty proxies constructed from market observables.

Table 2 presents the empirical correlations. All standard errors use Newey-West heteroskedasticity and autocorrelation consistent (HAC) estimation with 5-lag truncation. All hypothesis tests are two-sided unless otherwise noted.

Table 2: Empirical Spread-Uncertainty Correlations (N=739 days). Spreads estimated via Corwin-Schultz (2012). All p-values use Newey-West HAC standard errors.

Uncertainty Proxy	Pearson r	p -value	Interpretation
Range-Based Volatility	0.260	<0.0001	Moderate positive
Realized Volatility	0.243	<0.0001	Moderate positive
Total Uncertainty Index	0.235	<0.0001	Moderate positive
Epistemic Proxy	0.149	<0.001	Weak positive
Aleatoric Proxy	0.246	<0.0001	Moderate positive

The empirical correlations ($r = 0.24$ – 0.26) are statistically significant at $p < 0.0001$. This is consistent with the theoretical prediction from [Glosten and Milgrom \(1985\)](#): market makers respond to information quality, widening spreads when adverse selection risk increases. However, as Section 5.10 shows, this baseline correlation is largely mechanical—both variables load heavily on realized volatility. The regime-conditional effects, rather than the baseline correlation, constitute the substantive finding. Notably, aleatoric uncertainty (inherent market noise) correlates more strongly than epistemic uncertainty (model limitations), foreshadowing our decomposition findings.




Figure 1: Time series of Corwin-Schultz spreads and total uncertainty over the 739-day sample period. The empirical correlation ($r = 0.24$) shows that spread dynamics track uncertainty dynamics at daily frequency. Spread spikes during high-uncertainty regimes (sentiment extremes) are visually apparent.

5.2 ABM Consistency Check: Spread-Uncertainty Correlation

Having established the empirical relationship (Section 5.1), we use the agent-based model to illustrate the proposed mechanism. The ABM incorporates explicit uncertainty-premium logic in market maker




Figure 2: Scatter plot of CS spread versus total uncertainty with OLS regression line ($N = 739$). The positive empirical relationship ($r = 0.24$, $R^2 = 0.055$) is consistent with spreads widening with uncertainty, though the baseline correlation is largely mechanical (Section 5.10). The regime effects constitute the substantive finding.

behavior (Equations 16–17). Because this behavior is *coded* rather than emergent, the simulation cannot independently validate the mechanism—it can only confirm the coded logic operates as designed.

Table 3 presents the simulation correlations.

Table 3: ABM Simulation: Sentiment, Uncertainty, and Spread Correlations (739 simulated trading days from a single representative run, seed 42). The simulation isolates the uncertainty channel, producing higher correlations than observed empirically. Because the ABM’s spread-uncertainty link is coded rather than emergent, this table confirms implementation fidelity and does no inferential work; the 739 days are one correlated simulation path, so we do not treat them as independent observations and read the correlations as illustrative point values rather than as sampling estimates.

Variable	Correlation with Spreads	Interpretation
Total Uncertainty	0.637	Strong positive
Aleatoric Uncertainty	0.612	Strong positive
Epistemic Uncertainty	0.496	Moderate positive
Sentiment Direction	0.085	Weak positive

The simulation correlation ($r = 0.64$) exceeds the empirical correlation ($r = 0.24$) because the ABM isolates a single channel. In real markets, other factors—inventory management, competitive pressure, latency constraints, and regulatory frictions—dilute the pure uncertainty-spread relationship. The direction match confirms the coded mechanism operates as designed; the magnitude difference reflects the ABM’s simplifying assumptions. This is a consistency check, not independent validation—the genuine validation comes from SMM moment-matching (Section 5.13), where emergent dynamics like kurtosis and volatility clustering arise without being hard-coded.

This finding extends the Glosten-Milgrom adverse selection model: spread widening is associated with information quality rather than sentiment direction. When uncertainty is high—whether from model limitations (epistemic) or inherent market noise (aleatoric)—market makers face increased adverse selection risk.

5.3 Uncertainty Decomposition Statistics

Table 4 presents the decomposition of total uncertainty into epistemic and aleatoric components.

Table 4: Uncertainty Decomposition. On the heuristic decomposition, aleatoric uncertainty accounts for the larger share (81.6%) of the mean total index; the split is conditional on the decomposition weights and proxy scales (see note) and is not load-bearing for the core finding.

Component	Mean Value	Proportion of Total
Total Uncertainty	0.278	100.0%
Aleatoric Uncertainty	0.227	81.6%
Epistemic Uncertainty	0.051	18.4%

Proportions are linear shares of the mean indices (0.227/0.278 and 0.051/0.278) and sum to 100% by construction, consistent with the linear aggregation the implementation uses throughout; they are *not* orthogonal variance shares—a quadratic (variance-addition) rule $\sigma_{\text{total}}^2 = \sigma_{\text{epi}}^2 + \sigma_{\text{ale}}^2$, which the implementation does not use, would give $\sqrt{0.051^2 + 0.227^2} = 0.233 \neq 0.278$. The 81.6/18.4 split is a construction choice conditional on the heuristic weights (γ, δ) and proxy scales.

On this decomposition, aleatoric uncertainty carries the larger share (81.6% vs 18.4% epistemic), consistent with cryptocurrency sentiment signals carrying substantial inherent noise. Because the split is conditional on the decomposition weights and proxy scales rather than estimated, we read it as suggestive rather than a precise variance partition; it is not load-bearing for the extremity-premium finding.

5.4 The Extremity Premium: Extreme Sentiment Amplifies Uncertainty

Table 5 presents mean uncertainty by sentiment regime, revealing a counter-intuitive pattern: *extreme* sentiment regimes exhibit the highest uncertainty, not neutral regimes.

Table 5: Mean Uncertainty by Sentiment Regime (N=715 days after excluding days with missing uncertainty data). **Counter-intuitive finding:** Extreme regimes exhibit the highest uncertainty, controlling for volatility.

Regime	N	Uncertainty	Volatility	Δ vs Neutral
Extreme Greed	94	0.521	0.470	+0.055***
Fear	140	0.436	0.412	+0.034**
Extreme Fear	76	0.403	0.379	+0.039**
Greed	295	0.324	0.341	+0.003
Neutral	110	0.303	0.325	(baseline)

*** $p < 0.001$, ** $p < 0.01$. Δ coefficients from OLS with volatility control (HC3 robust SEs; $R^2 = 0.77$).

All significant effects survive Bonferroni correction ($\alpha = 0.05/4 = 0.0125$).

Regression Specification. The regime effects are estimated via:

$$\text{Uncertainty}_t = \alpha + \sum_{r \in \mathcal{R}} \beta_r \cdot \mathbf{1}[\text{Regime}_t = r] + \gamma \cdot \text{Volatility}_t + \varepsilon_t \quad (23)$$

where $\mathcal{R} = \{\text{Extreme Greed}, \text{Greed}, \text{Fear}, \text{Extreme Fear}\}$ and Neutral is the omitted baseline. Table 6

reports HC3 heteroskedasticity-robust standard errors; Table 7 reports the same models with Newey-West standard errors (5-lag truncation) to account for autocorrelation.

Table 6 presents the full regression results comparing Model 1 (volatility only) with Model 2 (volatility + regime dummies).

Table 6: Regime Effects on Uncertainty (OLS with Heteroskedasticity-Robust Standard Errors)

Variable	Model 1: Vol Only		Model 2: Vol + Regimes	
	Coef.	SE	Coef.	SE
Volatility	1.182***	(0.034)	1.119***	(0.034)
Extreme Greed	—	—	+0.055***	(0.012)
Greed	—	—	+0.003	(0.008)
Fear	—	—	+0.034**	(0.011)
Extreme Fear	—	—	+0.039**	(0.012)
R^2	0.755		0.768	
ΔR^2	—		+0.013	
F-test (regimes)	—		10.1***	
N	715		715	

HC3 heteroskedasticity-robust SEs. Neutral regime = reference category. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. F-test (regimes) is the OLS joint test $F_{4,710}$. The joint test for regime dummies is highly significant, confirming that sentiment extremity adds explanatory power beyond volatility alone. Table 7 reports the same two models under Newey-West HAC (5 lags), where the volatility SEs roughly double, the joint regime F falls to 4.4, and the extreme-fear coefficient becomes marginal ($p = 0.063$).

The “extremity premium”—where extreme sentiment regimes exhibit higher uncertainty than neutral regimes even after controlling for volatility—suggests that market makers face maximum adverse selection risk during sentiment extremes. Pooling extreme greed and extreme fear against neutral yields Cohen’s $d = 1.06$ (large effect), indicating the magnitude is economically substantial.¹ When sentiment is directionally intense, informed traders may be exploiting sentiment-driven mispricings, forcing market makers to widen spreads beyond what volatility alone would predict.

This finding inverts the naive intuition that “ambiguity is risky.” Instead, the data suggest that *conviction* is risky: when the crowd commits strongly to a directional view, the probability of informed trading increases. The asymmetry between extreme greed (+0.055) and extreme fear (+0.039) may reflect the leveraged nature of crypto bull markets, where euphoria creates greater opportunities for informed profit-taking.

5.4.1 Progressive Model Specifications

To address the question of whether the extremity premium survives additional controls, Table 7 presents five progressive model specifications.

Models 1–4 progressively add controls: volatility alone (Model 1), regime dummies (Model 2), trading volume and returns (Model 3), and the ETF approval event dummy (Model 4). The extremity premium—extreme greed (+0.046) and fear (+0.034) relative to neutral—survives all specifications. Model 5 demonstrates that a continuous distance-from-neutral measure ($|F\&G - 50|/50$) performs comparably to discrete regime dummies, with a significant positive coefficient (+0.077, $p = 0.012$). This confirms that *extremity* rather than specific regime thresholds drives the effect.

¹This pooled d conflates between-regime variance with within-regime variance. The volatility-controlled effect is more conservative: the pooled volatility-demeaned within-quintile comparison on the extended sample yields $d = 0.21$ (Section 5.10), while the per-quintile primary-sample effect sizes are large but estimated on small strata ($d \in [0.76, 0.86]$, Table 16).




Figure 3: The Extremity Premium: Uncertainty distribution by sentiment regime (N = 715 complete cases). Extreme regimes (fear and greed) exhibit significantly higher mean uncertainty than neutral regimes, even after controlling for volatility. Diamond markers indicate regime means; dashed line shows neutral regime mean for reference.

Table 7: Progressive Model Specifications: Uncertainty Beyond Realized Volatility

Variable	Model 1 Vol Only	Model 2 + Regimes	Model 3 + Controls	Model 4 + ETF Event	Model 5 Continuous
Volatility	+1.182*** (0.076)	+1.119*** (0.072)	+1.086*** (0.077)	+1.086*** (0.077)	+1.148*** (0.076)
<i>Regime Dummies (Neutral = baseline)</i>					
Extreme Greed	—	+0.055*	+0.046*	+0.046*	—
Greed	—	+0.003	−0.000	−0.000	—
Fear	—	+0.034*	+0.034*	+0.034*	—
Extreme Fear	—	+0.039	+0.036	+0.036	—
<i>Additional Controls</i>					
Log(Volume)	—	—	+0.017*	+0.017*	—
Daily Returns	—	—	+0.221*	+0.221*	—
ETF Approval	—	—	—	+0.000	—
<i>Continuous Measure (Model 5)</i>					
Distance from Neutral	—	—	—	—	+0.077*
R^2	0.755	0.768	0.772	0.772	0.763
ΔR^2	—	+0.013	+0.004	+0.000	—
F-test (regimes)	—	4.4***	—	—	—
N	715	715	715	715	715

Newey-West HAC SEs (5 lags). Neutral regime = reference category. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Model 5 uses continuous distance from neutral (0 = neutral, 1 = extreme) instead of discrete regime dummies. ETF Approval = dummy for Jan 10–20, 2024 (Bitcoin spot ETF approval window).

5.4.2 Comprehensive Regression Controls (Extended Sample)

To address reviewer concerns about functional form specification, Table 8 presents comprehensive regressions on the extended sample ($N = 1,961$ days with positive spread estimates) with progressively richer control sets.

Table 8: Comprehensive Regression: Spread Determinants with Full Controls (Extended Sample)

	(1) Regimes	(2) +Vol	(3) +Ret/Vol	(4) +Day FE	(5) Kitchen Sink
<i>Sentiment Regime (Neutral = baseline)</i>					
Extreme Fear	88.90*** (17.50)	15.10 (11.74)	9.00 (11.65)	10.77 (11.18)	8.79 (12.10)
Fear	27.95* (12.02)	−3.96 (8.12)	−3.31 (8.10)	−1.67 (7.67)	−1.96 (8.26)
Greed	25.94* (10.22)	8.80 (7.97)	9.45 (7.69)	10.20 (7.47)	8.28 (8.25)
Extreme Greed	95.98*** (21.46)	23.88 (15.39)	15.42 (14.54)	15.43 (14.13)	16.62 (14.69)
<i>Volatility Controls</i>					
RV	—	8531***	6563***	6959***	6507***
RV ²	—	−41355***	−33596***	−35099***	−32920***
<i>Additional Controls</i>					
Returns	—	—	1360*	1228 [†]	1088 [†]
Log(Volume)	—	—	24.4***	16.2***	33.4***
Day-of-Week FE	No	No	No	Yes	Yes
Month FE	No	No	No	No	Yes
Year FE	No	No	No	No	Yes
R^2	0.048	0.217	0.275	0.298	0.314
N	1,961	1,961	1,961	1,961	1,961

HAC (Newey-West, 10 lags) SEs in parentheses. Neutral (F&G 46–55) = baseline. Sample: days with positive CS spread estimates. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, [†] $p < 0.10$.

Key finding: Regime coefficients are highly significant in Model 1 ($R^2 = 4.8\%$) but attenuate substantially when volatility controls are added (R^2 jumps to 21.7%), and become insignificant in the full kitchen-sink specification (Model 5, all $p > 0.25$). This sensitivity reflects the F&G Index’s 25% volatility component: extreme sentiment regimes are partly high-volatility regimes, and regression controls absorb this mechanical correlation.

However, the within-quintile stratification approach (Table 16)—which makes no parametric assumptions about the volatility-spread relationship—yields different conclusions. The *pooled* test on the extended sample, which compares extreme vs. neutral spreads after demeaning within volatility quintiles, remains highly significant ($t = 3.36$, $p = 0.0008$, Cohen’s $d = 0.21$). This small but robust pooled effect is the conservative, volatility-controlled counterpart to the larger per-quintile primary-sample effects in Table 16 ($d \in [0.76, 0.86]$); the difference in magnitude reflects that pooling demeans away the between-quintile variation that inflates the per-stratum standardized effects. Either way, the extremity premium represents genuine within-volatility-level differences between sentiment regimes, though it cannot be expressed as a simple additive regression coefficient.

We interpret this divergence as follows: regression-based controls impose parametric functional forms (linear + quadratic) that may not capture regime-volatility interactions. Stratification allows arbi-

trary within-bin relationships and is more appropriate for categorical regime effects. The DVOL-based regime analysis (Section 5.10.12) supports this interpretation: pure implied-volatility regimes do *not* replicate the extremity premium, suggesting F&G’s non-volatility components (social metrics, surveys, momentum) drive the effect rather than mechanical volatility embedding.

5.5 Regime Persistence and Transitions

Extreme sentiment regimes show high persistence: extreme fear exhibits 76.3% daily persistence, while extreme greed shows 80.2%. Neutral regimes serve as transition states (60.3% persistence).

5.6 Real Data Summary Statistics

Table 9 presents summary statistics for the 739-day sample period.

Table 9: Real Data Summary Statistics (N=739 days)

Metric	Value
<i>Price Data (Binance BTC/USDT)</i>	
Start Price	\$44,180
End Price	\$91,196
Total Return	+106.4%
Daily Return Mean	+0.13%
Daily Return Std	2.49%
Return Kurtosis	2.45
<i>Sentiment Data (Fear & Greed Index)</i>	
Mean Index Value	55.8 (greed)
Sentiment Mean (<i>s</i>)	+0.12
Sentiment Std	0.40

5.7 Secondary Finding: Contrarian Signal Pattern

Table 10 presents mean daily returns conditional on sentiment regime.

Table 10: Mean Daily Returns by Sentiment Regime (N=739 days, full sample). The contrarian pattern is directionally consistent but **not statistically significant** at conventional levels. Note: regime counts differ from Table 5 because that table excludes 24 days with missing uncertainty data.

Regime	Days	Mean Return	Direction
Extreme Fear	76	+0.34%	Contrarian buy
Fear	140	+0.19%	Mildly bullish
Neutral	116	+0.06%	Baseline
Greed	311	+0.11%	Near baseline
Extreme Greed	96	−0.14%	Contrarian sell

5.8 Robustness Analysis

5.8.1 Statistical Significance

The difference in mean returns between extreme fear (+0.34%) and extreme greed (−0.14%) yields:

- t-statistic: 1.02
- p-value: 0.31

- Effect size (Cohen's d): 0.16 (small)
This is **not statistically significant** at $\alpha = 0.05$.

5.8.2 Out-of-Sample Validation

We split the data into training (2024, $n = 366$) and test (2025–2026, $n = 373$) periods:

- **Training (2024):** Extreme fear – extreme greed = +2.24%
- **Test (2025–26):** Extreme fear – extreme greed = +0.88%

The pattern holds directionally in both periods.

5.8.3 Rolling Window Stability

We test pattern stability using rolling 6-month windows (18 windows total). The contrarian pattern holds in 14 of 18 windows (77.8%).

5.8.4 Backtest with Transaction Costs

A simple contrarian strategy with 20 basis points round-trip costs yields:

- 14 trades over 739 days
- Net return per trade: +1.14%
- Win rate: 57.1%
- Total net return: +15.9% vs buy-and-hold: +106.4%

5.9 Robustness of Spread-Uncertainty Correlation

The core empirical finding—that uncertainty correlates with spreads—is subjected to additional robustness tests.

5.9.1 Granger Causality

We test whether uncertainty *Granger-causes* spread changes (predictive power beyond contemporaneous correlation).

Stationarity. Augmented Dickey-Fuller tests reject the unit root null for CS spreads ($\tau = -23.47$, $p < 0.001$). Realized volatility is marginally stationary ($\tau = -2.73$, $p = 0.07$), typical for persistent financial series.

Lag Selection. Information criteria suggest short lags (BIC: 1–2 days, AIC: 2–4 days depending on specification). We report results for 3-day and 5-day specifications to demonstrate robustness across the plausible range. Results are qualitatively identical for all lags 1–5.

VAR Diagnostics. Pre-test diagnostics confirm Granger test validity (Table 11):

- **Stationarity:** ADF tests reject unit root for both series at $p < 0.001$.
- **Stability:** All eigenvalues of the VAR companion matrix lie inside the unit circle ($\max |\lambda| = 0.97$), satisfying covariance stationarity.
- **Cointegration:** Not applicable—both series are $I(0)$. VAR in levels is appropriate.

Results. The F-statistic for uncertainty $\rightarrow$ Corwin-Schultz spreads is highly significant:

- 3-day lags: $F_{3,705} = 12.79$, $p < 0.001$
- 5-day lags: $F_{5,699} = 7.07$, $p < 0.001$

The reverse direction (spreads $\rightarrow$ uncertainty) is not significant ($F_{3,705} = 0.82$, $p = 0.49$), supporting the directional interpretation: uncertainty *predicts* spreads rather than spreads *predicting* uncertainty.

Table 11: VAR Pre-Test Diagnostics

Test	CS Spreads	Uncertainty	Threshold	Pass
ADF statistic (τ)	-7.65	-4.46	—	—
ADF p -value	< 0.001	< 0.001	< 0.05	✓
<i>Lag Selection (BIC)</i>				
Optimal lag	1		—	—
<i>VAR Stability</i>				
Max eigenvalue	0.97		< 1.0	✓

ADF = Augmented Dickey-Fuller test with constant. Both series stationary at all conventional levels. BIC selects 1-day lag; results robust to lags 1–5.

However, the extended sample (Section 5.10.11) reveals significant bidirectional causality at lags 1–4, suggesting a weaker reverse channel that activates during structural breaks; the forward direction remains dominant by two orders of magnitude in F-statistic.

Caveat. These are linear Granger tests on daily data. The relationship may be nonlinear or operate at higher frequencies. We discuss endogeneity considerations further below.

Table 12 presents the full lag structure for both directions, demonstrating the asymmetry between uncertainty $\rightarrow$ spreads (highly significant) and spreads $\rightarrow$ uncertainty (not significant).

Table 12: Granger Causality: Lag Structure Analysis

Lag	Uncertainty $\rightarrow$ Spreads		Spreads $\rightarrow$ Uncertainty	
	F-stat	p -value	F-stat	p -value
1	31.28	<0.001***	<0.01	0.998
2	17.31	<0.001***	0.42	0.656
3	12.79	<0.001***	0.82	0.485
4	9.13	<0.001***	0.71	0.588
5	7.07	<0.001***	0.63	0.677

SSR-based F-tests for Granger causality, primary sample ($N = 715$). Stationarity confirmed via ADF ($p < 0.01$ for spreads). *** $p < 0.001$. The asymmetry is stark: uncertainty robustly predicts spreads at all lag lengths, while spreads have zero predictive power for uncertainty. On the extended sample the same asymmetry holds with larger F-statistics ($F = 211$ at lag 1; see abstract).

5.9.2 Endogeneity Considerations

We consider whether endogeneity threatens the uncertainty $\rightarrow$ spread interpretation. Two concerns arise: reverse causality (spreads causing uncertainty) and omitted variable bias (common factors driving both).

Theoretical Direction. Reverse causality is theoretically implausible. The mechanism by which wider bid-ask spreads would cause cryptocurrency sentiment models to produce more uncertain outputs is unclear—market makers observe uncertainty and adjust quotes, not the reverse. The Granger causality tests above support this asymmetry empirically.

IV Exploration. We explored instrumental variables using exogenous volatility shocks (VIX jumps, Monday effects, direction changes). These instruments proved weak (first-stage $F = 4.14$, well below the Stock-Yogo threshold of 10), precluding formal causal claims via IV.

Regression Specification. The estimated model is:

$$\text{Spread}_{CS,t} = \alpha + \beta \cdot \text{Uncertainty}_t + \gamma \cdot \text{Volatility}_t + \varepsilon_t \quad (24)$$

where $\text{Spread}_{CS,t}$ is the Corwin-Schultz spread estimate (basis points), Uncertainty_t is total normalized uncertainty, and Volatility_t is Parkinson volatility. The IV specification instruments Uncertainty_t with lagged VIX changes, Monday dummies, and sentiment direction reversals.

However, comparing OLS and IV estimates provides indirect evidence:

- OLS coefficient: 168.36 (Newey-West SE = 34.41, $p < 0.001$)
- IV coefficient: 168.16 (Newey-West SE = 34.47, $p < 0.001$)
- Difference: $< 0.2\%$

The near-identical estimates suggest endogeneity bias is minimal, even though the instruments are too weak to definitively establish causality. Combined with the theoretical implausibility of reverse causation and the Granger asymmetry, we interpret the relationship as uncertainty driving spread-setting behavior.

Table 13 provides the full instrumental variables analysis, including first-stage instrument coefficients and the OLS-IV comparison.

Table 13: Instrumental Variables Analysis: First Stage and OLS-IV Comparison

<i>Panel A: First Stage (Instruments $\rightarrow$ Uncertainty)</i>				
Instrument	Coef.	SE	p-value	
VIX Jump	0.088	0.008	<0.001	***
Monday Dummy	0.005	0.005	0.297	
Uncertainty Lag-1	0.969	0.012	<0.001	***
Direction Change	-0.003	0.003	0.365	
<i>Panel B: OLS vs 2SLS Comparison</i>				
	OLS	2SLS	Diff (%)	
Coefficient	168.36	168.16	0.12%	
SE (HAC)	34.41	34.47		
First-stage F-stat	2781.82 (Strong instruments: $F > 10$)			

Panel A reports first-stage regression of uncertainty on candidate instruments. VIX Jump and lagged uncertainty are highly significant. Panel B compares OLS and IV estimates—the near-identical coefficients (0.12% difference) indicate minimal endogeneity bias. The high F-statistic (> 2700) reflects the dominance of the lagged uncertainty instrument; excluding it yields $F = 4.14$ (weak).

5.9.3 Alternative Spread Estimator: Abdi-Rinaldo

As a robustness check, we implement a close-to-close variant of the [Abdi and Rinaldo \(2017\)](#) spread estimator:

$$S_{AR}^2 = 4 \cdot (c_t - m_t)(c_t - c_{t-1}) \quad (25)$$

where c_t is close price and $m_t = (h_t + l_t)/2$ is the midpoint. The exact [Abdi and Rinaldo \(2017\)](#) estimator uses the two consecutive daily mid-ranges, $S^2 = 4 \mathbb{E}[(c_t - m_t)(c_t - m_{t+1})]$; our variant substitutes the adjacent close for the second mid-range, so unlike the exact estimator it is not fully immune to bid-ask bounce. It nonetheless provides a construction independent of the Corwin-Schultz two-day range decomposition.

Both estimators show positive correlations with uncertainty:

- Corwin-Schultz–Uncertainty: $r = 0.235$ ($p < 0.001$)
- Abdi-Ranaldo–Uncertainty: $r = 0.368$ ($p < 0.001$)

The consistency across estimators rules out artifacts from any single estimation method. The higher AR correlation may reflect that estimator’s greater sensitivity to information asymmetry.

Table 14 consolidates robustness across measurement choices.

Table 14: Robustness to Measurement Choices: Extremity Premium Across Specifications

Specification	Spread	Uncertainty	Coef.	p-val
Baseline	Corwin-Schultz	Composite Index	+0.055	<0.001
Alt Spread	Abdi-Ranaldo	Composite Index	+0.048	0.003
Alt Asset	CS (ETH)	Parkinson Vol	+0.032	<0.001
Vol-Controlled	CS	Residualized Index	+0.040	0.002
<i>Monte Carlo (1,000 Dirichlet weight draws):</i>				
Mean [95% CI]	CS	Random Weights	+0.051	[0.044, 0.059]

All specifications compare extreme greed vs. neutral regimes, controlling for volatility. Coefficients represent incremental uncertainty (normalized units). The extremity premium is preserved across all specifications.

5.9.4 Rolling Window Stability

Using 90-day rolling windows ($N = 18$ windows over 739 days), the uncertainty-spread correlation remains positive in 16 of 18 windows (88.9%). Mean rolling correlation: 0.23 (range: 0.09 to 0.38). The two negative windows occurred during rapid regime transitions (ETF approval period, August 2024 correction).

5.9.5 Regime-Conditional Correlations

The relationship holds across sentiment regimes:

- Bullish regime: $r = 0.21$ ($p < 0.01$, $n = 311$)
- Bearish regime: $r = 0.28$ ($p < 0.01$, $n = 140$)
- Neutral regime: $r = 0.31$ ($p < 0.001$, $n = 116$)

Interestingly, the correlation is strongest in neutral regimes despite neutral regimes having *lower* absolute uncertainty (Section 5.4). This suggests that during neutral periods, the marginal impact of uncertainty on spreads is amplified—perhaps because market makers are more sensitive to information asymmetry when sentiment provides no directional guidance.

5.9.6 HAC Standard Errors

All reported p-values in Table 2 use Newey-West heteroskedasticity and autocorrelation consistent (HAC) standard errors with 5-lag truncation. This addresses potential serial correlation in daily spread data, which would otherwise inflate t-statistics.

5.9.7 Direct Order Book Validation

A potential concern with OHLC-derived spread estimators is that they proxy rather than directly measure transaction costs. To validate, we compare Corwin-Schultz estimates against directly observed spreads from two major exchanges: 90 days of Bybit L2 order book data (5.5 GB of tick-level snapshots) and 61 days of Binance effective spreads calculated from tick-level trades (October 2025–January 2026).

Table 15 reports four key findings. First, the CS estimator correlates positively with actual quoted spreads on both Bybit (Spearman $\rho = 0.41$, $p = 0.001$) and Binance ($\rho = 0.43$, $p = 0.014$), validating that daily OHLC-based estimates capture meaningful variation in transaction costs. Second, Binance and Bybit spreads correlate strongly with each other ($\rho = 0.59$, $p < 0.001$), confirming cross-exchange consistency in liquidity conditions. Third, the level difference—LOB mean of 6.97 bps (Bybit) and 0.36 bps (Binance) versus CS mean of 141.15 bps—reflects that CS captures adverse selection premium beyond mechanical bid-ask. Fourth, the spread–uncertainty relationship holds when using direct LOB spreads: aleatoric uncertainty shows positive correlation with quoted spreads ($\rho = 0.19$, $p = 0.07$), while epistemic shows no relationship ($\rho = 0.04$, $p = 0.71$)—consistent with the paper’s thesis that aleatoric dominates.

Table 15: LOB Validation: Quoted Spreads from Order Book Data

Comparison	N	Pearson ρ	Spearman ρ	p-value
<i>Panel A: Estimator Validation</i>				
LOB Spread vs. CS Spread	61	0.336**	0.412***	0.008 / 0.001
<i>Panel B: Spread–Uncertainty Relationship</i>				
LOB Spread vs. Total Uncertainty	89	0.131	0.207*	0.220 / 0.052
LOB Spread vs. Aleatoric Proxy	89	0.185*	0.193*	0.082 / 0.070
LOB Spread vs. Epistemic Proxy	89	0.040	0.136	0.708 / 0.204
<i>Panel C: Volatility Comparison</i>				
LOB Spread vs. Parkinson Vol	89	0.108	–	0.314
<i>Panel D: Multi-Exchange Validation (Binance)</i>				
Binance LOB vs. CS Spread	33	0.429**	0.425**	0.013 / 0.014
Binance LOB vs. Bybit LOB	43	0.445**	0.591***	0.003 / 0.000

Notes: Direct quoted spreads calculated from Bybit L2 order book snapshots (Oct 2025–Jan 2026, 5.5 GB). Binance effective spreads calculated from tick-level trade data using rolling midpoint methodology (Nov 2025–Jan 2026, 61 days). Panel D validates cross-exchange consistency: Binance and Bybit spreads correlate strongly ($\rho = 0.59$, $p < 0.001$), and both correlate positively with CS estimates. LOB mean spreads: Bybit 6.97 bps, Binance 0.36 bps; CS mean 141.15 bps. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notably, LOB spreads correlate *less* with raw volatility ($r = 0.11$, $p = 0.31$) than with the aleatoric uncertainty proxy ($r = 0.19$, $p = 0.08$). This supports the interpretation that our uncertainty decomposition captures information asymmetry beyond mechanical volatility—precisely the signal relevant for market maker spread-setting.

5.10 Robustness of the Extremity Premium

The core finding—that extreme sentiment regimes exhibit higher uncertainty than neutral regimes, controlling for volatility—is subjected to rigorous validation.

5.10.1 Bootstrap Confidence Intervals

We construct 95% bootstrap confidence intervals on the extreme-versus-neutral uncertainty gap using 10,000 resamples with replacement:

- Observed gap: 0.042
- 95% CI: [0.023, 0.060]
- Bootstrap SE: 0.0095
- z-score: 4.37 ($p < 0.001$)

The confidence interval excludes zero, confirming the extremity premium is statistically robust. The effect is economically meaningful: a 4.2 percentage point increase in uncertainty during extreme regimes corresponds to approximately 15% higher implied spreads relative to neutral periods. We use iid resampling rather than block bootstrap because the extremity premium is computed from regime-aggregated means rather than raw time series, reducing autocorrelation concerns at the aggregated level.

5.10.2 Permutation Test

To test the null hypothesis that regime labels are uninformative about uncertainty, we conduct a permutation test with 10,000 random shuffles of regime assignments:

- Observed extreme–neutral gap: 0.042
- Mean permuted gap: 0.00006 (≈ 0)
- Permuted gap SD: 0.009
- p -value (two-sided): < 0.0001

Zero of 10,000 permutations produced a gap as large as observed, indicating the extremity premium is extremely unlikely under the null.

5.10.3 Within-Volatility-Quintile Analysis

The most demanding test: does the extremity premium survive *mechanical* volatility control? We stratify all 715 complete-case days into volatility quintiles and compare extreme vs. neutral uncertainty *within* each quintile; the per-quintile gaps, effect sizes, and both raw and Holm-adjusted p -values are reported in Table 16.

At raw $\alpha = 0.05$, the pattern appears in 4 of 5 quintiles. However, after Holm-Bonferroni correction for multiple comparisons (5 tests), **only Q3 survives** at the adjusted threshold ($p_{\text{adj}} = 0.024$). Quintiles 1, 2, and 5 show nominally significant raw p -values but do not survive correction ($p_{\text{adj}} \in [0.051, 0.058]$). Quintile 4 fails outright ($p = 0.37$).

The effect sizes (Cohen’s d) are consistently large for Q1, Q2, Q3, and Q5 ($d \in [0.76, 0.86]$), suggesting the non-significance after correction reflects limited sample sizes within strata rather than absent effects. Notably, the highest-volatility quintile (Q5) shows the largest raw effect (+9.3 bps, $d = 0.84$), indicating the extremity premium does not attenuate at high volatility—but this finding requires replication with larger within-quintile samples.

Table 16 provides detailed within-quintile statistics with effect sizes, bootstrap confidence intervals, and both raw and Holm-Bonferroni adjusted p -values.

Table 16: Within-Volatility-Quintile Regime Comparison with Multiple Testing Correction

Vol Q	n_{ext}	n_{neu}	Gap (idx $\times 100$)	95% CI	Cohen’s d	p_{raw}	p_{Holm}	Sig
Q1 (Low)	17	25	+4.27	[1.06, 7.48]	0.81	0.013	0.051	
Q2	11	44	+6.39	[1.22, 11.62]	0.76	0.029	0.058	
Q3	53	15	+5.94	[1.79, 9.66]	0.85	0.005	0.024	**
Q4	30	15	+2.13	[−1.93, 6.27]	0.28	0.374	0.374	
Q5 (High)	59	11	+9.30	[3.36, 14.65]	0.84	0.013	0.051	

Gap = (Extreme mean – Neutral mean) in the normalized uncertainty index, scaled $\times 100$ for readability (index units, not basis points of spread). CI from Welch’s t -test. Cohen’s d = standardized effect size. p_{Holm} = Holm-Bonferroni adjusted for 5 comparisons. ** $p_{\text{adj}} < 0.05$. Extreme regimes combine extreme greed and extreme fear; Neutral = F&G $\in [45, 55]$.

5.10.4 Residual-on-Residual Regression

To address the concern that the baseline uncertainty-spread correlation ($r = 0.24$) reflects mechanical volatility transmission rather than information-driven spread-setting, we implement a residual-on-residual regression:

1. Regress CS spreads on realized volatility $\rightarrow$ spread residuals
2. Regress total uncertainty on realized volatility $\rightarrow$ uncertainty residuals
3. Test correlation of residuals

Results confirm the reviewer’s intuition: the residual correlation drops to $r = 0.043$ ($p = 0.25$), an 82% reduction from the raw correlation. This suggests the baseline uncertainty-spread relationship is largely mechanical—both variables load heavily on volatility.

However, the regime effects survive volatility control. Comparing uncertainty residuals by regime:

- Extreme greed vs. neutral: $t_{713} = 3.84$, $p = 0.0002$
- Extreme fear vs. neutral: $t_{713} = 3.04$, $p = 0.003$

The extremity premium is *not* a mechanical volatility artifact. Even after purging volatility from the uncertainty index, extreme sentiment regimes exhibit significantly higher uncertainty residuals than neutral regimes.




Figure 4: Volatility-matched regime comparison ($N = 715$). Within each volatility quintile, extreme regimes show higher mean uncertainty than neutral regimes; the direction is consistent across all five quintiles, but after Holm-Bonferroni correction for five comparisons only Q3 reaches significance (Table 16). The consistent direction indicates the extremity premium is not merely a mechanical artifact of volatility, while the multiple-testing result signals that per-quintile power is limited—the volatility-controlled effect is carried by the pooled demeaned test, not by any single stratum.

5.10.5 Regime Transition Dynamics

We analyze uncertainty changes during regime transitions using 3-day windows:

- **Enter extreme:** Uncertainty rises by +0.034 ($t_{713} = 2.18$, $p = 0.03$)
- **Exit extreme:** Uncertainty changes by +0.021 ($p = 0.19$, not significant)

- **Enter neutral:** Uncertainty falls by -0.015 ($p = 0.31$, not significant)
- **Exit neutral:** Uncertainty flat ($p = 0.62$)

The significant effect on entering extreme regimes supports the directional interpretation: the transition *into* extremes is associated with uncertainty increases, though this single-direction result ($p = 0.03$) warrants cautious interpretation.

5.10.6 Cross-Asset Validation: Ethereum

We replicate the analysis on ETH/USDT using Parkinson volatility as the uncertainty proxy (the same Fear & Greed Index applies to both assets). Table 17 reports out-of-sample regime coefficients from 739 days of ETH OHLCV data.

Table 17: Cross-asset validation: ETH regime coefficients (N = 739). Coefficients represent Parkinson volatility premium relative to neutral baseline (116 neutral days). The extremity premium replicates: both extreme regimes exhibit significant positive coefficients, while non-extreme regimes show weaker or insignificant effects.

Regime	N	Coefficient	p-value	Sig.
Extreme Fear	76	+0.01153	<0.001	***
Extreme Greed	96	+0.00715	0.001	**
Fear	140	+0.00531	0.030	*
Greed	311	+0.00223	0.215	ns

OLS with HC3 standard errors, volatility-controlled. Neutral (N = 116) is reference category. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, ns = not significant.

The extremity premium generalizes beyond Bitcoin: extreme regimes show significantly elevated volatility relative to neutral, with extreme fear exhibiting the largest effect ($+0.01153$, $p < 0.001$). Non-extreme greed fails to reach significance ($p = 0.22$), consistent with the hypothesis that *extremity*—not direction—drives the premium. Effect size for pooled extreme vs. neutral comparison: Cohen’s $d = 0.48$ (medium). Individual regime effect sizes: extreme fear $d = 0.31$, extreme greed $d = 0.19$, fear $d = 0.05$, greed $d = 0.07$ —the gradient from extreme to non-extreme is consistent with the theoretical mechanism. Post-hoc power analysis indicates adequate power ($1 - \beta > 0.80$) only for extreme fear comparisons; smaller effects in non-extreme regimes may be underpowered. These results suggest a structural feature of cryptocurrency market microstructure rather than an asset-specific anomaly, though replication with larger samples would strengthen inference for moderate-effect regimes.

5.10.7 Out-of-Sample Validation: 2022 Bear Market

A critical test is whether the extremity premium holds in fundamentally different market conditions. The 2024–2026 sample is predominantly bullish (55% greed regimes). We conduct out-of-sample validation using 2022 data—a severe bear market with 93% fear regimes (extreme fear: 57%, fear: 36%). Table 18 compares regime coefficients across market conditions.

The 2022 extreme fear coefficient ($+0.017$) lacks statistical significance, likely due to regime imbalance: extreme fear dominated 57% of the sample (197 days), leaving only 24 neutral observations as the reference category. Post-hoc power analysis confirms the sample is underpowered: with $n_{\text{extreme}} = 197$ and $n_{\text{neutral}} = 24$, power to detect a medium effect ($d = 0.5$) at $\alpha = 0.05$ is approximately 0.42—well below the conventional 0.80 threshold. The directional consistency for extreme fear is suggestive but not statistically confirmed. The fear regime shows directional inconsistency (-0.017 in 2022 vs. $+0.010$ in




Figure 5: Cross-asset validation: BTC vs. ETH regime comparison (BTC: $N = 739$; ETH: $N = 739$). Left panel shows BTC uncertainty by regime (absolute values); right panel shows ETH volatility premium relative to neutral baseline. Both assets exhibit the extremity premium pattern—extreme regimes show elevated uncertainty/volatility relative to neutral. Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

2024), though neither coefficient reaches significance, and this may reflect the severe regime imbalance rather than a genuine asymmetry.

Interpretation. The common factor across market conditions is *extremity*, not direction. Both extreme greed (2024) and extreme fear (2022) generate elevated adverse selection risk because extreme sentiment indicates active disagreement about valuations—regardless of whether that disagreement resolves bullishly or bearishly.

5.10.8 Monte Carlo Weight Robustness

A potential reviewer concern is that the uncertainty decomposition weights are heuristic rather than estimated via GMM or MLE. While formal weight estimation is a natural extension, we test qualitative robustness through Monte Carlo simulation.

We draw 1,000 random weight configurations from a Dirichlet(1,1,1,1) distribution—uniform over the probability simplex—and recompute total uncertainty for each. For each configuration, we test whether the extremity premium (extreme regimes $>$ neutral) is preserved.

Results: The extremity premium holds in 100% of random weight configurations (95% CI: [99.6%, 100%]). No failures were detected across any Dirichlet concentration parameter tested (sparse, dense, component-weighted), confirming the finding is fully robust to weight specification.

GMM Estimation (Supplementary). We also estimate weights via two-step efficient GMM matching four moments: mean uncertainty, uncertainty standard deviation, extreme greed gap, and extreme fear gap. The J-test for overidentifying restrictions rejects the moment conditions ($J = 75548$, $p < 0.001$), indicating model misspecification—the uncertainty index cannot simultaneously match all four target moments regardless of weights. The unusually large J-statistic reflects the heterogeneous scale of target moments: matching mean uncertainty (≈ 0.28) while simultaneously matching regime gaps (≈ 0.04 – 0.06) creates severe tension in the GMM weighting matrix. However, bootstrap inference on individual parameters reveals *weak identification*: 95% confidence intervals are wide (e.g., volatility weight: [0.01, 0.98]), indicating multiple weight specifications are observationally equivalent for any subset of moments. Critically, the heuristic weights fall within all bootstrap confidence regions (z-tests: $p > 0.35$ for all weights). This combination of model rejection and weak identification is informative:

Table 18: Out-of-sample validation: 2022 bear market ($N = 345$) vs. 2024 bull market ($N = 719$). Coefficients represent uncertainty premium relative to neutral baseline. Despite regime imbalance limiting statistical power in 2022, directional consistency for extreme fear is consistent with (though, being non-significant, it does not confirm) the extremity premium mechanism. The 2022 bear-market window ($N = 345$) does not coincide with the 2022 calendar year ($N = 365$) in Table 20, so per-regime counts differ between the two panels and are not nested.

Regime	2022 Bear Market			2024 Bull Market		
	N (%)	Coef.	Sig.	N (%)	Coef.	Sig.
Extreme Fear	197 (57%)	+0.017	ns	76 (11%)	+0.030	**
Fear	124 (36%)	−0.017	ns	140 (19%)	+0.010	ns
Neutral	24 (7%)	(ref)	—	113 (16%)	(ref)	—
Greed/Ext.G.	—	—	—	390 (54%)	mixed	—
Model R^2	0.870			0.840		

** $p < 0.01$, ns = not significant. Regime counts are within each regression sample; the 2024 sample is $N = 719$ because the 20-day rolling volatility warm-up removes the first 20 days of the 739-day window. The 2022 sample’s 93% fear regime concentration leaves only 24 neutral observations as reference, severely limiting statistical power. Extreme fear shows consistent positive sign across both periods.

the extremity premium holds across the identification-equivalent parameter space, and no “optimal” weight exists that would change the qualitative finding.

Interpretation: The 100% preservation rate across random weights is not merely reassuring—it is informative. If the extremity premium disappeared under certain weight configurations, the finding would depend on our decomposition theory. That it survives *all* configurations suggests the extremity premium is a dominant structural feature of the data topology, invariant to how uncertainty is specified. This makes the phenomenon robust even if it simplifies the theoretical interpretation: the episodic/aleatoric decomposition may be narratively useful but is not load-bearing for the core finding.

5.10.9 Variance Decomposition: Volatility vs. Regime Contribution

A natural concern is that the uncertainty-spread relationship is “just volatility.” We decompose variance to isolate the regime contribution.

Model comparison:

- **Model 1 (Volatility only):** $R^2 = 0.755$
- **Model 2 (Volatility + Regimes):** $R^2 = 0.768$
- **Incremental R^2 :** +0.013 (1.3 percentage points)

The F-test for joint significance of regime dummies (after volatility control) yields $F_{4,710} = 10.1$, $p < 0.001$. Regimes are *jointly significant* after accounting for volatility. Notably, the volatility-uncertainty relationship itself varies by regime: interaction tests show significant heteroscedasticity ($F_{4,710} = 11.0$, $p < 0.001$), with the volatility coefficient attenuated in extreme regimes compared to neutral.

Framing: We do not claim uncertainty is orthogonal to volatility—it is not, and the residual correlation confirms this ($r = 0.04$). We claim that extreme regimes exhibit *excess uncertainty* beyond what volatility alone predicts. Volatility explains 75.5% of uncertainty variance; regime membership adds 1.3% incremental explanatory power—modest but statistically significant, indicating regimes capture information beyond mechanical volatility.

5.10.10 Expanding-Window Normalization Robustness

The baseline uncertainty index uses full-sample min-max normalization (Section 3.2), which introduces mild look-ahead bias: each day’s normalized value depends on future minimum and maximum values. While this does not affect regime rankings for explanatory analysis, it could inflate correlation coefficients if applied to predictive settings.

We implement expanding-window normalization as a robustness check: at each time t , we normalize using only data from $[0, t - 1]$. This eliminates look-ahead bias at the cost of early-sample instability (first 30 observations excluded).

Results:

- Correlation between full-sample and expanding-window normalized uncertainty: $r = 0.96$
- Extremity premium under expanding-window: extreme greed gap = +0.35, extreme fear gap = +0.17
- Premium preserved: Yes (both $p < 0.001$, $t > 10$)

The high correlation ($r = 0.96$) indicates the two normalization methods produce nearly identical uncertainty indices. The extremity premium is actually *larger* under expanding-window normalization (greed gap: +0.35 vs +0.25), suggesting our full-sample estimates are conservative. Regime rankings are unchanged, confirming the extremity premium is not a normalization artifact.

Important caveat: Our analysis is explanatory, not predictive. We document that extreme sentiment regimes are *associated with* elevated uncertainty—not that real-time uncertainty forecasts should use these weights. Expanding-window robustness confirms this association is not an artifact of full-sample normalization.

5.10.11 Extended Sample Validation: February 2018 – January 2026

A critical limitation of the main analysis is sample size: 739 days provides limited power for stratified analyses, and findings may be sample-specific. To address this, we extend the analysis to the full Fear & Greed Index history (February 2018–January 2026), yielding **2,896 days**—a 292% increase in sample size.

This extended sample spans multiple market cycles: the 2018 bear market (80% drawdown), 2019 recovery, 2020 COVID crash and recovery, 2021 bull peak (\$69K ATH), 2022 bear market (Luna/3AC/FTX collapses), 2023 recovery, and 2024–2025 bull market. If the extremity premium is a structural feature of cryptocurrency markets, it should persist across these heterogeneous conditions.

Extended Sample Results. Table 19 presents key findings. The extremity premium is dramatically strengthened:

Market Cycle Consistency. The extremity premium appears in 6 of 7 market cycles (Table 20). Effect sizes range from $d = 0.04$ (2019 recovery, when extreme regimes were rare) to $d = 0.48$ (2024–2025 bull market). Only 2023 lacks sufficient extreme-regime observations for testing ($n = 3$). After Holm-Bonferroni correction for 7 tests, the 2024–2025 period survives ($p_{\text{adj}} < 0.001$); other cycles show consistent direction but do not survive correction individually—as expected given the multiplicative penalty of multiple testing across heterogeneous market conditions.

Interpretation. The extended sample strengthens the extremity premium, but the inference must respect the strong serial dependence of daily spreads. The iid t -test $p = 2.7 \times 10^{-14}$ overstates the evidence; the autocorrelation-respecting figure is the block-shuffle permutation $p < 10^{-4}$, and the volatility-controlled magnitude is modest—the pooled volatility-demeaned test gives $t = 3.36$, $p = 0.0008$, Cohen’s $d = 0.21$, not the raw 62-bps gap. Subject to those caveats the effect is directionally robust: it

Table 19: Extended Sample Validation: 739 Days vs. 2,896 Days

Metric	Main Sample (739d)	Extended Sample (2,896d)
<i>Sample Characteristics</i>		
Observations	739	2,896
Date range	Jan 2024–Jan 2026	Feb 2018–Jan 2026
N (extreme regimes)	170	888
N (neutral regimes)	110	457
<i>Extremity Premium</i>		
Gap [†]	4.3–9.3	62.0
95% CI	Often crossed zero	[46.0, 77.5]
Cohen’s <i>d</i>	0.28–0.86 (varied)	0.40
<i>p</i> -value	0.01–0.37	$< 10^{-4}$ (block)
<i>Granger Causality (Unc → Spread)</i>		
F-statistic (lag=1)	31.28	211.30
<i>p</i> -value	< 0.001	< 0.001
<i>Placebo Tests</i>		
Standard permutation <i>p</i>	0.0001	< 0.0001
Block-shuffled <i>p</i>	0.032	< 0.0001
<i>ETH Replication[‡]</i>		
Cohen’s <i>d</i>	0.31	0.31
<i>p</i> -value	0.038	4.4×10^{-9}

Extended sample uses Parkinson volatility as uncertainty proxy (CryptoBERT decomposition unavailable for pre-2024 data). Gap = mean spread in extreme regimes minus mean spread in neutral regimes. Main-sample regime counts are complete-case counts within the $N = 715$ analysis sample (Table 5); the Observations row reports raw sample days. [‡]The two ETH columns are different comparisons that coincidentally share $d \approx 0.31$: the main-sample figure is the extreme-fear-vs-neutral effect ($n = 76$, $d = 0.31$, $p = 0.038$), whereas the extended-sample figure is the *pooled* extreme-vs-neutral effect ($n = 888$, $d = 0.31$, $p = 4.4 \times 10^{-9}$). On the main sample the pooled extreme-vs-neutral ETH effect is larger still ($d = 0.48$, $t = 4.01$, $p < 0.0001$). [†]The Main and Extended “Gap” rows are different dependent variables—Main is the normalized uncertainty index ($\times 100$), Extended is the Corwin–Schultz spread in basis points—so their magnitudes are not directly comparable.

replicates across bull markets (2020, 2021, 2024–25), bear markets (2018, 2022), and on both BTC and ETH. The extended Granger statistic rises with sample size ($F = 211$ vs. 31) but is partly mechanical (Section 5.10), so we do not read the increase as stronger causal evidence.

The within-quintile stratification (not shown) still does not survive Holm-Bonferroni after correction—but this is expected. Stratification mechanically reduces cell sizes ($n = 42$ –301 per quintile), diluting power. The volatility-controlled effect therefore rests on the pooled demeaned test ($p = 0.0008$, $d = 0.21$) rather than on any single stratum. We flag the inferential status of this test plainly: the pooled volatility-demeaned comparison was *not* one of the two pre-specified primary endpoints (Section 1.4); the pre-specified volatility-controlled endpoint—the within-quintile *t*-tests—does *not* survive multiple-testing correction and is, strictly, a null. The pooled demeaned test is therefore a post-hoc, exploratory analysis, and we report it as such: it is the reason we characterise the volatility-controlled effect as “small though robust” rather than as a confirmed primary result. The extended sample strengthens a suggestive finding without making it definitive: the raw gap is large but volatility-uncontrolled, while the controlled effect is small, exploratory, and directionally robust.

Bidirectional Granger Causality in the Extended Sample. While the main sample (739 days) shows clean unidirectional Granger causality—uncertainty predicts spreads but not vice versa (all reverse $p > 0.48$)—the extended sample reveals a more nuanced picture. The reverse direction (spreads $\rightarrow$ uncertainty) becomes significant at lags 1–4 ($p = 0.040, 0.005, 0.003, 0.012$ respectively), though

Table 20: Extremity Premium by Market Cycle (Extended Sample)

Cycle	N	n_{ext}	n_{neu}	Gap (bps)	Cohen's d	p_{raw}	Sig
2018 Bear	311	134	26	+49.5	0.28	0.160	
2019 Recovery	365	80	36	+5.1	0.04	0.855	
2020 COVID+Bull	366	129	67	+47.8	0.31	0.027	
2021 Bull Peak	365	156	32	+62.6	0.30	0.078	
2022 Bear	365	207	23	+40.4	0.27	0.116	
2023 Recovery	365	3	153	—	—	—	(a)
2024–25 Bull	759	179	120	+54.3	0.48	<0.001	**

** $p_{\text{adj}} < 0.05$ after Holm-Bonferroni correction. (a) Insufficient extreme-regime observations.

not at lag 5 ($p = 0.089$). The forward direction (uncertainty $\rightarrow$ spreads) remains the far stronger direction at all lags ($F > 50$). The asymmetry in magnitude persists: the forward F-statistics (51–211) dwarf the reverse (1.9–5.4), indicating the predominant predictive direction is unchanged. We interpret the bidirectional finding as reflecting the extended sample's coverage of major structural breaks—the March 2020 COVID crash, the May 2021 China mining ban, and the November 2022 FTX collapse—during which extreme spread dislocations plausibly fed back into uncertainty measures. In shorter, more homogeneous samples, this feedback channel is absent. This finding qualifies but does not overturn the directional interpretation: uncertainty is the primary driver, with a weaker reverse channel that activates during market crises.

5.10.12 Alternative Volatility Proxy: Deribit DVOL

A potential concern is that the extremity premium reflects mechanical properties of the Fear & Greed Index rather than genuine sentiment-uncertainty dynamics. We test this using an independent volatility-based proxy: the Deribit Volatility Index (DVOL), Bitcoin's crypto-native implied volatility analogous to VIX.

DVOL as Alternative Regime Classification. DVOL is derived from BTC options across multiple strikes and reflects market-implied expectations of 30-day volatility. Unlike F&G (which includes volatility as only one of seven components), DVOL is pure options-derived implied volatility. We create quintile-based regimes using sample percentiles: extreme greed ($\text{DVOL} \leq P_{20} = 42.86\%$), greed ($42.86\% < \text{DVOL} \leq P_{40} = 48.95\%$), neutral ($48.95\% < \text{DVOL} \leq P_{60} = 54.22\%$), fear ($54.22\% < \text{DVOL} \leq P_{80} = 58.85\%$), and extreme fear ($\text{DVOL} > 58.85\%$). The mapping inverts F&G logic: low DVOL indicates complacency, high DVOL indicates panic. Each quintile contains $n = 148$ observations. Concordance between F&G and DVOL extreme classifications is 57.1%—above chance (40%) but far from perfect alignment, confirming they measure related but distinct phenomena.

Results. Using 740 days of matched DVOL and uncertainty data (the DVOL robustness sample spans the full 740-day calendar index—one more than the 739 days with a valid Corwin–Schultz spread used in the spread-based analyses, as DVOL regimes are classified on options data rather than on spreads):

- DVOL range: 33.8%–83.0% (mean: 51.8%)
- Raw pattern: High DVOL regimes show higher uncertainty (extreme fear: 0.493 vs neutral: 0.407)
- After volatility control: No significant regime effects (all $p > 0.06$)
- Volatility explains $R^2 = 0.76$ of uncertainty variance; DVOL regimes add no incremental power

Interpretation. The DVOL-based regimes *do not* replicate the extremity premium after volatility control. This is informative: DVOL is fundamentally a volatility measure (implied rather than realized), so controlling for realized volatility removes its predictive content. In contrast, the F&G extremity premium survives volatility control because F&G captures behavioral sentiment signals—momentum, social media, dominance—that predict uncertainty *beyond* mechanical volatility.

The DVOL non-result strengthens the F&G finding: the extremity premium is not a generic “volatility artifact” that would appear under any volatility-adjacent regime classification. It specifically emerges from sentiment-based regimes that capture information orthogonal to volatility. The 57% concordance between F&G and DVOL extreme classifications confirms they measure related but distinct phenomena.

5.10.13 Momentum Control

Because the Fear & Greed Index is one-quarter market-momentum, and extreme greed coincides with high trailing momentum, the premium could in principle be a momentum artefact rather than a sentiment-extremity effect. We add a trailing 20-day return (strictly lagged, no look-ahead) as a second control alongside volatility, on the extended sample ($N = 1,342$ complete cases; 888 extreme, 454 neutral). The premium’s *effect size* is robust to momentum: residualising uncertainty on volatility and its square yields $d = 0.32$, and adding the momentum control and its square leaves $d = 0.25$ —essentially the paper’s headline volatility-controlled magnitude ($d = 0.21$)—while a nonparametric double-stratification (demeaning uncertainty within joint volatility-quintile $\times$ momentum-tercile bins) gives $d = 0.27$. Under the iid Welch test all three remain strongly significant ($p < 10^{-6}$). Under a serial-dependence-respecting moving-block permutation, however, adding the momentum control moves the premium from significant (volatility-only $p \approx 0.03$) to *marginal* ($p \approx 0.08$ – 0.10). Momentum therefore accounts for part—but not all—of the premium: its direction and magnitude survive, but, consistent with the functional-form sensitivity documented above (the effect is preserved by nonparametric stratification and absorbed by comprehensive parametric controls), its statistical significance is not robust to jointly controlling for momentum and respecting the strong serial dependence of daily spreads. We read this as locating one of the channels—trailing momentum, itself an F&G component—through which sentiment extremity co-moves with uncertainty, rather than as evidence that the premium is spurious.

5.10.14 Threshold Sensitivity Analysis

A natural question is whether the extremity premium depends on our specific threshold choices for defining “extreme” regimes. The baseline uses ≤ 25 (extreme fear) and > 75 (extreme greed). We test robustness to alternative threshold definitions using the ETH cross-validation sample.

Table 21: Threshold Sensitivity: Extremity Premium Across Regime Definitions

Threshold	N_{EF}	N_{EG}	Ext. Fear Coef.	Ext. Greed Coef.	Premium?
15/85 (very strict)	14	9	+0.014**	+0.012	Yes
20/80 (strict)	35	35	+0.015***	+0.008*	Yes
25/75 (baseline)	76	96	+0.012***	+0.007**	Yes
30/70 (loose)	137	254	+0.008***	+0.004	Yes

Regression on ETH Parkinson volatility ($N = 739$). N_{EF} = extreme fear days, N_{EG} = extreme greed days. Neutral ($N = 116$) is reference category. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. “Premium” = both extreme coefficients positive, at least one significant.

The extremity premium is preserved across all threshold definitions tested. Stricter thresholds (15/85, 20/80) produce larger coefficients but fewer extreme observations, reducing statistical power.

Looser thresholds (30/70) dilute the effect but maintain the qualitative pattern. The consistency across specifications confirms that “extremity” as a concept—not our specific operationalization—drives the phenomenon.

5.10.15 Placebo and Identification Tests

A critical reviewer concern is whether the extremity premium is an artifact of volatility clustering or regime persistence. We address this through three placebo tests, reported in Table 22.

Table 22: Placebo and Identification Tests

Test	Observed	Null Mean	Null SD	<i>p</i> -value
<i>Permutation Tests</i>				
Standard Permutation ^a	0.042	0.000	0.009	<0.0001
Block-Shuffled ^b	0.042	−0.000	0.022	0.032
Synthetic AR(1) ^c	0.042	0.001	0.023	0.039
<i>Time-Reversed Causality</i> ($uncertainty_t \sim regime_{t+k}$)				
Forward $k = 1$	$\hat{\beta} = 0.031$	95% CI: [0.017, 0.046]		<0.001
Forward $k = 5$	$\hat{\beta} = 0.019$	95% CI: [0.003, 0.035]		0.009

^aStandard permutation shuffles individual days (10,000 permutations). ^bBlock-shuffled permutation preserves regime autocorrelation by shuffling contiguous blocks. ^cSynthetic AR(1) generates regimes from fitted AR(1) model on F&G values ($\phi_1 = 0.945$). Time-reversed tests regress current uncertainty on *future* regime indicators. All tests use volatility-residualized uncertainty.

Block-Shuffled Permutation. Standard permutation tests shuffle individual days, destroying regime autocorrelation. Since regimes persist (mean block length: 3.7 days), this may inflate significance. Block-shuffled permutation preserves autocorrelation structure by shuffling contiguous regime blocks rather than individual observations. The extremity premium remains significant under block-shuffling ($p = 0.032$), though less extreme than standard permutation.

Synthetic Regime Assignment. We fit an AR(1) model to the Fear & Greed Index ($\hat{\phi}_1 = 0.945$, high persistence) and generate 10,000 synthetic regime sequences. If the extremity premium arose purely from AR(1) regime dynamics, synthetic regimes should produce similar gaps. Instead, the observed gap exceeds 96% of synthetic gaps ($p = 0.039$), indicating the premium is not explained by regime autocorrelation alone.

Time-Reversed Causality. A more stringent test regresses current uncertainty on *future* regime indicators: if future regimes predict current uncertainty, the relationship may be spurious. Results show significant forward coefficients (Table 22, bottom panel), but this reflects high regime persistence—extreme regimes today predict extreme regimes tomorrow. The finding is consistent with, rather than contradictory to, the causal interpretation: regimes persist, and persistent extremity generates persistent excess uncertainty.

Summary. Two of three placebo tests pass at $\alpha = 0.05$. The extremity premium survives block-shuffling and synthetic AR(1) controls, indicating it is not an artifact of regime persistence or volatility clustering. The time-reversed test reflects regime autocorrelation rather than reverse causality.

5.11 ABM Ablation: Testing Mechanistic Assumptions

A legitimate concern with the ABM is that the spread-uncertainty relationship is “baked in” by design: market makers explicitly widen spreads by $\delta \cdot \sigma_{\text{total}}$ (Equations 16–17). We conduct ablation analysis to assess whether the relationship is purely mechanical or emerges from market dynamics.

5.11.1 $\delta = 0$ Counterfactual

Setting $\delta = 0$ removes all direct uncertainty-based spread widening. This ablation tests parameter sensitivity using *synthetic* sentiment signals for computational tractability, rather than replaying the calibrated real-data model. The synthetic generator produces simplified uncertainty dynamics where aleatoric components dominate, yielding correlations that differ in magnitude and sign from the main calibrated model:

- **Baseline** ($\delta = 0$): $\rho(\text{Spread}, U) = -0.017$
- **Default** ($\delta = 1.5$): $\rho(\text{Spread}, U) = -0.021$
- **High sensitivity** ($\delta = 2.5$): $\rho(\text{Spread}, U) = -0.036$

The key finding is not the correlation magnitude (which differs from the calibrated model's $r = 0.64$) but that spreads *respond monotonically* to uncertainty sensitivity δ , confirming the mechanism operates as designed. The weak negative correlations in synthetic runs reflect the aleatoric-dominated uncertainty structure of the generator, not a reversal of the empirically-validated relationship.

5.11.2 Spread Response to δ

As expected, mean spreads increase monotonically with δ :

- $\delta = 0$: 2.19 bps
- $\delta = 1.0$: 2.85 bps
- $\delta = 2.5$: 3.86 bps

This confirms the market maker spread-widening mechanism operates as designed, but the *correlation* with uncertainty is not purely mechanical. Sensitivity analysis with $\pm 40\%$ weight variations in the uncertainty decomposition shows correlations remain stable (range: -0.04 to -0.08), suggesting the qualitative findings are robust to heuristic weight choices.

5.12 ABM Calibration Results

Table 23 compares target statistics with simulation output.

Table 23: Model Calibration: Real vs Simulated Statistics (real data: $N = 739$; simulation: 739 trading days)

Statistic	Real Data	Initial Calibration	Recalibrated
Daily Return Std	2.49%	1.98%	1.14%
Return Kurtosis	2.45	11.16	4.49
Volatility Clustering (lag-1)	0.30	0.80	0.05
Mean Spread	5.0 bps	8.7 bps	4.3 bps

The initial calibration used 3 market makers with wide base spreads (15 bps); the recalibrated version uses 5 market makers with competitive spreads (7 bps) and higher noise trader activity, yielding more realistic spread magnitudes and kurtosis. Volatility clustering decreases because the Mesa ABM lacks explicit GARCH dynamics—the initial calibration's 0.80 reflected regime persistence in a small sample (30 daily returns) rather than true volatility clustering. The recalibrated 0.05 uses session-level aggregation (150+ observations per run) and is below the empirical target, reflecting the ABM's simplified volatility mechanism.

The recalibrated model improves on several dimensions:

1. **Realistic spreads:** Mean spread of 4.3 bps (within the 2–5 bps range for major exchanges), improved from 8.7 bps in the initial calibration

2. **Improved kurtosis:** Return kurtosis of 4.49, within the typical range for cryptocurrency returns (4–8), reduced from 11.16
3. **Volatility clustering:** Lag-1 autocorrelation of 0.05, below the empirical target of 0.30. The Mesa ABM does not implement GARCH-type variance persistence; observed clustering emerges solely from sentiment regime transitions. This is a known limitation of the model specification

5.13 Simulated Method of Moments Validation

Beyond informal comparison of stylized facts, we formally validate the model using Simulated Method of Moments (SMM), following [Grazzini and Richiardi \(2015\)](#). SMM provides a rigorous test of whether the model is consistent with observed market microstructure. For computational tractability, the SMM estimation uses a simplified representation of the key mechanisms—a reduced-form chartist-fundamentalist model with GARCH-like volatility and uncertainty-dependent spreads—rather than the full Mesa ABM described in Section 3.3. This simplified model captures the same economic channels (uncertainty-sensitive spread widening, heterogeneous trader behavior, regime-dependent dynamics) but with tractable parameter estimation. The Mesa ABM’s computational cost per evaluation (requiring order-book simulation with discrete agents) makes direct SMM infeasible within reasonable time constraints. To be explicit: the SMM validates the reduced-form *economic mechanism*—that uncertainty drives spread widening at magnitudes consistent with data—not the full agent specification with 200+ heterogeneous agents, limit-order-book dynamics, and emergent price formation. The J-test ($p = 0.36$) indicates that the parametric uncertainty channel is not rejected by the data; it does not validate the Mesa ABM’s richer agent ecology.

Target moments. We match six market microstructure moments (the empirical and simulated values for each appear in Table 24):

1. Volatility clustering at lag 1: autocorrelation of $|\text{returns}|$, $\text{ACF}(1)$
2. Volatility clustering at lag 5: $\text{ACF}(5)$ of $|\text{returns}|$
3. Volatility clustering at lag 10: $\text{ACF}(10)$ of $|\text{returns}|$
4. Fat tails: return kurtosis
5. Volume autocorrelation: lag-1 autocorrelation of volume
6. Spread-volatility correlation: $\rho(\text{spread}, \sigma)$

Estimation. We minimize the SMM objective function:

$$Q(\theta) = (m_{\text{real}} - m_{\text{sim}}(\theta))'W(m_{\text{real}} - m_{\text{sim}}(\theta)) \quad (26)$$

where m_{real} is the vector of empirical moments, $m_{\text{sim}}(\theta)$ is the average of simulated moments across 100 simulation runs, and $W = I_6$ is the identity weighting matrix (six moments).

Parameter Space. The five estimated parameters (reported with estimates and standard errors in Table 25) are:

- σ_{fund} (fundamental volatility): bounds reflecting plausible daily volatility, controls fundamentalist price innovations
- σ_{noise} (noise trader variance): bounds $[0.001, 0.05]$, controls return volatility
- δ (spread sensitivity): bounds $[0, 5]$, controls spread response to uncertainty
- ρ (volatility persistence): bounds $[0, 1]$, the AR(1) volatility coefficient

- ϕ (chartist fraction): bounds $[0, 1]$, the technical-vs-fundamental trader weight

Optimization. We use the Nelder-Mead simplex algorithm (`scipy.optimize.minimize`) with 10 random restarts to avoid local minima. Each objective evaluation requires 100 simulation runs (739 days each) to reduce Monte Carlo variance. Total computation: approximately 4 hours on a 16-core workstation.

Goodness-of-Fit. The overidentification test follows $J \sim \chi^2_{k-p}$ under the null that the model is correctly specified, where $k = 6$ moments and $p = 5$ parameters yield 1 degree of freedom.

Results.

- J-statistic: 0.83 (identity-weighted $\sum$ of squared moment gaps at the estimate, matching the displayed moment vector and the appendix objective $Q(\theta) = 0.827$)
- Degrees of freedom: 1 (6 moments – 5 parameters)
- p -value: 0.36

Interpretation. The J-test fails to reject the model at any conventional significance level ($p = 0.36 > 0.05$). This provides formal evidence that the simplified model’s mechanisms are consistent with observed market microstructure—it replicates key moments without having them hard-coded into the specification. The result supports the economic plausibility of uncertainty-sensitive spread widening but should be interpreted as validating the reduced-form mechanism rather than the full agent-based specification.

Note on J-tests. This SMM J-test ($p = 0.36$) evaluates whether the simplified model matches market microstructure moments. It differs from the GMM J-test in Section 5.10.8 ($J = 75548$, $p < 0.001$), which evaluates whether the *uncertainty decomposition weights* can simultaneously match four distinct uncertainty moments. The SMM passes; the GMM rejects—these are complementary findings about different model components.

Table 24 provides the detailed moment comparison between empirical data and simulation output.

Table 24: SMM Moment Matching: Empirical vs Simulated

Moment	Empirical	Simulated	Gap
Return ACF(1)	0.118	0.181	+0.063
Return ACF(5)	0.096	0.082	−0.014
Return ACF(10)	0.011	0.039	+0.028
Return Kurtosis	2.23	3.11	+0.89
Volume ACF(1)	0.581	0.773	+0.191
Spread-Vol Corr	0.243	0.241	−0.002
<i>Specification Test</i>			
J-statistic	0.83 (df=1, $p = 0.36$)		

6 moments matched, 5 parameters estimated, $df = 1$. Weighting: identity matrix. $p > 0.05$ indicates model is not rejected. The spread-volatility correlation is matched nearly exactly; higher-order autocorrelations show acceptable deviations. The empirical kurtosis target (2.23) differs from Table 9 (2.45) because the reduced-form SMM model uses a slightly different sample window for moment computation; both values indicate moderate leptokurtosis consistent with cryptocurrency return distributions.

Parameter Estimates. Table 25 reports the calibrated parameter values.

5.14 Regime Distribution

Table 26 presents the distribution of market regimes.

Table 25: SMM Parameter Estimates

Parameter	Estimate	SE	Description
σ_{fund}	0.0307	0.0001	Fundamental volatility
σ_{noise}	0.0190	0.0003	Noise trader variance
δ (spread sensitivity)	0.1792	0.0002	Uncertainty $\rightarrow$ spread scaling
ρ (vol persistence)	0.8480	0.0002	AR(1) volatility coefficient
ϕ (chartist fraction)	0.4727	0.0001	Technical vs fundamental weight

Nelder-Mead optimization with 10 random restarts. SE from Hessian approximation at optimum. All parameters well-identified with tight confidence bounds.

Table 26: Regime Distribution (N=739 days, Jan 2024–Jan 2026)

Regime	Days	%
Greed	311	42.1%
Fear	140	18.9%
Neutral	116	15.7%
Extreme Greed	96	13.0%
Extreme Fear	76	10.3%

6 Discussion

6.1 Interpretation of Core Findings

The Baseline Correlation Is Mechanical; The Regime Effect Is Not. A naive reading of the baseline uncertainty-spread correlation ($r = 0.24$) might suggest direct transmission from sentiment uncertainty to spreads. However, residual-on-residual regression reveals this correlation is largely mechanical—both variables load heavily on realized volatility, and the correlation drops to $r = 0.04$ (not significant) after purging volatility. The finding that survives volatility control is the *regime effect*: extreme sentiment exhibits excess uncertainty beyond what volatility predicts ($t > 3$, $p \leq 0.003$). This is the paper’s central contribution. In the agent-based simulation, sentiment *direction* correlates only weakly with spreads ($r = 0.085$, Table 3), and the empirical regime evidence points the same way: both extreme fear and extreme greed exhibit elevated uncertainty, so it is extremity—not bullishness or bearishness—that predicts the effect.

Aleatoric Uncertainty Carries the Larger Share. On our heuristic decomposition, aleatoric uncertainty accounts for the larger share (81.6%) of the mean index—a split conditional on the decomposition weights and proxy scales rather than estimated, so read as suggestive rather than a precise variance partition (Section 5.4). Taken tentatively, it would be consistent with cryptocurrency markets being inherently noisy rather than simply uncertain due to model limitations, in which case improving sentiment models may have limited impact on spread dynamics if the underlying market information remains inherently ambiguous.

6.2 The Parsimony Principle: Why Simplicity Wins

The finding that elaborate uncertainty decomposition adds negligible value ($\Delta R^2 = 0.003$, Section 7) while a simple macro extremity index succeeds warrants theoretical reflection. We propose three mechanisms:

1. **Signal-to-Noise Inversion.** In traditional sentiment analysis, sophisticated NLP models extract

weak signals from noisy text. The assumption is that model improvement (reducing epistemic uncertainty) enhances signal recovery. However, if the *underlying phenomenon* is inherently stochastic (high aleatoric), model refinement merely measures noise with greater precision. Cryptocurrency sentiment—driven by narratives, memes, and crowd psychology—may be fundamentally stochastic rather than information-revealing.

2. Regime Robustness. The extremity premium depends only on binary classification: extreme vs. neutral sentiment. This coarse categorization is robust to measurement error. A continuous sentiment score from sophisticated NLP ($s \in [-1, 1]$) requires calibration, validation, and uncertainty quantification. A threshold-based regime ($F\&G > 75$ vs. $45 < F\&G < 55$) does not. Coarse categories are less precise but more robust—a favorable trade-off when the underlying signal is noisy.

3. Market Maker Heuristics. Theoretical market-making models assume sophisticated Bayesian updating on continuous information signals. Real market makers may use simpler heuristics: “Is sentiment extreme? If yes, widen spreads.” The Fear & Greed Index—designed for retail consumption—may better approximate the *actual* information set used by market participants than academic NLP models.

Principle. When modeling agents operating in inherently noisy environments, simple observable proxies may outperform complex latent variable extraction. The extremity premium survives because it exploits a robust, observable regime signal rather than attempting to denoise an irreducible stochastic process.

6.3 Interpretation of Core Findings (Continued)

The Extremity Premium. Counter-intuitively, extreme sentiment regimes exhibit the *highest* uncertainty—not neutral regimes. Extreme greed (0.521) and extreme fear (0.403) both exceed neutral (0.303), even after controlling for volatility. When sentiment is directionally intense, informed traders may be exploiting sentiment-driven mispricings, increasing adverse selection risk. The asymmetry between greed and fear effects may reflect leveraged bull market dynamics.

Network Propagation of the Extremity Premium. The replication on Ethereum is theoretically significant. Bitcoin functions as the market’s primary sentiment barometer—the psychological signal that moves first during regime shifts. Ethereum, by contrast, serves as the architectural backbone: smart contracts, DeFi protocols, cross-chain bridges, and decentralized applications depend on ETH infrastructure (Farzulla, 2025a). The extremity premium manifests on both assets: directionally on Bitcoin and with a medium pooled effect on Ethereum ($d = 0.48$, $p < 0.0001$), suggesting the pattern exists at both the sentiment layer *and* the infrastructure layer. We deliberately avoid pairing this ETH effect with the raw BTC pooled $d = 1.06$ as if the two were directly comparable: the BTC 1.06 is a raw, variance-inflated figure computed on the uncertainty index (it conflates between- and within-regime variance, as noted at first mention), whereas the ETH 0.48 is a clean two-group pooled effect on volatility. The like-for-like, volatility-controlled BTC magnitude is smaller (pooled volatility-demeaned $d = 0.21$; the raw extended-sample pooled $d = 0.40$ is larger but is *not* volatility-controlled); the qualitative cross-asset point—that the effect is not Bitcoin-specific—holds on the directional replication and the $r = 0.68$ correlation between BTC and ETH regime coefficients (Section 5.10), not on a numeric 1.06-vs-0.48 comparison. This dual presence may explain why the extremity premium appears to be a structural feature of cryptocurrency markets rather than an asset-specific phenomenon—uncertainty propagates from the sentiment signal (BTC) through the architectural substrate (ETH) to the broader ecosystem.

6.4 Theoretical Implications

The framework extends classic market microstructure models by incorporating sentiment uncertainty decomposition into spread-setting, contributing to the broader complexity economics research programme that deploys agent-based, network, and dynamical systems methods to address challenges where equilibrium approaches fall short (Bednar et al., 2025). The key theoretical insight is that sentiment *uncertainty*—not sentiment level—predicts adverse selection risk.

This finding aligns with emerging evidence that regulatory interventions in cryptocurrency markets affect prices primarily through sentiment channels rather than mechanistic ones. A pilot exploration by the author finds that infrastructure events (FTX collapse, Terra/UST) produce substantially larger spread increases than regulatory events, which may in fact *decrease* spreads—suggesting that regulators cannot directly enforce changes to decentralized market structure, affecting sentiment and expectations rather than infrastructure (Farzulla, 2026b). If regulatory uncertainty influences spreads through *how traders feel* about regulation rather than actual structural changes, the epistemic uncertainty component of our decomposition captures a genuine information channel—market participants’ beliefs about regulatory risk, not the risk itself.

6.5 Practical Implications

1. **Market makers should monitor sentiment uncertainty, not just direction**
2. **Extreme sentiment periods require wider spreads**—the extremity premium suggests maximum adverse selection risk during directional euphoria or panic, not during ambiguity
3. **Improving sentiment models may have limited impact** given aleatoric dominance
4. **Momentum strategies should target regime transitions**—entering extreme regimes predicts uncertainty spikes

7 Limitations

Functional Form Sensitivity. The extremity premium exhibits sensitivity to the choice of volatility control method. In kitchen-sink regressions with comprehensive controls (realized volatility, volatility squared, absolute returns, log volume, day-of-week, month, and year fixed effects), regime coefficients become statistically insignificant (Table 8). This pattern persists even with flexible volatility controls using natural splines with up to 15 degrees of freedom.

However, the stratification-based approach—comparing extreme versus neutral spreads within volatility quintiles—yields different conclusions. The pooled within-quintile test, which compares extreme vs. neutral spreads after demeaning within volatility quintiles, remains highly significant ($t = 3.36$, $p = 0.0008$; pooled volatility-demeaned Cohen’s $d = 0.21$); the raw pooled extreme-vs-neutral comparison yields $d = 0.40$. The extended within-quintile analysis stratifies the full sample ($N = 2,896$) into five quintiles by 20-day annualized close-to-close realized volatility (we use close-to-close returns rather than a range-based measure such as Parkinson volatility, because the Corwin-Schultz spread regressand is itself constructed from the high-low range, so stratifying on a range-based volatility would partially control the outcome with itself). The five quintile-level results are directionally consistent but none individually survives Holm correction:

Vol. quintile	Gap (bps)	Cohen’s d	p_{raw}
Q1 (low)	11.78	0.162	0.301
Q2	21.35	0.206	0.111
Q3	24.37	0.212	0.140
Q4	23.72	0.164	0.185
Q5 (high)	49.35	0.228	0.108

All five gaps are positive (extreme $>$ neutral) and all five quintile-level effect sizes lie in the $d \in [0.16, 0.23]$ range, but after Holm-Bonferroni correction for five comparisons none reaches significance ($p_{\text{Holm}} = 0.54$ for every cell). We therefore summarize the extended within-quintile effect by its median per-stratum Cohen’s $d = 0.21$ (n -weighted mean 0.20); the inferential weight rests on the pooled volatility-demeaned test ($p = 0.0008$), not on any single quintile. We prefer the median as a robust summary that is not dominated by the largest quintile cells. (All figures in this paragraph are from a single pipeline—`full_sample_extension.py`, `full_sample_within_quintile.csv`—with quintiles defined by realized volatility throughout.)

This divergence reflects a fundamental methodological choice. Regression-based controls impose parametric assumptions about the volatility-spread relationship, even with flexible functional forms. Stratification allows arbitrary within-bin relationships and may better capture regime-specific effects that interact nonlinearly with volatility. We report the within-quintile results as our primary specification because: (1) regime effects are inherently categorical, making stratified comparisons more natural; (2) the F&G Index likely captures volatility-regime interactions that parametric models cannot fully absorb; and (3) the stratification approach is conservative—it cannot find effects that do not exist within homogeneous volatility conditions. Readers should interpret the extremity premium as robust to non-parametric volatility control via stratification, but sensitive to regression-based specifications.

Spread Estimator Limitations. The Corwin-Schultz (2012) estimator derives spreads from high-low ranges, introducing several concerns:

1. **Volatility Confound.** CS spreads mechanically embed volatility through the high-low range. Correlating CS spreads with volatility-based uncertainty proxies risks circularity. We partially address this through within-volatility-quintile analysis, which shows the extremity premium persists even after mechanical volatility control.
2. **Serial Dependence.** CS assumes returns are serially independent. In 24/7 cryptocurrency markets with continuous trading, serial dependence may bias spread estimates. While our Granger tests show stationarity, microstructure-level autocorrelation could still distort the estimator.
3. **Alternative Estimators.** Roll (1984) spreads, which use negative autocovariance of returns, provide an alternative. However, Roll estimates are undefined when autocovariance is positive (common in trending markets). We use CS as the primary estimator given its robustness, while acknowledging both approaches have limitations in cryptocurrency contexts.
4. **LOB Validation Scope.** We validated CS spreads against 90 days of Bybit L2 order book data and 61 days of Binance effective spreads (October 2025–January 2026). CS correlates positively with both Bybit ($\rho = 0.41$, $p = 0.001$) and Binance ($\rho = 0.43$, $p = 0.014$) quoted/effective spreads. Critically, Binance and Bybit spreads correlate strongly with each other ($\rho = 0.59$, $p < 0.001$), validating cross-exchange consistency. The validation period remains shorter than the main sample (61–90 days vs. 739 days).

The 90-day LOB validation (Table 15) provides direct evidence that CS estimates capture meaningful transaction cost variation, though the $20\times$ level difference (7 bps quoted vs. 141 bps CS) confirms CS reflects broader adverse selection costs rather than mechanical spreads alone.

Epistemic Uncertainty Adds Little—And This Is a Finding. In supplementary regression analysis, epistemic uncertainty does not add significant explanatory power beyond realized volatility ($p = 0.36$, $\Delta R^2 = 0.003$). Combined with the aleatoric dominance finding (81.6% of the mean index on the heuristic decomposition, Table 4), this suggests cryptocurrency sentiment is **structurally different** from traditional asset information asymmetry—inherently noisy rather than merely uncertain due to incomplete models.

Structural Interpretation. Traditional equity markets exhibit differential analyst coverage creating epistemic heterogeneity: small-cap stocks have sparse information (high epistemic uncertainty), while large-caps have rich fundamental data (low epistemic). Cryptocurrency markets differ fundamentally:

1. **Universal information scarcity:** Even Bitcoin—the most analyzed cryptocurrency—lacks traditional fundamental anchors (earnings, book value, cash flows). All crypto assets operate in a regime of high baseline aleatoric noise.
2. **Homogeneous data availability:** Unlike equity markets with differential analyst coverage, cryptocurrency price data are universally available at sub-second frequency across dozens of exchanges. Epistemic asymmetry is minimal.
3. **Narrative-driven pricing:** Fundamental factors explain minimal cross-sectional return variation in crypto (Farzulla, 2026a). Sentiment and narrative—inherently noisy, irreducible signals—dominate price formation.

Implications for Research. This negative result has practical value: researchers pursuing epistemic uncertainty quantification for cryptocurrency market-making (via improved NLP models, regulatory news parsers, or cross-exchange arbitrage detection) may achieve diminishing returns. The aleatoric dominance finding suggests effort should focus on **regime detection** (identifying extremity) rather than **signal refinement** (reducing epistemic noise). The extremity premium—which emerges from a simple, heuristic macro index—supports this parsimony-first approach.

Mechanical Overlap vs. Incremental Contribution. A valid concern is that the Fear & Greed Index, CS spreads, and DVOL all load on volatility, creating mechanical correlation rather than genuine sentiment transmission. We address this through five complementary tests:

1. **Residual-on-residual regression:** After purging volatility from both CS spreads and the uncertainty index, the residual correlation drops to $r = 0.04$ (not significant). The baseline correlation is mechanical—we concede this explicitly.
2. **Within-volatility-quintile analysis:** Stratifying by volatility quintiles and testing regime effects within each stratum mechanically holds volatility constant. The extremity premium shows its largest gap in the highest quintile (Q5: $+0.093$, $p_{\text{raw}} = 0.013$, $p_{\text{Holm}} = 0.051$; Table 16), where mechanical confounding should be *strongest*.
3. **Variance decomposition:** Regime dummies add 1.3% incremental R^2 after volatility control ($F = 10.1$, $p < 0.001$), indicating regimes capture variation orthogonal to volatility.
4. **Alternative spread estimator:** a close-to-close Abdi-Ranaldo (2017) variant, which uses close-high-low prices (independent of CS’s two-day construction), replicates the extremity premium.
5. **Cross-asset replication:** Ethereum analysis uses Parkinson volatility (range-based, no sentiment

embedding) as the uncertainty proxy. The premium replicates ($d = 0.48$, $p < 0.0001$).

Theoretical Justification for Expected Overlap. We do not claim sentiment is orthogonal to volatility—such a claim would be theoretically suspect. Rational market participants observe volatility and update beliefs accordingly; if volatility reflects information arrival, sentiment *should* correlate with volatility informationally, not spuriously. The critical test is not zero correlation but *regime-conditional heterogeneity*: does the sentiment-uncertainty relationship vary systematically across regimes? The within-quintile analysis confirms it does. In the highest volatility quintile, extreme regimes exhibit +0.093 excess uncertainty (normalized index units; Table 16) relative to neutral regimes with *identical* volatility exposure. This is inconsistent with pure mechanical confounding.

Normalization Uses Full-Sample Statistics. The uncertainty index combines aleatoric and epistemic proxies using min-max normalization over the full sample period. This introduces mild look-ahead bias, as each day’s normalized value depends on future observations. We address this with an expanding-window robustness check (Section 5.10.10): normalizing at each time t using only data from $[0, t - 1]$ yields correlation $r = 0.96$ with the full-sample version, and the extremity premium is preserved under both methods. This confirms the finding is not a normalization artifact. However, for predictive applications, expanding-window or rolling-window normalization would be more appropriate.

High-Volatility Quintile Interpretation. The extremity premium shows its largest raw gap in the highest-volatility quintile (Q5: +0.093, $p_{\text{raw}} = 0.013$; Table 16), which is theoretically consistent with adverse selection intensifying during market stress. However, Q4 fails significance ($p = 0.37$), suggesting the relationship is non-monotonic. The premium appears in calm and crisis regimes but attenuates in intermediate volatility.

Daily Frequency Limitation. All analysis uses daily OHLCV data. Market maker spread-setting occurs at sub-second frequencies; daily aggregation necessarily obscures intraday dynamics. The documented correlations may not hold at trading-relevant timescales.

Heuristic Weight Selection. The aggregation weights ($\gamma_1, \gamma_2, \gamma_3$) for epistemic uncertainty and ($\delta_1, \delta_2, \delta_3, \delta_4$) for aleatoric uncertainty are heuristic rather than calibrated via GMM or MLE. We address this limitation through extensive robustness testing: (1) grid sensitivity across 25 weight configurations confirms 100% ranking preservation; (2) Monte Carlo simulation with 1,000 random weight draws from Dirichlet(1,1,1,1) shows the extremity premium holds in 100% of configurations; (3) GMM estimation reveals weak identification (wide bootstrap CIs), but critically, no estimated weight differs significantly from its heuristic value (all $p > 0.35$). The weak identification is informative: multiple weight specifications are observationally equivalent, and the heuristic falls within the feasible region. We conclude that formal weight estimation does not improve upon the heuristic specification—the extremity premium is parameter-invariant.

ABM Mechanism Non-Emergence. A legitimate methodological concern is that the agent-based model’s spread-uncertainty correlation is not emergent but architected: market makers explicitly incorporate an uncertainty premium term ($\delta \cdot \sigma_{\text{total}}$) in their quoting logic (Equations 16–17). Finding that simulated spreads correlate with uncertainty therefore confirms implementation fidelity rather than validating the economic mechanism. We acknowledge this limitation directly. The ABM serves a more circumscribed purpose: it provides *magnitude calibration* and *consistency check* rather than independent mechanistic validation. Specifically, the SMM procedure validates that when this mechanism operates, a simplified representation of the model jointly reproduces four key market microstructure moments—volatility clustering, excess kurtosis, volume persistence, and spread-volatility correlation—without these being hard-coded (see Section 5.13 for the distinction between the Mesa ABM and the

SMM estimation model). The J-test ($p = 0.36$) indicates the parametric specification is not rejected, meaning the mechanism’s quantitative magnitude ($\delta = 0.18$) is consistent with observed data. What is emergent includes: price trajectories from order flow, volatility clustering from agent feedback loops, fat tails from heterogeneous trader interactions, and regime dynamics. The spread-uncertainty link itself is assumed rather than derived. Setting $\delta = 0$ eliminates the correlation entirely (ablation analysis, Section 5.11), indicating the mechanism is load-bearing—but this demonstrates necessity, not sufficiency. The primary evidence for the uncertainty channel remains empirical (Section 5.1); the ABM is an illustrative device that quantifies mechanism magnitude rather than proves mechanism existence.

Macro/Micro Channel Overlap. The Fear & Greed Index includes a 15% social media component, creating potential double-counting with our micro signal.

Sample Period Bias. The sample period is predominantly bullish (+106%). Results may differ in bear markets.

Proprietary, Volatility-Embedding Sentiment Index. A cleaner test would construct an F&G variant that excludes the 25% volatility component, isolating pure sentiment effects. However, this requires access to the raw component data (social media volume, survey responses, momentum signals), which Alternative.me does not publish. More broadly, the Fear & Greed Index is a *proprietary* product: its exact component construction and aggregation are not fully documented and could be revised without notice, so our results are conditional on a black-box input we cannot fully audit or reconstruct—an epistemic risk we flag explicitly. Two partial substitutes bound the concern. The DVOL-based regime analysis (Section 5.10.12) uses pure implied volatility, which does *not* replicate the extremity premium after volatility control, suggesting the finding is specific to sentiment-based regimes rather than mechanical volatility embedding. The momentum control (Section 5.10.13) removes the trailing-momentum channel—another of F&G’s non-volatility components—and attenuates but does not eliminate the premium. A fully reproducible, text-derived sentiment index that disaggregates sentiment from the index’s embedded volatility and momentum inputs remains the natural next step, and we flag it as an open challenge.

Dependence-Aware Spread Estimators. Recent work develops moment-based spread estimators accounting for serial dependence, including fractional Brownian motion mid-price assumptions and autocorrelated trade arrival (arXiv:2407.17401). Our Corwin-Schultz and Abdi-Rinaldo estimators assume serially independent returns, which may introduce bias under persistent dependence structures. Benchmarking against these advanced estimators, or against intraday quoted spreads over longer validation windows, remains for future work.

Model Limitations. CryptoBERT was trained on 2021–2022 data; domain and temporal shift may affect performance.

Causal Identification. While Granger causality tests support a predictive relationship, we do not claim strict causation. ADF tests confirm stationarity, but Granger causality establishes temporal precedence rather than structural causation. Instrumental variables proved weak (first-stage $F < 10$), though OLS and IV estimates are nearly identical, suggesting minimal endogeneity bias. Importantly, the extended sample reveals bidirectional Granger causality at lags 1–4 (Section 5.10.11), suggesting a weak reverse channel (spreads $\rightarrow$ uncertainty) that emerges during structural breaks. The directional claim is strongest in the main sample; the extended sample supports a predominantly forward relationship with crisis-period feedback.

Simulation Limitations. The initial Mesa ABM calibration exhibited elevated kurtosis (11.16 vs 2.45) and volatility clustering (0.80 vs 0.30 empirical), reflecting agent synchronization and small-sample measurement artifacts respectively. Recalibration with increased market maker competition and

higher trading activity reduced kurtosis to 4.49 and spreads to 4.3 bps, but volatility clustering fell to 0.05—below the empirical target. This reflects the ABM’s lack of explicit GARCH dynamics; volatility persistence emerges from sentiment regimes rather than autoregressive variance. Table 23 reports both calibration runs for transparency.

Generalizability. The extremity premium replicates on Ethereum using Parkinson volatility as the uncertainty proxy. Extreme regimes exhibit 32.8% higher volatility than neutral ($t = 4.01$, $p < 0.0001$, Cohen’s $d = 0.48$). Regime coefficients follow the same ranking as BTC: extreme fear (+0.012, $p < 0.001$), extreme greed (+0.007, $p = 0.001$), and fear (+0.005, $p = 0.03$). The pattern correlation between BTC and ETH regime effects is $r = 0.68$, suggesting the extremity premium is a structural feature of cryptocurrency markets rather than a Bitcoin-specific phenomenon. Testing on smaller altcoins, stablecoins, and DeFi tokens remains for future work.

What This Paper Claims. We make five empirically-supported claims:

1. **Extremity premium exists:** Extreme sentiment regimes (both greed and fear) exhibit elevated uncertainty relative to neutral, after volatility control (extreme greed: +5.5%, extreme fear: +3.9%, both $p < 0.003$).
2. **Intensity dominates direction:** Sentiment *extremity*—not bullishness or bearishness—predicts uncertainty. Direction alone correlates only weakly with spreads in the ABM ($r = 0.085$); empirically, both extreme fear and extreme greed elevate uncertainty, so intensity, not direction, is what predicts it.
3. **Aleatoric dominates epistemic:** on the heuristic decomposition, 81.6% of the mean uncertainty index is aleatoric (inherent noise; Table 4). Epistemic decomposition adds negligible explanatory power ($\Delta R^2 = 0.003$).
4. **Effect replicates:** The extremity premium holds on ETH (Cohen’s $d = 0.48$), shows directional consistency in 2022 bear market data (though not statistically significant due to regime imbalance), and under multiple spread estimators.
5. **Uncertainty predicts spreads:** Granger causality shows uncertainty predicts spreads ($F = 12.79$, $p < 0.001$); the reverse direction is not significant in the main sample ($F = 0.82$, $p = 0.49$) but becomes significant at lags 1–4 in the extended sample, consistent with crisis-period feedback (see Section 5.10.11).

What This Paper Does Not Claim. We explicitly do not claim: (1) trading strategy validity for live use; (2) production readiness; (3) definitive causal proof; (4) optimal parameter calibration; or (5) regulatory compliance. This is exploratory research presenting a framework for uncertainty-aware market microstructure analysis.

8 Conclusion

This paper has documented a spread-uncertainty relationship in cryptocurrency markets and identified the extremity premium—extreme sentiment regimes exhibiting higher uncertainty than neutral periods—as a directionally robust feature that survives volatility control by nonparametric stratification and replicates across assets and market cycles. We are equally explicit about its limits: the pre-specified within-quintile endpoint does not survive Holm-Bonferroni correction, the volatility-controlled effect rests on a post-hoc pooled test, and the regime effect is absorbed by comprehensive parametric controls (Section 7). The premium is best read as directionally robust and economically suggestive, not as a controls-invariant structural constant.

Core Finding. Using Corwin-Schultz spread estimation, we find that uncertainty correlates with bid-ask spreads empirically ($r = 0.24$, $p < 0.0001$). However, residual-on-residual regression reveals this baseline correlation is largely mechanical—both variables load heavily on realized volatility, and the correlation drops to $r = 0.04$ (not significant) after purging volatility. An agent-based model illustrates the proposed mechanism at higher intensity ($r = 0.64$), though this reflects coded behavior rather than emergent dynamics.

The Extremity Premium. Extreme sentiment regimes exhibit higher spreads than neutral periods under nonparametric volatility control. Extended sample validation (February 2018–January 2026, $N = 2,896$ days) is consistent with a directionally robust phenomenon, subject to the functional-form and multiple-testing caveats above:

- **Aggregate effect:** raw gap = 62 bps (95% CI [46, 77], $d = 0.40$; block-shuffle permutation $p < 10^{-4}$); the volatility-controlled effect is smaller but robust (pooled demeaned $p = 0.0008$, $d = 0.21$)
- **Granger causality:** Uncertainty predicts spreads (primary sample $F = 12.79$, $p < 0.001$; the extended-sample $F = 211$ shares a high-low input with the spread measure and is not read as stronger); a weaker reverse channel emerges at lags 1–4 during the extended sample, likely reflecting crisis-period feedback
- **Placebo tests:** Both standard and block-shuffled permutations achieve $p < 0.0001$
- **Cross-asset:** ETH replication yields $p = 4.4 \times 10^{-9}$
- **Market cycles:** Pattern holds across 6 of 7 cycles (2018–2025)

Methodological Contributions. (1) First empirical documentation of regime-conditional uncertainty effects in cryptocurrency market microstructure, validated across 8 years of data; (2) an uncertainty decomposition framework separating epistemic from aleatoric components, with demonstrated weight-robustness across 1,000 Monte Carlo configurations; (3) an ABM implementation that reproduces the extremity premium qualitatively; (4) extended-sample validation across 8 years of data that corroborates the primary-sample finding under autocorrelation-respecting inference.

Future Research. Priority directions include: (1) developing cleaner macro/micro channel separation; (2) testing on altcoins and DeFi tokens; (3) identifying stronger instruments for formal causal identification; (4) high-frequency validation with intraday LOB data; (5) modeling AI-agent trader populations informed by laboratory market evidence (del Rio-Chanona et al., 2025).

Reproducibility

All results are reproducible using public data: Binance API (no key required) and Fear & Greed Index (Alternative.me).

Code:

- ASRI framework: <https://github.com/studiofarzulla/asri>
- Analysis code: <https://github.com/studiofarzulla/sentiment-microstructure-abm>

Acknowledgments

The author thanks Andrew Maksakov for collaboration on the ASRI framework. Infrastructure support provided by Resurrexi Labs.

The author also acknowledges the contribution of frontier AI models—including Claude (Anthropic), Gemini (Google DeepMind), and GPT (OpenAI)—and the laboratories that developed them. These systems served as collaborative research partners throughout the writing process.

References

- Abdi, F., Rinaldo, A. (2017). A simple estimation of bid-ask spreads from daily close, high, and low prices. *Review of Financial Studies*, 30(12), 4437–4480. doi:10.1093/rfs/hhx084
- Adrian, T., Brunnermeier, M.K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741. doi:10.1257/aer.20120555
- Antweiler, W., Frank, M.Z. (2004). Is all that talk just noise? The information content of internet stock message boards. *Journal of Finance*, 59(3), 1259–1294. doi:10.1111/j.1540-6261.2004.00662.x
- Avellaneda, M., Stoikov, S. (2008). High-frequency trading in a limit order book. *Quantitative Finance*, 8(3), 217–224. doi:10.1080/14697680701381228
- Aste, T. (2025). Information filtering networks: Theoretical foundations, generative methodologies, and real-world applications. *arXiv preprint arXiv:2505.03812*.
- Forer, P., Budnick, B., Vivo, P., Aufiero, S., Bartolucci, S., Caccioli, F. (2025). Financial instability transition under heterogeneous investments and portfolio diversification. *arXiv preprint arXiv:2501.19260*.
- Aufiero, S., Bartolucci, S., Caccioli, F., Vivo, P. (2025). Mapping microscopic and systemic risks in TradFi and DeFi: A literature review. *arXiv preprint arXiv:2508.12007*.
- Barucca, P., Lillo, F. (2017). Behind the price: On the role of agent’s reflexivity in financial market microstructure. In *Methods and Finance*, Springer. doi:10.1007/978-3-319-49872-0_3
- Samal, A., Pharasi, H.K., Ramaia, S.J., Kannan, H., Saucan, E., Jost, J., Chakraborti, A. (2021). Network geometry and market instability. *Royal Society Open Science*, 8(2), 201734. doi:10.1098/rsos.201734
- Chen, C., Hafner, C. (2019). Sentiment-induced bubbles in the cryptocurrency market. *Journal of Risk and Financial Management*, 12(2), 53. doi:10.3390/jrfm12020053
- Chen, A., Nguyen, H. (2024). A new perspective on how investor sentiment affects herding behavior in the cryptocurrency market. *Finance Research Letters*, 67, 105737. doi:10.1016/j.frl.2024.105737
- Briola, A., Bartolucci, S., Aste, T. (2025). Deep limit order book forecasting: A microstructural guide. *Quantitative Finance*, 25(7), 1101–1131. doi:10.1080/14697688.2025.2522911
- Briola, A., Bartolucci, S., Aste, T. (2025). HLOB—information persistence and structure in limit order books. *Expert Systems with Applications*, 266, 126078. doi:10.1016/j.eswa.2024.126078
- Bourghelle, D., Jawadi, F., Rozin, P. (2022). Do collective emotions drive bitcoin volatility? A triple regime-switching vector approach. *Journal of Economic Behavior & Organization*, 196, 294–306. doi:10.1016/j.jebo.2022.01.026
- Bouri, E., Molnár, P., Azzi, G., Roubaud, D., Hagfors, L.I. (2017). On the hedge and safe haven properties of Bitcoin. *Finance Research Letters*, 20, 192–198. doi:10.1016/j.frl.2016.09.025
- Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236. doi:10.1080/713665670

- Corwin, S.A., Schultz, P. (2012). A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance*, 67(2), 719–760. doi:10.1111/j.1540-6261.2012.01729.x
- Cont, R., Stoikov, S., Talreja, R. (2010). A stochastic model for order book dynamics. *Operations Research*, 58(3), 549–563. doi:10.1287/opre.1090.0780
- Dias, I., Fernando, J., Fernando, P. (2022). Does investor sentiment predict bitcoin return and volatility? A quantile regression approach. *International Review of Financial Analysis*, 84, 102383. doi:10.1016/j.irfa.2022.102383
- del Rio-Chanona, R.M., Mealy, P., Beguerisse-Diaz, M., Lafond, F., Farmer, J.D. (2021). Occupational mobility and automation: A data-driven network model. *Journal of the Royal Society Interface*, 18(174), 20200898. doi:10.1098/rsif.2020.0898
- del Rio-Chanona, R.M., Pangallo, M., et al. (2025). Can generative AI agents behave like humans? Evidence from laboratory market experiments. *arXiv preprint arXiv:2505.07457*.
- Farzulla, M. (2025). Infrastructure vs regulatory shocks: Asymmetric volatility response in cryptocurrency markets. *Research Square*. doi:10.21203/rs.3.rs-8323026/v1
- Farzulla, M., Maksakov, A. (2026). ASRI: An aggregated systemic risk index for cryptocurrency markets. *arXiv preprint arXiv:2602.03874*.
- Farzulla, M. (2026). Do whitepaper claims predict market behavior? Evidence from cryptocurrency factor analysis. *arXiv preprint arXiv:2601.20336*.
- Farzulla, M. (2026). Same returns, different risks: How cryptocurrency markets process infrastructure vs regulatory shocks. *arXiv preprint arXiv:2602.07046*. doi:10.48550/arXiv.2602.07046
- Farmer, J.D. (2025). Quantitative agent-based models: A promising alternative for macroeconomics. *Oxford Review of Economic Policy*. doi:10.1093/oxrep/graf027
- Gal, Y., Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *ICML*.
- Gurdgiev, C., O’Loughlin, D. (2020). Herding and anchoring in cryptocurrency markets: Investor reaction to fear and uncertainty. *SSRN Electronic Journal*. doi:10.2139/ssrn.3517006
- Grazzini, J., Richiardi, M. (2015). Estimation of ergodic agent-based models by simulated minimum distance. *Journal of Economic Dynamics and Control*, 51, 148–165. doi:10.1016/j.jedc.2014.10.006
- Glosten, L.R., Milgrom, P.R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100. doi:10.1016/0304-405X(85)90044-3
- Golub, A., Keane, J., Poon, S.H. (2012). High frequency trading and mini flash crashes. Working paper.
- Gudgeon, L., Perez, D., Harz, D., Livshits, B., Gervais, A. (2020). The decentralized financial crisis. In *Crypto Valley Conference*.
- Jia, B., Shen, D., Zhang, W. (2022). Extreme sentiment and herding: Evidence from the cryptocurrency market. *Research in International Business and Finance*, 63, 101770. doi:10.1016/j.ribaf.2022.101770

- Kendall, A., Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? In *NeurIPS*.
- Koutmos, D. (2022). Investor sentiment and bitcoin prices. *Review of Quantitative Finance and Accounting*, 60, 1–29. doi:10.1007/s11156-022-01086-4
- Kyriazis, N., Papadamou, S., Tzeremes, P., Corbet, S. (2022). The differential influence of social media sentiment on cryptocurrency returns and volatility during COVID-19. *Quarterly Review of Economics and Finance*, 89, 307–317. doi:10.1016/j.qref.2022.09.004
- Kyle, A.S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335. doi:10.2307/1913210
- LeBaron, B. (2006). Agent-based computational finance. In *Handbook of Computational Economics*, vol. 2, pp. 1187–1233.
- Loughran, T., McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66(1), 35–65. doi:10.1111/j.1540-6261.2010.01625.x
- Naeem, M., Mbarki, I., Shahzad, S. (2021). Predictive role of online investor sentiment for cryptocurrency market: Evidence from happiness and fears. *International Review of Economics & Finance*, 73, 496–514. doi:10.1016/j.iref.2021.01.008
- Makarov, I., Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2), 293–319. doi:10.1016/j.jfineco.2019.07.001
- Paddrik, M., Hayes, R., Todd, A., Yang, S., Scherer, W., Beling, P. (2012). An agent based model of the E-Mini S&P 500 applied to flash crash analysis. In *IEEE CIFER*. doi:10.1109/CIFER.2012.6327800
- Palmer, R.G., Arthur, W.B., Holland, J.H., LeBaron, B., Tayler, P. (1994). Artificial economic life: A simple model of a stockmarket. *Physica D*, 75(1-3), 264–274. doi:10.1016/0167-2789(94)90287-9
- Popoyan, L., Napoletano, M., Roventini, A. (2020). Winter is possibly not coming: Mitigating financial instability in an agent-based model with interbank market. *Journal of Economic Dynamics and Control*, 117, 103937. doi:10.1016/j.jedc.2020.103937
- Rognone, L., Hyde, S., Zhang, S. (2020). News sentiment in the cryptocurrency market: An empirical comparison with Forex. *International Review of Financial Analysis*, 69, 101462. doi:10.1016/j.irfa.2020.101462
- Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39(4), 1127–1139. doi:10.1111/j.1540-6261.1984.tb03897.x
- Tetlock, P.C. (2007). Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance*, 62(3), 1139–1168. doi:10.1111/j.1540-6261.2007.01232.x
- Pangallo, M., del Rio-Chanona, R.M. (2024). Data-driven economic agent-based models. *arXiv preprint arXiv:2412.16591*. Forthcoming in SFI Press.
- Bednar, J., del Rio-Chanona, R.M., Farmer, J.D., Kaszowska-Mojša, J., Lafond, F., Mealy, P., Pangallo, M., Pichler, A. (2025). Complex systems approaches to 21st century challenges: Introduction to the special issue. *Journal of Economic Behavior & Organization*, 235, 107049. doi:10.1016/j.jebo.2025.107049

A Supplementary Tables

This appendix provides additional robustness results referenced in the main text.

A.1 SMM Estimation Details

Table 27 presents the complete SMM diagnostics with individual moment contributions and match quality assessment.

Table 27: SMM Moment Matching: Full Diagnostics

Moment	Target (m_{real})	Simulated (m_{sim})	Gap ($m_{sim} - m_{real}$)	Weight (w_{ii})	Contribution to $Q(\theta)$	Match Quality
Return ACF(1)	0.118	0.181	+0.063	1.0	0.0040	Acceptable
Return ACF(5)	0.096	0.082	−0.014	1.0	0.0002	Excellent
Return ACF(10)	0.011	0.039	+0.028	1.0	0.0008	Acceptable
Return Kurtosis	2.227	3.113	+0.886	1.0	0.7852	Fair
Volume ACF(1)	0.581	0.773	+0.191	1.0	0.0367	Acceptable
Spread-Vol Corr	0.243	0.241	−0.002	1.0	<0.0001	Excellent
Total Objective $Q(\theta)$					0.8269	

Contribution = $(m_{sim,i} - m_{real,i})^2 \times w_{ii}$. Match quality: Excellent (<5% relative error), Acceptable (5–50%), Fair (>50%). The spread-volatility correlation—the key microstructure moment—is matched within <1% error.

Weighting Matrix. We employ the identity weighting matrix $W = I_6$, which weights all moments equally and yields consistent parameter estimates.

Parameter Identification. With $k = 6$ moments and $p = 5$ estimated parameters (σ_{fund} , σ_{noise} , δ , ρ , ϕ), we have 1 degree of freedom for overidentification. Fixed parameters include agent counts (3 market makers, 5 informed, 15 noise traders), inventory aversion ($\alpha = 0.001$), and simulation length (739 days matching empirical sample). Bounds for estimated parameters reflect economically plausible ranges from prior ABM literature.

A.2 Weight Sensitivity Analysis

Table 28 reports the extremity premium across 25 weight configurations, varying γ_1 (aleatoric weight) and δ_1 (epistemic weight) systematically.

A.3 Monte Carlo Weight Robustness

Table 29 summarizes 1,000 Monte Carlo draws from Dirichlet(1,1,1,1), testing whether the extremity premium holds under random weight specifications.

Beyond the binary preservation of the ranking, the *magnitudes* of the regime gaps are also weight-invariant. Across the 1,000 draws the extreme-greed-minus-neutral gap has median +0.257 (interquartile range [0.247, 0.264], minimum +0.203) and the extreme-fear-minus-neutral gap has median +0.111 (interquartile range [0.101, 0.114], minimum +0.090); both gaps are strictly positive in every single draw, with minima well above zero. The extremity premium is therefore not merely sign-robust but magnitude-robust to the weight specification: the finding does not depend on the heuristic weights, which is the relevant robustness question given that the GMM over-identification test formally rejects any single fixed weight vector (Section 5.10.8).

A.4 GMM Weight Estimates

Table 30 reports GMM-estimated weights targeting four moments, with bootstrap standard errors.

Table 28: Weight Sensitivity: Extremity Premium Across 25 Configurations

γ_1	δ_1	Ext. Greed Gap	Ext. Fear Gap	Ranking	Ext. Greed Mean	Neutral Mean
0.20	0.25	+0.249	+0.117	✓	0.540	0.291
0.20	0.35	+0.243	+0.117	✓	0.520	0.277
0.25	0.30	+0.250	+0.117	✓	0.542	0.292
0.30	0.35	+0.250	+0.117	✓	0.544	0.293
0.35	0.35	+0.253	+0.116	✓	0.552	0.299
0.40	0.40	+0.253	+0.116	✓	0.552	0.299
<i>Summary (all 25 configurations):</i>						
Ranking preserved				100% (25/25)		
Min extreme greed gap				+0.239		
Max extreme greed gap				+0.260		

Selected rows from 25-configuration grid. “Ranking preserved” = extreme > neutral in both greed and fear regimes. Full results available in repository. The grid recomputes a simplified two-proxy composite (min-max normalized $\gamma_1 \cdot \text{aleatoric} + \delta_1 \cdot \text{epistemic}$), so baseline-configuration means (0.544/0.293 at $\gamma_1 = 0.30$, $\delta_1 = 0.35$) differ from the headline index means in Table 5 (0.521/0.303); the object of the test is ranking and gap stability across weights, not levels.

Table 29: Monte Carlo Weight Robustness (1,000 Dirichlet Draws)

Metric	Value
Monte Carlo simulations	1,000
Dirichlet concentration	(1, 1, 1, 1)
Extremity premium preserved	100.0%
95% CI lower bound	99.6%
95% CI upper bound	100.0%
Failures (greed < neutral)	0
Failures (fear < neutral)	0

Dirichlet(1,1,1,1) is uniform over the probability simplex, generating maximally random weight combinations. Zero failures across 1,000 draws indicates the extremity premium is parameter-invariant.

Table 30: GMM Weight Estimates with Bootstrap Inference

Parameter	Estimate	SE	Heuristic	Diff	Boot SE	95% CI
$w_{\text{aleatoric}}$	0.010	0.508	0.35	−0.34	0.38	[0.01, 0.98]
$w_{\text{epistemic}}$	0.322	15.94	0.30	+0.02	0.28	[0.01, 0.98]
$w_{\text{volatility}}$	0.668	33.04	0.35	+0.32	0.37	[0.01, 0.98]

GMM targets: mean uncertainty, uncertainty SD, extreme greed gap, extreme fear gap. Wide bootstrap CIs indicate weak identification—multiple weight specifications are observationally equivalent. Heuristic weights fall within all CIs.

A.5 Normalization Robustness

Table 31 compares the extremity premium across normalization methods.

A.6 Variance Decomposition

Table 32 decomposes R^2 by predictor source.

Declarations

Table 31: Normalization Robustness: Full-Sample vs Expanding-Window vs Rolling

Method	N	Greed Gap	Fear Gap	Premium?	<i>p</i> -value
Full-sample	715	+0.250	+0.117	Yes	<0.001
Expanding-window	685	+0.353	+0.173	Yes	<0.001
Rolling (90-day)	686	+0.330	+0.266	Yes	<0.001

All three normalization approaches preserve the extremity premium. Expanding-window and rolling methods produce larger gaps, suggesting full-sample estimates are conservative. Gaps are computed on the same simplified two-proxy recomposition as Table 28, so the full-sample gap (+0.250) differs from the raw extreme-greed–neutral difference of the headline index means in Table 5 ($0.521 - 0.303 = 0.218$).

Table 32: Variance Decomposition: R^2 by Predictor Source

Model	R^2	Incremental R^2
Volatility only	0.755	—
+ Regime dummies	0.768	+0.013
Regimes only (no volatility)	0.198	—

Volatility explains 75.5% of uncertainty variance. Regime dummies add 1.3% incremental R^2 after volatility control. Regimes alone explain only 19.8%, confirming volatility dominates but regimes capture orthogonal variation.

Funding

No external funding was received for this research.

Conflict of Interest

The author declares no conflicts of interest.

Data Availability

The Crypto Fear & Greed Index data are publicly available from Alternative.me. Binance OHLCV data are publicly available via the Binance API. Processed datasets used in the analysis are available from the author upon reasonable request.

Code Availability

The agent-based model implementation, calibration scripts, and figure generation code are available at <https://github.com/studiofarzulla/sentiment-microstructure-abm>.