BETA
Bureau
d'économie
théorique
et appliquée

Documents de travail

« Political Communication and Cryptocurrency Volatility.
Level Effects and Regime Transitions at High Frequency»

Auteurs

Sami Es-snibi, Mehdi Guelmamen

Document de Travail n° 2026 – 27

Août 2026

Bureau d'Économie Théorique et Appliquée BETA

https://www.beta-economics.fr/

Contact :
jaoulgrammare@beta-cnrs.unistra.fr

Université de Strasbourg
cnrs
Université de Lorraine
INRAE
AgroParisTech
Université Haute-Alsace

Political Communication and Cryptocurrency Volatility

Level Effects and Regime Transitions at High Frequency

Sami Es-snibi* Mehdi Guelmamen†

August 18, 2026

Working Paper

Abstract

We ask whether high-frequency political communication shapes the level of conditional volatility, its persistence, or the probability of transiting between volatility states. The question is addressed with hourly Bitcoin prices and the universe of Donald Trump's social media posts from August 2017 to February 2026, scoring the directional tone of market-relevant posts and entering it as an exogenous regressor in single-regime and Markov-switching variance equations. Three results follow. First, the tone coefficient is recovered only in log-linear specifications, where more positive communication coincides with lower conditional volatility and adversarial communication with higher volatility; additive specifications place the estimate at the boundary of the admissible parameter space. Second, a two-state Markov-switching EGARCH identifies a calm and a turbulent state with expected durations of eight and four hours, so that regime alternation is an intraday phenomenon that daily aggregation cannot resolve. Third, the tone effect is proportionally identical across the two states, and a fourteen-month interruption in communication leaves regime alternation essentially unchanged. Political communication thus modulates volatility intensity within whichever state prevails, without governing transitions between states.

JEL Classification: C58; E44; G15; G17.

Keywords: Political communication; cryptocurrency volatility; sentiment analysis; Markov-switching GARCH; regime transitions; high-frequency data

*Université de Lorraine.

†Corresponding author. Université de Lorraine, Université de Strasbourg, AgroParisTech, CNRS, INRAE, BETA. E-mail: mehdi.guelmamen@univ-lorraine.fr.

1 Introduction

Financial markets absorb information from a heterogeneous set of sources, and those sources have been studied unevenly. Scheduled macroeconomic releases and central bank communication have received sustained attention, and their transmission to prices and to volatility is by now well characterised. Political communication delivered directly to the public through social media has not. It arrives without a release calendar, passes through no institutional filter, and carries a directional charge that formal announcements are drafted to avoid. Its market consequences remain correspondingly less well understood.

The question is not only whether such communication matters, but through which margin of the volatility process it operates. A signal may raise or lower the intensity of conditional variance within the market state that prevails when it arrives, or it may alter the probability that the market moves from one state to another. The two margins are distinct objects with distinct implications: the first describes a market absorbing news within a stable dynamic, the second a market whose structure is itself displaced by political events. They also carry different implications for financial stability, since only the second constitutes a source of structural risk. The empirical literature has rarely separated them, in part because doing so requires a volatility model in which states and within-state dynamics are parameterised independently.

Cryptocurrency markets offer a setting in which the question can be posed sharply. Bitcoin is only weakly anchored to conventional cash-flow fundamentals, which makes its volatility a natural candidate for transmission from sentiment, attention and narrative shocks (Liu and Tsyvinski, 2021); it trades continuously, so no observation is lost to market closure; and the second Trump administration, marked by an explicit strategic embrace of digital assets that includes executive initiatives such as the Strategic Bitcoin Reserve and broader efforts at crypto-market deregulation, has made political communication about the asset both frequent and consequential. Our evidence comes from a single communicator, and we return in Section 6 to what that implies for the scope of the conclusions.

The macro-finance literature has long established that political uncertainty affects asset prices and volatility. The economic policy uncertainty index of Baker et al. (2016) is associated with higher stock-market volatility, lower investment and reduced employment growth; Pástor and Veronesi

(2013) provide the theoretical foundation in a general-equilibrium model where political uncertainty generates time-varying risk premia and amplifies return volatility; and Kelly et al. (2016) isolate a political uncertainty premium embedded in option prices around electoral events. These contributions concern relatively structured institutional channels, however: formal policy announcements, legislative decisions, central bank communication. Their informational content is scheduled, filtered and directionally explicit.

Social media communication by elected officials introduces a signal of a different kind, one that is immediate, unfiltered, highly visible and often directionally charged (Tetlock, 2007; Born et al., 2021). For traditional financial assets this channel is already documented: Colonescu (2018) shows that tweet sentiment drives abnormal returns in major US equity indices, and Ge et al. (2019) document statistically significant equity and exchange-rate responses to Trump's tweets. The implications for cryptocurrency markets remain comparatively underexplored, despite the sensitivity of those markets to narrative shocks noted above, and the absence of a scheduled release calendar means that the timing of such communication is itself informative.

Two methodological features of the existing evidence help explain why the channel has proved difficult to detect. First, most studies rely on daily data, which mechanically smooth over the intraday process through which information shocks are absorbed (Tetlock, 2007). Second, political communication is usually reduced to a binary indicator of occurrence, which discards the directional content of the message. A market-supportive message from a sitting president is informationally distinct from a regulatory threat, and conflating the two within a dummy variable will attenuate any detectable effect toward zero.

Given this backdrop, we ask three questions. Does the directional content of political communication co-vary with conditional variance, and does the answer depend on how the variance equation is specified? Is hourly volatility governed by a single stable process or by distinct market states, and if the latter, how long do those states last? And does communication modulate the intensity of volatility within a state, shift the probability of transiting between states, or neither?

We address these questions using an hourly dataset spanning August 2017 to February 2026. We combine the universe of social media posts published by Donald Trump, across Twitter and Truth Social, with high-frequency Bitcoin price data sourced from the Binance API. Market-relevant posts are identified through a keyword taxonomy grounded in the cryptocurrency and political risk literatures, and each post is scored with the VADER compound index (Hutto and Gilbert, 2014), which yields a continuous measure of communication tone. This score enters the conditional variance equation as an exogenous regressor in a family of GARCH-X specifications and, subsequently, in the regime-specific log-variance equations of a two-state Markov-switching EGARCH.

Our study yields three findings, one on each margin of the volatility process. On the level of conditional variance, the tone coefficient is negative and statistically significant in the log-linear EGARCH framework, at -0.043 contemporaneously and -0.060 at the one-hour horizon, and is constrained to the boundary of the admissible parameter space in the additive specifications: whether the channel is detected at all depends on the functional form of the variance equation. On the structure of the volatility process, the two-state MS-EGARCH identifies a calm state and a turbulent state with expected durations of approximately eight and four hours and realised volatilities differing by a factor of nearly three-and-a-half, so that regime alternation occurs roughly twice a day and cannot be resolved by daily aggregation. On the transition margin, the tone effect survives the separation of regimes at -0.037, a likelihood-ratio test cannot reject the restriction that it takes the same value in both states, and a fourteen-month period during which no market-relevant post was published on either platform leaves the regime structure essentially unchanged. Communication operates on the intensity of volatility within a state rather than on the probability of transiting between states.

Our study contributes to the literature in three respects. The first is measurement. Political communication is measured here at the frequency at which markets absorb it and by the direction of its content rather than by its mere occurrence, which is what allows a channel that daily binary designs attenuate toward zero to be recovered. The second is the separation of margins. Because the Markov-switching specification parameterises within-state dynamics and state transitions independently, the level channel is estimated free of regime persistence rather than conflated with it within a single parameter; this connects to the distinction drawn by Filardo (1994) and, to our knowledge, has not been applied to political communication. The third is evidence on the transition margin. The fourteen-month interruption in communication that follows the January 2021 platform suspension removes the signal entirely rather than merely varying its intensity, and provides a natural experiment on the regime channel that does not rely on the functional form of the sentiment term.

These contributions sit within a growing literature on the volatility dynamics of cryptocurrency markets. Cryptocurrencies exhibit features that strain the standard GARCH assumptions, including extreme excess kurtosis, contested asymmetry and pronounced structural instability (Katsiampa, 2017; Chu et al., 2017; Baur et al., 2018). Ardia et al. (2019b) test for regime changes in Bitcoin GARCH volatility and find that a two-state asymmetric Student specification delivers the best in-sample fit, while Caporale and Zekokh (2019) fit more than a thousand GARCH specifications across four cryptocurrencies and conclude that two-regime Markov-switching models produce better value-at-risk forecasts than their single-regime counterparts. Both studies work at daily frequency and neither incorporates an exogenous communication regressor. On the political side, Pankratz (2019) shows that political risk indices predict cryptocurrency return volatility and Smales (2021) documents that investor attention Granger-causes cryptocurrency returns, again at daily frequency and without directional content.

These findings carry practical relevance for regulators and market participants. For regulators concerned with financial stability in digital asset markets, the evidence suggests that the systemic risk posed by political communication is limited: what warrants monitoring is the market dynamics that drive regime change rather than the day-to-day flow of political messaging. For market participants, the directional content of communication is more informative than its mere occurrence, which implies that sentiment-based hedging strategies should condition on tone rather than on the volume of posts.

The remainder of the paper is organised as follows. Section 2 describes the cryptocurrency and social media data. Section 3 presents the volatility specifications and the estimation strategy. Section 4 reports the empirical findings and Section 5 the robustness checks. Section 6 discusses the results and Section 7 concludes.

2 Data

2.1 Cryptocurrency dataset

We source Bitcoin prices from the Binance API (`BTCUSDT` spot pair) at hourly resolution, covering 17 August 2017 to 16 February 2026 and yielding 74,389 hourly observations. Returns are computed as the difference between the logarithms of consecutive hourly closing prices, multiplied by 100.

Unlike regular currency markets, cryptocurrency exchanges operate continuously, so no observation is excluded for weekends or holidays. All timestamps are expressed in Coordinated Universal Time.

The choice of hourly frequency is motivated by two considerations. Theoretically, markets absorb publicly available information within minutes to a few hours, which implies that daily aggregation introduces measurement error in the dependent variable (Tetlock, 2007). Empirically, hourly Bitcoin returns exhibit strong ARCH effects whereas sub-minute data introduce excessive microstructure noise. The one-hour interval thus balances informational richness against statistical tractability.

2.2 Social media posts

We compile the universe of posts published by Donald Trump between 2015 and 2026, encompassing his Twitter activity until his suspension in January 2021 and, from Truth Social's public launch onwards, his posts on that platform. The longitudinal dataset covers more than 90,000 individual messages over twelve years.

Because the raw corpus contains posts unrelated to digital assets or macroeconomic policy, a systematic filtering strategy is required. Following Tetlock (2007) and Hassan et al. (2019), we construct a keyword dictionary grounded in the cryptocurrency market and political risk literatures, organised into seven economically motivated categories (Table 1). The first five categories capture direct references to the crypto ecosystem. The regulatory category isolates posts referencing enforcement authorities such as the SEC and CFTC and actions such as bans and crackdowns that may shift the perceived regulatory environment. The macroeconomic category captures broader policy signals, including trade tensions, monetary policy and inflation, that transmit to Bitcoin through risk-premium and capital-flow channels, consistent with the frameworks of Pástor and Veronesi (2013) and Liu and Tsyvinski (2021).

The breadth of the taxonomy is deliberate and defines the object of study. The set we wish to capture is not communication about cryptocurrency, but political communication capable of moving the cryptocurrency market, and the two are not the same. A statement on tariffs, on the dollar or on the conduct of monetary policy carries no reference to digital assets and yet reaches Bitcoin through the risk-premium and capital-flow channels described by Pástor and Veronesi (2013) and Liu and Tsyvinski (2021); conversely, a passing mention of blockchain in an unrelated context need not carry market content at all. The taxonomy is therefore built around potential market relevance rather than around subject matter, and combines direct references to digital assets with the regulatory and macro-policy vocabulary through which political communication transmits to the asset indirectly. We refer to the selected messages as market-relevant posts throughout, and the sentiment score constructed below measures the tone of that communication as a whole.

Table 1: Keyword taxonomy for identifying market-relevant posts

Category	Keywords
Cryptocurrency terms	`bitcoin, btc, crypto, cryptocurrency, blockchain`
Digital assets	`digital currency, digital asset`
Major tokens	`eth, ethereum, stablecoin, tether, usdt`
Exchange platforms	`coinbase, binance`
Mining infrastructure	`mining, miner, hashrate`
Regulatory terminology	`sec, cftc, regulation, ban, banned, crackdown`
Macroeconomic signals	`tariff, sanction, china, inflation, interest rate, fed, dollar`

Notes: A post is classified as market-relevant if it contains at least one whole-word match with the above taxonomy, matching being case-insensitive. The first five categories capture direct references to digital assets; the last two capture regulatory and macro-policy content that transmits to Bitcoin indirectly.

2.3 Sentiment measurement

We score each post with the VADER compound index (Hutto and Gilbert, 2014), $S_t \in [-1, +1]$, a lexicon-based measure of valence developed for social media text. Hourly windows containing more than one market-relevant post are aggregated by simple averaging, and S_t is set to zero in hours without such a post. The aggregation rule binds for a minority of observations only, since the average post window contains 1.5 posts.

Two properties motivate this instrument for the present corpus. VADER was constructed and validated on microblog-like text rather than on editorial prose, and it supplements its lexicon with rules governing amplification through punctuation, capitalisation, degree modifiers and negation.[1] This second feature matters because the communication we study relies heavily on such devices: a bag-of-words dictionary such as that of Loughran and McDonald (2011) assigns identical scores to a message and to its capitalised, exclamation-laden variant, discarding precisely the variation

[1] Benchmarked against eleven alternatives including LIWC, ANEW, the General Inquirer and SentiWordNet, VADER attains a correlation of 0.881 with aggregated human ground truth on tweets, against 0.888 for an individual annotator (Hutto and Gilbert, 2014).

in intensity we seek to measure. VADER is also deterministic and estimates no parameter on the sample, which rules out leakage of later-period information into earlier scores over a nine-year window. The instrument has been applied to cryptocurrency data at both daily and hourly frequency (Kraaijeveld and De Smedt, 2020).

Two measurement caveats apply. VADER is a general-purpose lexicon, known to misread irony and sarcasm and lacking the crypto-native vocabulary whose valence is unambiguous to market participants; and our results rest on a single sentiment instrument and a single keyword taxonomy. We treat the VADER score as a noisy but unbiased proxy for tone, so that measurement error of this kind is likely to attenuate the estimated coefficient toward zero.

2.4 Descriptive statistics

Table 2 summarises the post corpus. Of 74,389 hourly windows, 4,938 (6.64%) contain at least one market-relevant post, corresponding to 7,467 individual posts. The sentiment distribution is directionally charged: positive-sentiment hours account for 54.0% of post windows with a mean VADER score of +0.688, while negative-sentiment hours account for 33.9% with a mean score of −0.622. Only 12.2% carry a neutral tone. This asymmetry motivates the use of a continuous score in place of a binary indicator.

Yearly coverage varies substantially. Crypto-relevant posts occupy 3.6% of hourly windows in 2017 and 6.1% in 2018, reflecting the limited political salience of cryptocurrencies prior to their mainstreaming. Coverage falls to 0.1% in 2021, corresponding to Trump's indefinite Twitter suspension following the January 6 Capitol events, and recovers to 3.9% in 2022, a year in which posting resumes only in April. It reaches 10.6% in 2025, coinciding with the second presidential campaign and an increasingly explicit pro-crypto policy agenda. We exploit the 2021–2022 interruption as a natural experiment in Section 4.

Table 3 reports distributional statistics for hourly log-returns over the full sample and by posting regime. The return distribution displays the hallmarks of cryptocurrency data: a mean of 0.004% consistent with near-zero drift at this frequency, a standard deviation of 0.784%, and an excess kurtosis of 35.67. These features motivate the use of Student-t distributed innovations in all subsequent specifications.

Two contrasts across posting regimes are worth noting. Mean log-returns are modestly but

Table 2: Trump posts and market-relevant content

Panel A. Post frequency		
Total hourly windows	74,389	obs.
Windows with market-relevant post	4,938	obs.
Total market-relevant posts	7,467	count
Avg. posts per post window	1.51	count
Share of post windows	6.64	%
Panel B. Sentiment of post windows		
	Windows	Share (%)
Positive ($S_t > 0.05$)	2,665	54.0
Neutral ($-0.05 \leq S_t \leq 0.05$)	601	12.2
Negative ($S_t < -0.05$)	1,672	33.9

Notes: Sentiment regimes defined by the VADER compound score S_t, using the ± 0.05 classification thresholds proposed by Hutto and Gilbert (2014).

significantly higher during post windows, at 0.024% against 0.002% ($p = 0.046$), while excess kurtosis falls from 36.44 to 22.45, suggesting lower unconditional tail risk during post hours. Average Bitcoin prices are also substantially higher during post windows, at \$51,240 against \$34,981 (Welch $t = -14.80$, $p < 0.001$), which is consistent with the concentration of market-relevant communication during bull-market phases. This selection effect is not corrected by the econometric strategy below, which models the conditional variance rather than price levels, and we return to it in Section 6. Finally, mean returns across the positive and negative sentiment regimes are statistically indistinguishable ($t = 0.150$, $p = 0.881$), which indicates that the directional content of communication does not translate mechanically into return-level differentials and sharpens the focus on the second moment.

3 Methodology

3.1 Conditional mean

Let P_t denote the Bitcoin closing price at the end of hour t and $r_t = 100 \times (\ln P_t - \ln P_{t-1})$ the log-return. Although cryptocurrency returns are typically close to serially uncorrelated in levels, weak linear dependence can contaminate inference on second-moment dynamics. We therefore filter

Table 3: Distributional statistics of hourly Bitcoin log-returns (%)

Statistic	All hours	No post	Post
Mean	0.0037	0.0023	0.0244
Maximum	16.0280	16.0280	8.9576
Minimum	−20.103	−20.103	−9.518
Std. deviation	0.7839	0.7861	0.7518
Skewness	−0.467	—	—
Excess kurtosis	35.671	36.440	22.447
N	74,388	69,450	4,938
Welch t-test on mean returns, post vs. no post			
t-statistic		1.997**	
p-value		0.046	

Notes: Returns are computed as $100 \times [\ln(P_t) - \ln(P_{t-1})]$. A window is classified as a post window if it contains at least one market-relevant post. ** denotes significance at the 5% level. The Jarque–Bera statistic for the full sample is 3,946,532 ($p < 0.001$).

the series with an ARMA(p, q) process,

$$r_t = \sum_{i=1}^{p} \phi_i r_{t-i} + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j} + \varepsilon_t, \qquad \mathbb{E}[\varepsilon_t \mid \mathcal{F}_{t-1}] = 0, \tag{1}$$

where $\mathcal{F}_{t-1}$ denotes the information set at time $t-1$ and the innovation is decomposed as $\varepsilon_t = \sigma_t z_t$. Here σ_t^2 is the variance of ε_t conditional on $\mathcal{F}_{t-1}$ augmented with any same-hour exogenous covariate entering the variance equation, and z_t follows a standardised Student-t distribution with ν degrees of freedom to accommodate heavy tails. Among specifications with $p, q \in \{0, 1, 2\}$, the ARMA$(1, 1)$ without intercept minimises the BIC, and its residuals retain pronounced ARCH effects, with an ARCH-LM statistic of 9,555 at lag 24 ($p < 0.001$).[2]

3.2 Single-regime specifications

We begin with a GARCH(1,1) augmented by the VADER sentiment score as an exogenous regressor,

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 + \delta S_t, \tag{2}$$

[2] Augmented Dickey–Fuller, Phillips–Perron and KPSS tests confirm that hourly log-returns are stationary in levels without drift or deterministic trend. Detailed results are available on request.

subject to $\omega > 0$, $\alpha_1 \geq 0$, $\beta_1 \geq 0$ and $\alpha_1 + \beta_1 < 1$. The coefficient δ is the parameter of interest, with a negative estimate indicating that more positive communication coincides with lower conditional variance. The sum $\alpha_1 + \beta_1$ measures shock persistence.

The exponential GARCH of Nelson (1991) models the log-conditional variance, which guarantees non-negativity without explicit parameter constraints and allows for differential responses to positive and negative innovations:

$$\ln \sigma_t^2 = \omega + \beta_1 \ln \sigma_{t-1}^2 + \alpha_1 z_{t-1} + \gamma_1 \left(|z_{t-1}| - \mathbb{E}\,|z_{t-1}|\right) + \delta S_t, \qquad z_{t-1} \equiv \frac{\varepsilon_{t-1}}{\sigma_{t-1}}. \tag{3}$$

The sign parameter α_1 captures asymmetry: $\alpha_1 < 0$ implies that negative innovations raise conditional variance by more than positive innovations of equal magnitude, the standard leverage pattern. The parameter γ_1 governs the magnitude effect, with the centring term ensuring that the innovation enters in deviation from its mean. In the log-linear framework, $\delta < 0$ implies that a positive unit increase in the score reduces $\ln \sigma_t^2$ and therefore compresses conditional variance.

We complement these with two further specifications. The Glosten et al. (1993) model introduces a threshold indicator $I_{t-1} = \mathbf{1}[\varepsilon_{t-1} < 0]$ in an additive variance structure,

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \gamma_1 I_{t-1} \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 + \delta S_t, \tag{4}$$

which allows us to verify whether a common sentiment effect survives once the asymmetric response to negative innovations is controlled for. The sentiment term enters additively and without interaction, so that $\partial \sigma_t^2 / \partial S_t = \delta$ irrespective of the sign of ε_{t-1}; the specification does not test whether the transmission of tone is itself sign-dependent. The integrated GARCH restricts $\alpha_1 + \beta_1 = 1$, implying permanent volatility shocks, which is relevant given the near-unit-root behaviour frequently documented in high-frequency cryptocurrency data. Finally, to assess whether single-regime specifications are structurally adequate, we apply the Nyblom (1989) parameter-constancy test, whose joint null hypothesis is that all parameters are time-invariant against an alternative allowing a random walk in the coefficient vector.

3.3 Markov-switching EGARCH

Let $s_t \in \{1, 2\}$ denote an unobserved regime indicator governed by a first-order homogeneous Markov chain with constant transition probabilities $P_{ij} = \Pr(s_t = j \mid s_{t-1} = i)$, so that the expected duration of regime i is $(1 - P_{ii})^{-1}$. Conditional on regime k, the log-conditional variance follows an EGARCH(1,1):

$$\ln \sigma^2_{t,k} = \alpha_{0,k} + \beta_k \ln \sigma^2_{t-1,k} + \alpha_{1,k} \left(|z_{t-1,k}| - \mathbb{E}\,|z_k| \right) + \alpha_{2,k}\, z_{t-1,k} + \delta_k S_t, \qquad z_{t-1,k} \equiv \frac{\varepsilon_{t-1}}{\sigma_{t-1,k}}, \quad (5)$$

which allows the intercept, magnitude coefficient, asymmetry parameter and persistence to differ freely across regimes.[3] We retain the exponential form for the regime-specific variance equations on the evidence of the single-regime results: it is the only specification in which the sentiment coefficient is recovered, and it attains the lowest information criteria among the four reported in Table 4. The switching estimates below are therefore conditional on this functional form. The innovation distribution within each regime is Student-t with regime-specific degrees of freedom ν_k, which provides a cleaner separation between heavy-tailed shocks and structural shifts in the volatility level.

This specification is what allows the two margins of interest to be separated. The first is within-regime persistence, governed by β_k and reflecting the decay rate of shocks within a stable market state. The second is regime persistence, governed by P_{kk} and reflecting the stickiness of the states themselves. The distinction matters both substantively, since a communication channel operating on one margin has different implications from one operating on the other, and econometrically, since a high β in a single-regime model may conflate genuine shock persistence with unmodelled structural breaks (Lamoureux and Lastrapes, 1990).

We follow standard practice in this literature and fix the number of states at two. Ardia et al. (2019b) and Caporale and Zekokh (2019) both report that two-state Markov-switching specifications deliver the best in-sample fit for Bitcoin returns and outperform their single-regime counterparts out of sample, without formally selecting the number of states. A standard likelihood-ratio test is unavailable here, since the second-regime parameters and the transition probabilities are unidentified

[3] The two packages we use order the asymmetry terms differently. In `rugarch` the coefficient on the signed innovation is labelled α and that on the centred absolute innovation γ, whereas `MSGARCH` labels them α_2 and α_1 respectively. We report each package's native names and indicate in the table notes which coefficient governs the sign effect.

under the null of a single regime (Hansen, 1992).

The specification remains tractable because the K variance processes are updated in parallel rather than through a collapsed aggregate (Haas et al., 2004): each recursion depends only on ε_{t-1} and on its own lagged variance, so no path dependence arises and an exogenous regressor may enter each state without rendering the likelihood infeasible. This is the structure implemented by Ardia et al. (2019a) and the one we adopt.

3.4 Estimation and inference

The Markov-switching model imposes no structure on the conditional mean, so returns are pre-filtered by the ARMA$(1, 1)$ of Equation (1) and the switching specifications are estimated on the resulting residual sequence with the mean parameters held fixed. Information criteria are therefore comparable within the switching block but not across it and the single-regime block, which is obtained from joint estimation of the mean and variance equations. Because the likelihood of Markov-switching GARCH models is known to be multimodal (Ardia et al., 2019a), estimation is repeated from a set of perturbed starting values and the highest-likelihood solution is retained.

Inference on δ_k relies on likelihood-ratio rather than Wald statistics. Wald inference requires a well-conditioned Hessian, which is not guaranteed in this class of models: the near-unit persistence typical of high-frequency volatility series flattens the likelihood in the direction of β_k, and the degrees-of-freedom parameter is weakly identified whenever a regime approaches Gaussianity. Likelihood-ratio statistics are invariant to reparameterisation and better behaved near a boundary. We test three restrictions: the joint restriction $\delta_1 = \delta_2 = 0$, which asks whether communication is associated with conditional variance at all; the individual restrictions $\delta_k = 0$, which locate the association across states; and the equality $\delta_1 = \delta_2$, which asks whether the effect operates with the same intensity in both states.

Each specification is estimated with the contemporaneous score S_t and with its one-hour lag S_{t-1}. Only the latter is $\mathcal{F}_{t-1}$-measurable, since S_t aggregates posts published during hour t; the contemporaneous specification therefore conditions on a realised same-hour covariate rather than delivering a one-step-ahead variance forecast, and we place the interpretive weight on the lagged results accordingly.

4 Empirical results

4.1 Single-regime estimates

Table 4 reports maximum-likelihood estimates for the four GARCH-family specifications, with the sentiment coefficient shown for both the contemporaneous and the lagged score. The conditional mean is an ARMA$(1,1)$ throughout.

Across symmetric specifications, the ARCH parameter ranges from 0.087 in the IGARCH to 0.134 in the GARCH-X, indicating substantial short-term shock amplification, while $\hat{\beta}_1 \approx 0.908$ implies slow volatility decay. The sum $\hat{\alpha}_1 + \hat{\beta}_1 = 1.042$ in the baseline GARCH-X exceeds unity. Lamoureux and Lastrapes (1990) attribute this pattern theoretically to the conflation of genuine shock persistence with unmodelled regime shifts, which pre-announces the need for a switching specification. A further symptom appears in the tail parameter: the single-regime degrees of freedom, at $\hat{\nu} \approx 2.71$, lie below four, so the standardised innovation has no finite fourth moment.

The sentiment coefficient yields a sharply differentiated picture. In the symmetric GARCH-X, the IGARCH-X and the asymmetric GJR-GARCH-X, $\hat{\delta}$ is numerically zero under both timings. This is not a rejection: the estimates lie at the lower bound of the admissible parameter space imposed by the estimation software for exogenous variance regressors in additive specifications, and re-estimating with $-S_t$ in place of S_t, which renders a negative effect on S_t attainable, leaves the estimate at zero.

In the EGARCH-X the coefficient is negative and highly significant in both the contemporaneous version ($\hat{\delta} = -0.043$, $p = 0.009$) and the lagged version ($\hat{\delta} = -0.060$, $p < 0.001$). Hours of more positive political communication are associated with lower hourly conditional variance, and hours of adversarial or regulatory rhetoric with higher conditional variance. We interpret this association cautiously: because posting concentrates in bull-market phases (Table 3), the estimate conditions on, rather than identifies, the causal effect of communication. In the log-linear framework δ scales the log-variance, so the implied response is a constant proportional change in conditional variance, whereas the additive specifications impose a constant absolute change. We note that this contrast does not by itself isolate the functional form of the sentiment term, since the two families differ in the whole variance equation and not only in the way S_t enters it. The finding is nonetheless consistent with Silva et al. (2024), who report that within the class of asymmetric specifications the

EGARCH systematically outperforms alternative formulations for cryptocurrencies.

Two features of the EGARCH estimates merit attention. First, the effect is present at both horizons and is somewhat larger at the one-hour lag than contemporaneously. Part of this gap is mechanical: a post published mid-hour affects only the remainder of the current window but the entirety of the next. Either way, the association is concentrated within a two-hour window that daily aggregation would average away. Second, to contextualise the magnitude, the VADER score has a standard deviation of approximately 0.55 among post hours, so a one-standard-deviation increase reduces $\ln \sigma_t^2$ by approximately 0.033 at the one-hour horizon. The effect is modest in absolute terms but precisely estimated, and directionally consistent with the prediction of Pástor and Veronesi (2013) that policy-supportive signals lower uncertainty premia.

The bottom panel of Table 4 reports Nyblom parameter-constancy statistics. The joint statistics range from 19.29 to 27.94 across specifications and timings, systematically exceeding the 1%-level critical values of 2.12 to 2.59. This rejection is consistent with parameter instability over the sample. We read it as suggestive rather than decisive: with $T = 74{,}388$ observations the test has considerable power against even minor departures from constancy, and it is uninformative about the form the instability takes. It motivates a specification in which parameters may vary, without by itself pointing to a Markov-switching one. Examining individual statistics, the persistence parameter is the most unstable in symmetric models, at 5.63 in the GARCH-X, while in the EGARCH it is the magnitude parameter, at 8.47 contemporaneously and 8.60 at the lag.

4.2 Regime identification

Table 5 reports estimates for the two-state MS-EGARCH-X. The regimes are sharply distinguished. The log-variance intercept is negative in Regime 1 ($\hat{\alpha}_{0,1} = -0.026$) and positive in Regime 2 ($\hat{\alpha}_{0,2} = +0.037$). Because $\hat{\beta}_1$ lies at the integrated boundary these intercepts do not map into unconditional variances and cannot be used to rank the states directly; we therefore label the regimes on the basis of realised volatility within each state, at 0.428% and 1.465% respectively, a factor of nearly three-and-a-half.

Tail behaviour differs by a wider margin still. The degrees-of-freedom parameter is $\hat{\nu}_1 = 7.78$ against $\hat{\nu}_2 = 3.38$, so that the regime-specific standardised innovation has a finite fourth moment in the calm state and none in the turbulent one. We stress that this is a statement about the

Table 4: GARCH-family models with the VADER sentiment score

Parameter	GARCH-X	IGARCH-X	GJR-GARCH-X	EGARCH-X
Mean equation				
$\hat{\phi}$ (AR1)	0.491^{***}	0.490^{***}	0.489^{***}	0.468^{***}
$\hat{\theta}$ (MA1)	-0.575^{***}	-0.572^{***}	-0.573^{***}	-0.552^{***}
Variance equation				
$\hat{\omega}$	0.003^{***}	0.003^{***}	0.003^{***}	-0.006^{***}
$\hat{\alpha}_1$	0.134^{***}	0.087^{***}	0.122^{***}	-0.016^{***}
$\hat{\beta}_1$	0.908^{***}	0.913	0.908^{***}	0.992^{***}
$\hat{\gamma}_1$	—	—	0.025^{***}	0.213^{***}
$\hat{\nu}$	2.711^{***}	3.154^{***}	2.711^{***}	2.760^{***}
Sentiment coefficient				
$\hat{\delta}$ (S_t)	$0.000^{\dagger}$	$0.000^{\dagger}$	$0.000^{\dagger}$	$-\mathbf{0.043}^{***}$
$\hat{\delta}$ (S_{t-1})	$0.000^{\dagger}$	$0.000^{\dagger}$	$0.000^{\dagger}$	$-\mathbf{0.060}^{***}$
Model fit and diagnostics				
Log-likelihood (S_t)	−54,468	−54,602	−54,456	−53,881
AIC (S_t)	1.4646	1.4682	1.4643	1.4489
BIC (S_t)	1.4655	1.4690	1.4653	1.4499
ARCH-LM [5] (p)	0.065	0.084	0.064	0.333
Nyblom joint (S_t)	19.29	19.98	20.12	27.61
Nyblom joint (S_{t-1})	19.53	20.27	20.35	27.94
Nyblom 1% critical value	2.35	2.12	2.59	2.59

Notes: ML estimation with Student-t innovations; mean equation ARMA$(1,1)$ without intercept, estimated jointly with the variance equation, which accounts for the small differences in $\hat{\phi}$ and $\hat{\theta}$ across columns. Parameter names follow each package's convention: in the EGARCH column α_1 loads on the signed innovation, a negative value denoting standard leverage, and γ_1 on its centred absolute value; elsewhere α_1 is the ARCH term and γ_1 the GJR threshold. Variance and distribution parameters are unchanged to three decimal places between the contemporaneous and the lagged specification and are reported once.

† Estimate at the lower bound of the admissible parameter space; re-estimating with $-S_t$ in place of S_t leaves it at zero.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

conditional innovation within each regime: the unconditional return distribution is a probability-weighted mixture across states with time-varying scale, and its moments are not determined by $\hat{\nu}_1$ and $\hat{\nu}_2$ alone. The contrast does indicate, however, that the extreme single-regime tail reported in Section 4.1 partly reflects regime pooling rather than a property of either state taken separately.

Within-regime persistence differs across states. In the turbulent regime the estimate is interior, at $\hat{\beta}_2 = 0.991$, so volatility shocks decay within the state. In the calm regime the likelihood increases monotonically up to the boundary and $\hat{\beta}_1 = 1$: shocks have no internal mean-reverting force and the state behaves as an integrated EGARCH. Separating regimes therefore does not lower within-state persistence relative to the single-regime EGARCH-X, where $\hat{\beta}_1 = 0.992$.

The estimated chain is nonetheless ergodic: since $\hat{P}_{11} = 0.874 < 1$, the calm state is exited almost surely and its expected duration is finite. We stress, however, that ergodicity of the regime process does not by itself deliver stationarity of the joint state–variance process. Under the parallel-recursion structure of Haas et al. (2004), each regime-specific log-variance is propagated at every date whether or not its regime is active, so leaving the calm state neither resets nor contracts the recursion. We therefore report $\hat{\beta}_1 = 1$ as a boundary estimate and do not claim that regime switching restores stationarity. The substantive implication is negative and informative in its own right: contrary to the mechanism of Lamoureux and Lastrapes (1990), whereby unmodelled structural change inflates single-regime persistence, separating regimes here leaves persistence unchanged. The near-unit-root behaviour of hourly Bitcoin volatility appears to be a property of the volatility process itself rather than an artefact of regime pooling.

The self-transition probabilities are $\hat{P}_{11} = 0.874$ and $\hat{P}_{22} = 0.757$, implying expected durations of 7.9 and 4.1 hours. The corresponding ergodic probabilities are 0.658 and 0.342, so the calm state prevails roughly two-thirds of the time. Classification on filtered probabilities assigns a larger share still to the calm state, at 78.1% against 21.9%; the gap between the two figures indicates that the filter spends a non-trivial share of the sample near the classification threshold, so the separation between states is less clear-cut than a hard assignment suggests. Both durations are short, which is the substantive point at this frequency: a full cycle lasts around twelve hours, so regime alternation is an intraday phenomenon that daily aggregation cannot resolve. The two states are best understood not as crisis and tranquillity but as distinct speeds of adjustment.

The asymmetry estimates are consistent across specifications. Both regime-specific sign coefficients are negative ($\hat{\alpha}_{2,1} = -0.006$, $\hat{\alpha}_{2,2} = -0.016$), as is the corresponding coefficient in the single-regime EGARCH-X, and the GJR threshold term loads positively on negative innovations. All four therefore indicate standard leverage. This is worth noting because the crypto literature has often reported the opposite pattern at daily frequency (Baur et al., 2018; Chu et al., 2017): at hourly frequency, Bitcoin behaves conventionally in this respect.

Table 5: Two-state Markov-switching EGARCH-X: ML estimates

Parameter	Regime 1 (calm)	Regime 2 (turbulent)
$\hat{\alpha}_{0,k}$ (Intercept)	-0.026^{***}	0.037^{***}
	(0.001)	(0.003)
$\hat{\alpha}_{1,k}$ (Magnitude)	0.072^{***}	0.148^{***}
	(0.003)	(0.012)
$\hat{\alpha}_{2,k}$ (Sign)	-0.006^{***}	-0.016^{***}
	(0.002)	(0.005)
$\hat{\beta}_k$ (Persistence)	$1.000^{\dagger}$	0.991
		(0.001)
$\hat{\nu}_k$ (Student-t)	7.78	3.38
	(0.71)	(0.11)
$\hat{\delta}$ (Sentiment S_t)	$-\mathbf{0.037}^{***}$	(0.012)
$\hat{P}_{kk}$	0.874 (0.008)	0.757 (0.018)
Expected duration (hours)	7.9	4.1
Ergodic probability	0.658	0.342
Classified frequency	78.1%	21.9%
Realised std. dev.	0.428%	1.465%
Log-likelihood	−53,373.63	
AIC	1.4354	
BIC	1.4370	

Notes: ML estimation on ARMA$(1,1)$ residuals by direct evaluation of the Hamilton filter under the parallel-recursion structure of Haas et al. (2004). Standard errors in parentheses, from the numerical Hessian with β_1 fixed at its boundary value. $\alpha_{1,k}$ loads on the centred absolute innovation and $\alpha_{2,k}$ on the signed innovation, so a negative $\alpha_{2,k}$ denotes standard leverage. δ is restricted to be common across regimes; Table 6 reports the unrestricted and lagged specifications.

† Likelihood increases monotonically to the boundary; the calm regime is integrated in log-variance. Because the estimate lies on the boundary, the reference distribution of the likelihood-ratio statistics reported below is non-standard and the associated p-values should be read as indicative.

*** $p < 0.01$.

4.3 The level channel

Following the distinction drawn by Filardo (1994), political communication may either alter the intensity of volatility within a regime, through a level channel, or shift the probability of transiting between regimes. We estimate the former by entering the sentiment score in the regime-specific log-variance equations and holding the transition matrix constant, and turn to the latter in Section 4.4.

Table 6 reports the sentiment coefficient across specifications. In the preferred model the estimate is $\hat{\delta} = -0.037$ with a likelihood-ratio statistic of 9.78 ($p = 0.002$), so the restriction $\delta = 0$ is rejected at the one percent level. The estimate is close to its single-regime counterpart of -0.043, which matters for interpretation: separating regimes overturns neither the sign nor the order of magnitude of the association. Evaluated at the standard deviation of the score among post hours, a one-standard-deviation increase in tone lowers $\ln \sigma_t^2$ by roughly 0.020, or about one percent of conditional volatility. The effect is precisely estimated and economically modest.

Allowing δ to differ across states leaves the estimates almost unchanged, at -0.036 in the calm regime and -0.039 in the turbulent one, and the restriction $\delta_1 = \delta_2$ is not rejected ($LR = 0.01$, $p = 0.935$). Political communication is therefore associated with the intensity of Bitcoin volatility to the same proportional degree whichever state the market occupies. The uniformity is itself informative. Because δ scales the log-variance, a common coefficient implies an absolute effect on σ_t^2 more than an order of magnitude larger in the turbulent state than in the calm one, so a constant proportional response is consistent with a channel that amplifies whatever volatility dynamics are already in train rather than adding a fixed increment to them.

The two states are far enough apart for this common coefficient to bear on functional form. Their realised standard deviations, at 0.428% and 1.465%, differ by a factor of 3.4, so their variances differ by a factor of roughly twelve. A single proportional coefficient describes both states over that range, whereas a constant absolute shift common to the two would have to be either negligible relative to the turbulent state or large relative to the calm one. The switching estimates thus corroborate the functional-form result of Section 4.1 from a second direction, though they do not test it: no additive switching specification is estimated here, and what the evidence establishes is that a proportional coefficient survives a twelvefold change in the level of variance, not that an additive one fails.

The lagged specification yields a larger coefficient, at -0.047, with a higher likelihood-ratio statistic (15.82, $p < 0.001$) and marginally lower information criteria. The ordering replicates that of the single-regime estimates and is partly mechanical for the reason given above. On either reading, the association is concentrated within a two-hour window.

Table 6: The level channel across specifications

	(1) No sentiment $\delta = 0$	(2) Regime-specific S_t	(3) Common δ S_t	(4) Common δ S_{t-1}
$\hat{\delta}_1$ (Regime 1)	—	−0.036	—	—
$\hat{\delta}_2$ (Regime 2)	—	−0.039	—	—
$\hat{\delta}$ (common)	—	—	-0.037^{***}	-0.047^{***}
			(0.012)	(0.012)
Parameters k	12	14	13	13
Log-likelihood	−53,378.52	−53,373.63	−53,373.63	−53,370.61
AIC	1.4355	1.4354	1.4354	**1.4353**
BIC	1.4370	1.4371	1.4370	**1.4369**
$H_0: \ \delta = 0$	—	9.78^{***}	9.78^{***}	15.82^{***}
		[0.008]	[0.002]	[< 0.001]
$H_0: \ \delta_1 = \delta_2$	—	0.01	[0.935]	—

Notes: All specifications share the structure of Table 5 and are estimated on the same ARMA(1, 1) residuals. Column (2) allows δ to differ across regimes; columns (3) and (4) impose $\delta_1 = \delta_2$. The LR test of $\delta = 0$ has two degrees of freedom in column (2) and one in columns (3) and (4); the test of $\delta_1 = \delta_2$ compares columns (2) and (3). p-values in brackets. Individual restrictions in column (2) yield $LR = 6.76$ [0.009] for δ_1 and $LR = 3.00$ [0.083] for δ_2. *** $p < 0.01$.

4.4 Conditional volatility dynamics and the platform interruption

Figure 1 plots the conditional volatility implied by the MS-EGARCH-X, aggregated monthly over the full sample in Panel A and at hourly resolution over three weeks spanning the March 2020 liquidity crisis in Panel B.

Panel B clarifies what the two states capture. The turbulent regime marks abrupt transitions rather than sustained stress: the sharp ascents of 12–13 March are assigned to it, whereas the subsequent decline from 13 to 18 March remains in the calm state despite conditional volatility of three to four percent, far above its sample mean. With an expected duration of four hours, the turbulent state identifies the intraday clustering of large innovations rather than macroeconomic crisis episodes, which unfold over weeks and are absorbed within the calm state once the initial adjustment has passed.

Panel A establishes a first fact about the relation between communication and volatility, and it runs against the view that political messaging is a source of market turbulence. The share of hours containing a market-relevant post rises from three to ten percent between 2018 and 2025, reflecting the growing political salience of digital assets, while mean conditional volatility declines over the same period, from above two percent during the 2017–2018 bubble to roughly one half of one percent from 2023 onward. The two series are negatively correlated at monthly frequency (-0.45), and the association is not merely a common trend: removing a linear time trend from both leaves a correlation of -0.34. Months of intense market-relevant political communication are thus, on average, calmer months.

The platform interruption provides a sharper test of the transition channel, because it removes political communication entirely rather than merely varying its intensity. Trump was suspended from Twitter in January 2021, and the first market-relevant post on Truth Social in our corpus appears only in April 2022. Between February 2021 and March 2022 — fourteen months covering 13.7% of the sample, marked by the shaded band — not a single market-relevant post was published on either platform. The regime structure of the market was nevertheless essentially unaffected. The mean probability of the turbulent state is 0.360 during the interruption against 0.339 outside it: a difference of two percentage points, and one whose sign is the opposite of what a transition channel would imply. Had political communication been driving the market into its turbulent state, its complete absence for over a year should have produced fewer such episodes, not marginally more.

The window was also, in market terms, an eventful one. It spans the April and November 2021 peaks and the May 2021 collapse, during which Bitcoin lost roughly half its value in six weeks, and mean conditional volatility is accordingly higher during the interruption than outside it, at 0.843% against 0.660%. The market therefore produced both its regime alternation and some of its largest volatility episodes with no political input whatsoever. Regime transitions appear to be driven by market dynamics that require none, which is consistent with the level channel estimated above: tone is associated with the intensity of volatility within whichever state prevails, and not with the movement between states.

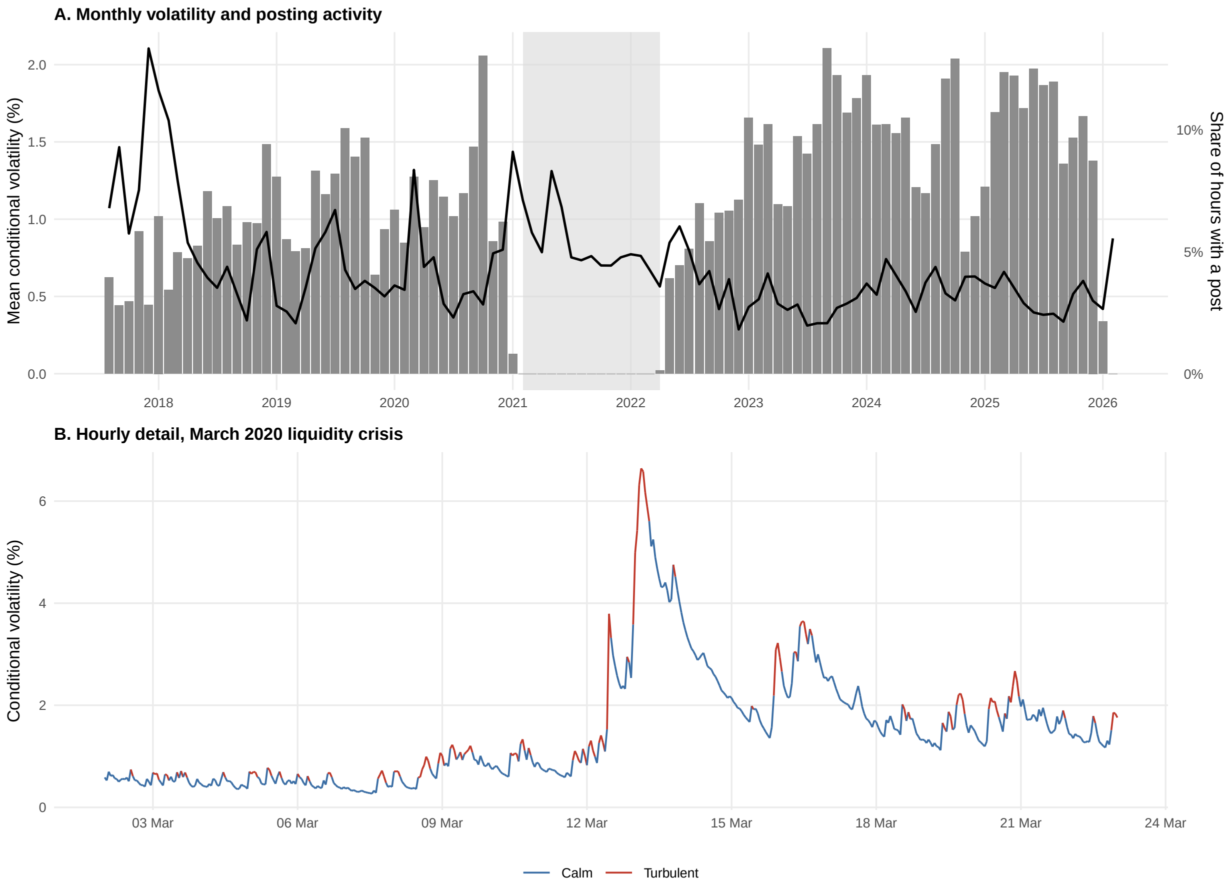

Figure 1: Conditional volatility from the MS-EGARCH-X. *Panel A*: monthly mean conditional volatility (black line, left axis) and share of hourly windows containing at least one market-relevant post (grey bars, right axis). The shaded band marks February 2021 to March 2022, during which no market-relevant post was published on either platform: Trump had been suspended from Twitter in January 2021 and the first market-relevant Truth Social post in our corpus dates from April 2022. *Panel B*: hourly conditional volatility over three weeks spanning the March 2020 liquidity crisis, coloured by the regime with the higher filtered probability. Volatility is the probability-weighted mixture of the two regime-specific conditional standard deviations.

5 Robustness checks

This section subjects the level result to three perturbations of the empirical design, taken in turn: the instrument used to measure tone, the asset on which the model is estimated, and the distributional assumption on the innovations. The second of these is as much an extension of scope as a check, since it asks whether the channel is specific to Bitcoin. In each case only the element under scrutiny changes; the taxonomy, the hourly aggregation rule, the mean equation and the switching specification are those of Section 4 throughout. The quantity reported is the common coefficient δ of column (3) of Table 6, and Table 7 collects the estimates.

5.1 An alternative sentiment lexicon

Because the level result rests on a single measurement instrument, we re-estimate the switching specification with the finance dictionary of Loughran and McDonald (2011) in place of VADER. Every post retained by the taxonomy is rescored, the hourly aggregation rule is unchanged, and the model is the common-δ MS-EGARCH-X of column (3) in Table 6. The two lexicons differ in scale, since the Loughran–McDonald score normalises by document length, so both are standardised to unit variance among post hours and δ is read as the response to a one-standard-deviation change in tone.

The association survives, as the first two rows of Table 7 show. The coefficient is -0.024 under VADER and -0.016 under Loughran–McDonald, and the restriction $\delta = 0$ is rejected in both cases ($LR = 9.78$, $p = 0.002$; $LR = 4.45$, $p = 0.035$). A one-standard-deviation increase in tone lowers conditional volatility by 1.20% under the first measure and 0.81% under the second. The regime structure is unaffected: expected durations, tail parameters and transition probabilities are unchanged to the reported precision, which confirms that the switching estimates identify the two states from the return process and not from the communication regressor.

The attenuation under Loughran–McDonald is itself informative and follows from the construction of the two instruments. As noted in Section 2, a bag-of-words dictionary assigns identical scores to a message and to its capitalised, exclamation-laden variant, discarding the amplification devices on which this style of communication relies; VADER scores them differently. Under the reading of the sentiment score as a noisy but unbiased proxy for tone, the instrument that discards more of the relevant variation should return a coefficient closer to zero, which is what we observe. That the sign, the significance and the order of magnitude are preserved is the substantive point.

5.2 An alternative cryptocurrency

The taxonomy selects political communication for its potential relevance to the cryptocurrency market rather than to a single asset, which invites the question of whether the association we document is specific to Bitcoin. We therefore repeat the exercise on Ethereum, the second asset by capitalisation, using hourly `ETHUSDT` prices from the same exchange over the same window. The communication measure is unchanged — the same posts, the same scores, the same hourly aggregation — so the design holds the signal fixed and varies the asset that receives it. Everything downstream of the price series is re-estimated: among specifications with $p, q \in \{0, 1, 2\}$ the BIC selects an ARMA$(0, 2)$ for Ethereum rather than the ARMA$(1, 1)$ retained for Bitcoin, and the switching parameters are initialised from a two-state fit on Ethereum residuals rather than carried over.

The coefficient is -0.016 and the restriction $\delta = 0$ is rejected ($LR = 4.36$, $p = 0.037$), against -0.024 for Bitcoin. The point estimate is smaller, but it lies well inside the confidence interval of the baseline, so the evidence does not support a difference between the two assets; what it supports is that the association is present in both. A one-standard-deviation increase in tone lowers conditional volatility by 0.81 percent on Ethereum against 1.20 percent on Bitcoin.

The regime structure replicates more closely still, and in a way that is informative about what the two states capture. Ethereum is the more volatile asset, with an hourly standard deviation of 0.97 percent against 0.78 percent, and its realised volatilities within the two states are correspondingly higher, at 0.62 and 2.25 percent. Their ratio, however, is 3.6, against 3.4 for Bitcoin: the two states stand in almost the same relative position on both assets. Expected durations are 7.5 and 2.6 hours, against 7.9 and 4.1, so the calm state is essentially indistinguishable across assets and the turbulent state is shorter on Ethereum. The intraday alternation between a slow and a fast state is therefore a property of this market rather than of Bitcoin alone, and the level channel operates within it on both assets.

5.3 An alternative innovation distribution

The regime-specific degrees of freedom reported in Table 5 motivate a third check. In the turbulent state $\hat{\nu}_2 = 3.38$ lies below four, so the standardised innovation there has no finite fourth moment, and the single-regime estimate of Section 4.1 is more extreme still. A specification whose tail parameter sits in that region invites the question of whether the result is an artefact of the Student assumption. We therefore re-estimate the model with two alternative innovation distributions: the generalised error distribution, which departs from the Student in the shape of the tail rather than its thickness alone, and the skewed Student-t of Fernández and Steel (1998), which relaxes the symmetry of the conditional density. Both are standardised to zero mean and unit variance, as the exponential specification requires, and the centring term $\mathbb{E}\,|z_k|$ in Equation (5) is recomputed for each.

The coefficient is unchanged. The final rows of Table 7 report -0.024 under the skewed Student-t and -0.025 under the generalised error distribution, against -0.024 in the baseline, and $\delta = 0$ is rejected at the one percent level in every case. Expressed in conditional volatility the three estimates span 1.20 to 1.23 percent per standard deviation of tone. The level result is therefore insensitive to the distributional assumption within this family.

Two features of the fit are worth recording, since they bear on the baseline specification rather than on the sentiment coefficient. Allowing for skewness raises the log-likelihood by 38.2 points for two additional parameters, so the restriction to a symmetric conditional density is decisively rejected ($LR = 76.4$, $p < 0.001$); the tone coefficient nonetheless does not move, which indicates that the asymmetry of the innovation and the transmission of communication are separate features of the data. The generalised error distribution, by contrast, fits substantially worse than either Student specification, losing 162 points of log-likelihood at equal parameter count. The heavy-tailed Student family is thus the appropriate one for hourly returns at this frequency, and the symmetric version we retain in the main text is a restriction of the better-fitting member rather than a departure from it.

Table 7: Robustness of the level coefficient

	Variant	$\hat{\delta}$	LR	Effect on σ_t
Baseline	VADER, BTC, Student-t	-0.024^{***}	9.78	-1.20%
		(0.008)	[0.002]	
Lexicon	Loughran–McDonald	-0.016^{**}	4.45	-0.81%
		(0.008)	[0.035]	
Asset	Ethereum	-0.016^{**}	4.36	-0.81%
			[0.037]	
Distribution	Skewed Student-t	-0.024^{***}	9.97	-1.21%
			[0.002]	
	GED	-0.025^{***}	10.07	-1.23%
			[0.002]	

Notes: Each row re-estimates the common-δ MS-EGARCH-X of column (3) of Table 6, changing one element of the design and holding the taxonomy, the hourly aggregation rule, the mean equation and the switching specification fixed. Sentiment scores are standardised to unit variance among post hours, so $\hat{\delta}$ is the response of $\ln \sigma_t^2$ to a one-standard-deviation change in tone and is comparable across rows. The final column converts that response into a percentage change in conditional volatility. Standard errors in parentheses where available, LR p-values in brackets; the test is of $H_0 : \delta = 0$ with one degree of freedom, and inference is likelihood-ratio based throughout for the reasons given in Section 3. Each distribution row compares the free model to a constrained model re-estimated under the same distribution, so the statistics are internally consistent but the LR values are not comparable across rows; the log-likelihoods are. The Ethereum row is estimated on its own returns with its own mean equation and is not nested in the baseline, so its log-likelihood is not comparable with the others either.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6 Discussion

The results describe a channel that operates on the second moment, proportionally to the prevailing level of volatility, and within rather than across market states. Each of these three features carries an interpretation, and taken together they suggest a specific account of how political communication enters a cryptocurrency market.

That the channel is confined to the second moment is itself informative. Mean returns are statistically indistinguishable across sentiment regimes ($t = 0.150$), so the directional content of communication does not translate into a return differential even though it is associated with conditional variance. This is the pattern the mechanism of Pástor and Veronesi (2013) predicts at this frequency: political signals revise the uncertainty attached to future policy rather than the expected cash flows of the asset, and in a market with no cash flows to revise the first channel is the only one available. Our findings extend the equity-market results of Colonescu (2018) and Ge et al. (2019) to an asset whose weak fundamental anchoring should, in principle, make it more rather than less responsive to high-frequency political signals, and the direction of the association is the one their framework implies.

That the effect is proportional rather than additive bears on how the signal is absorbed. Because δ scales the log-variance, a coefficient common to both states implies an absolute effect on σ_t^2 more than an order of magnitude larger in the turbulent state than in the calm one. Communication does not add a fixed increment of volatility to the market; it amplifies whatever dynamics are already in train. A market already adjusting quickly reacts more, in absolute terms, to the same message than a market at rest. This reading is supported from two directions: the coefficient is statistically indistinguishable across states whose variances differ by a factor of roughly twelve, and it is stable at -0.024 to -0.025 across two alternative innovation distributions and at -0.016 on a second asset with a different volatility level.

That the effect is confined to the level margin is the finding with the clearest implication outside the paper. A signal that shifts the level of variance within a prevailing state is absorbed by the market's existing dynamics; a signal that shifts transition probabilities displaces those dynamics themselves. Only the second would constitute a source of structural risk, and the evidence points to the first. The fourteen-month interruption is the sharpest observation here: during a period when no market-relevant post was published on either platform, the market produced both its usual regime alternation and some of its largest volatility episodes, and the mean probability of the turbulent state was, if anything, marginally higher than outside the window. Regime change in this market is driven by dynamics that require no political input. For regulators concerned with financial stability in digital asset markets, what warrants monitoring is therefore the mechanics of regime change rather than the day-to-day flow of political messaging; for market participants, the directional content of communication is more informative than its mere occurrence, so sentiment-based strategies should condition on tone rather than on volume.

The regime structure deserves comment in its own right, independently of the communication channel. The two states are not crisis and tranquillity but distinct speeds of adjustment: with expected durations of eight and four hours a full cycle lasts around half a day, and Panel B of Figure 1 shows the turbulent state marking the abrupt ascents of March 2020 while the subsequent decline, at conditional volatility several times its sample mean, is absorbed within the calm state. That this structure replicates on Ethereum — durations of 7.5 and 2.6 hours, and realised volatilities standing in almost the same ratio, 3.6 against 3.4 — indicates a property of the market rather than of a single asset. A further point runs against the received reading of high persistence in cryptocurrency volatility. Lamoureux and Lastrapes (1990) attribute near-integrated GARCH estimates to unmodelled structural change, and the single-regime estimate of $\hat{\alpha}_1 + \hat{\beta}_1 = 1.042$ invites exactly that diagnosis. Yet separating regimes leaves within-state persistence at the integrated boundary. The near-unit-root behaviour of hourly Bitcoin volatility appears to be a property of the volatility process itself, not an artefact of regime pooling.

A methodological lesson follows from the pattern of what is and is not detected. The channel is recovered only in the log-linear specifications, only at hourly frequency, and only when communication is measured by the direction of its content. Each of these is a design choice that the literature has often resolved the other way, and each is sufficient on its own to hide the effect: an association that scales with the prevailing level of volatility is invisible to a specification imposing a constant absolute shift, an association concentrated within a two-hour window is averaged away by daily aggregation, and a binary indicator of occurrence conflates a market-supportive message with a regulatory threat. That the effect has proved difficult to detect in earlier work may owe as much to these choices as to its modest size.

What the design establishes should be stated precisely, since the three margins are not identified in the same way. The timing of market-relevant communication is not exogenous: posting concentrates in bull markets, and at monthly frequency the share of hours containing a post is negatively correlated with mean conditional volatility, at -0.45 in levels and -0.34 once both series are detrended. Because communication on cryptocurrency is predominantly positive in tone, the sentiment score is mechanically higher in calmer periods, in the same direction as the estimated $\hat{\delta}$. The hourly specification conditions on lagged conditional variance and absorbs much of this low-frequency co-movement, but the level coefficient is best read as a conditional association net of own-volatility dynamics rather than as an identified treatment effect. The platform interruption is stronger on this count, since it removes the signal entirely rather than varying its intensity, but it disciplines the conclusion about regime dynamics and not the magnitude of the level effect: the window spans a distinct market era, and it is a weaker counterfactual for volatility levels than for regime alternation. A design-based identification of the tone channel remains open.

Finally, the evidence comes from a single communicator. Trump's posting is distinctive in volume, tone and market salience. Section 5.2 shows that the channel is not confined to a single asset, but the corresponding question on the sending side is not addressed here. What we take to travel beyond this setting is the structure of the result — a proportional level effect, uniform across states, without an accompanying transition effect — and the design that recovers it, rather than the size of the coefficient.

7 Conclusion

This paper asks through which margin of the volatility process high-frequency political communication operates: the level of conditional variance within a prevailing market state, or the probability of transiting between states. We answer the question with an hourly dataset spanning August 2017 to February 2026 that combines the universe of Donald Trump's social media posts with Binance price data, and with a Markov-switching EGARCH framework in which the two margins are parameterised independently. Three findings follow.

First, the association depends on the functional form of the variance equation. In the EGARCH-X the sentiment coefficient is negative and statistically significant, indicating that hours of more positive political communication coincide with lower conditional volatility and hours of more adverse rhetoric with higher conditional volatility, though the effect is economically modest. It is not recovered in the additive specifications, where the estimates lie at the boundary of the admissible parameter space and remain at zero once that boundary is made non-binding.

Second, Bitcoin volatility is strongly regime-dependent. The MS-EGARCH-X identifies two structurally distinct states with expected durations of eight and four hours, so that a full cycle lasts around twelve hours and alternation cannot be resolved by daily aggregation. The states correspond to distinct speeds of adjustment rather than to crisis and tranquillity, and near-integrated behaviour within the calm state persists once regimes are separated, which indicates that the high persistence of hourly Bitcoin volatility is not an artefact of unmodelled regime shifts.

Third, the association operates through a level effect that is uniform across states: a likelihood-ratio test cannot reject the restriction that the sentiment coefficient takes the same value in both regimes. A fourteen-month interruption in communication leaves regime alternation essentially unchanged, which suggests that communication modulates volatility intensity within whichever state prevails without governing transitions between states. The level coefficient is stable across an alternative sentiment lexicon, an alternative cryptocurrency and two alternative innovation distributions.

The broader implication is that political communication is best understood as a high-frequency informational input that alters uncertainty conditions within the prevailing market state, rather than as a determinant of the structural changes that reorganise it. The distinction matters for how such communication should be treated as a risk factor: a channel confined to the level margin is absorbed by the market's existing dynamics, whereas one operating on transitions would displace those dynamics themselves. Several extensions bear directly on how far this pattern generalises: separating the occurrence of a message from its tone, entering directly crypto-related and broader macro-political messages as distinct regressors, estimating an additive switching specification to test the functional-form result within the regime-switching framework itself, allowing time-varying transition probabilities to test the transition channel directly, and extending the design to other communicators and platforms.

Data availability statement

Hourly Bitcoin and Ethereum prices are publicly available through the Binance API. The corpus of social media posts was compiled from publicly available archives; the constructed hourly dataset and the replication code are available upon request.

References

Ardia, D., Bluteau, K., Boudt, K., Catania, L., and Trottier, D.-A. (2019a). Markov-switching garch models in r: The msgarch package. *Journal of Statistical Software*, 91:1–38.

Ardia, D., Bluteau, K., and Rüede, M. (2019b). Regime changes in Bitcoin GARCH volatility dynamics. *Finance Research Letters*, 29:266–271.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Baur, D. G., Hong, K., and Lee, A. D. (2018). Bitcoin: Medium of exchange or speculative assets? *Journal of International Financial Markets, Institutions and Money*, 54:177–189.

Born, B., Müller, G. J., Schularick, M., and Sedláček, P. (2021). The costs of economic nationalism: Evidence from the brexit experiment. *Economic Journal*, 131(635):1042–1078.

Caporale, G. M. and Zekokh, T. (2019). Modelling volatility of cryptocurrencies using Markov-Switching GARCH models. *Research in International Business and Finance*, 48:143–155.

Chu, J., Chan, S., Nadarajah, S., and Osterrieder, J. (2017). Garch modelling of cryptocurrencies. *Journal of Risk and Financial Management*, 10(4):17.

Colonescu, C. (2018). The effects of president trump's tweets on financial markets. *Finance Research Letters*, 25:188–192.

Fernández, C. and Steel, M. F. J. (1998). On Bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association*, 93(441):359–371.

Filardo, A. J. (1994). Business-cycle phases and their transitional dynamics. *Journal of Business & Economic Statistics*, 12:299–308.

Ge, Q., Kurov, A., and Wolfe, M. H. (2019). Political uncertainty and stock market volatility: The case of twitter. *Journal of Financial Markets*, 45:1–19.

Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The journal of finance*, 48(5):1779–1801.

Haas, M., Mittnik, S., and Paolella, M. S. (2004). A new approach to Markov-switching GARCH models. *Journal of Financial Econometrics*, 2(4):493–530.

Hansen, B. E. (1992). The likelihood ratio test under nonstandard conditions: Testing the Markov switching model of GNP. *Journal of Applied Econometrics*, 7(S1):S61–S82.

Hassan, T. A., Hollander, S., van Lent, L., and Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics*, 134(4):2135–2202.

Hutto, C. J. and Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1):216–225.

Katsiampa, P. (2017). Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters*, 158:3–6.

Kelly, B., Pástor, Ł., and Veronesi, P. (2016). The price of political uncertainty: Theory and evidence from the option market. *Journal of Finance*, 71(5):2417–2480.

Kraaijeveld, O. and De Smedt, J. (2020). The predictive power of public Twitter sentiment for forecasting cryptocurrency prices. *Journal of International Financial Markets, Institutions and Money*, 65:101188.

Lamoureux, C. G. and Lastrapes, W. D. (1990). Persistence in variance, structural change, and the GARCH model. *Journal of Business & Economic Statistics*, 8:225–234.

Liu, Y. and Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, 34(6):2689–2727.

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66(1):35–65.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the econometric society*, pages 347–370.

Nyblom, J. (1989). Testing for the constancy of parameters over time. *Journal of the American Statistical Association*, 84(405):223–230.

Pankratz, N. (2019). Political risk and cryptocurrency returns. *Journal of International Money and Finance*, 99:102056.

Pástor, Ł. and Veronesi, P. (2013). Political uncertainty and risk premia. *Journal of Financial Economics*, 110(3):520–545.

Silva, Á. J. B., Caldeira, J. F., and Silva, W. V. (2024). Modeling and forecasting cryptocurrency volatility: evidence from garch-type models. *RAUSP Management Journal*, 59(1):45–63.

Smales, L. A. (2021). Investor attention and cryptocurrency returns. *Journal of International Financial Markets, Institutions and Money*, 73:101374.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance*, 62(3):1139–1168.