

Rolling Shannon Entropy and the Structural Complexity of Bitcoin Market Dynamics

AUTHORS: Markoulis S.*, University of Cyprus

SUMMARY: This study investigates whether rolling Shannon entropy can capture structural complexity and provide predictive information in Bitcoin's return dynamics. Using normalized rolling Shannon entropy over 15-, 30-, and 60-day windows, the research finds a strong positive association between entropy and rolling volatility, with high-entropy regimes exhibiting higher risk and lower average returns. Correlation and regression analyses show that entropy significantly enhances volatility forecasting models, doubling the explanatory power compared to an AR(1) benchmark. Out-of-sample tests indicate that entropy-augmented models reduce mean squared forecast errors by 5% to 17%, supporting the use of complexity-based measures in financial market modeling.

METHOD: The study employs rolling Shannon entropy calculated over different time windows and evaluates its predictive power in volatility forecasting. Correlation and regression analyses are used to assess the relationship between entropy and volatility, while out-of-sample tests are conducted using both recursive estimation and a fixed 2/3–1/3 split.

KEY FINDINGS:

- High-entropy regimes are associated with higher risk and lower average returns.
- Entropy remains strongly associated with forward volatility measures.
- Entropy inclusion in volatility regressions approximately doubles explanatory power.
- Entropy-augmented models reduce mean squared forecast errors by 5% to 17%.
- Tests of predictive accuracy reject the null hypothesis of equal forecast performance.

IMPLICATION FOR TRADING: These findings suggest that rolling Shannon entropy can provide valuable insights into Bitcoin's structural complexity, offering traders a tool to better predict market volatility. By incorporating entropy measures into their models, traders can potentially improve their forecasting accuracy and risk management strategies, leading to more informed trading decisions.

Rolling Shannon Entropy and the Structural Complexity of Bitcoin Market Dynamics

Markoulis S.*

Abstract

This study examines whether rolling entropy captures structural complexity and provides predictive information in the return dynamics of Bitcoin. Using normalized rolling Shannon entropy calculated over 15-, 30-, and 60-day windows, we examine whether entropy conveys information beyond conventional volatility measures. We first document a strong positive association between entropy and rolling volatility, with high-entropy regimes exhibiting substantially higher risk and weaker average returns. Correlation and regression analyses show that entropy remains strongly associated with forward volatility measures, and its inclusion in volatility regressions approximately doubles explanatory power relative to an AR(1) benchmark based solely on volatility persistence. Building on the in-sample correlation and regression evidence, we evaluate the out-of-sample predictive performance of entropy-augmented models, using both recursive estimation and a fixed 2/3–1/3 split. The entropy model consistently reduces mean squared forecast errors relative to the benchmark, with improvements ranging from 5% to 17% depending on horizon and estimation approach. Tests of predictive accuracy reject the null hypothesis of equal forecast performance. Overall, our findings suggest that rolling entropy captures persistent structural features of Bitcoin's return-generating process and provides economically meaningful incremental information for volatility forecasting, thus supporting the relevance of complexity-based and information-theoretic measures in financial market modelling.

Keywords: Rolling entropy; Bitcoin; Volatility forecasting; Market complexity; Information theory; Econophysics.

JEL Classification: G17, C22, C53, D83

*: University of Cyprus

Email: markoulis.stelios@ucy.ac.cy

Rolling Shannon Entropy and the Structural Complexity of Bitcoin Market Dynamics

Abstract

This study examines whether rolling entropy captures structural complexity and provides predictive information in the return dynamics of Bitcoin. Using normalized rolling Shannon entropy calculated over 15-, 30-, and 60-day windows, we examine whether entropy conveys information beyond conventional volatility measures. We first document a strong positive association between entropy and rolling volatility, with high-entropy regimes exhibiting substantially higher risk and weaker average returns. Correlation and regression analyses show that entropy remains strongly associated with forward volatility measures, and its inclusion in volatility regressions approximately doubles explanatory power relative to an AR(1) benchmark based solely on volatility persistence. Building on the in-sample correlation and regression evidence, we evaluate the out-of-sample predictive performance of entropy-augmented models, using both recursive estimation and a fixed 2/3–1/3 split. The entropy model consistently reduces mean squared forecast errors relative to the benchmark, with improvements ranging from 5% to 17% depending on horizon and estimation approach. Tests of predictive accuracy reject the null hypothesis of equal forecast performance. Overall, our findings suggest that rolling entropy captures persistent structural features of Bitcoin's return-generating process and provides economically meaningful incremental information for volatility forecasting, thus supporting the relevance of complexity-based and information-theoretic measures in financial market modelling.

Keywords: Rolling entropy; Bitcoin; Volatility forecasting; Market complexity; Information theory; Econophysics.

JEL Classification: G17, C22, C53, D83

1. Introduction

Financial markets are widely regarded as complex adaptive systems that are characterized by nonlinear dynamics, heterogeneous agents, and evolving distributional properties [1]. In such environments, traditional second-moment-based risk measures may often not be sufficient to fully capture the structure of market uncertainty. This concern is particularly relevant for cryptocurrency markets and in particular Bitcoin, which exhibits characteristics such as volatility clustering, heavy-tailed return distributions, regime shifts, and incidents of sharp price movements [2,3]. These features suggest that in such markets, risk may be better understood not only through standard variance-based measures, but also through alternative metrics that capture broader distributional complexity.

Entropy, originally introduced by Claude Shannon [4] in information theory, provides a model-free measure of the expected information content of a probability distribution. When applied to returns of financial assets, Shannon entropy captures the overall probabilistic organization of returns, reflecting changes in tail thickness, concentration, and structural disorder that are not fully summarized by second-moment measures.

As such, in markets like Bitcoin, entropy-based measures may capture dimensions of uncertainty that standard second-moment approaches overlook. In this context, an increase in entropy would reflect a flattening or redistribution of return probabilities across states, which is consistent with heightened structural uncertainty and reduced predictability in market dynamics.

Despite the growing literature on cryptocurrency volatility using GARCH-type models, realised volatility measures, and regime-switching frameworks, relatively little attention has been given to the use of rolling entropy as a time-varying indicator of structural market complexity. Existing studies primarily focus on modelling persistence in second moments, but they do not directly focus on quantifying distributional disorder or evolving probability structures. A systematic assessment of whether entropy provides additional explanatory and predictive content, beyond volatility persistence, remains limited.

This paper investigates whether rolling entropy provides meaningful information about Bitcoin market risk and whether it possesses incremental predictive power beyond traditional return-based measures. Using daily Bitcoin price data from January 2014 to February 2026, we construct normalized rolling entropy measures over 15-, 30-, and 60-day windows. To facilitate interpretation and regime classification entropy is then normalised by bounding the measure between zero and one, thus ensuring comparability across rolling windows. Within this framework, we examine the relationship between entropy and rolling volatility, explore entropy-based regime classification, and evaluate the predictive performance of entropy through both in-sample regressions and out-of-sample forecasting experiments.

The empirical results reveal a number of interesting findings. First, rolling entropy exhibits substantial variation over time and aligns closely with periods of heightened market turbulence. Second, entropy is strongly and positively correlated with rolling volatility across all window

lengths, with correlations remaining economically meaningful at forward horizons. Third, entropy-based regime classification produces a clear monotonic ordering; high-entropy periods are associated with elevated volatility and lower average returns, while low-entropy periods correspond to calmer and more favourable market conditions. Fourth, regression analysis shows that entropy explains a statistically and economically significant share of volatility variation, retaining significance even after controlling for return persistence. Finally, out-of-sample forecasting results demonstrate that incorporating entropy reduces mean squared prediction error by approximately 5–17% relative to benchmark volatility models, depending on horizon and estimation method, hence confirming its incremental predictive value. When taken together, these findings suggest that entropy captures structural features of return distributions that are not fully summarized by volatility persistence, particularly during periods of heightened market instability.

The contribution of this study is threefold. Firstly, it develops and systematically evaluates normalized rolling entropy as a measure of distributional complexity in Bitcoin returns. Secondly, it demonstrates that entropy provides economically meaningful regime differentiation and substantial incremental explanatory power beyond volatility persistence. Finally, it establishes the out-of-sample forecasting relevance of entropy using nested-model comparisons. By linking information-theoretic measures of disorder with financial volatility dynamics, our paper brings together concepts from statistical physics and econometric modelling and offers a distribution-sensitive perspective on market complexity.

The remainder of the paper unfolds as follows. Section 2 provides a review of the related literature on entropy in financial markets, volatility modelling, and cryptocurrency dynamics. Section 3 describes the data and methodology used to construct rolling entropy and related measures. Section 4 presents and discusses the empirical results and section 5 concludes.

2. Literature Review

This paper is related to three connected strands of literature: (i) the stylized facts and regime dynamics of the Bitcoin market, (ii) entropy and information-theoretic measures in the analysis of financial time series, and (iii) the application of entropy-based measures as dynamic indicators of risk and complexity in cryptocurrency markets.

Starting with the first one, several of the stylized empirical regularities documented in traditional financial time series, such as heavy-tailed return distributions, volatility clustering, and scaling behavior, are also observed in cryptocurrency markets [5,6]. Empirical studies show that Bitcoin returns tend to display substantial excess kurtosis and pronounced tail thickness [2,7], as well as strong volatility persistence and evidence of long memory in higher-order moments [8]. These features are consistent with the broader econophysics literature, which emphasises power-law behavior and nonlinear dependence structures in financial markets [9–11].

Beyond second-moment persistence, a growing body of work has also documented structural shifts and regime changes in Bitcoin dynamics. Specifically, Markov-switching and hidden-

state models identify distinct volatility regimes and time-varying risk states [3], thus suggesting that Bitcoin operates in an environment that is characterized by distributional properties that are evolving. Such findings lend support to the belief that cryptocurrency markets can be considered to be complex adaptive systems with nonlinear feedback mechanisms and regime-dependent behavior.

Parametric volatility frameworks, and in particular GARCH-type model specifications [12,13], capture conditional variance dynamics given explicit assumptions about the evolution of variance. However, although these models are successful in describing second-moment persistence, they are designed to model conditional variance and as such may not fully capture shifts in the overall distributional structure of returns, particularly during episodes of extreme market turbulence or structural breaks. These distributional and regime-dependent features motivate the use of other non-parametric measures that are capable of capturing broader structural uncertainty beyond conditional variance dynamics.

The second strand of literature is related to the use of entropy-based and information-theoretic approaches to analyse financial markets. In particular, Shannon entropy was introduced in information theory and subsequently incorporated into economic measurement frameworks [14], where it provides a model-free measure of uncertainty that is defined over the entire probability distribution. While permutation entropy and ordinal-pattern approaches focus on temporal ordering and complexity, the present study employs distribution-based Shannon entropy that is constructed from return histograms, thus directly quantifying probabilistic dispersion rather than ordinal structure. Unlike variance, which summarizes dispersion around the mean, entropy captures the full distributional structure of outcomes, thereby incorporating higher-order characteristics such as skewness, tail thickness, and probability concentration across states. As such, Shannon entropy provides a natural tool that can be utilised to examine how the probabilistic structure of the distribution of returns evolves over time and to assess whether these structural changes contain meaningful information about market risk.

Within the context of financial markets, entropy has been used to assess market efficiency, predictability, and structural complexity. Permutation entropy and related complexity-based approaches have been applied to equity markets to identify regime changes and variations in informational efficiency [15–18]. Time-varying entropy measures have also been used to examine the evolution of financial complexity across different market conditions [19]. These approaches reflect a broader movement within econophysics towards the use of disorder- and complexity-based measures to characterize nonlinear and non-equilibrium financial dynamics beyond descriptions offered by traditional second-moment measures.

The appeal of entropy-based methods lies in their non-parametric nature and in their ability to quantify distributional disorder, without imposing any assumptions on conditional variance. This can be particularly useful in environments that are characterized by heavy tails, nonlinearities, and regime shifts. As already noted, these are features widely documented in cryptocurrency markets therefore making entropy a theoretically coherent and empirically flexible measure of structural uncertainty.

In addition to Shannon entropy, the econophysics literature also provides generalized entropy measures such as Rényi [20] and Tsallis [21] entropy to capture different scaling patterns and heavy-tailed behavior in complex systems. These generalized formulations extend the standard Shannon framework by allowing different sensitivities to tail events and probability concentration. Although such approaches have been applied in the analysis of financial markets and complex economic systems, Shannon entropy remains the most widely used benchmark measure due to its strong theoretical foundation and straightforward interpretability as a measure of probabilistic disorder.

The third strand of literature focuses on cryptocurrency volatility modelling and the measurement of risk dynamics. A substantial body of work applies GARCH-type models and regime-switching specifications to Bitcoin and other digital assets [8,12,13]. These studies find evidence of persistent volatility clustering and regime dependence, confirming that second-moment dynamics play a central role in cryptocurrency risk modelling.

At the same time, nevertheless, relatively less attention has been given to the evaluation of entropy as a time-varying measure of market risk in cryptocurrency markets, in particular with respect to its additional explanatory and forecasting value in comparison to established volatility benchmarks. While entropy has been used to assess informational efficiency and complexity in Bitcoin and other financial assets [7,22,23], most applications remain descriptive in nature or focus on efficiency testing rather than risk prediction.

In markets that are characterised by structural breaks, evolving tail behavior, and nonlinear feedback effects, a normalized rolling entropy measure may provide meaningful complementary insight into distributional dynamics that are not fully summarized by conditional variance models. Despite this, however, a systematic evaluation of whether entropy delivers incremental explanatory power and statistically significant out-of-sample forecasting improvements when compared to benchmark volatility models remains limited.

This study seeks to address this gap by constructing normalized rolling Shannon entropy measures for Bitcoin returns and systematically evaluating their empirical relevance within a unified analytical framework. In particular, the analysis examines (i) the relationship between entropy and volatility across different market regimes, (ii) the incremental explanatory content of entropy in regression settings, and (iii) its out-of-sample forecasting performance relative to established second-moment volatility models. By combining distribution-based entropy measures with standard volatility benchmarks and formal forecast comparison procedures [24–26], the analysis provides a systematic assessment of whether entropy captures structural features of return distributions that are not fully summarized by conventional variance-based approaches.

3. Methodology

The empirical analysis uses daily closing prices of Bitcoin denominated in U.S. dollars covering the period 2 February 2014 to 12 February 2026. The data were obtained from Investing.com, which provides a continuous historical price series constructed as a volume-

weighted composite across major cryptocurrency exchanges. Given that Bitcoin trades 7 days per week, our data set includes prices over weekends.

Daily returns are computed as logarithmic returns:

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) \quad (\text{Equation 1})$$

where P_t refers to the closing price at day t .

As already noted, to quantify distributional uncertainty in returns, we utilise the Shannon entropy. Let r_t denote the daily log return at time t . For a given estimation window (W), the range of returns is partitioned into K discrete bins. Let $p_{i,t}^{(W)}$ denote the empirical probability that a return falls into bin i ($i = 1, \dots, K$), so that:

$$\sum_{i=1}^K p_{i,t}^{(W)} = 1 \quad (\text{Equation 2})$$

The Shannon entropy of the empirical return distribution within the window is then defined as follows:

$$H_t^{(W)} = -\sum_{i=1}^K p_{i,t}^{(W)} \ln(p_{i,t}^{(W)}) \quad (\text{Equation 3})$$

To facilitate comparability across estimation windows, entropy is then normalised as follows:

$$H_t^{norm, (W)} = \frac{H_t^{(W)}}{\ln(K)} \quad (\text{Equation 4})$$

where K is the number of bins. $H_t^{norm, (W)}$ ranges between 0 and 1. Higher entropy values would suggest a more dispersed and less predictable return distribution, whereas lower entropy ones reflect greater concentration of probability mass and relatively more predictable market dynamics.

For all calculations, the range of returns within each estimation window is divided into equally spaced bins ($K = 20$). The bin boundaries are defined using equally spaced intervals over the range of observed returns within each estimation window. In order to ensure comparability of entropy values and to avoid spurious time variation resulting from changing discretization schemes, the number of bins has been fixed across time. This approach balances statistical reliability since too few bins may understate distributional heterogeneity, while too many bins could produce sparse probabilities and unstable entropy estimates. This discretization approach is widely used in empirical entropy estimation for financial time series, where the underlying distribution is unknown and must be approximated non-parametrically.

To capture the time-varying uncertainty in the Bitcoin market, we also compute rolling entropy. For each time $t \geq W$, we consider the rolling window of log returns $\{r_{t-W+1}, \dots, r_t\}$. Within this

window, returns are assigned to bins and empirical probabilities $p_{i,t}^{(W)}$ are estimated for $i = 1, \dots, K$. The normalised entropy $H_t^{norm,(W)}$ is then computed using Equations (3) and (4). The superscript W indicates that both probabilities and entropy are computed over a rolling window of length W . This procedure is applied to multiple window lengths ($W = 15, 30, 60$) days to assess robustness to the choice of horizon.

It should be noted that in addition to entropy, the empirical analysis employs rolling volatility measures, regime classification based on entropy quantiles, predictive regression models, and out-of-sample forecasting procedures. To ensure comparability, rolling volatility is also computed as the sample standard deviation of returns within the same rolling window W . To maintain clarity of exposition and avoid unnecessary repetition of notation, the formal econometric specifications for the regression and forecasting exercises are presented in the corresponding empirical sections.


4. Results

In this section we present and discuss the empirical findings of Bitcoin price dynamics based on rolling entropy measures. We start by presenting the time variation of normalized rolling entropy relative to the price and volatility of Bitcoin and then elaborate by carrying out a rolling entropy-based regime analysis over multiple window lengths. We then turn our attention to correlation and regression analysis and in particular on the relationship between rolling entropy and volatility and finally assess the predictive power of entropy through out-of-sample forecasting experiments.

4.1. Rolling Entropy and Market Dynamics

To evaluate the evolving structural complexity of the Bitcoin market, we compute normalized rolling Shannon entropy of daily returns over 15-, 30-, and 60-day windows. Figure 1 depicts the 30-day normalized entropy together with the logarithm of Bitcoin prices. Entropy varies between near zero and about 0.55, indicating that market complexity changes over time; values close to zero would be indicative of predictable periods, while values hovering around 0.55 would suggest moderate random behaviour. In addition, it also appears that periods of elevated entropy tend to coincide with sharp price movements and episodes of heightened market turbulence.

Figure 1. Rolling 30-Day Normalized Entropy and Bitcoin Prices



Notes: The above figure plots the 30-day normalized Shannon rolling entropy (left axis) of Bitcoin returns alongside the logarithm of Bitcoin daily prices (right axis) for the period January 2014 to February 2026. Entropy values are normalized between 0 and 1, and the log price series is included to illustrate the relationship between market complexity and price movements.

Figure 2 elaborates on the relationship between entropy and market risk by plotting the 30-day normalized rolling entropy against the 30-day rolling volatility, which is estimated over the same window (Equation 5).


$$\sigma_t^{(W)} = \sqrt{\frac{1}{W-1} \sum_{j=0}^{W-1} (r_{t-j} - \bar{r}_t^{(W)})^2} \quad (\text{Equation 5})$$

where $\bar{r}_t^{(W)}$ denotes the mean return within the rolling window of length W (equation 6).

$$\bar{r}_t^{(W)} = \frac{1}{W} \sum_{j=0}^{W-1} r_{t-j} \quad (\text{Equation 6})$$

A clear positive relationship between entropy and volatility is observed; periods of high volatility are associated with higher entropy, while calmer market conditions correspond to lower entropy values. Taken together, the two figures suggest that rolling entropy captures the complexity of the Bitcoin market and its alignment with market risk. Similar patterns are observed for 15- and 60-day windows, thus confirming the robustness of these dynamics across different time scales.

Figure 2. Rolling 30-Day Normalized Entropy and Rolling Volatility of Bitcoin



Notes: The above figure plots the 30-day normalized rolling Shannon entropy (left axis) and the 30-day rolling volatility (right axis) of Bitcoin returns over the sample period. Entropy values are normalized between 0 and 1, and the rolling volatility is presented to capture market uncertainty.

4.2. Entropy-Based Regime Analysis

To further investigate the economic relevance of rolling entropy, we next classify each trading day into three regimes - Low, Medium, and High entropy - based on the 33rd and 66th percentiles of the 30-day normalized rolling entropy distribution. This percentile-based classification ensures approximately balanced regime sizes while allowing for clear differentiation between low- and high-uncertainty market states. We then report the average volatility and returns across the three entropy regimes, along with pairwise t-tests for differences in mean values. Results are summarised in table 1.

Table 1: Volatility and Returns across Entropy Regimes (30-day rolling window)

Panel A: Descriptive Statistics					
Regime	Obs.	Avg. Volatility	Std. Dev. Volatility	Avg. Returns	Std. Dev. Returns
Low	1394	0.016371	0.127948	0.002677	0.019674
Medium	1419	0.029539	0.005129	0.000403	0.031404
High	1579	0.053983	0.009051	-0.000062	0.048905
Panel B: Mean Comparison Tests					
Comparison	Volatility t-stat	p-value	Returns t-stat	p-value	
High vs Low	82.616321	0.000000	-10.673876	0.000000	

High vs Medium	50.574183	0.000000	-8.241623	0.000000	
Low vs Medium	-50.347460	0.000000	2.305563	0.000000	

Notes: The above table presents average volatility and returns across entropy-defined regimes based on the 30-day normalized rolling entropy. Regimes are determined using the 33rd and 66th percentiles of the entropy distribution. Panel A presents descriptive statistics. Panel B reports pairwise t-tests for equality of means. Volatility is measured as the 30-day rolling standard deviation of daily log returns. Returns refer to daily log returns. Given that volatility is computed using rolling windows, observations may exhibit some overlap. Nevertheless, the magnitude of the differences across regimes remains economically large.

The results reveal a clear monotonic ordering pattern across entropy regimes. High-entropy periods are associated with substantially higher volatility, medium-entropy periods exhibit intermediate volatility levels, and low-entropy periods correspond to the calmest market conditions. The magnitude of the differences is economically large and statistically significant across all pairwise comparisons. In economic terms, average volatility in the high-entropy regime is more than three times larger than in the low-entropy regime, thus highlighting the substantial risk differential associated with elevated structural disorder. This monotonic relationship suggests that entropy effectively captures shifts between calm and turbulent market regimes.

By contrast, average returns display an inverse ordering. High-entropy regimes are associated with weaker (actually negative) mean return performance, whereas low-entropy regimes exhibit comparatively stronger average returns. Differences in mean returns are statistically significant across pairwise comparisons for the 30-day specification. This pattern suggests that periods of elevated structural disorder are associated not only with increased risk but also with weaker return performance.

To assess robustness, we replicate the regime classification analysis using 15- and 60-day rolling entropy measures (Table 2). The monotonic ordering of volatility across regimes remains consistent across window lengths, verifying that high-entropy periods systematically correspond to higher market risk. Although differences in returns are less pronounced for the 60-day window, the overall pattern remains broadly stable. Taken together, these findings indicate that entropy-based regimes provide a robust and economically meaningful segmentation of Bitcoin market conditions, consistent with entropy capturing shifts in the underlying structural organisation of returns.

Table 2: Robustness of Entropy Regime Effects Across Rolling Window Lengths

Window	Comparison	Volatility t-stat	p-value	Return t-stat	p-value
15	High vs Low	32.51228489	0.0000	-2.749329429	0.0060
15	High vs Medium	21.06033997	0.0000	-2.193381204	0.0284
15	Low vs Medium	-42.05234081	0.0000	0.447172026	0.6548

60	High vs Low	32.58113281	0.0000	-1.025883938	0.3051
60	High vs Medium	21.32171825	0.0000	0.474982561	0.6348
60	Med vs Low	-54.06244808	0.0000	2.002394318	0.0453

Notes: The above table reports pairwise t-tests for differences in mean volatility and returns across entropy-defined regimes for alternative rolling window lengths (15 and 60 days). Regimes are constructed using the 33rd and 66th percentiles of the respective rolling entropy distributions. Volatility is measured as the 30-day rolling standard deviation of daily log returns. Returns refer to daily log returns. Given that volatility is computed using rolling windows, observations may exhibit some overlap. Nevertheless, the magnitude of the differences across regimes remains economically large.

4.3 Correlation and Regression Evidence

To further examine the relationship between entropy and market risk, we turn our attention to correlation and regression analysis. Table 3 reports Pearson correlations between normalized rolling entropy and rolling volatility, including forward volatility measures up to five days ahead.

Table 3: Correlation Between Rolling Entropy and Volatility

Window	$\text{Corr}(H_t, \sigma_t)$	$\text{Corr}(H_t, \sigma_{t+1})$	$\text{Corr}(H_t, \sigma_{t+2})$	$\text{Corr}(H_t, \sigma_{t+3})$	$\text{Corr}(H_t, \sigma_{t+5})$
15	0.686	0.679	0.668	0.653	0.613
30	0.720	0.718	0.716	0.711	0.696
60	0.736	0.736	0.735	0.734	0.731

Notes: The above table reports Pearson correlations between normalized rolling entropy (H) and rolling volatility (σ), including forward volatility measures ($\sigma + k$), with $k = 1, 2, 3$, and 5 days. All correlations are statistically significant at the 1% level.

The contemporaneous correlation between entropy and volatility is strongly positive across window lengths, ranging from 0.686 for the 15-day window to 0.736 for the 60-day window. The strength of the association increases with the window length, thus suggesting that longer entropy measures more effectively capture persistent structural uncertainty in the Bitcoin market.

Importantly, this positive relationship persists when volatility is measured at future horizons (+1, +2, +3, and +5 days), with correlations remaining economically large and statistically significant. For instance, under the 60-day specification, the correlation between entropy and five-day-ahead volatility remains above 0.73. Although correlation does not imply causality, the persistence of strong positive associations at forward horizons is consistent with entropy containing forward-looking information about volatility dynamics.

To further examine the above relationship, we also estimate linear regressions of absolute Bitcoin returns on lagged absolute returns and rolling entropy measures across different window lengths. Following a large empirical literature, absolute returns are commonly used as a simple proxy for daily volatility in empirical studies [27, 28]. We estimate two models, the

first one is an AR(1) baseline specification that includes only lagged absolute returns (equation 7), while the second one is an augmented specification that includes, in addition to lagged absolute returns, normalised rolling entropy calculated over the 15-, 30-, and 60-day windows (equation 8). Table 4 presents our results.

$$|r_t| = \beta_0 + \beta_1|r_{t-1}| + \varepsilon_t \quad (\text{Equation 7})$$

$$|r_t| = \beta_0 + \beta_1|r_{t-1}| + \beta_2H_t + \varepsilon_t \quad (\text{Equation 8})$$

Panel A reports the baseline AR(1) specification. The coefficient on lagged absolute returns is positive and highly statistically significant across all windows, with values ranging between 0.214 and 0.237. This confirms the presence of volatility persistence, consistent with the well-documented phenomenon of volatility clustering in financial markets [11,12]. The explanatory power of the AR(1) model, nevertheless, remains modest with adjusted R² values between 4.6% and 5.7%.

Table 4. Regressions of Absolute Returns on Lagged Returns and Rolling Entropy

Panel A: Model $|r_t| = \beta_0 + \beta_1|r_{t-1}| + \varepsilon_t$

Window	$\beta_1(r_{t-1})$	t-stat	p-value	Adjusted R ²
15	0.2141	5.725	0.000	0.0456
30	0.2141	5.720	0.000	0.0456
60	0.2374	8.671	0.000	0.0565

Panel B: Model $|r_t| = \beta_0 + \beta_1|r_{t-1}| + \beta_2H_t + \varepsilon_t$

Window	$\beta_1(r_{t-1})$	t-stat	p-value	$\beta_2(H_t)$	t-stat	p-value	Adjusted R ²
15	0.1318	3.940	0.000	0.0922	4.344	0.000	0.0917
30	0.1462	4.269	0.000	0.0957	4.682	0.000	0.0871
60	0.1721	5.431	0.000	0.0719	10.809	0.000	0.0969

Notes: The dependent variable is the absolute daily log return of Bitcoin ($|r_t|$). Panel A reports AR(1) regressions including only lagged absolute returns. Panel B augments the specification with normalized rolling entropy (H_t) computed over 15-, 30-, and 60-day windows. All models include a constant (not reported). Coefficients are estimated using OLS. t-statistics are based on Newey–West heteroskedasticity and autocorrelation consistent (HAC) standard errors (lag = 5). Adjusted R² is reported.

Panel B augments the model by incorporating normalized rolling entropy. The entropy coefficient is positive and statistically significant at the 1% level across all window lengths. Moreover, the inclusion of entropy increases the adjusted R² from roughly 4–6% in the baseline specification to approximately 9–10%, indicating substantial incremental explanatory power beyond simple volatility persistence. Interestingly, although the coefficient of lagged absolute returns remains significant, it declines once entropy is included, thus suggesting that entropy captures additional distributional information not fully reflected in past volatility. Across window lengths, the 60-day entropy measure yields the strongest statistical significance, with shorter windows remaining robust. These findings provide consistent evidence that rolling entropy can function as a meaningful dynamic risk indicator, offering complementary information to traditional volatility-based predictors.

Beyond statistical significance, the interpretation of these findings is noteworthy also from an economic perspective. While volatility measures capture the magnitude of return fluctuations, entropy reflects the degree of distributional disorder in the return-generating process. In this sense, entropy provides a distribution-based measure of uncertainty that may complement traditional second-moment volatility indicators. In a market such as Bitcoin that is characterized by structural breaks, liquidity shocks, and rapid shifts in investor sentiment, elevated entropy likely signals increasing instability in the underlying distribution rather than merely higher dispersion.

Overall, the strong and persistent correlations between entropy and forward volatility, along with the substantial increase in adjusted R^2 when entropy is included, suggest that entropy captures evolving regime conditions beyond simple volatility persistence. The decline in the AR(1) coefficient once entropy is introduced further indicates that part of observed volatility clustering may indeed reflect underlying distributional complexity. In this sense, rolling entropy appears to function as a dynamic indicator of regime transitions and concealed risk accumulation in the Bitcoin market.

4.4 Out-of-Sample Forecasts

To assess whether rolling entropy contains predictive information beyond the in-sample relation documented in the previous section, we carry out a comprehensive out-of-sample forecast experiment. More specifically, we compare a benchmark model that includes lagged absolute returns, capturing volatility persistence in absolute Bitcoin returns, with an augmented model specification that also incorporates lagged entropy. Forecast accuracy is evaluated at 15-, 30-, and 60-day horizons using mean squared forecast errors (MSE).

Two complementary estimation approaches are utilised. Firstly, we employ recursive estimation where the model parameters are re-estimated sequentially using only information available at each forecast origin. This approach mirrors real-time forecasting and allows the parameters to adapt to evolving market conditions. Secondly, we use a fixed 2/3 - 1/3 sample split, where the first two-thirds of the sample are used for parameter estimation and the remaining one-third for the evaluation of the forecast. This alternative design provides a robustness check that allows for the assessment of the stability of the entropy effect when parameters are estimated over a stable historical window.

The results are reported in Table 5. Under recursive estimation, the entropy-augmented model consistently outperforms the benchmark across all forecast horizons. MSE reductions range from 5.23% to 6.8% thus indicating that entropy provides incremental predictive information that spans beyond volatility persistence. Although the magnitude of these improvements is moderate, they are uniform across horizons and estimation windows. Under the fixed 2/3 - 1/3 split, the predictive gains become substantially larger with the entropy-augmented model reducing forecast errors by 14.19% to 17.27% across horizons. The larger gains under the fixed estimation approach suggest that entropy captures relatively persistent structural characteristics of return dynamics, which are effectively exploited when parameters are estimated over a stable historical window.

Given that the entropy model nests the benchmark specification, we evaluate the statistical significance of improvement by using the Clark–West test [26], which corrects for the bias that arises in nested forecast comparisons. The Clark–West statistics are strongly significant at the 1% level across all horizons and under both estimation approaches. These results confirm that the observed reductions in forecast errors reflect genuine incremental predictive content rather than sampling variation.

Table 5: Out-of-Sample Forecasting Performance - Models With, and Without Entropy

Horizon	Method	MSE (Benchmark)	MSE (Entropy)	Improvement (%)	CW t-stat	
15-day	Recursive	0.000631	0.000598	5.23%	11.05	***
30-day	Recursive	0.000632	0.000589	6.80%	12.46	***
60-day	Recursive	0.000620	0.000584	5.81%	13.97	***
15-day	Fixed 2/3 - 1/3	0.000473	0.000403	14.80%	15.99	***
30-day	Fixed 2/3 - 1/3	0.000469	0.000388	17.27%	16.69	***
60-day	Fixed 2/3 - 1/3	0.000437	0.000375	14.19%	17.39	***

Notes: The above table reports out-of-sample forecast performance for a baseline volatility model including lagged absolute returns and an augmented model that additionally incorporates lagged entropy. Forecast accuracy is evaluated using mean squared forecast errors (MSE) at 15-, 30-, and 60-day horizons. Two estimation schemes are implemented. Under recursive estimation, model parameters are re-estimated sequentially using all information available up to each forecast origin. Under the fixed 2/3 - 1/3 split, the first two-thirds of the sample are used for parameter estimation and the remaining one-third for forecast evaluation. “Improvement (%)” denotes the percentage reduction in MSE of the entropy-augmented model relative to the benchmark. Statistical significance is assessed using the Clark–West (2007) test for nested forecast models, which corrects for the upward bias in the larger model’s prediction error. Reported statistics correspond to one-sided tests of equal predictive accuracy against the alternative that the entropy model yields lower forecast errors.

*** indicates significance at the 1% level.

From a complexity perspective, entropy may be interpreted as a quantification of the degree of disorder and heterogeneity in the underlying return dynamics. Periods of high entropy correspond to increased structural uncertainty, fragmented information, and heterogeneous trading behavior, which in financial markets may translate as increased volatility risk, unstable expectations formation, and weaker coordination among traders. The consistent out-of-sample gains therefore suggest that entropy captures latent structural features of market organization that are not fully reflected in past volatility alone.

Overall, the out-of-sample evidence provides strong support for the additional predictive value of entropy in modelling Bitcoin volatility dynamics. The improvements are economically meaningful, robust from a statistical perspective and consistent across horizons and estimation methods. These findings reinforce the interpretation that entropy captures persistent structural complexity in the Bitcoin return-generating process and lend support to the usefulness of complexity-based measures in a financial forecasting framework.

5. Conclusion

This paper investigates whether normalized rolling entropy can provide a meaningful measure of structural complexity in Bitcoin returns, as well as whether it may offer predictive information beyond traditional volatility-based indicators. Using daily Bitcoin prices from February 2014 to February 2026, we construct 15-, 30-, and 60-day normalized rolling entropy measures and systematically evaluate their empirical relevance through descriptive analysis, regime classification, regression models, and out-of-sample forecasting experiments.

Our results reveal a number of important insights. First, rolling entropy exhibits substantial temporal variation, with peaks coinciding with periods of sharp price movements and market turbulence, thus indicating its ability to capture evolving market structure. Second, entropy is strongly and positively correlated with both contemporaneous and forward-looking volatility, hence suggesting that it reflects persistent hidden risk, beyond that captured by conventional variance measures. Third, entropy-based regime classification produces economically and statistically distinct market states; specifically, high-entropy periods correspond to elevated volatility and weaker average returns, whereas low-entropy periods coincide with calmer market conditions and stronger returns. Fourth, regression analysis shows that entropy provides significant incremental explanatory power spanning beyond AR(1) volatility models, with the inclusion of entropy approximately doubling the adjusted R^2 across all window lengths examined. Finally, out-of-sample forecasting experiments demonstrate that entropy-augmented models reduce mean squared forecast errors by 5–17% relative to benchmark models, with improvements being robust across horizons and estimation methods. Taken together, these findings indicate that entropy provides complementary information about market complexity beyond traditional second-moment volatility measures.

This study contributes to the literature in several ways. Firstly, it develops and evaluates normalized rolling entropy as a measure of the distributional complexity in Bitcoin returns. Secondly, it demonstrates that entropy provides incremental explanatory and predictive value relative to standard volatility benchmarks. Thirdly, by integrating information-theoretic measures with volatility modelling and forecast evaluation, the analysis connects concepts from statistical physics and econophysics with empirical applications in financial markets.

Although our results provide evidence that entropy-based measures capture meaningful aspects of market dynamics, there are certain limitations that should be acknowledged. First, the analysis focuses exclusively on Bitcoin returns, and therefore the extent to which the findings generalize to other cryptocurrencies or traditional financial assets remains an open question. Second, the entropy measures are constructed using rolling-window estimates, which may introduce sensitivity to window length and sampling choices. Finally, despite the fact that the forecasting results indicate improvements relative to benchmark volatility models, the magnitude of predictive gains may vary across alternative model specifications and market conditions.

Looking towards future research avenues, these could extend this analysis to other cryptocurrencies and financial assets, explore alternative entropy measures such as Rényi [20] or Tsallis [21] entropy, or combine entropy with high-frequency realized volatility for a more

granular assessment of market dynamics. Overall, the evidence suggests that rolling entropy captures persistent structural features of Bitcoin return distributions that are not fully summarized by conventional volatility measures, thus offering a robust tool for better understanding and quantifying the complex dynamics of digital asset markets.

References:

- [1] R.N. Mantegna, H.E. Stanley, *An Introduction to Econophysics: Correlations and Complexity in Finance*, Cambridge University Press, Cambridge, 2000.
- [2] D. Begušić, Z. Kostanjčar, H.E. Stanley, B. Podobnik, Scaling properties of extreme price fluctuations in Bitcoin markets, *Physica A* 510 (2018) 400–406.
- [3] G.M. Caporale, A. Zekokh, Regime switches and commonalities of the cryptocurrencies asset class, *North Am. J. Econ. Finance* 58 (2021) 101489.
- [4] C.E. Shannon, A mathematical theory of communication, *Bell Syst. Tech. J.* 27 (1948) 379–423, 623–656.
- [5] R. Cont, Empirical properties of asset returns: Stylized facts and statistical issues, *Quant. Finance* 1 (2001) 223–236.
- [6] A. Phillip, J.S.K. Chan, S. Peiris, A new look at cryptocurrencies, *Econ. Lett.* 163 (2018) 6–9.
- [7] A.F. Bariviera, The inefficiency of Bitcoin revisited: A dynamic approach, *Econ. Lett.* 161 (2017) 1–4.
- [8] T. Bollerslev, Generalized autoregressive conditional heteroskedasticity, *J. Econometrics* 31 (1986) 307–327.
- [9] D. Sornette, *Why Stock Markets Crash: Critical Events in Complex Financial Systems*, Princeton University Press, Princeton, 2003.
- [10] X. Gabaix, P. Gopikrishnan, V. Plerou, H.E. Stanley, A theory of power-law distributions in financial market fluctuations, *Nature* 423 (2003) 267–270.
- [11] R.N. Mantegna, H.E. Stanley, Scaling behaviour in the dynamics of an economic index, *Nature* 376 (1995) 46–49.
- [12] C. Katsiampa, Volatility estimation for Bitcoin: A comparison of GARCH models, *Econ. Lett.* 158 (2017) 3–6.
- [13] J. Chu, S. Chan, S. Nadarajah, J. Osterrieder, GARCH modelling of cryptocurrencies, *J. Risk Financial Manag.* 10 (2017) 17.
- [14] E. Maasoumi, A compendium to information theory in economics and econometrics, *Econometric Rev.* 12 (1993) 137–181.
- [15] J. Alvarez-Ramirez, E. Rodriguez, J.C. Echeverria, Permutation entropy analysis of financial time series, *Physica A* 389 (2010) 5228–5236.

- [16] L. Zunino, M. Zanin, B.M. Tabak, D.G. Pérez, O.A. Rosso, Complexity-entropy causality plane: A useful approach to quantify the stock market inefficiency, *Physica A* 389 (2010) 1891–1901.
- [17] P. Fiedor, Information-theoretic approach to lead-lag effect on financial markets, *Physica A* 413 (2014) 518–531.
- [18] W.A. Risso, The informational efficiency and the financial crashes, *Res. Int. Bus. Finance* 22 (2008) 396–408.
- [19] Z.-Q. Jiang, W.-J. Xie, M.-X. Li, B. Podobnik, W.-X. Zhou, H.E. Stanley, Measuring complexity and information flow in financial markets, *Entropy* 23 (2021) 1–20.
- [20] A. Rényi, On measures of entropy and information, in: *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, Berkeley, 1961, pp. 547–561.
- [21] C. Tsallis, Possible generalization of Boltzmann–Gibbs statistics, *J. Stat. Phys.* 52 (1988) 479–487.
- [22] A. Sensoy, The inefficiency of Bitcoin revisited: A high-frequency analysis with alternative currencies, *Finance Res. Lett.* 28 (2019) 68–73.
- [23] A.K. Tiwari, R.K. Jana, D. Das, D. Roubaud, Informational efficiency of Bitcoin—An extension, *Econ. Lett.* 163 (2018) 106–109.
- [24] F.X. Diebold, R.S. Mariano, Comparing predictive accuracy, *J. Bus. Econ. Stat.* 13 (1995) 253–263.
- [25] K.D. West, Asymptotic inference about predictive ability, *Econometrica* 64 (1996) 1067–1084.
- [26] T.E. Clark, K.D. West, Approximately normal tests for equal predictive accuracy in nested models, *J. Econometrics* 138 (2007) 291–311.
- [27] G.W. Schwert, Stock volatility and the crash of '87, *Review of Financial Studies* 3 (1990) 77–102.
- [28] Z. Ding, C.W.J. Granger, R.F. Engle, A long memory property of stock market returns and a new model, *Journal of Empirical Finance* 1 (1993) 83–106.