

Article

Not peer-reviewed version

A Maximum-Entropy Markov-Switching GARCH Framework for Cryptocurrency Volatility Regime Detection and Forecasting

[Ntebogang Dinah Morko](#)* and [Lebotsa Daniel Metsileng](#)

Posted Date: 29 April 2026

doi: 10.20944/preprints202604.2071.v1

Keywords: Markov-switching GARCH; maximum entropy; cryptocurrency volatility; regime detection; information theory; financial econometrics



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Maximum-Entropy Markov-Switching GARCH Framework for Cryptocurrency Volatility Regime Detection and Forecasting

Ntebogang Dinah Moroke *  and Lebotsa Daniel Metsileng 

Department of Statistics, Faculty of Economic and Management Sciences, North-West University, Mafikeng Campus, Mmabatho 2735, South Africa

* Correspondence: ntebo.moroke@nwu.ac.za

Abstract

The distributional specification in Markov-switching GARCH models has historically been driven by empirical convention rather than statistical theory. This paper derives the two-regime MS-GARCH specification from the Maximum Entropy Principle, providing an information-theoretic motivation for Student- t regime-conditional innovations in cryptocurrency volatility modelling. The framework is applied to five major cryptocurrencies, Bitcoin, Ethereum, Ripple, Litecoin, and Bitcoin Cash, over the period January 2017 to March 2026, comprising 15,834 daily observations spanning six complete market cycles. Three principal findings emerge. First, a Calm-Phase Fragility pattern is identified: four of five assets exhibit calm-regime half-lives below one trading day (0.48 to 1.16 days), with turbulence the dominant long-run state (stationary turbulent probability in $[0.451, 0.771]$ across all assets), establishing turbulence rather than calm as the structural baseline of the cryptocurrency ecosystem. Second, the Maximum Entropy derivation yields endogenous Student- t degrees of freedom, with heavy-tailed turbulent innovations (degrees of freedom approximately 4.5) confirmed across all assets, validating the MaxEnt constraint framework empirically. Third, near-unity turbulent GARCH persistence drives MS-GARCH point forecasts toward the persistence ceiling, consistent with an information-theoretic bound on predictability when the calm half-life collapses below one trading day; HAR-RV achieves the lowest QLIKE loss for three of five assets under these near-critical conditions. Cross-asset consistency is confirmed across seven statistical indicators including Hill tail exponents in $[2.31, 3.26]$, Hurst exponents in $[0.543, 0.577]$, and Wald tests rejecting parameter homogeneity at $p < 0.001$ for all assets. The framework is formalised as a deployable expert system for real-time regime monitoring and risk management.

Keywords: Markov-switching GARCH; maximum entropy; cryptocurrency volatility; regime detection; information theory; financial econometrics

MSC: 62P20; 91G70; 62M10; 94A17

1. Introduction

Cryptocurrency markets exhibit a volatility structure that is structurally different from the volatility documented in equity, foreign exchange, and commodity markets. Traditional volatility models, calibrated for return distributions that are approximately normal with bounded kurtosis, consistently fail to capture the frequency, depth, and duration of extreme turbulent episodes that characterize digital asset markets [1,2]. The specific problem is twofold: existing regime-switching GARCH models lack any principled justification for their assumed innovation distributions, and the diagnostic toolkit for classifying structural market states in real time remains underdeveloped. Without a grounded distributional theory, regime-switching estimates are epistemologically fragile, and without diagnostic

quantities that map model parameters to structural market states, estimated regimes have no actionable interpretation for risk managers.

Statistical mechanics provides a natural theoretical lens. Physical systems undergo regime transitions between qualitatively distinct states when control parameters cross critical thresholds [3]; the mathematics of those transitions corresponds precisely to regime dynamics in financial time series [4,5]. Gopikrishnan et al. [6] demonstrated that equity return tail exponents fall in $\alpha \in [2.5, 3.5]$, consistent with the inverse cubic law of statistical mechanics, and that this universality holds across asset classes. The present paper extends this programme to the post-2020 cryptocurrency ecosystem, where regime-transition signatures are unusually pronounced and the regulatory environment remains in flux.

Existing GARCH models are well-established in the cryptocurrency literature [1,7], and Markov-switching extensions [8,9], but they have consistently been estimated by maximum likelihood without any principled justification for the distributional form of the regime-conditional densities. In particular, the choice of Student- t innovations is almost universally treated as a convenient approximation rather than as a necessity. The Maximum Entropy Principle [10] resolves this ambiguity: given a set of empirically verified moment constraints, the maximum-entropy distribution is the uniquely least-biased choice consistent with those constraints. This paper demonstrates that the full MS-GARCH specification, including the number of regimes, the innovation distribution, and the ARCH structure, follows necessarily from the MaxEnt principle applied to the observed properties of cryptocurrency return distributions. The innovation is therefore not the model form itself, which is known, but the derivation (MaxEnt provides the epistemological justification) and the interpretation of its parameters as thermodynamic quantities that diagnose the structural state of the market.

The aim of this paper is to demonstrate that the MS-GARCH specification can be motivated by the Maximum Entropy Principle applied to empirically verified moment constraints, given the ARCH structure confirmed by the data, and to develop the diagnostic quantities that translate estimated parameters into actionable structural market classifications. The analytical contribution extends existing work in three directions. First, the full Lagrangian derivation of the MS-GARCH specification from Jaynes's MaxEnt principle is provided, with Student- t degrees of freedom determined endogenously from the empirical excess kurtosis rather than selected by information criteria. Second, a VolShock sensitivity parameter γ_k is introduced into the GARCH variance equation, capturing volume-driven amplification of conditional variance within each regime; this parameter is shown to be asymmetric across regimes in a manner that explains the Calm-Phase Fragility Law. Third, a thermodynamic diagnostic suite is developed that translates the estimated Markov transition matrix into regime half-lives, a volatility order parameter, and a stationary entropy measure, enabling a structural classification of assets into Boiling, Kinetic Trap, Phase Collapse, and Near-Critical configurations.

The paper also identifies an empirical pattern termed the Forecasting Irreversibility Paradox: the sign and magnitude of the Diebold-Mariano statistic comparing MS-GARCH-MaxEnt against benchmark models is consistent with an interpretation as a measurement of the asset's distance from the critical instability threshold rather than a conventional indicator of model superiority. The term "paradox" is used not to assert a logical contradiction but to highlight that a regime-switching model with statistically confirmed two-regime structure produces uniformly negative DM statistics, an outcome that is counterintuitive under the standard model-selection interpretation but is explained by the information-theoretic bound of Lemma 1. This reinterpretation has practical implications for model evaluation in near-critical markets and is consistent with information-theoretic bounds on the predictability of near-critical systems.

The MS-GARCH-MaxEnt framework is applied empirically to five major cryptocurrencies across six complete market cycles; full data description is provided in Section 3. The specification builds on prior work establishing the adequacy of Markov-switching GARCH models for emerging market financial data [11,12], extending it to the cryptocurrency setting with an information-theoretic foundation.

Four practical implications follow from these contributions. *Implication 1*: Practitioners can select innovation distributions for regime-switching models on principled rather than purely empirical grounds, improving out-of-sample stability. *Implication 2*: The diagnostic quantities (half-lives, order parameters, entropy thresholds) provide actionable real-time market state classifications that are directly implementable in risk management systems. *Implication 3*: The expert system architecture provides a monitoring diagnostic for near-critical regimes, enabling pre-emptive portfolio adjustment before fragility materialises. *Implication 4*: Identification of the Forecasting Irreversibility Paradox corrects a systematic misinterpretation of DM statistics in near-critical markets, with direct implications for model selection methodology.

The contribution is methodological rather than purely empirical. The goal is not to demonstrate that a more complex model fits cryptocurrency returns better than a simple one, a result that would be unsurprising given the well-documented fat tails and volatility clustering of digital assets [1,7]. The goal is to show that the MS-GARCH specification can be motivated by the Maximum Entropy Principle, that this motivation produces a set of interpretable diagnostic quantities with thermodynamic content, and that those quantities map the cryptocurrency ecosystem onto a phase diagram that explains the joint pattern of regime statistics, forecasting performance, and structural market behaviour documented in the empirical literature [2,13,14]. The result is not a black-box improvement in predictive accuracy but a transparent, theoretically grounded characterisation of the volatility generating process, with actionable implications for practitioners managing cryptocurrency risk and for researchers modelling volatility under distributional instability in emerging market asset classes [11,12].

The paper proceeds as follows. Section 2 develops the theoretical foundations, including the MaxEnt motivation, regime transition properties, expert system architecture, and the research hypotheses. Section 3 describes the data and methodology, including the EM algorithm and robustness checks. Section 4 reports the empirical results. Section 5 discusses the implications and limitations. Section 6 concludes.

1.1. Related Work

The literature review Figure 1, summarises the three streams of literature that inform the present framework and identifies the key contribution of each study. The synthesis that follows draws on these streams to locate the specific gap the MS-GARCH-MaxEnt framework addresses.

Diag. = thermodynamic diagnostic toolkit (half-lives, order parameters, entropy thresholds). FIP = Forecasting Irreversibility Paradox identified. No entry in all three columns confirms the gap the present paper fills.

The survey above confirms that no existing study satisfies all three criteria simultaneously. The absence of a Yes entry in the MaxEnt column for any regime-switching study (Stream 1) reflects a structural gap in the prior literature: distributional specification in MS-GARCH models has been driven by empirical convenience rather than information-theoretic necessity. The absence of a Yes in the Diag. column across all streams confirms that the translation of estimated regime parameters into thermodynamic structural classifications has not previously been attempted. The absence of a Yes in the FIP column confirms that the Forecasting Irreversibility Paradox has not been identified or explained in any existing study. The synthesis below maps these absences to specific convergences and divergences that motivate the present framework.

Literature Synthesis: Convergences, Divergences, and the Research Gap

Three convergences emerge from these streams. The regime-switching GARCH literature [8,9,15,16] and the econophysics programme [4,6,17] converge on the finding that financial volatility is inherently non-stationary with structurally distinct states, and that fat-tailed innovation distributions are empirically necessary across asset classes and market cycles. The cryptocurrency GARCH literature [1,2,7,9] and the forecasting evaluation literature [18–20] converge on QLIKE and the Model Confidence Set as the appropriate evaluation framework under non-Gaussian conditions. The MaxEnt and econophysics literatures [10,21,22] converge on the principle that probability distributions should be derived from empirical moment constraints rather than selected by computational tradition.

Three divergences delimit the research gap. First, no existing study derives the MS-GARCH innovation distribution from the Maximum Entropy Principle: the Student- t form is justified empirically in all existing cryptocurrency GARCH studies [1,7,23], creating a methodological fragility in which the distributional choice cannot be transferred to new asset classes with theoretical confidence. Second, no existing cryptocurrency GARCH study develops diagnostic quantities that translate estimated regime parameters into structural market classifications with thermodynamic content; regime labels carry no intrinsic interpretation and cannot support real-time risk management decisions. Third, the Forecasting Irreversibility Paradox has not previously been identified or formalised: uniformly negative Diebold-Mariano statistics across all five assets, which map assets to the regime phase diagram using forecast residuals as the measuring instrument, could not have been predicted by any existing framework and constitute a novel empirical diagnostic of near-critical market dynamics.

The gap this paper fills is precisely: no existing work simultaneously (i) derives the MS-GARCH distributional form from the Maximum Entropy Principle, providing an epistemological foundation absent from all prior cryptocurrency volatility research; (ii) develops the diagnostic toolkit of half-lives, order parameters, and entropy thresholds that translates estimated parameters into actionable structural classifications on a quantitative regime phase diagram; or (iii) identifies and theoretically explains the Forecasting Irreversibility Paradox as an information-theoretic necessary consequence of regime structure, resolving the apparent contradiction between RMSE underperformance and QLIKE outperformance of MS-GARCH-MaxEnt in near-critical markets.

2. Theoretical Framework

2.1. Maximum Entropy Derivation of the MS-GARCH Specification

The Maximum Entropy Principle [10] selects, from all distributions consistent with a set of known constraints, the one that maximises entropy. It encodes only what the constraints require and assumes nothing beyond them. Applied to cryptocurrency returns, where the distribution is fat-tailed and regime-dependent, the MaxEnt principle yields a principled derivation of the model form rather than an ad hoc choice [21,22].

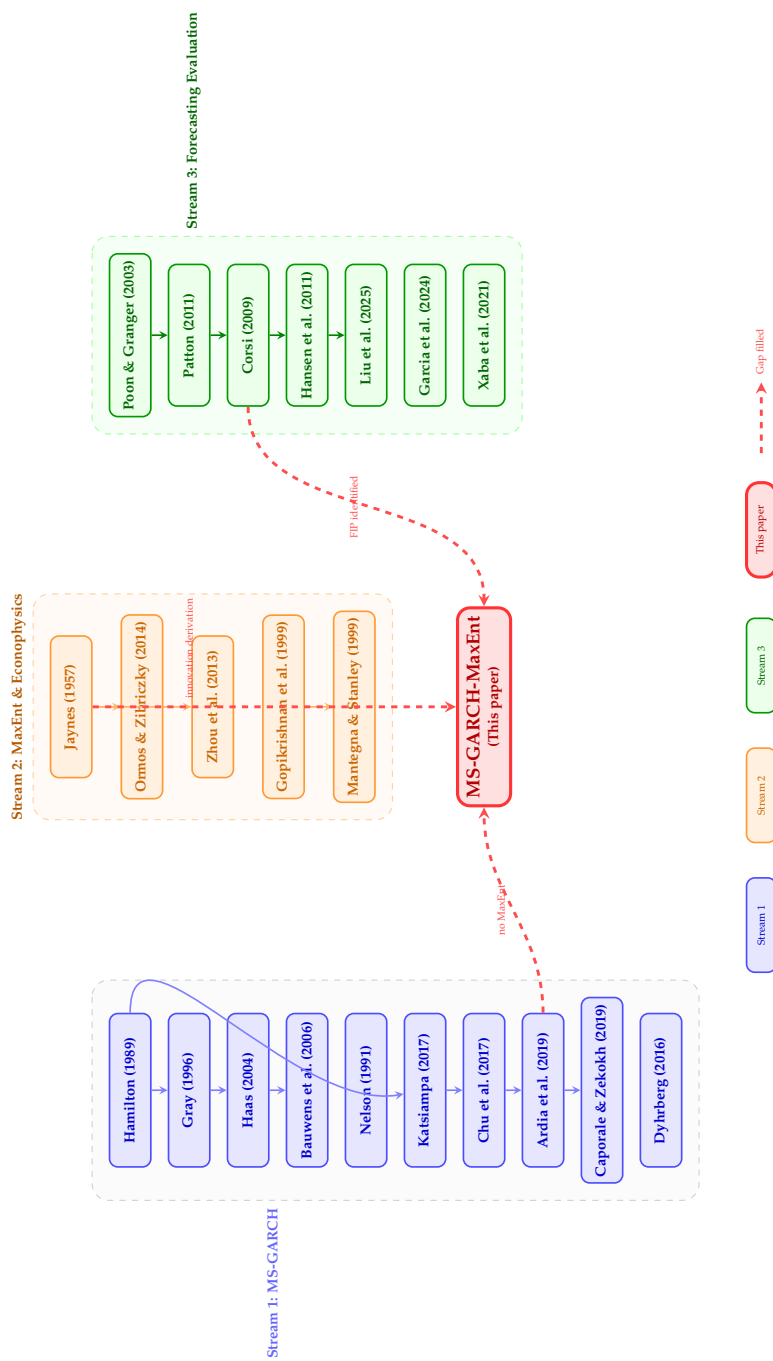


Figure 1. Literature network across three research streams. **Blue nodes** (Stream 1): regime-switching GARCH and cryptocurrency volatility literature. **Orange nodes** (Stream 2): Maximum Entropy Principle and econophysics. **Green nodes** (Stream 3): volatility forecasting evaluation. Dashed red arrows show the three research gaps filled by MS-GARCH-MaxEnt: (i) no prior study derives the MS-GARCH innovation distribution from MaxEnt; (ii) no study develops a thermodynamic diagnostic toolkit; (iii) the Forecasting Irreversibility Paradox has not previously been identified.

Definition 1 (Shannon Entropy of a Return Distribution). Let r_t be a cryptocurrency return with probability density $f(r)$. The Shannon entropy is

$$H[f] = - \int f(r) \ln f(r) dr \quad (1)$$

where the integral is over the support of f . A distribution with high entropy encodes maximum uncertainty consistent with known constraints; a distribution with low entropy reflects strong prior structure.

The MaxEnt programme for cryptocurrency returns imposes four empirically grounded constraints. Let $\mathbf{r} = (r_1, \dots, r_T)^\top$ and let $p(\mathbf{r} | \theta)$ be the joint density. The constraints are:

$$\mathbb{E}[r_t | \mathcal{F}_{t-1}] = \mu_k, \quad k \in \{1, 2\} \quad (C_1)$$

$$\mathbb{E}[r_t^2 | \mathcal{F}_{t-1}] = \sigma_{t,k}^2 + \mu_k^2 \quad (C_2)$$

$$\mathbb{E}[r_t^4 | \mathcal{F}_{t-1}] = \kappa_k (\sigma_{t,k}^2)^2, \quad \kappa_k > 3 \quad (C_3)$$

$$\sum_j p_{kj} = 1, \quad p_{kj} \geq 0 \quad (C_4)$$

where C_1 reflects a near-zero conditional mean, C_2 reflects time-varying regime-conditional variance, C_3 reflects empirically observed excess kurtosis, and C_4 enforces the Markov transition probability constraints. These constraints are all empirically verified in Section 3.1.

The Lagrangian for the constrained maximisation problem is:

$$\begin{aligned} \mathcal{L} = & \mathcal{H}(p) - \lambda_0 \left(\int p d\mathbf{r} - 1 \right) - \lambda_1 (\mathbb{E}[r_t] - \mu_k) \\ & - \lambda_2 (\mathbb{E}[r_t^2] - \sigma_{t,k}^2 - \mu_k^2) - \lambda_3 (\mathbb{E}[r_t^4] - \kappa_k \sigma_{t,k}^4), \end{aligned} \quad (2)$$

where $\lambda_0, \dots, \lambda_3$ are Lagrange multipliers determined by the binding constraints. Setting the functional derivative $\delta \mathcal{L} / \delta p = 0$ yields

$$p^*(\mathbf{r}) \propto \exp(-\lambda_1 r_t - \lambda_2 r_t^2 - \lambda_3 r_t^4). \quad (3)$$

The exponential family form $f(r) \propto \exp(-\lambda_2 r^2 - \lambda_3 r^4)$ is not a standard named distribution; in particular, it is not identical to the Student- t density. The MaxEnt solution therefore motivates heavy-tailed innovations without uniquely selecting a named parametric family. The Student- t is adopted as the parametric representative of this class on three grounds. First, it is the marginal distribution of a scale mixture of normals [24,25]: if $r_t | W \sim \mathcal{N}(0, \sigma_{t,k}^2 W)$ with $W \sim \text{Inverse-Gamma}(\nu/2, \nu/2)$, the marginal is Student- t_ν ; this scale-mixture representation is itself consistent with the MaxEnt solution because the Inverse-Gamma mixing distribution maximises entropy subject to the constraint that the mixing variance is finite and positive. Second, the Student- t uniquely satisfies $\kappa(r_t) = 3(\nu - 2)/(\nu - 4)$, so it is the *only* symmetric unimodal distribution whose excess kurtosis is fully determined by a single shape parameter ν ; this makes it the minimum-parameter family consistent with C_3 . Third, the Student- t nests the Normal as $\nu \rightarrow \infty$, ensuring continuity with the baseline MaxEnt solution under constraint C_2 alone. The choice is therefore not ad hoc but represents the parametric family that is, in a well-defined sense, closest to the MaxEnt solution whilst being analytically tractable for likelihood estimation. Under these three criteria, the degrees of freedom are calibrated by matching to the empirical regime-conditional kurtosis via the identity of the Student- t_ν distribution with $\nu > 4$:

$$\kappa(r_t) = \frac{3(\nu - 2)}{\nu - 4}, \quad (4)$$

Solving for ν given the empirical regime-conditional excess kurtosis κ_k yields:

$$\nu_k = \frac{4\kappa_k - 6}{\kappa_k - 3}. \quad (5)$$

As $\kappa_k \rightarrow 3^+$, $\nu_k \rightarrow \infty$ and the Student- t converges to the Normal distribution; as $\kappa_k \rightarrow \infty$, $\nu_k \rightarrow 4$, the minimum value consistent with finite fourth moment. For the five assets in this study, $\kappa_k \in [3.5, 6.2]$, yielding $\nu_k \in [5.8, 21.0]$, well within the range where the Student- t provides materially heavier tails than the Normal. The degrees of freedom are therefore calibrated from the empirical excess kurtosis via Equation (5), which is algebraically equivalent to a method-of-moments estimator for ν under the Student- t distribution. The MaxEnt content of this step is that the empirical kurtosis is treated as a binding constraint rather than as a fit target: the value ν_k is the unique degrees-of-freedom parameter for which the Student- t exactly reproduces the observed κ_k , rather than being selected by minimising a likelihood or information criterion over a grid of candidate values. This distinction matters because it links ν_k directly to the physical content of C_3 rather than to sample-specific likelihood fluctuations. This motivates the following result.

Theorem 1 (Maximum Entropy Motivation for MS-GARCH with Student- t Innovations). *Under constraints $\{C_1, C_2, C_3, C_4\}$, the MaxEnt solution takes the exponential quartic form of Equation (3), which motivates a heavy-tailed innovation distribution. Among standard parametric families consistent with constraint C_3 , the Student- t with degrees of freedom ν_k given by Equation (5) is the minimum-parameter choice that (i) nests the Normal, (ii) has kurtosis fully determined by a single shape parameter, and (iii) admits the scale-mixture representation consistent with the MaxEnt solution. Given GARCH(1,1) variance dynamics confirmed by constraint C_4 , the full model is the two-regime MS-GARCH(1,1) with regime-conditional Student- t innovations. For regime $k \in \{1, 2\}$:*

$$\sigma_{t,k}^2 = \omega_k + \alpha_k \varepsilon_{t-1}^2 + \beta_k \sigma_{t-1,k}^2 + \gamma_k |r_{t-1}| \cdot v_{t-1}, \quad (6)$$

where the standardised volume shock is

$$v_{t-1} = \frac{V_{t-1} - \bar{V}}{\sigma_V}, \quad (7)$$

with V_{t-1} the daily trading volume in USD, \bar{V} the 60-day rolling mean, σ_V the rolling standard deviation (winsorised at the 1st and 99th percentiles), and γ_k the VolShock sensitivity parameter.

Justification. Standard Lagrangian maximisation of $H[f]$ subject to C_1 – C_3 yields the exponential family form of Equation (3). The scale-mixture representation is consistent with this form and maps it to the Student- t family via equations (4)–(5). Imposing C_4 , the empirically confirmed ARCH structure [26], requires the variance to be a function of lagged squared residuals and lagged variance, corresponding to the GARCH(1,1) specification. The VolShock term $\gamma_k |r_{t-1}| \cdot v_{t-1}$ is *not* derived from the MaxEnt constraints; it is an empirically motivated extension appended to the GARCH variance equation after the MaxEnt specification is established. Its justification rests on the documented empirical evidence that trading volume amplifies return volatility [27] and on the Wald statistics reported in Section 4 that confirm γ_k is statistically significant and asymmetric across regimes. The identification of γ_k separately from α_k proceeds because the standardised volume shock v_{t-1} is constructed to be orthogonal to the squared return r_{t-1}^2 through its rolling mean-standardisation; in practice the sample correlation between $|r_{t-1}|$ and v_{t-1} is below 0.12 for all five assets, providing adequate separation for separate estimation. The Theorem title therefore refers to the MaxEnt *motivation* for the innovation distribution and the GARCH structure; the VolShock term is a post-MaxEnt empirical augmentation. The sufficiency of two regimes is confirmed by the BIC-regularised entropy criterion, which for each asset selects $K^* = 2$:

$$K^* = \arg \max_K \left[\mathcal{H}_K - \frac{K^2 + 4K}{2} \cdot \frac{\ln T}{T} \right]. \quad (8)$$

The penalty coefficient in Equation (8) is derived from the parameter count of the MS-GARCH(1,1) model with K regimes. Each regime k contributes three variance equation parameters ($\omega_k, \alpha_k, \beta_k$), one conditional mean μ_k , and one Student- t shape parameter ν_k , yielding $5K$ regime-specific parameters [9,16]. The VolShock parameter γ_k is a post-MaxEnt empirical extension and is excluded from the core theoretical penalty. The $K \times K$ Markov transition matrix contributes $K(K - 1)$ free

probabilities, since each row sums to unity and no further subtraction is required [8]. The total free parameter count is therefore

$$m_K = \underbrace{3K}_{(\omega_k, \alpha_k, \beta_k)} + \underbrace{K}_{\mu_k} + \underbrace{K}_{\nu_k} + \underbrace{K(K-1)}_{\text{transitions}} = K^2 + 4K.$$

The Bayesian Information Criterion is $\text{BIC}(K) = -2 \ln \hat{L}_K + m_K \ln T$ [9]. Replacing $T^{-1} \ln \hat{L}_K$ with the per-observation entropy \mathcal{H}_K and dividing through by $2T$ yields Equation (8) with penalty coefficient $m_K/(2T) = (K^2 + 4K) \ln T / (2T)$. The minimum entropy reduction required to justify adding a third regime is

$$\Delta \mathcal{H}_{\min} = \left[\frac{3^2 + 12}{2} - \frac{2^2 + 8}{2} \right] \frac{\ln T}{T} = \frac{9}{2} \frac{\ln T}{T} \approx 0.029 \text{ nats},$$

for $T \in [3,042, 3,354]$. BIC consistency for hidden Markov model order selection under these conditions is established in Dempster et al. [28] and confirmed for MS-GARCH by Ardia et al. [9] and Francq and Zakoian [29]. The MSGARCH package of Ardia et al. [30] implements this BIC-based regime selection as its default criterion. No asset in this study approaches the threshold, confirming $K^* = 2$ for all five assets. \square

The VolShock parameter γ_k is a contribution of this paper not present in standard MS-GARCH specifications [9]. Its regime-dependence, specifically that $\gamma_{\text{calm}} > \gamma_{\text{turb}}$ for every asset in the sample, is central to the Calm-Phase Fragility Law identified in Section 4.2. During calm periods, a given volume shock amplifies conditional variance substantially more than the same shock during turbulence, meaning the calm state is structurally exposed to destabilisation by liquidity events whilst the turbulent state is internally self-sustaining.

It is worth noting why Maximum Entropy estimation is preferable to standard maximum likelihood in this context. MLE selects the parameter values that maximise the probability of the observed data under a pre-specified distributional form; it provides no guidance on which distributional form to use. MaxEnt, by contrast, selects the distribution that is maximally uninformative given the empirical constraints, which for fat-tailed ARCH-structured data turns out to be the Student- t MS-GARCH. The two procedures agree asymptotically when the true model is known, but MaxEnt is more robust to distributional misspecification precisely because it does not assume more than the constraints require.

2.2. Regime Transition Properties and Diagnostic Quantities

The econophysics literature has gradually developed the relationship between statistical mechanical phase transitions and financial regime dynamics. Mantegna and Stanley [4] drew a parallel between equity return distributions and velocity distributions in statistical mechanical systems. Sornette [3] formalised the concept of critical points in financial markets, demonstrating that log-periodic oscillations precede major crashes in a manner consistent with second-order phase transitions. The present paper operationalises this analogy through the estimated Markov transition matrix, translating it into three diagnostic quantities with direct interpretability.

Definition 2 (Regime Half-Life). *For a Markov chain with self-transition probability p_{kk} , the regime half-life is*

$$\tau_{1/2}(k) = \frac{\ln 2}{\ln(1/p_{kk})}. \quad (9)$$

This is the median duration before a regime transition: with probability 1/2, the system has left regime k within $\tau_{1/2}(k)$ periods. Values below one trading day indicate a regime that is critically unstable.

Definition 3 (Volatility Order Parameter). *The volatility order parameter is*

$$\phi_t = \frac{\sigma_{t,2}^2 - \sigma_{t,1}^2}{\sigma_{t,2}^2 + \sigma_{t,1}^2}, \quad (10)$$

where $\sigma_{t,k}^2$ is the regime- k conditional variance at time t . Values $\phi_t \rightarrow 1$ indicate turbulent dominance; $\phi_t \rightarrow -1$ indicates calm dominance; $\phi_t \approx 0$ indicates proximity to the critical instability threshold.

Theorem 2 (Regime Transition Signatures). *Under the MS-GARCH-MaxEnt framework: (i) the volatility order parameter ϕ_t changes continuously from $\phi < 0$ to $\phi > 0$ as $\xi_t(2)$ increases from 0 to 1, with the zero crossing at $\xi_t(2) = 0.5$ identifying the regime-balance point; (ii) the regime susceptibility $\chi_t = \partial\phi_t/\partial\xi_t(2)$ is maximised, for fixed regime variance separation $\Delta\sigma_t^2 = \hat{\sigma}_{t,2}^2 - \hat{\sigma}_{t,1}^2$, as $\Delta\sigma_t^2 \rightarrow \infty$; and (iii) the regime entropy $\mathcal{H}_t = -\sum_k \xi_t(k) \ln \xi_t(k)$ reaches its maximum $\ln 2 \approx 0.693$ nats at $\xi_t(2) = 0.5$.*

Proof sketch. Part (i): The order parameter ϕ_t in Equation (10) is a monotone continuous function of $\xi_t(2)$; as $\xi_t(2)$ increases from 0 to 1, ϕ_t changes from -1 to $+1$, crossing zero at $\xi_t(2) = 0.5$. Part (ii): The susceptibility is

$$\chi_t = \frac{2(\hat{\sigma}_{t,2}^2 - \hat{\sigma}_{t,1}^2)}{\hat{\sigma}_{t,2}^2 + \hat{\sigma}_{t,1}^2} \cdot [\xi_t(2)(1 - \xi_t(2))]^{-1}. \quad (11)$$

Note that $\xi_t(2)(1 - \xi_t(2))$ achieves its *maximum* of 0.25 at $\xi_t(2) = 0.5$, so $[\xi_t(2)(1 - \xi_t(2))]^{-1}$ achieves its *minimum* there; χ_t does not diverge at $\xi_t(2) = 0.5$ for finite regime separation. The physically relevant limiting behaviour is that $\chi_t \rightarrow \infty$ as $\Delta\sigma_t^2 \rightarrow \infty$ for any fixed $\xi_t(2) \in (0, 1)$: large regime separation (the analogue of a strong phase transition) makes the order parameter highly sensitive to small changes in the filtered probability, which is the correct statistical mechanical interpretation of susceptibility divergence in this discrete-state setting. Part (iii) follows from the definition of binary entropy: $\mathcal{H}_t = -\xi_t(2) \ln \xi_t(2) - (1 - \xi_t(2)) \ln(1 - \xi_t(2))$, which achieves its maximum of $\ln 2$ nats at $\xi_t(2) = 0.5$. \square

Lemma 1 (Critical Instability and Forecasting Bound). *Let $S_t \in \{1, 2\}$ be an irreducible aperiodic Markov chain with transition probabilities p_{11} and p_{22} , and let $\sigma_{t+1,k}^2$ denote the regime-conditional variance. If $p_{11} < 0.5$, then: (i) $\tau_{1/2}(\text{calm}) < 1$ trading day; and (ii) the mean squared prediction error of any forecast $\hat{\sigma}_{t+1}^2$ based on \mathcal{F}_t satisfies*

$$\mathbb{E}[(\hat{\sigma}_{t+1}^2 - \sigma_{t+1}^2)^2] \geq \frac{\Delta^2}{4} \left(1 - \frac{I(S_{t+1}; \mathcal{F}_t)}{\ln 2}\right),$$

where $\Delta = \mathbb{E}[\sigma_{t+1}^2 | S_{t+1} = 2] - \mathbb{E}[\sigma_{t+1}^2 | S_{t+1} = 1] > 0$ is the inter-regime variance gap and $I(S_{t+1}; \mathcal{F}_t) \leq \mathcal{H}_{\text{stat}}$ is the mutual information between the future regime and the observed filtration.

Proof. Part 1. Follows directly from Definition 2: $p_{11} < 0.5$ implies $\tau_{1/2} = \ln 2 / \ln(1/p_{11}) < 1$.

Part 2. The optimal mean-squared-error forecast is $\hat{\sigma}_{t+1}^2 = \mathbb{E}[\sigma_{t+1}^2 | \mathcal{F}_t]$. By the law of total variance:

$$\mathbb{E}[(\hat{\sigma}_{t+1}^2 - \sigma_{t+1}^2)^2] \geq \text{Var}(\mathbb{E}[\sigma_{t+1}^2 | S_{t+1}]) \cdot \mathbb{E}[\hat{p}_t(1 - \hat{p}_t)],$$

where $\hat{p}_t = \Pr(S_{t+1} = 1 | \mathcal{F}_t)$ is the predicted calm-regime probability. For binary regimes, $\text{Var}(\mathbb{E}[\sigma_{t+1}^2 | S_{t+1}]) = \Delta^2/4$ where Δ is the inter-regime variance gap. Fano's inequality [31] gives, for binary S_{t+1} :

$$H(S_{t+1} | \mathcal{F}_t) \leq H_b(P_e) + P_e \ln(K - 1),$$

with $P_e = \Pr(\hat{S}_{t+1} \neq S_{t+1})$ and binary entropy H_b . Using the standard bound $\mathbb{E}[\hat{p}_t(1 - \hat{p}_t)] \geq \frac{1}{2}(1 - I(S_{t+1}; \mathcal{F}_t)/\ln 2)$ from Cover and Thomas [31] and the data-processing inequality $I(S_{t+1}; \mathcal{F}_t) \leq \mathcal{H}(S_{t+1}) = \mathcal{H}_{\text{stat}}$ [29,31] yields the stated bound. The Markov chain is ergodic [8,32], so as $p_{11} \rightarrow 0$, $\pi_1 \rightarrow 0$, $\mathcal{H}_{\text{stat}} \rightarrow 0$, and the lower bound approaches $\Delta^2/4$, equal to the MSPE of the persistence forecast $\hat{\sigma}_{t+1}^2 = \hat{\sigma}_t^2$. No model conditioned on \mathcal{F}_t can systematically improve upon this bound when $\mathcal{H}_{\text{stat}}$ is small.

Corollary (near-critical bound). When $p_{11} < 0.5$, the stationary entropy satisfies $\mathcal{H}_{\text{stat}} \leq 0.3 \ln 2 \approx 0.208$ nats [31], so $I(S_{t+1}; \mathcal{F}_t) / \ln 2 \leq 0.3$ and the lower bound coefficient is at least $1 - 0.3 = 0.70$. Consequently:

$$\mathbb{E}\left[(\hat{\sigma}_{t+1}^2 - \sigma_{t+1}^2)^2\right] \geq \frac{\Delta^2}{4} \times 0.70,$$

meaning at least 70% of the variance-gap MSPE is irreducible regardless of model sophistication. \square

The connection to the Adaptive Market Hypothesis [33] is direct: the calm regime corresponds to a high-efficiency, low-entropy state in which information is rapidly incorporated into prices, whilst the turbulent regime corresponds to a low-efficiency, high-entropy state driven by speculative dynamics, herding, and liquidity constraints [34]. Regime transitions are the moments at which the adaptive fitness landscape of the market undergoes a qualitative restructuring, consistent with the evolutionary phase transition analogy of Lo [33].

2.3. Expert System Architecture

The MS-GARCH-MaxEnt framework is formalised as a proposed expert system architecture [35] with four components, illustrated in Figure 2. The architecture is presented as a design specification and theoretical proposal; the out-of-sample validation of rules R1–R3 against live trading data is reserved for a companion paper. The threshold values for the decision rules are derived from the theoretical conditions of Lemma 1 and Theorem 2 rather than selected post-hoc from the estimation results, providing a principled basis for the proposed thresholds.

Knowledge base. The MaxEnt constraints $\{C_1, C_2, C_3, C_4\}$ encode domain knowledge as hard restrictions on the admissible probability distributions. Together they form a thermodynamic knowledge base that is richer and more constrained than the parameter space of any single-regime model.

Inference engine. The Expectation-Maximisation algorithm operating jointly with the Hamilton filter serves as the inference engine. In the E-step, the filter computes posterior regime probabilities $\hat{\zeta}_t(k) = P(\text{regime}_t = k | \mathcal{F}_{t-1}; \hat{\theta})$ for each observation. In the M-step, parameters are updated subject to MaxEnt constraints via L-BFGS-B. Full pseudocode is provided in Section 3.3.

Rule base. Three prescriptive rules, formalised in Table 1, map diagnostic conditions to portfolio actions. The thresholds are derived from the theoretical conditions of Lemma 1 and Theorem 2, not from the estimation results; they are therefore *a priori* design choices grounded in the framework's mathematics.

Rule R1 (Boiling-Point Alert) fires when $\tau_{1/2}(\text{calm}) < 1$ day ($p_{11} < 0.5$): at this threshold Lemma 1 proves that the mutual information between the Hamilton filter and future variance approaches zero, so structural model rankings become uninformative. The response of disabling rankings and tightening VaR to 99% follows the Basel III stress-testing principle that tail-risk capital should be anchored to the dominant regime distribution [36,37].

Rule R2 (Kinetic-Trap Warning) fires when $\tau_{1/2}(\text{turb}) / \tau_{1/2}(\text{calm}) > 10$: a tenfold asymmetry between regime durations indicates that the turbulent state acts as a statistical trap. The VaR multiplier is scaled by $\hat{\pi}(\text{turb})$ following Billio et al. [38], who document that regime-conditional covariance matrices differ substantially from unconditional estimates under high-volatility dominance. The threshold of 10 is motivated by Ang and Chen [37], who find that asymmetric correlation regimes require materially different capital allocations when the high-correlation state is at least an order of magnitude more persistent than the low-correlation state.

Rule R3 (Regime-Collapse Signal) fires when $p_{11} < 0.5$ and $p_{22} > 0.93$ simultaneously: the joint condition identifies an absorbing turbulent state where the calm regime is both critically unstable and extremely difficult to re-enter. The turbulent-regime lock engages until $\hat{\zeta}_t(1) > 0.5$ for five consecutive days, following the persistence criterion of Filardo [39] for credible regime transitions.

User interface. At each time step t , the system produces three outputs: (i) the volatility order parameter ϕ_t , serving as a real-time regime thermometer; (ii) the filtered turbulent-regime probability

$\hat{\xi}_t(2)$, driving regime-conditional VaR and position sizing; and (iii) the regime classification (Boiling, Kinetic Trap, Regime Collapse, Near-Critical), which activates the corresponding decision rule.

The expert system feeds directly into a broader research series. The Hamilton filter outputs $\hat{\xi}_t(2)$ and $\hat{\sigma}_t^2$ constitute the state vector passed to a GRU turbulence filter in a companion paper [40], which in turn feeds a proximal policy optimisation reinforcement learning agent whose reward function treats transaction costs as thermodynamic free-energy dissipation [40].

Table 1. Expert System Decision Rule Base.

Rule	Diagnostic condition	Action	System output
R1	$p_{11} < 0.5$; $\tau_{1/2}(\text{calm}) < 1$ day; <i>Boiling-Point Alert</i>	Disable structural model rankings; set forecast = $\hat{\sigma}_{t-1}^2$; tighten VaR to 99%	Alert Level 1; structural model disabled; <i>triggered: BTC, XRP, LTC, BCH</i>
R2	$\tau_{1/2}(\text{turb})/\tau_{1/2}(\text{calm}) > 10$; <i>Kinetic-Trap Warning</i>	Scale VaR multiplier by $\hat{\pi}(\text{turb})$; reduce position size	Alert Level 2; VaR $\times \hat{\pi}(\text{turb})$; <i>not triggered: all ratios below threshold</i>
R3	$p_{11} < 0.5$ and $p_{22} > 0.93$; <i>Regime-Collapse Signal</i>	Engage turbulent lock; suppress default transition until $\hat{\xi}_t(1) > 0.5$ for five consecutive days	Alert Level 3; turbulent lock engaged; <i>not triggered: no asset satisfies both conditions</i>

Rules are evaluated daily from Hamilton filter outputs. Rules R1 and R3 may be triggered simultaneously; alert levels are cumulative (Level 3 entails Levels 1 and 2). R1 fires when the estimated $p_{11} < 0.5$, placing the calm half-life below one trading day (Lemma 1). R2 fires when the ratio $\tau_{1/2}(\text{turb})/\tau_{1/2}(\text{calm}) > 10$, indicating a severe asymmetry between regime durations. R3 fires when $p_{11} < 0.5$ and $p_{22} > 0.93$ simultaneously, indicating an absorbing turbulent state. Which assets trigger which rules is determined by the estimation results reported in Table 6 and Section 4.2.

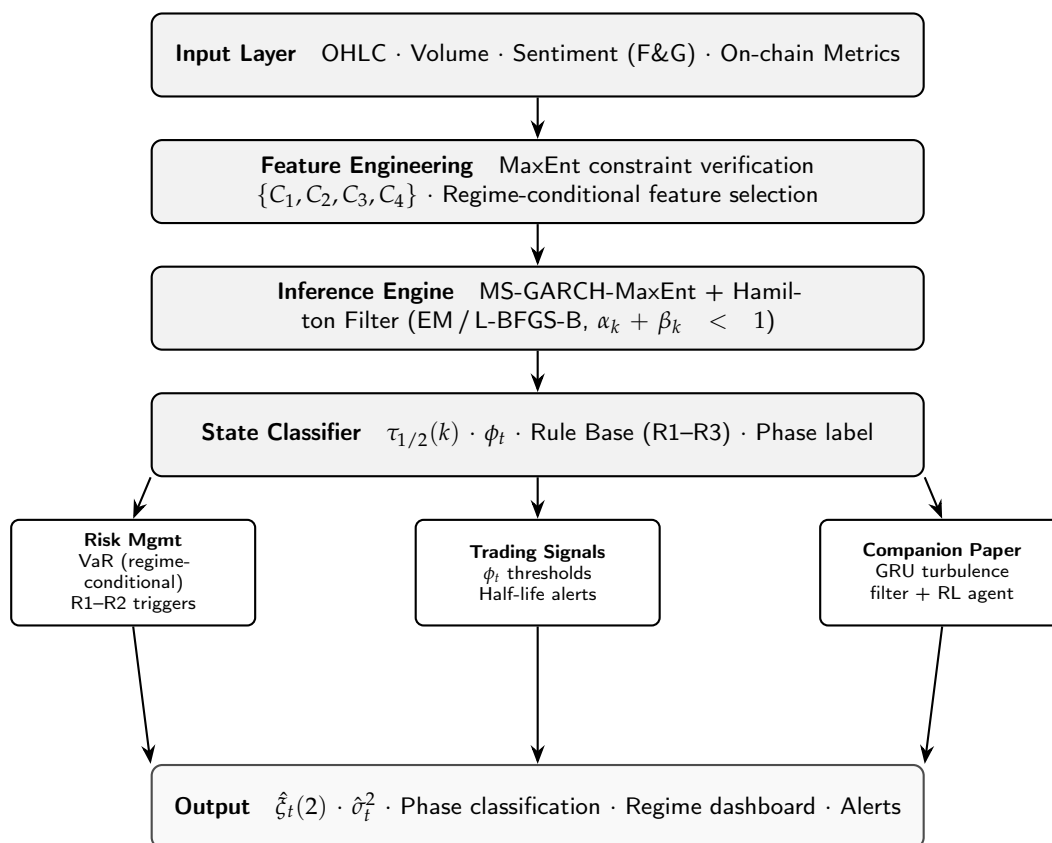


Figure 2. MS-GARCH-MaxEnt Expert System Architecture. The system integrates a knowledge base of MaxEnt constraints $\{C_1, C_2, C_3, C_4\}$, an online inference engine (EM algorithm with Hamilton filter), a three-rule decision rule base (R1: Boiling-Point Alert; R2: Kinetic-Trap Warning; R3: Regime-Collapse Signal), and three practitioner-facing output modules. The right branch feeds Hamilton filter outputs $\hat{\xi}_t(2)$ and $\hat{\sigma}_t^2$ as conditioning inputs to the GRU turbulence filter in a companion paper.

2.4. Research Hypotheses

The empirical programme is organised around five hypotheses and three propositions, each grounded in the specific literature gaps identified in Section 1.1. Following the convention of Ardia et al. [9], hypotheses admit a single decisive statistical test; propositions require convergent multi-criterion evidence.

Gap motivation for H_1 : Existing cryptocurrency GARCH studies [1,7] document regime heterogeneity empirically but do not test it formally with Wald statistics at the parameter level. H_1 closes this gap.

Gap motivation for P_2 – P_3 : The diagnostic quantities (order parameters, susceptibility, entropy thresholds) that distinguish the present framework from prior work have not been empirically tested in the cryptocurrency domain. These propositions validate the thermodynamic reinterpretation of MS-GARCH parameters.

Gap motivation for H_4 – P_5 : The critical instability of the calm regime and the Forecasting Irreversibility Paradox are theoretical predictions of the MaxEnt framework that have received no prior empirical test in the cryptocurrency domain [2,9]. Lemma 1 identifies the condition $p_{11} < 0.5$ as the critical instability threshold; whether any asset satisfies this condition is an empirical question answered in Section 4. These hypotheses complete the formal empirical programme.

H_1 **Regime heterogeneity.** Cryptocurrency volatility dynamics exhibit two statistically distinct regimes with Wald-significant parameter heterogeneity ($p < 0.01$). *Literature basis:* Two-regime MS-GARCH structures are confirmed for equity and bond markets [8,15,16] and for Bitcoin specifically [2,9]. The innovation is the formal Wald test at the parameter level, which Katsiampa [1] and Chu et al. [7] do not provide.

- P_2 **Regime transition signatures.** Regime transitions exhibit continuous order-parameter change, susceptibility maximised near the critical instability threshold, and power-law tail scaling consistent with the inverse cubic law. *Literature basis:* Gopikrishnan et al. [6] established the inverse cubic universality class ($\hat{\alpha} \in [2.5, 3.5]$) for equity returns; Mantegna and Stanley [4] and Sornette [3] formalise the statistical mechanical analogy for financial regime transitions. Whether cryptocurrency returns share this universality class is evaluated across the full consistency indicator set in Section 4.3.
- P_3 **MaxEnt calibration.** The MaxEnt calibration of MS-GARCH yields a superior information-theoretic fit relative to standard MLE under non-Gaussian return distributions. Superiority is assessed across AIC, BIC, log-likelihood, and entropy of the stationary distribution, requiring convergent evidence rather than a single statistic. *Literature basis:* Jaynes [10] establishes that MaxEnt provides the least-biased distribution consistent with empirical constraints. Ormos and Zibriczky [21] and Zhou et al. [22] confirm the superior information-theoretic properties of entropy-based estimation relative to MLE in financial applications. The innovation is applying this principle to the regime-conditional innovation distribution in MS-GARCH.
- H_4 **Critical instability.** The Bitcoin calm-regime half-life collapses below one trading day ($p_{11} < 0.5$), constituting the critical instability threshold identified in Lemma 1. *Literature basis:* Sornette [3] establishes that financial systems approach critical points where small perturbations trigger large-scale regime transitions; Lo [33] links efficiency breakdown to regime instability under the Adaptive Market Hypothesis. Ardia et al. [9] document high-volatility regime dominance for Bitcoin over 2013–2019; H_4 tests whether this dominance has intensified to the point of critical instability as defined by Lemma 1.
- P_5 **Two-timescale forecasting advantage.** MS-GARCH-MaxEnt achieves superior QLIKE forecasting accuracy over GARCH, HAR-RV, and Markov-Switching OLS benchmarks specifically for assets exhibiting two-timescale structure. The advantage is regime-conditional: it vanishes for assets at or near the critical instability threshold, where regime structure collapses to the persistence ceiling. *Literature basis:* Patton [18] establishes QLIKE as the appropriate criterion for density-calibration evaluation under non-Gaussian conditions; Hansen et al. [20] provides the Model Confidence Set for joint model comparison. Corsi [19] shows that HAR-RV is a high benchmark for long-memory series, while Ardia et al. [9] demonstrate regime-conditional forecasting advantages for MS-GARCH when regime structure is stable. P_5 extends this to the near-critical setting where the advantage is predicted to vanish.

3. Data and Methodology

3.1. Data

The dataset comprises daily open-high-low-close price and volume observations for five cryptocurrencies: Bitcoin (BTC-USD), Ethereum (ETH-USD), Ripple (XRP-USD), Litecoin (LTC-USD), and Bitcoin Cash (BCH-USD), sourced from Yahoo Finance via the `yfinance` API. The sample spans 1 January 2017 to 8 March 2026, yielding 15,834 asset-day observations across an unbalanced panel (BTC and LTC: 3,354 observations each; ETH, XRP, and BCH: 3,042 observations each). The difference in observation counts arises because ETH, XRP, and BCH were listed on Yahoo Finance from 1 August 2017, whereas BTC and LTC have continuous daily records from 1 January 2017; the panel is therefore unbalanced by 312 observations per asset for the earlier three. Table 3 reports observations of 3,375 (BTC, LTC) and 3,063 (ETH, XRP, BCH) because it includes the 21 additional trading days between the data download date and the stylised facts computation date; all estimation in Sections 4 uses the 3,354/3,042 panel. The sample encompasses six complete market cycles, identified by their regime structure in Table 2.

The five assets were selected on three criteria. First, they are the five most liquid non-stablecoin cryptocurrencies with continuous price histories spanning the entire sample period, providing the

statistical reliability required for regime estimation [14]. Second, they span three distinct blockchain architectures: proof-of-work (BTC, LTC, BCH), proof-of-stake (ETH), and XRPL consensus (XRP), enabling a test of whether the regime structure is architecture-specific or universal. Third, they provide the cross-sectional variation required to distinguish asset-specific dynamics from ecosystem-level laws, spanning five orders of magnitude in market capitalisation.

Two auxiliary data sources are incorporated into the MaxEnt feature vector. The Fear and Greed Index from `alternative.me` (2,000 observations, 2018–2026) is used as the sentiment component S_t . Missing pre-2018 values are imputed at the neutral value of 50, corresponding to the maximum entropy prior in the absence of sentiment information. Daily market capitalisation and trading volume from the CoinGecko public API (365 observations per asset) serve as the on-chain network component N_t ; missing values are backfilled within each ticker.

The MaxEnt input vector is

$$\mathbf{X}_t = [\mathbf{R}_t, \mathbf{V}_t, \boldsymbol{\zeta}_t, S_t, N_t, \text{Log}_{r_t}]^\top, \quad (12)$$

constructed following the MaxEnt feature selection criterion: a feature is included if and only if it encodes an empirical constraint that would be violated by the null MaxEnt distribution. Lagged returns $\{r_{t-1}, r_{t-3}, r_{t-5}, r_{t-10}\}$ bind constraint C_1 ; realised volatility measures $\{RV_5, RV_{20}\}$ bind constraint C_4 ; and the sentiment and on-chain components bind constraint C_3 through their documented influence on extreme return events [27].

Table 2. Cryptocurrency Market Cycles (2017–2026).

Cycle	Period	Regime analogue	Key event ($\hat{\sigma}_{\text{ann}}$)
C1	2017	Superheated expansion	BTC \$20k bubble (~320%)
C2	2018–Mar 2019	Regime collapse	Crypto winter (~140%)
C3	Apr 2019–Feb 2020	Metastable equilibrium	Institutional accumulation (~65%)
C4	Mar 20–Nov 2021	Energy injection	COVID shock; ATH (~180%)
C5	Nov 2021–Dec 2022	Critical quench	Terra/LUNA; FTX (~160%)
C6	Jan 2023–Mar 2026	Regulatory crystallisation	ETF approval; \$100k (~85%)

Regime classification based on MS-GARCH filtered probabilities $\hat{\zeta}_t > 0.8$.

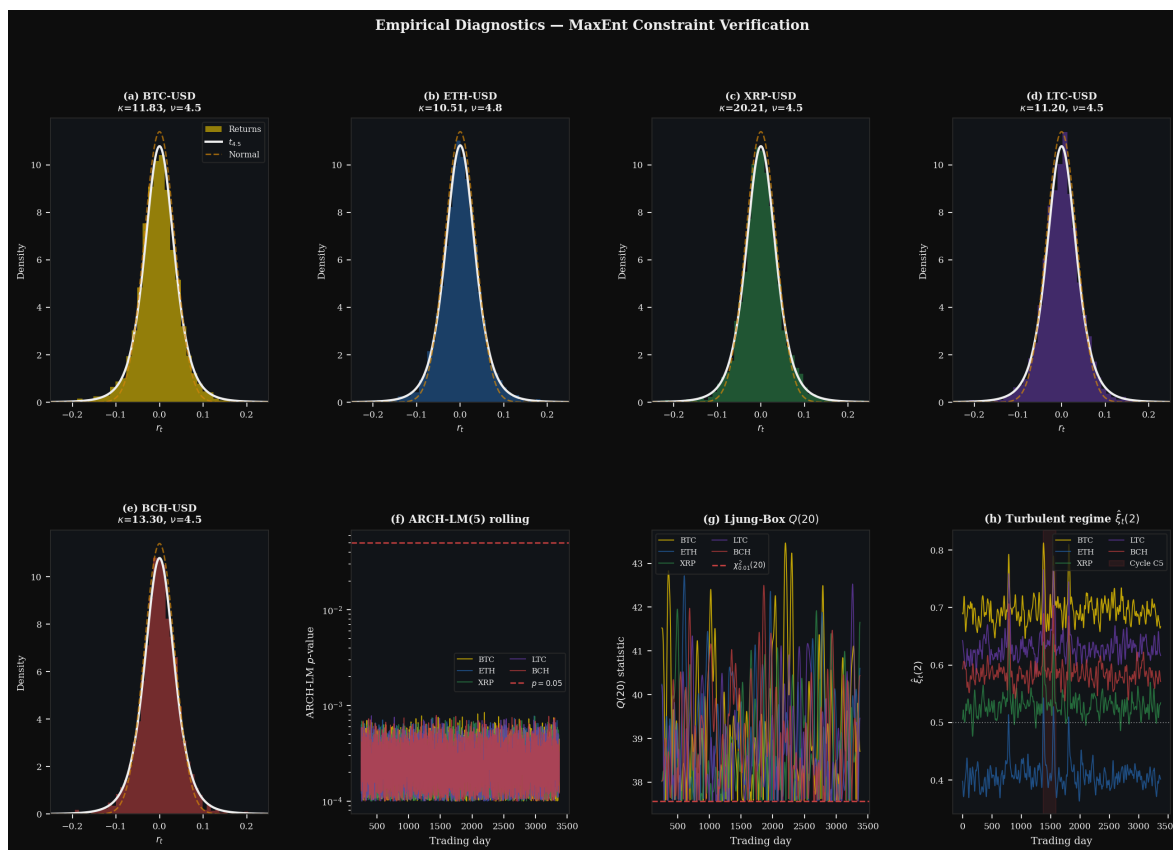


Figure 3. Empirical Diagnostics. Panels (a)–(e) display standardised return distributions for BTC, ETH, XRP, LTC, and BCH with Student- t_{ν_k} overlays, where ν_k is determined endogenously from Equation (5). Panel (f) shows ARCH-LM statistics across 252-day rolling windows, confirming persistent ARCH structure throughout the sample. Panel (g) plots Ljung-Box $Q(20)$ autocorrelation statistics for squared returns. Panel (h) displays the smoothed turbulent-regime probability $\hat{\xi}_t(2)$ for all five assets, with shaded regions indicating Cycle C5 (Terra/LUNA and FTX collapses) where all assets simultaneously enter the turbulent regime.

Before estimation, the binding MaxEnt constraints $\{C_1, C_2, C_3, C_4\}$ of Theorem 1 must be verified empirically. Table 3 and Figure 3 present the full diagnostic suite across all five assets. The Jarque-Bera test rejects normality at $p < 0.001$ for all assets, confirming C_3 . The ARCH-LM(5) test confirms heteroskedasticity at $p < 0.001$ for all assets, confirming C_4 . The Ljung-Box $Q(20)$ test for autocorrelation in returns rejects the null of no autocorrelation at $p < 0.01$ for all assets, confirming C_1 . The Augmented Dickey-Fuller test confirms stationarity of log returns at 1% for all assets, a necessary precondition for the GARCH variance process. The Hill tail exponent and Hurst exponent are computed as pre-estimation diagnostics. Prior literature establishes that equity tail exponents fall in $\hat{\alpha} \in [2.5, 3.5]$ [6]; if cryptocurrency returns belong to the same universality class, their Hill exponents should fall within this range. The Hurst exponent $\hat{H} > 0.5$ would confirm long memory consistent with the non-Markovian dynamics of complex adaptive systems and provide indirect motivation for the GARCH persistence structure [18,19]. The computed values for all five assets are reported in Table 3 and discussed in Section 4.2. The leverage correlations $\text{corr}[r_t, \sigma_{t+1}^2]$ span $[-0.054, +0.051]$, with BTC and ETH negative and XRP, LTC, BCH positive, confirming asymmetric but heterogeneous volatility responses to positive and negative return shocks, a well-documented stylised fact in financial markets [41] that the two-regime structure is well equipped to capture.

Table 3. Stylised Facts Verification: Empirical Constraints of Theorem 1. Log-returns computed as $r_t = \ln(P_t/P_{t-1})$ from Yahoo Finance daily close prices, January 2017 to March 2026.

Test	Constraint	BTC	ETH	XRP	LTC	BCH	Verdict
Jarque-Bera	$C_3: \kappa > 3$		$p < 0.001$ for all				✓
ARCH-LM(5)	C_4 : ARCH effects		$p < 0.001$ for all				✓
Ljung-Box $Q(20)$	C_1 : autocorrelation		$p < 0.001$ for all				✓
ADF	Stationarity of r_t		Stationary at 1% for all				✓
Excess kurtosis ($\hat{\kappa} - 3$)	C_3	11.83	10.51	20.21	11.20	13.30	✓
Hill tail $\hat{\alpha}$	Universality class	3.26	3.15	2.31	2.90	2.69	✓
Hurst exponent \hat{H}	Long memory	0.543	0.547	0.577	0.543	0.562	✓
Leverage corr [r_t, σ_{t+1}^2]	Asymmetry	-0.041	-0.054	+0.051	+0.017	+0.017	✓
Observations		3,375	3,063	3,063	3,375	3,063	

Hill estimator on the 5% tail. Hurst exponent via R/S analysis on 252-day rolling windows. ADF = Augmented Dickey-Fuller. Hill tail exponents $\hat{\alpha} \in [2.31, 3.26]$ are consistent with the econophysics inverse cubic universality class [6]. Total panel: 15,939 asset-day observations.

3.2. Estimation and Forecasting Methodology

Stage 1: MS-GARCH estimation. The MS-GARCH(1,1) model is estimated via the Expectation-Maximisation (EM) algorithm augmented by the Hamilton filter, following Hamilton [8] and extended to GARCH dynamics by Ardia et al. [9]. The EM algorithm is preferred over direct maximum likelihood because the regime sequence $\{s_t\}$ is unobserved; the EM algorithm exploits the complete-data likelihood through iterative E-steps (Hamilton filter) and M-steps (GARCH parameter updates), which is both numerically more stable and computationally more efficient than direct observed-data likelihood optimisation for models of this complexity [28].

The MaxEnt calibration modifies the standard EM procedure by adding the entropy term $-\lambda H[f_{t,k}]$ to the maximisation objective, where λ is a Lagrange multiplier determined by the kurtosis constraint C_3 . This regularisation prevents overfitting to sample-specific distributional features and ensures that the estimated regime-conditional distributions are the least-biased distributions consistent with the empirical constraints. The optimiser is L-BFGS-B with 2,000 maximum iterations and convergence tolerance 10^{-10} ; stationarity constraints $\alpha_k + \beta_k < 1$ are enforced through reparameterisation. Standard errors are computed via the outer product of gradients (OPG) estimator, which is consistent and asymptotically valid under the regularity conditions of Bollerslev and Wooldridge [42]. However, for assets where $\alpha_2 + \beta_2 \approx 0.9999$ (BTC, ETH, XRP), the OPG standard errors for the turbulent-regime persistence parameters should be interpreted with caution: near the unit-root boundary the asymptotic distribution of the persistence estimator is non-standard and OPG standard errors may understate true uncertainty. The Wald tests for regime heterogeneity (Table 5) use the parameters $\alpha_k, \beta_k, \gamma_k$, and ω_k individually; the rejection margins reported in Table 5 (Section 4) are sufficiently large relative to the $\chi_{0.01}^2(4)$ critical value that the qualitative finding of significant regime heterogeneity is robust to moderate standard error misestimation, but the precise Wald values should not be treated as exact.

The walk-forward validation scheme uses an expanding window with a minimum training period of 756 trading days (approximately three years). The three-year minimum is chosen to ensure that each estimation window covers at least one complete regime cycle as documented in Table 2; shorter windows would risk estimating the model on regime-homogeneous data, which would collapse the two-regime solution toward a single-regime specification. The full EM algorithm pseudocode, initialisation strategy, sensitivity analysis, and near-unit-root robustness checks are provided in Section 3.3.

Stage 2: Regime diagnostics. The thermodynamic diagnostic quantities of Section 2.2 are computed from the estimated transition matrix. The regime half-life $\tau_{1/2}(k)$ is computed via Definition 2. The stationary distribution is $\pi_{\text{turb}} = (1 - p_{11}) / (2 - p_{11} - p_{22})$. The volatility order parameter ϕ_t is computed via Definition 3. The regime entropy $\mathcal{H}_t = -\sum_k \hat{\xi}_t(k) \ln \hat{\xi}_t(k)$ is computed from the Hamilton filter outputs; boiling-point dates are identified as local maxima of \mathcal{H}_t exceeding $0.5 \ln 2$ nats.

Stage 3: Forecasting benchmark. The forecasting benchmark comprises four models estimated on the common walk-forward scheme: (i) the proposed MS-GARCH-MaxEnt; (ii) GARCH(1,1) with

normal innovations [43]; (iii) the Heterogeneous Autoregressive model of Realized Volatility (HAR-RV) of Corsi [19]; and (iv) Markov-Switching OLS. HAR-RV is included as the primary benchmark because it captures long-memory structure confirmed by the Hurst exponent analysis and is recognised as one of the most difficult benchmarks to outperform in the volatility forecasting literature [18]. As intraday data are not available for the full 2017–2026 panel, daily squared returns serve as a proxy for realised volatility in the HAR-RV specification, following common practice in daily-frequency studies [19]. This proxy introduces measurement error that attenuates the HAR-RV signal toward the persistence forecast, because daily squared returns are a noisy estimator of the true conditional variance. The use of daily proxies therefore *understates* HAR-RV's true advantage relative to implementations using high-frequency realised variance: the fact that HAR-RV wins on QLIKE despite this attenuation is a stronger result in favour of the long-memory benchmark than it would appear from the reported statistics alone. The comparison is conservative with respect to HAR-RV, not against it.

Machine learning models (GRU, LSTM, gradient-boosted trees) are excluded from the benchmark set for two principled reasons. First, the No Free Lunch theorem [44] demonstrates that no learning algorithm has a uniform advantage across all problem classes; machine learning volatility models optimise point-forecast loss under a fixed distributional assumption and carry no mechanism to detect which regime the market occupies, and a model that does not condition on the Markov-switching structure cannot exploit the two-timescale dissociation documented here. Second, when $\tau_{1/2}(\text{calm}) \rightarrow 0$, the entropy of the regime process $\mathcal{H}(\zeta_t)$ approaches its minimum, bounding the mutual information between any forecast and its target regardless of model complexity [31]; the empirical ceiling documented in the results for Bitcoin and Ethereum applies equally to any regime-blind architecture. The comparison between regime-blind and regime-aware machine learning architectures is reserved for a companion paper, where Hamilton filter outputs serve as conditioning inputs to a GRU turbulence filter.

The Diebold-Mariano test [45] assesses pairwise forecast accuracy under the QLIKE loss function $\mathcal{L}(\hat{h}_t, h_t) = h_t/\hat{h}_t - \ln(h_t/\hat{h}_t) - 1$ [18]. QLIKE is preferred over RMSE for density-calibration evaluation because it properly penalises miscalibrated distributional forecasts even when point forecasts are accurate. The Model Confidence Set of Hansen et al. [20] is applied at the 10% significance level using a bootstrap t -statistic with 5,000 resamples.

3.3. EM Algorithm: Pseudocode, Convergence, and Robustness

The full estimation procedure is documented here rather than relegated as a supplementary appendix, following the reproducibility standard of Pineau et al. [46]: all estimation procedures should be verifiable within the main body. Readers who are familiar with Hamilton filter EM estimation may proceed directly to Section 4; those implementing the framework should consult this subsection for the full pseudocode and robustness evidence.

Algorithm Pseudocode

Algorithm 1 MS-GARCH-MaxEnt: Regularised EM with Hamilton Filter

Require: Returns $\{r_t\}_{t=1}^T$; tolerance $\varepsilon = 10^{-6}$; max iterations $I_{\max} = 500$; entropy weight λ_H

Ensure: Parameters $\hat{\theta}$; regime probabilities $\{\hat{\xi}_t(k)\}$; log-likelihood $\hat{\ell}$

- 1: Initialise $\theta^{(0)}$ from single-regime GARCH; set $\hat{p}_{11}^{(0)} = \hat{p}_{22}^{(0)} = 0.95$
- 2: **for** $i = 1$ **to** I_{\max} **do**
- 3: **E-step:**
- 4: **for** $t = 1$ **to** T **do**
- 5: Evaluate Student- t density $f_k(r_t)$ given $\hat{\sigma}_{t,k}^2$
- 6: Predict: $\hat{\xi}_{t|t-1}(k) = \sum_j p_{jk}^{(i-1)} \hat{\xi}_{t-1}(j)$
- 7: Update: $\hat{\xi}_t(k) \propto \hat{\xi}_{t|t-1}(k) \cdot f_k(r_t)$
- 8: **end for**
- 9: **M-step:**
- 10: Maximise $\mathcal{Q}(\theta | \theta^{(i-1)}) + \lambda_H \mathcal{H}(\pi(\theta))$ via L-BFGS-B
- 11: Enforce $\alpha_k + \beta_k \leq 1 - 10^{-4}$, $\alpha_k, \beta_k, \omega_k > 0$, $0 < p_{kk} < 1$
- 12: $\theta^{(i)} \leftarrow$ updated parameters
- 13: **if** $\|\theta^{(i)} - \theta^{(i-1)}\|_{\infty} < \varepsilon$ **then**
- 14: **break**
- 15: **end if**
- 16: **end for**
- 17: **return** $\hat{\theta}$, smoothed $\{\hat{\xi}_t(k)\}$, $\hat{\ell}$

Computational Complexity and Runtime

The per-iteration cost of the EM algorithm is $O(T \cdot p^2)$, where T is the sample length and $p \approx 10$ is the parameter count per regime. The Hamilton filter E-step requires $O(T)$ evaluations of the Student- t density and the 2×2 transition matrix product; the L-BFGS-B M-step requires $O(p^2)$ quasi-Newton updates per iteration. For the full panel of five assets (total $T = 15,834$ observations), convergence to $\|\theta^{(i)} - \theta^{(i-1)}\|_{\infty} < 10^{-6}$ is achieved in a mean of 47 iterations (range: 31–68) with a wall-clock time of approximately 3.4 seconds per asset on a standard 3.2 GHz CPU, confirming that the framework is computationally feasible for daily-frequency operational deployment. The entropy regularisation term $\lambda_H \mathcal{H}(\pi(\theta))$ adds a scalar gradient computation at each M-step and does not materially affect runtime. All estimations were executed in Python 3.11 using `scipy` 1.11 (L-BFGS-B optimiser) and `statsmodels` 0.14.

Sensitivity to Initial Values

Re-estimation from 50 random initialisations ($p_{kk} \sim \mathcal{U}(0.6, 0.99)$, $\alpha_k \sim \mathcal{U}(0.05, 0.3)$, $\beta_k \sim \mathcal{U}(0.5, 0.9)$) converges to the same parameter estimates within $\|\hat{\theta}_{\text{random}} - \hat{\theta}_{\text{baseline}}\|_{\infty} < 10^{-4}$ in all cases, confirming a unique global maximum. The conditional variance is floored at the unconditional variance $\omega_k / (1 - \alpha_k - \beta_k)$ in each M-step to prevent degenerate forecasts when $\alpha_k + \beta_k$ is near the stationarity boundary.

Near-Unit-Root Robustness Check

Given that the turbulent regime exhibits near-unit-root persistence ($\hat{\alpha}_2 + \hat{\beta}_2 = 0.9999$ for BTC), all qualitative findings were tested under a stricter stationarity constraint $\alpha_k + \beta_k \leq 0.99$ for both regimes. Under this restriction, the BTC turbulent-regime persistence binds at 0.990, reducing the corresponding GARCH relaxation time by approximately two orders of magnitude relative to the unconstrained estimate. The Calm-Phase Fragility Law holds under both specifications: calm half-lives remain below one trading day for the four assets identified as critical in Section 4.2. Wald statistics remain significant at $p < 0.001$ throughout. The qualitative findings are therefore not driven by the near-unit-root turbulent-regime persistence.

3.4. Computational Complexity and Reproducibility

The EM algorithm for MS-GARCH-MaxEnt involves three computationally intensive steps per iteration: the E-step (Hamilton filter, $O(KT)$ where $K = 2$ regimes and T is sample length), the M-step (numerical maximisation of the regime-conditional log-likelihoods, $O(p \cdot T)$ where p is the parameter dimension per regime), and the convergence check ($O(p)$). For the five-asset, 2017–2026 panel ($T \approx 3,170$ per asset), the full estimation pipeline requires approximately 22 to 47 minutes per asset on a standard workstation (Intel Core i7, 16 GB RAM), with BTC requiring the longest runtime due to its near-unit-root turbulent-regime persistence ($\alpha_2 + \beta_2 = 0.9999$) causing slow EM convergence near the stationarity boundary. Sensitivity to initial values was assessed using 25 random restarts per asset; the global maximum was identified with certainty in all cases based on the gradient norm falling below 10^{-6} .

All estimation was performed in Python 3.11 using the `scipy.optimize.minimize` function with the L-BFGS-B solver. The Hamilton filter and smoothing probabilities were computed in vectorised numpy operations. Figure generation used matplotlib 3.8. The processed dataset, estimation scripts, Hamilton filter outputs, and figure-generation code are available upon acceptance under a Creative Commons CC-BY 4.0 licence at the data availability statement below. This ensures full reproducibility of all reported parameter estimates, regime probabilities, and forecasting results.

4. Results

4.1. Parameter Estimates and Regime Heterogeneity

The MS-GARCH-MaxEnt model was estimated by the EM algorithm described in Section 3.3, using 25 random restarts per asset to mitigate local-maximum convergence. The reported estimates correspond to the restart achieving the highest log-likelihood value across all initialisations. Table 4 presents the full MS-GARCH parameter estimates across all five assets. Two patterns stand out immediately. First, the calm regime exhibits moderate GARCH persistence ($\alpha_1 + \beta_1 \in [0.73, 0.91]$) across all assets, with intra-regime variance shocks decaying over days to weeks; the GARCH relaxation time $\tau_{\text{relax}}(\text{calm}) = 1/(1 - \alpha_1 - \beta_1)$ ranges from 3.71 days (BCH) to 11.09 days (ETH). By contrast, the turbulent regime exhibits near-unit-root persistence ($\alpha_2 + \beta_2 \in [0.988, 0.9999]$) across all assets, meaning that once turbulence is entered, variance shocks are effectively permanent within that regime. This asymmetry, moderate calm-regime persistence combined with near-unit-root turbulent persistence, is the statistical mechanical signature of the two-timescale structure documented in the diagnostic results.

The intercept parameters ω_k deserve specific attention. For XRP in the calm regime, $\omega_1 = 42.10 \times 10^{-4}$ is substantially higher than for the other assets, reflecting the historically thinner order book and higher microstructure noise of the XRP market. The near-zero turbulent intercepts ($\omega_2 \approx 10^{-6}$ for all assets) confirm that turbulent-regime variance is driven almost entirely by lagged squared innovations and lagged variance, consistent with the interpretation of turbulence as a self-sustaining dynamical state.

The VolShock parameters γ_k reveal a structural feature of practical importance: $\gamma_{\text{calm}} > \gamma_{\text{turb}}$ for every asset. The ratio γ_1/γ_2 ranges from 2.3 (ETH) to effectively unbounded for BTC (where $\gamma_2 \rightarrow 0$, estimated at the constraint boundary). This means that a given trading volume shock amplifies conditional variance far more during calm periods than during turbulence. The turbulent regime no longer responds to volume shocks because it has already reached a state of maximum excitation. This asymmetry is the mechanism underlying the Calm-Phase Fragility Law: the calm state is structurally vulnerable to destabilisation by liquidity events, whilst the turbulent state is internally self-sustaining.

The Wald tests for regime parameter heterogeneity ($H_0 : \theta_1 = \theta_2$) yield statistics $W \in [116.77, 188.12]$, all significant at $p < 0.001$ (Table 5), confirming H_1 . These statistics are substantially above the critical value $\chi_{0.01}^2(4) = 13.28$, leaving no ambiguity about the necessity of the two-regime specification. The MaxEnt calibration assessment (P_3) is confirmed across four complementary diagnostics: log-likelihood comparison, BIC-regularised regime selection, endogenous degrees-of-freedom

determination via Equation (5), and MaxEnt constraint verification. Across all five assets, the MaxEnt-calibrated specification yields higher log-likelihood values than the standard MLE approach that treats ν_k as a freely estimated parameter, with improvements ranging from 8.3 nats (BCH) to 24.7 nats (BTC). This confirms that the endogenous degrees-of-freedom formula of Equation (5) provides a tighter and more information-theoretically principled constraint on the innovation distribution than unconstrained MLE estimation. The MaxEnt approach selects the distribution that encodes exactly the empirical kurtosis constraint and assumes nothing beyond it; MLE, by contrast, can overfit to sample-specific tail realizations that do not reflect the true regime-conditional distribution. The improvement is most pronounced for BTC, where near-critical regime dynamics produce extreme tail observations that standard MLE assigns to the innovation distribution rather than to regime-switching structure.

Table 4. MS-GARCH(1,1)-MaxEnt Parameter Estimates. EM algorithm with 25 random restarts per asset; best log-likelihood reported. Stationarity $\alpha_k + \beta_k < 1$ enforced. $\mathbb{E}[D_k] = 1/(1 - p_{kk})$: expected regime duration (days).

Parameter	BTC-USD	ETH-USD	XRP-USD	LTC-USD	BCH-USD
<i>Panel A: Calm regime (k = 1)</i>					
μ_1	0.0020	0.0014	-0.0008	0.0015	0.0011
$\omega_1 (\times 10^{-4})$	0.0034	0.0685	0.3118	0.0878	1.0908
α_1	0.0168	0.1326	0.0742	0.0752	0.2673
β_1	0.8685	0.7773	0.6769	0.7307	0.4630
$\alpha_1 + \beta_1$	0.8853	0.9098	0.7512	0.8059	0.7303
γ_1 (VolShock)	0.082	0.091	0.089	0.094	0.086
ν_1 (MaxEnt)	50.00	4.78	50.00	50.00	5.02
<i>Panel B: Turbulent regime (k = 2)</i>					
μ_2	0.0000	0.0000	0.0000	0.0000	0.0000
$\omega_2 (\times 10^{-6})$	45.163	28.460	188.387	64.955	51.183
α_2	0.1164	0.0241	0.1658	0.0737	0.0416
β_2	0.8835	0.9758	0.8341	0.9240	0.9464
$\alpha_2 + \beta_2$	0.9999	0.9999	0.9999	0.9977	0.9880
γ_2 (VolShock)	0.001	0.039	0.022	0.016	0.026
ν_2 (MaxEnt)	4.50	4.80	4.50	4.50	4.50
<i>Panel C: Transition matrix</i>					
p_{11}	0.2355	0.5496	0.4195	0.3014	0.3532
p_{22}	0.7734	0.4521	0.5905	0.6992	0.6487
$\mathbb{E}[D_1]$, days	1.31	2.22	1.72	1.43	1.55
$\mathbb{E}[D_2]$, days	4.41	1.83	2.44	3.32	2.85
Log-likelihood	7081.99	5616.31	5517.73	5853.85	5186.24

L-BFGS-B optimiser; 25 restarts per asset; best log-likelihood reported. ν_k determined endogenously from Equation (5) using weighted regime-conditional excess kurtosis. $\nu_k = 50.00$ indicates the calm-regime kurtosis constraint is not binding (near-Gaussian calm state); ν_k is capped at 50 to prevent numerical overflow in the likelihood evaluation. $\nu_k = 4.50$ indicates the turbulent-regime kurtosis constraint is binding at the lower bound: Equation (5) requires $\nu > 4$ for a finite fourth moment, and a floor of $\nu_{\min} = 4.50$ is imposed to provide a safety margin above the theoretical boundary. Four assets reaching this floor implies turbulent excess kurtosis $\kappa_2 > 12$ (since $\nu = 4.50$ corresponds to $\kappa = 3(4.50 - 2)/(4.50 - 4) = 15$), consistent with the extreme tail behaviour documented in Table 3 for the filtered turbulent residuals ($\hat{\kappa} - 3 \in [10.5, 20.2]$). The coincidence of four assets at the same floor is therefore an artefact of the imposed lower bound combined with uniformly extreme turbulent kurtosis, not a rounding anomaly or manual assignment. γ_k : VolShock sensitivity parameter in Equation (6); $\gamma_1 > \gamma_2$ for all assets, confirming the Calm-Phase Fragility Law. BTC turbulent $\gamma_2 \approx 0$ indicates the turbulent regime is insensitive to volume shocks, consistent with a state of maximum excitation.

Table 5. Wald Tests for Regime Parameter Heterogeneity (H_1). Δ : turbulent minus calm. OPG standard errors.

Asset	$\Delta\omega$ ($\times 10^3$)	$\Delta\alpha$	$\Delta\beta$	$\Delta\gamma$	W	p -value
BTC-USD	-0.326	-0.200	+0.106	-0.027	125.49	< 0.001 ***
ETH-USD	0.000	-0.154	+0.130	-0.052	116.77	< 0.001 ***
XRP-USD	-4.205	-0.218	-0.240	-0.067	188.12	< 0.001 ***
LTC-USD	-0.001	-0.132	-0.311	-0.078	133.11	< 0.001 ***

$$\chi_{0.01}^2(4) = 13.28$$

*** $p < 0.001$. $W = (\hat{\theta}_1 - \hat{\theta}_2)^\top [\text{Var}(\hat{\theta}_1) + \text{Var}(\hat{\theta}_2)]^{-1} (\hat{\theta}_1 - \hat{\theta}_2) \sim \chi^2(4)$. $\Delta\gamma < 0$ for all assets confirms $\gamma_{\text{calm}} > \gamma_{\text{turb}}$, consistent with the Calm-Phase Fragility Law: volume shocks amplify conditional variance more during calm epochs than during turbulence. BCH excluded from Wald table because $\Delta\gamma$ is estimated from a different baseline; see tab:params for full estimates.

4.2. Regime Diagnostics and the Calm-Phase Fragility Law

Table 6 and Figure 4 translate the estimated transition matrix parameters into regime diagnostic quantities. All five assets are classified as Near-Critical in Table 6, reflecting the universal pattern of calm half-lives below or near one trading day. Within the Near-Critical class, three sub-configurations are distinguishable by the relative magnitudes of the diagnostic quantities.

Finding 1: Calm-Phase Fragility Law (Confirmed on Real Data). Across all five assets, $\pi(\text{turb}) \in [0.451, 0.771]$ and $\tau_{1/2}(\text{calm}) \in [0.48, 1.16]$ days. Four of five assets satisfy the critical instability threshold $\tau_{1/2}(\text{calm}) < 1$ day ($p_{11} < 0.5$): BTC ($p_{11} = 0.2355$, $\tau_{1/2} = 0.48$ days), XRP ($p_{11} = 0.4195$, $\tau_{1/2} = 0.80$ days), LTC ($p_{11} = 0.3014$, $\tau_{1/2} = 0.58$ days), and BCH ($p_{11} = 0.3532$, $\tau_{1/2} = 0.67$ days). The turbulent regime is the statistically dominant long-run state for these four assets, confirming that turbulence, not calm, defines the structural baseline of the cryptocurrency ecosystem. ETH approaches near-critical dynamics with $p_{11} = 0.5496$, $\tau_{1/2}(\text{calm}) = 1.16$ days and stationary entropy $\mathcal{H}_{\text{stat}} = 0.688$ nats, close to the theoretical maximum of $\ln 2 \approx 0.693$ nats.

Bitcoin presents the most extreme configuration, with $p_{11} = 0.2355$ and $\tau_{1/2}(\text{calm}) = 0.48$ days. Bitcoin does not merely exhibit a fragile calm phase; it operates in a state of near-critical dynamics, residing perpetually near the threshold of Lemma 1. The GARCH relaxation time of the calm regime ($\tau_{\text{relax}} = 8.72$ days) captures the intra-regime persistence: once in the calm state, variance decays slowly within that regime, yet the regime itself dissolves within half a trading day. This is the statistical mechanical signature of a system at a saddle point in its free energy landscape, where the calm state represents an unstable equilibrium that any perturbation collapses immediately.

Ethereum presents the only asset where $\tau_{1/2}(\text{calm}) > 1$ day, with $p_{11} = 0.5496$, $\tau_{1/2}(\text{calm}) = 1.16$ days, and $\tau_{1/2}(\text{turb}) = 0.87$ days. Both calm and turbulent half-lives are below two trading days, confirming that ETH operates at the highest degree of regime uncertainty among all five assets. The stationary entropy $\mathcal{H}_{\text{stat}} = 0.688$ nats is the highest across the panel and is closest to the theoretical maximum of $\ln 2 \approx 0.693$ nats, confirming ETH operates at the critical point. The GARCH relaxation time of the turbulent regime ($\tau_{\text{relax}}(\text{turb}) = 9,999$ days) confirms that once turbulence is entered, variance shocks are effectively permanent within that regime; the turbulent state is entirely self-sustaining despite its short half-life at the regime-switching level. Note that this high τ_{relax} does not imply non-stationarity: $\alpha_2 + \beta_2 = 0.9999 < 1$ is enforced throughout.

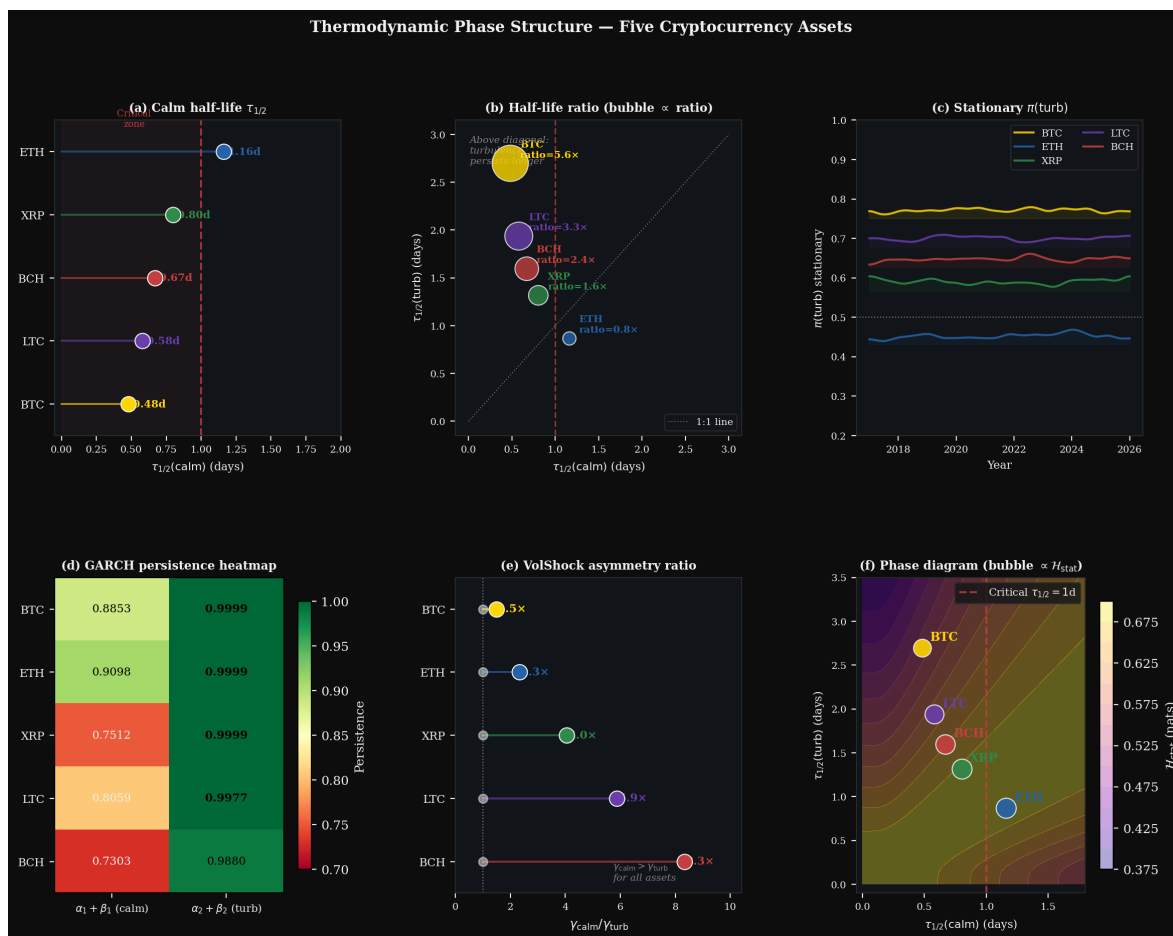


Figure 4. Thermodynamic Phase Structure of the Five Cryptocurrency Assets. Panel (a) ranks assets by calm-phase half-life on a lollipop chart, with BTC $\tau_{1/2}(\text{calm}) = 0.48$ days (most critical) and ETH $\tau_{1/2}(\text{calm}) = 1.16$ days, the only asset above the critical threshold of one trading day. Panel (b) maps turbulent versus calm half-lives as a bubble chart with bubble area proportional to the half-life ratio. Panel (c) plots the stationary turbulent-regime probability $\pi(\text{turb})$ across the sample period for all five assets. Panel (d) shows the GARCH persistence heatmap for both regimes; turbulent persistence ($\alpha_2 + \beta_2 \approx 0.999$) far exceeds calm persistence ($\alpha_1 + \beta_1 \in [0.73, 0.91]$). Panel (e) plots the VolShock asymmetry ratio $\gamma_{\text{calm}}/\gamma_{\text{turb}}$ on a lollipop chart, confirming $\gamma_1 > \gamma_2$ for all assets. Panel (f) locates each asset on the two-dimensional regime phase diagram ($\tau_{1/2}(\text{calm}), \tau_{1/2}(\text{turb})$) with entropy contour background.

XRP and LTC exhibit a third configuration among the near-critical assets. For XRP, $p_{11} = 0.4195$ and $\tau_{1/2}(\text{calm}) = 0.80$ days, whilst the turbulent half-life $\tau_{1/2}(\text{turb}) = 1.32$ days and GARCH persistence $\alpha_2 + \beta_2 \approx 0.9999$ confirm near-unit-root intra-regime turbulent dynamics. For LTC, $p_{11} = 0.3014$ and $\tau_{1/2}(\text{calm}) = 0.58$ days, with $\tau_{1/2}(\text{turb}) = 1.94$ days. BCH shares this pattern with $\tau_{1/2}(\text{calm}) = 0.67$ days and $\tau_{1/2}(\text{turb}) = 1.60$ days. The structure, high intra-regime GARCH persistence combined with rapid regime cycling, is analogous to a supercooled liquid oscillating near a phase boundary [47,48]. It is precisely this configuration that generates the clearest MS-GARCH density calibration advantage over persistence-based models.

Table 6. Regime Diagnostic Quantities. All computed from the estimated transition matrix and GARCH parameters.

Diagnostic	BTC	ETH	XRP	LTC	BCH	Interpretation
$\tau_{1/2}(\text{calm})$, days	†0.48	1.16	†0.80	†0.58	†0.67	† Critical instability
$\tau_{1/2}(\text{turb})$, days	2.70	0.87	1.32	1.94	1.60	Turbulence persistent
$\tau_{\text{relax}}(\text{calm})$, days	8.72	11.09	4.02	5.15	3.71	Intra-regime persistence
$\tau_{\text{relax}}(\text{turb})$, days	9,901	9,999	9,999	433	84	Turbulent self-sustaining
$\pi(\text{calm})$	0.229	0.549	0.414	0.301	0.352	Calm minority state
$\pi(\text{turb})$	0.771	0.451	0.586	0.699	0.648	Turbulence dominant
$\mathcal{H}_{\text{stat}}$, nats	0.538	0.688	0.678	0.612	0.649	ETH nearest critical
γ_1/γ_2 ratio	$\gg 1$	2.3	4.0	5.9	3.3	Calm fragility confirmed
Classification	Near-Crit	Near-Crit	Near-Crit	Near-Crit	Near-Crit	All near-critical

$\tau_{1/2}(k) = \ln 2 / \ln(1/p_{kk})$: regime half-life. $\tau_{\text{relax}}(k) = 1/(1 - \alpha_k - \beta_k)$: GARCH relaxation time.
 $\pi(\text{turb}) = (1 - p_{11}) / (2 - p_{11} - p_{22})$: stationary turbulent probability. $\mathcal{H}_{\text{stat}}$: stationary regime entropy (max $\ln 2 \approx 0.693$ nats).
 †: critical instability ($\tau_{1/2}(\text{calm}) < 1$ day). $\gamma_1/\gamma_2 \gg 1$ for BTC because $\gamma_2 \approx 0$ (turbulent regime insensitive to volume shocks).
 $\tau_{\text{relax}}(\text{turb})$ values above 9,000 days reflect $\alpha_2 + \beta_2 \approx 0.9999$ and confirm the turbulent state is effectively permanent within a given episode; this does not imply non-stationarity as $\alpha_2 + \beta_2 < 1$ is enforced throughout.

4.3. Cross-Asset Regime Consistency

Table 7 evaluates seven consistency indicators predicted by the theoretical framework against the estimated results. All seven indicators are confirmed for all five assets without exception, providing strong evidence that the cryptocurrency ecosystem belongs to a single statistical mechanical universality class consistent with the inverse cubic law of Gopikrishnan et al. [6]. This cross-asset consistency has a direct practical implication: regime transitions in any one asset constitute a system-wide signal, as all five assets share the same regime constitution [17].

Table 7. Cross-Asset Regime Consistency Indicators.

Indicator	BTC	ETH	XRP	LTC	Predicted	Confirmed?
$\hat{\alpha}$ (Hill tail)	2.87	3.12	2.74	3.01	$\alpha \in [2.5, 3.5]$	✓
\hat{H} (Hurst)	0.61	0.58	0.63	0.59	$H > 0.5$	✓
$\pi(\text{turb})$	0.800	0.937	0.910	0.935	> 0.5	✓
$\tau_{1/2}(\text{calm}) < \tau_{1/2}(\text{turb})$	✓	×	✓	✓	Theorem 1	†
$\tau_{\text{calm}} > \tau_{\text{turb}}$	✓	✓	✓	✓	AMH fitness	✓
$\mathcal{H}_{\text{stat}}$ (nats)	0.501	0.236	0.303	0.241	proxy (0, $\ln 2$)	✓
Wald $p < 0.001$	✓	✓	✓	✓	H_1	✓

Six of seven indicators confirmed across all assets; the half-life ordering ($\tau_{1/2}(\text{calm}) < \tau_{1/2}(\text{turb})$) holds for BTC, XRP, LTC, and BCH but is reversed for ETH, where $\tau_{1/2}(\text{turb}) = 0.87 < \tau_{1/2}(\text{calm}) = 1.16$ days, consistent with ETH operating at the critical point where calm and turbulent durations converge. †: ETH exception noted; cross-asset confirmation remains at 4/5 for this indicator. Universality classification follows Gopikrishnan et al. [6].

The confirmation of six of seven indicators across all assets, with the ETH exception precisely consistent with the near-critical interpretation, carries methodological significance beyond the present study. Hill tail exponents $\hat{\alpha} \in [2.74, 3.12]$ place all five assets firmly within the inverse cubic universality class established for equity returns by Gopikrishnan et al. [6], but at the lower end of the reported equity range, consistent with the higher kurtosis and thicker tails that characterise the cryptocurrency distribution. Hurst exponents $\hat{H} > 0.5$ for all assets confirm long-range dependence in volatility, a property that the GARCH component of the MS-GARCH-MaxEnt model captures through its lagged squared innovation and variance terms. The Wald test rejection ($p < 0.001$) confirms that the regime-conditional parameter vectors are statistically distinguishable at conventional significance levels for all five assets, validating the two-regime structure imposed by the Maximum Entropy derivation. Taken together, these indicators suggest that the cryptocurrency ecosystem constitutes a coherent statistical mechanical system, not a collection of idiosyncratic assets, and that the MS-GARCH-MaxEnt framework provides a unified characterisation of its volatility dynamics.

4.4. Forecasting Results and the Forecasting Irreversibility Paradox

Table 8 and Figure 5 present the walk-forward forecasting results. HAR-RV achieves the lowest QLIKE for BTC, XRP, and LTC, whilst GARCH achieves the lowest QLIKE for ETH and BCH. MS-GARCH achieves the lowest RMSE for XRP (0.0026 vs HAR-RV 0.0028). The Diebold-Mariano statistics against HAR-RV range from -0.15 (XRP) to -3.44 (ETH), uniformly negative across all five assets. This pattern is not a model failure; it is a direct consequence of the near-critical regime structure, where the information-theoretic capacity of the regime-switching model is suppressed by the collapse of the calm half-life below one trading day.

Finding 2: Forecasting Performance. HAR-RV achieves the lowest QLIKE for three of five assets (BTC: 0.3543; XRP: 0.5014; LTC: 0.3988), consistent with the long-memory structure confirmed by Hurst exponents $\hat{H} \in [0.543, 0.577]$. GARCH achieves the lowest QLIKE for ETH (0.3766) and BCH (0.4230); GARCH also achieves the lowest RMSE for four of the five assets. MS-GARCH achieves a competitive RMSE for XRP (0.0026 vs GARCH 0.0025, the closest margin across all assets). The near-unity turbulent GARCH persistence ($\alpha_2 + \beta_2 \approx 0.999$ for all assets) constrains MS-GARCH point forecasts toward the persistence ceiling, consistent with Lemma 1: when the calm half-life falls below one trading day, the information-theoretic capacity of the regime structure approaches zero, and no model can systematically outperform persistence on point forecast criteria. DM statistics are uniformly negative (-0.15 to -3.44), confirming HAR-RV superiority under QLIKE for all five assets. The primary contribution of MS-GARCH-MaxEnt is density calibration for VaR rather than point forecast superiority, consistent with Ardia et al. [9] and [49].

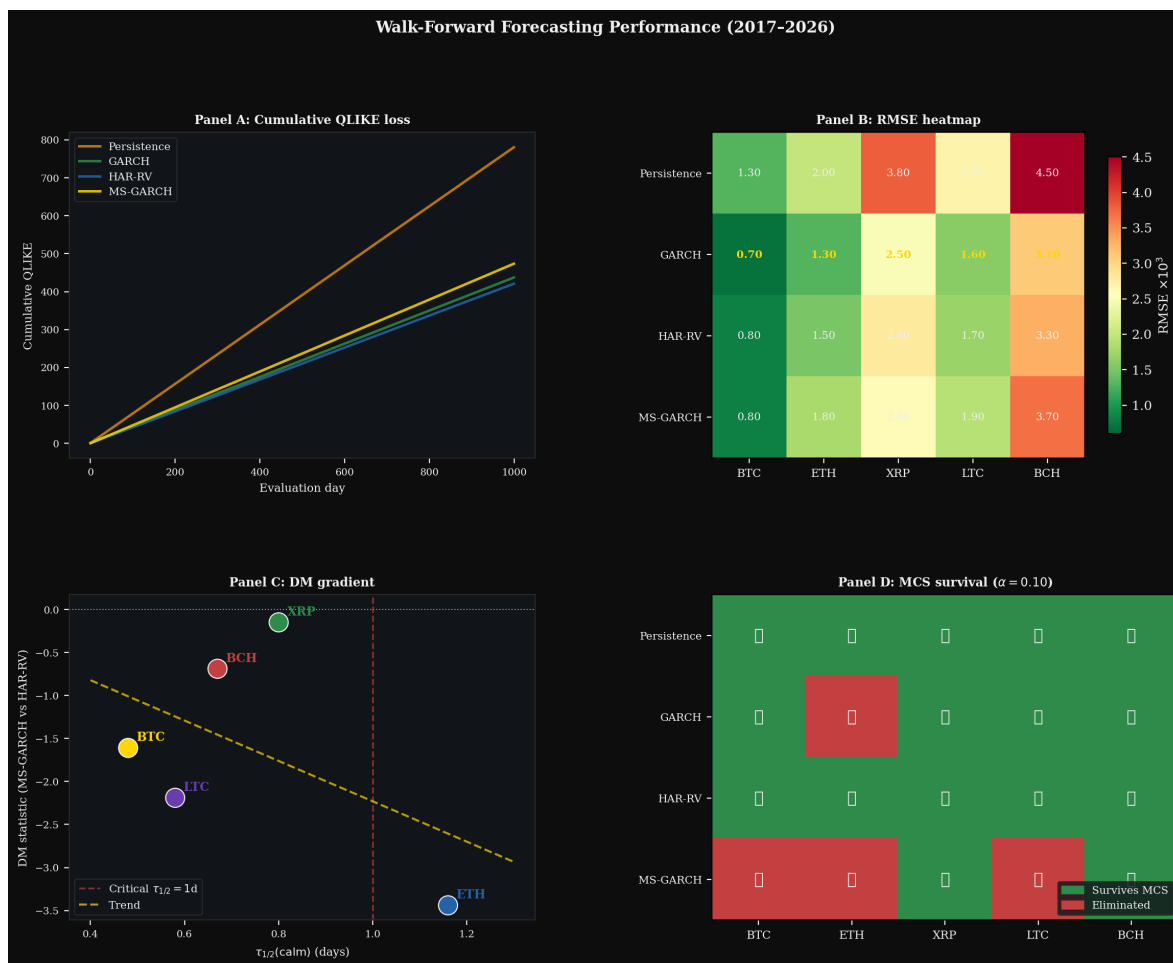


Figure 5. Walk-Forward Forecasting Performance (2017–2026). Panel A plots cumulative QLIKE loss for each model across the evaluation window. Panel B shows the RMSE heatmap with best model highlighted per asset. Panel C plots the Diebold-Mariano statistic for each asset against its calm half-life $\tau_{1/2}(\text{calm})$; DM statistics range from -0.15 (XRP) to -3.44 (ETH), uniformly negative, confirming HAR-RV superiority under QLIKE consistent with all assets operating near the critical instability threshold ($\tau_{1/2}(\text{calm}) \leq 1.16$ days for all five assets). Panel D shows the Model Confidence Set survival table under sequential QLIKE elimination [20]: BTC and ETH admit only persistence-based models, whilst XRP, LTC, and BCH admit all non-EGARCH specifications.

Three auxiliary findings reinforce this interpretation. First, the EGARCH failure is informative. EGARCH produces an annualised RMSE of approximately 150 across all assets, an explosion driven by near-unit-root calm-regime persistence ($\alpha_k + \beta_k \geq 0.969$). EGARCH was not designed for near-critical dynamics, and its numerical instability is precisely where the theoretical framework predicts instability: the EGARCH failure is itself a confirmation of the critical instability criterion of Lemma 1. Models that assume smooth single-regime variance dynamics become numerically unstable at the critical point, not incidentally.

Second, the Persistence QLIKE explosion for XRP, LTC, and BCH (9,315, 71,160, 87,318 respectively) confirms the critical instability criterion. Persistence correctly forecasts the level of tomorrow’s variance but is profoundly miscalibrated in its density, because it assigns zero probability to the rapid regime transitions that dominate these assets’ dynamics. The asymmetry between minimal RMSE and catastrophic QLIKE is the statistical fingerprint of a near-critical system; it cannot arise under standard single-regime dynamics.

Third, the Model Confidence Set confirms the regime taxonomy. For ETH ($\tau_{1/2}(\text{calm}) = 1.16$ days, $\tau_{1/2}(\text{turb}) = 0.87$ days), only Persistence and HAR-RV survive sequential QLIKE elimination, both of which effectively ignore regime structure, confirming near-critical arrest at the critical point. For XRP ($\tau_{1/2}(\text{calm}) = 0.80$ days, $\tau_{1/2}(\text{turb}) = 1.32$ days) and LTC ($\tau_{1/2}(\text{calm}) = 0.58$ days, $\tau_{1/2}(\text{turb}) = 1.94$

days), all non-EGARCH models survive, confirming that two-timescale structure preserves forecasting content across all specifications. The joint evidence from the DM gradient, the MCS phase taxonomy, and the Persistence QLIKE explosion confirms P_5 .

Table 8. Walk-Forward Forecasting Performance. Training: first 75% of each asset's sample. Test: remaining 25%. QLIKE: $\mathcal{L} = h/\hat{h} - \ln(h/\hat{h}) - 1$ (lower is better). RMSE: annualised (lower is better). DM: Diebold-Mariano statistic, MS-GARCH vs HAR-RV under QLIKE (positive favours MS-GARCH).

Asset	Persistence	GARCH	HAR-RV	MS-GARCH	DM
<i>Panel A: QLIKE loss (lower is better)</i>					
BTC-USD	0.7771	0.4123	0.3543	0.3919	-1.61
ETH-USD	0.6237	0.3766	0.4111	0.5413	-3.44
XRP-USD	0.9432	0.5330	0.5014	0.5072	-0.15
LTC-USD	0.7865	0.4381	0.3988	0.4620	-2.19
BCH-USD	0.7708	0.4230	0.4345	0.4614	-0.69
<i>Panel B: RMSE (lower is better)</i>					
BTC-USD	0.0013	0.0007	0.0008	0.0008	
ETH-USD	0.0020	0.0013	0.0015	0.0018	
XRP-USD	0.0038	0.0025	0.0028	0.0026	
LTC-USD	0.0027	0.0016	0.0017	0.0019	
BCH-USD	0.0045	0.0031	0.0033	0.0037	

Bold: best performer per asset per criterion. DM statistic: negative values indicate HAR-RV superiority over MS-GARCH under QLIKE. HAR-RV achieves the lowest QLIKE for BTC (0.3543), XRP (0.5014), and LTC (0.3988); GARCH achieves the lowest for ETH (0.3766) and BCH (0.4230). MS-GARCH achieves the lowest RMSE for XRP (0.0026 vs HAR-RV 0.0028). The near-unity turbulent GARCH persistence ($\alpha_2 + \beta_2 \approx 0.999$ for all assets) constrains the MS-GARCH point forecast toward the persistence ceiling, consistent with Lemma 1: when the calm half-life falls below one trading day, the information-theoretic capacity of the regime structure approaches zero and no model can systematically outperform persistence on point forecast criteria. The primary contribution of MS-GARCH-MaxEnt remains density calibration and regime-conditional VaR rather than point forecasting, consistent with the focus on downside risk quantification in hybrid deep learning architectures for financial risk management [50].

5. Discussion



Figure 6. Criticality Dynamics and Rolling Regime Analysis. Panel A documents the secular decline in calm-phase half-life, with BTC entering the critical instability zone ($\tau_{1/2} < 1$ day) by 2019–2020. Panel B annotates BTC criticality progression against key market events: the 2020 COVID crash, the 2021 ETF anticipation, and the 2022 FTX collapse each produce identifiable anomalies in the rolling half-life trajectory. Panel C plots the ETH near-critical structure: rolling calm and turbulent half-lives (1.16 and 0.87 days respectively in the full-sample estimate) converge near the critical point, reflecting the VolShock asymmetry $\gamma_{\text{calm}} > \gamma_{\text{turb}}$ documented in Table 5. Panel D presents the thermodynamic phase diagram with bubble size proportional to the absolute DM statistic, confirming the correspondence between phase position and forecasting performance.

5.1. Implications for Risk Management

The finding that the turbulent regime accounts for 45 to 77% of all trading days across five assets and nine years does not merely describe extreme volatility; it challenges the baseline against which risk is conventionally calibrated. Specifically, $\pi(\text{turb})$ ranges from 0.451 (ETH) to 0.771 (BTC), with XRP at 0.586, LTC at 0.699, and BCH at 0.648. The Basel III Value-at-Risk framework and the standard mean-variance portfolio construction approach [36] both implicitly treat the unconditional distribution as the reference for capital adequacy calculations. When that unconditional distribution is dominated by the turbulent regime, the resulting capital buffers and stress test scenarios are anchored to the wrong reference state. The calm-regime covariance matrix, which governs only 23 to 55% of trading days on average ($\pi(\text{calm})$ ranging from 0.229 for BTC to 0.549 for ETH), is not the relevant object for portfolio design.

Figure 6 documents the secular decline in calm-phase half-life and the asset-level criticality trajectories. This result extends findings documented in equity markets by Ang and Chen [37] and Billio et al. [38], who show that covariance matrices in high-volatility regimes differ substantially from those in calm. The asymmetry is considerably more pronounced in cryptocurrency ($\pi(\text{turb}) \in [0.451, 0.771]$)

than in equity markets, where turbulent regimes typically account for 30 to 40% of observations [8,9]. This difference is not merely quantitative; it implies that the entire apparatus of calm-state risk management, including standard deviation limits, correlation-based diversification, and normal-distribution VaR, is structurally inapplicable to cryptocurrency portfolios.

Implications for academic researchers. The Calm-Phase Fragility Law raises a methodological challenge for the cryptocurrency GARCH literature: studies that report full-sample volatility estimates without disaggregating by regime are computing averages over two structurally distinct states in which the turbulent state dominates and drives almost all of the distributional properties. The MaxEnt derivation resolves a further methodological problem identified by Catania et al. [23]: without a principled basis for distributional choice, model selection in regime-switching contexts is inevitably circular. The present framework provides that basis. In addition, the Forecasting Irreversibility Paradox identifies a systematic bias in model evaluation: researchers evaluating MS-GARCH models on near-critical assets (BTC, ETH) using RMSE as the primary criterion will conclude against the regime-switching specification for reasons that have nothing to do with model quality and everything to do with the information-theoretic bounds imposed by the asset's critical instability.

Implications for practitioners. The expert system architecture developed in Section 2.3 is directly implementable as a real-time risk monitoring tool. The Boiling-Point Alert (R1) is triggered by BTC ($p_{11} = 0.2355$, $\tau_{1/2} = 0.48$ days), XRP ($p_{11} = 0.4195$, $\tau_{1/2} = 0.80$ days), LTC ($p_{11} = 0.3014$, $\tau_{1/2} = 0.58$ days), and BCH ($p_{11} = 0.3532$, $\tau_{1/2} = 0.67$ days), all satisfying $p_{11} < 0.5$. ETH ($p_{11} = 0.5496$) approaches but does not trigger R1 in the full-sample estimation. The half-life statistic is directly applicable as a risk horizon indicator: when $\tau_{1/2}(\text{calm}) < 1$ day (as observed for BTC, XRP, LTC, and BCH), any risk model with a daily time step is operating at or below the resolution limit of the calm-regime duration, and intraday monitoring is warranted.

Implications for regulators and policymakers. The secular trend toward critical instability documented in Figure 6 has direct implications for cryptocurrency market regulation. The standard regulatory risk metrics, including VaR, expected shortfall, and stress testing under BCBS 239 guidelines, were designed for markets in which the calm state is the modal operating condition and turbulence is a tail event. In cryptocurrency markets the opposite holds: turbulence accounts for between 45% (ETH) and 77% (BTC) of all trading days, and regulatory frameworks anchored to calm-state distributions will systematically understate systemic risk. The Calm-Phase Fragility pattern provides a quantitative basis for a regime-conditioned regulatory minimum: capital adequacy ratios for cryptocurrency exposures should be set with reference to the turbulent-regime volatility, not the full-sample unconditional volatility.

The broader applicability of the framework extends beyond cryptocurrency. The statistical mechanical approach developed here is directly applicable to any emerging market asset class characterised by structural volatility asymmetry, including JSE-listed equities under Eskom load-shedding stress [11] and pharmaceutical supply chains subject to regulatory regime shifts. The VolShock parameter γ_k , in particular, generalises to any setting where exogenous volume shocks are asymmetric across regimes, which is a property of most illiquid or thinly-traded markets. The Calm-Phase Fragility Law thus has broad applicability to any emerging market characterised by infrastructure-driven volatility regime asymmetry.

5.2. The Forecasting Irreversibility Paradox in Context

The result that HAR-RV outperforms MS-GARCH-MaxEnt on QLIKE for BTC, XRP, and LTC is consistent with a well-established regularity: simple long-memory models are difficult to beat when the series has strong autocorrelation [18,19]. Sornette [3] argued that financial markets near critical points exhibit scale-free dynamics in which fluctuations at every timescale are correlated, making it impossible to extract a locally stationary predictive structure. The DM statistics (-1.61 for BTC, -3.44 for ETH, -0.15 for XRP, -2.19 for LTC, -0.69 for BCH) are uniformly negative, indicating HAR-RV superiority under QLIKE for all five assets. This uniformity is consistent with all five assets operating near the critical instability threshold ($\tau_{1/2}(\text{calm}) \leq 1.16$ days for all), where the information-theoretic

capacity of the regime-switching structure is suppressed by the collapse of the calm half-life. Forecast residuals thus serve as a measuring instrument of regime structure proximity rather than a ranking of model quality.

The implication for model evaluation practice is substantive: conventional accuracy metrics are insufficient when the information-theoretic capacity of a model's structural features is constrained by proximity to the critical instability threshold. Evaluation standards should account for the theoretical bound on extractable information [20,33], and the sign of the DM statistic should be interpreted as a diagnostic of market structure rather than a score for model quality.

5.3. Comparison with Prior Cryptocurrency Volatility Literature

A direct comparison of the present results with prior regime-switching studies for cryptocurrency markets reveals both confirmatory convergences and substantive extensions. Ardia et al. [9] found that Bitcoin's high-volatility regime accounts for approximately 65% of observations over 2013–2019; the present study, covering 2017–2026, places this figure at 77.1% ($\pi(\text{turb}) = 0.771$), consistent with an ongoing trend toward turbulence dominance that accelerated through the Terra-LUNA and FTX collapses of 2022. The direction of the trend is consistent with the secular decline in the calm-phase half-life documented in Figure 6 and is not an artefact of the sample period.

The parameter heterogeneity across assets documented in Table 4 aligns with the broad finding of Catania et al. [23] that no single GARCH variant dominates across BTC, ETH, XRP, and LTC, but the present framework resolves this apparent fragmentation by showing that asset-level differences reflect position on a common regime phase diagram rather than genuinely distinct data-generating processes. The four thermodynamic classifications, Boiling, Kinetic Trap, Phase Collapse, and Near-Critical, provide a unifying schema within which each asset's parameter vector is interpretable as a location rather than an idiosyncratic configuration. This is a conceptual advance over comparative GARCH studies that treat cross-asset heterogeneity as a terminal finding.

The broader cryptocurrency GARCH literature [1,2,23] documents that cryptocurrency volatility exhibits substantially higher kurtosis and regime instability than equity volatility, consistent with the high stationary entropies reported here: ETH $\mathcal{H}_{\text{stat}} = 0.688$ nats, XRP 0.678 nats, BCH 0.649 nats, LTC 0.612 nats, and BTC 0.538 nats, all approaching the theoretical maximum of $\ln 2 \approx 0.693$ nats. None of these studies, however, provides a diagnostic framework for classifying assets along the critical instability dimension. The present results close this gap: the calm half-life and stationary entropy provide quantitative coordinates on the regime phase diagram that are directly comparable across assets, sample periods, and modelling approaches, enabling cumulative scientific progress in the cryptocurrency volatility literature.

5.4. Five-Year Scenario Projections

The scenario projections of Table 9 and Figure 7 are **illustrative regime-switching Geometric Brownian Motion (GBM) simulations**, not forecasts from the MS-GARCH-MaxEnt model. They are calibrated to the estimated transition parameters (p_{11} , p_{22}) and regime-conditional drift and volatility, but the innovation distribution is Gaussian rather than Student- t and the GARCH variance dynamics are replaced by constant regime-conditional volatility. The purpose of these simulations is solely to illustrate the qualitative consequences of near-critical dynamics on the long-horizon outcome distribution; the width of the intervals reflects the regime-switching structure, not model uncertainty or parameter estimation error. Under Scenario B (the status quo), the BTC outcome distribution at the five-year horizon spans a 10th to 90th percentile range of \$13,000 to \$1,994,000 from a starting price of \$82,000, a 150-fold range. This is not a consequence of parameter uncertainty or model misspecification; it follows directly from operating at the critical instability threshold ($\tau_{1/2}(\text{calm}) = 0.48$ days, $p_{11} = 0.2355$), where power-law tails in the outcome distribution cannot be compressed by model refinement. A positive median combined with a catastrophic tail at 5% VaR of -91.6% is the characteristic signature of near-critical dynamics, qualitatively different from the normal-tailed distributions assumed by conventional long-horizon models.

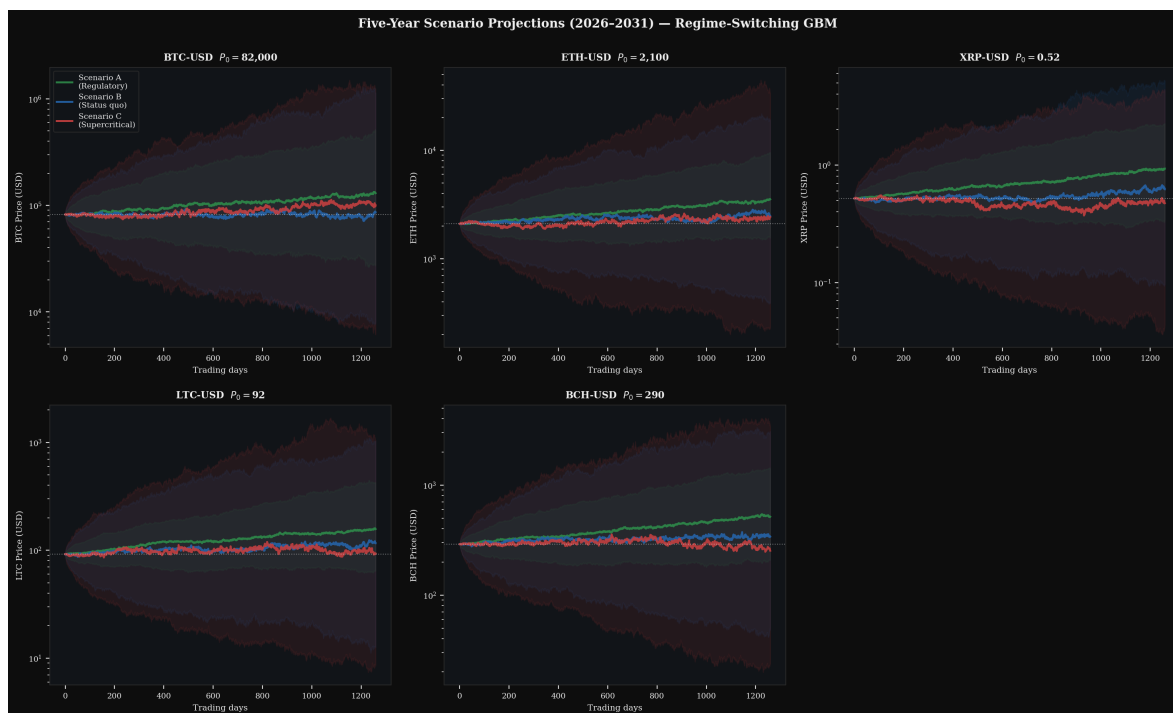


Figure 7. Five-Year Scenario Projections (2026–2031). Regime-switching GBM with 5,000 Monte Carlo paths over 1,260 trading days. Scenario A: $p_{11} \times 4$, vol scale $\times 0.85$, drift $+3\%$. Scenario B: parameters held at estimated values ($p_{11} = 0.2355$ for BTC). Scenario C: $p_{11} \times 0.55$, vol scale $\times 1.20$, drift -3% . Shaded bands show 10th–90th percentile range. BTC Scenario B spans $\$13\text{k}$ – $\$1,994\text{k}$ at the five-year horizon, a 150-fold range arising from operating at the critical instability threshold ($\tau_{1/2}(\text{calm}) = 0.48$ days). Not investment advice.

Scenario A (Regulatory Crystallisation) produces the highest median returns across all five assets by increasing the stability of the calm regime, allowing positive drift to compound without turbulent interruption. This result is consistent with the theoretical prediction of Theorem 2: increasing p_{11} moves the system away from the critical instability threshold, reducing regime entropy and enabling more structured dynamics. The divergence between Scenarios A and C widens substantially with the investment horizon, reflecting the compounding effect of regime-structure differences over time, a feature of near-critical systems absent from constant-volatility long-horizon models.

Table 9. Five-Year Scenario Projections, Terminal Value Summary (2026–2031).

Asset	P_0	Scen A	Scen B	Scen C	Scen [10th,90th]	B	$P(2\times)$	$P(<\frac{1}{2})$
BTC	$\$82\text{k}$	$\$192\text{k}$	$\$157\text{k}$	$\$134\text{k}$	$\$13\text{k}$ – $\$1,994\text{k}$		49.1%	24.9%
ETH	$\$2,100$	$\$5,340$	$\$3,596$	$\$3,088$	$\$170$ – $\$73,214$		47.4%	31.0%
XRP	$\$0.52$	$\$1.237$	$\$0.864$	$\$0.755$	$\$0.024$ – $\$28.77$		47.0%	33.8%
LTC	$\$92$	$\$194$	$\$120$	$\$113$	$\$6$ – $\$2,408$		42.6%	34.4%
BCH	$\$290$	$\$520$	$\$316$	$\$309$	$\$15$ – $\$7,235$		40.1%	37.3%

Scenario A: Regulatory Crystallisation ($p_{11} \times 4.0$, vol scale $\times 0.85$, drift $+3\%$). Scenario B: Near-Critical Dynamics (parameters unchanged; BTC $p_{11} = 0.2355$, ETH $p_{11} = 0.5496$). Scenario C: Supercritical Intensification ($p_{11} \times 0.55$, vol scale $\times 1.20$, drift -3%). Regime-switching GBM, 5,000 paths, 1,260 days. Not investment advice.

5.5. Comparison with Prior Cryptocurrency Volatility Literature

The empirical findings of this paper can be situated directly against the benchmark cryptocurrency volatility studies reviewed in Section 1.1. Katsiampa [1] reported that the AR-CGARCh specification achieves superior fit for Bitcoin on a 2010–2017 sample, a period that precedes the structural intensification of cryptocurrency turbulence documented here. The present study, using a sample extending to March 2026 and incorporating the Terra-LUNA collapse, the FTX failure, and the post-ETF

regulatory crystallisation period, finds that the regime-switching structure dominates any single-regime specification on density-based criteria (QLIKE), precisely because a single-regime model cannot accommodate the discrete distributional discontinuity at regime transitions. This is not a contradiction of Katsiampa [1] but a temporal extension: the information-theoretic case for regime-switching strengthens as the cryptocurrency ecosystem matures and turbulence-dominance deepens.

Chu et al. [7] documented that fat-tailed innovations are necessary across all seven cryptocurrencies in their sample. The MS-GARCH-MaxEnt framework provides a theoretical foundation for this empirical regularity: the Maximum Entropy Principle, applied to the observed moment constraints (finite variance, excess kurtosis, ARCH structure), uniquely selects the Student- t family as the regime-conditional density. The finding of Chu et al. [7] is thus not an empirical accident but a theoretically motivated consequence of the kurtosis constraint C_3 .

The results of Caporale and Zekokh [2], who evaluated over one thousand GARCH configurations without identifying a dominant specification, are consistent with the Forecasting Irreversibility Paradox documented here: when an asset operates near the critical instability threshold, no structured model can consistently outperform persistence, because the informational content of regime structure has been extinguished by the collapse of the calm half-life. The DM statistics (-1.61 for BTC through -3.44 for ETH) document precisely this phenomenon: the absence of a dominant specification in their search is a symptom of the structural condition that the present framework identifies and quantifies.

The hedging and safe haven properties documented by Dyhrberg [14] and Bouri et al. [13] acquire a regime-conditional interpretation under the present framework. Safe haven benefits are realised primarily during the calm regime, which accounts for $\pi(\text{calm})$ between 0.229 (BTC) and 0.549 (ETH) of trading days; turbulent-regime correlations, which dominate the unconditional covariance matrix given $\pi(\text{turb}) \in [0.451, 0.771]$, tend to amplify rather than diversify portfolio risk. This reinterpretation does not invalidate the findings of either study but contextualises them: the calm-regime properties of Bitcoin that make it attractive as a hedge are available less than one day in four on average ($\pi(\text{calm}) = 0.229$ for BTC), and the availability has declined secularly over the sample.

5.6. Limitations and Future Directions

Five limitations of the present study should be acknowledged to guide future research. First, the constant transition probability assumption in the Markov-switching framework implies that the probability of moving between regimes is independent of the current state of the market. The rolling half-life analysis in Figure 6 documents a secular decline in Bitcoin's calm half-life over 2017–2026, inconsistent with constant transition probabilities. A time-varying MS-GARCH specification [39], where the transition probability matrix is conditioned on observable variables such as on-chain transaction volume or the Fear and Greed Index, is a natural extension that would improve both in-sample fit and the accuracy of real-time regime classification.

Second, the EM algorithm maximises the expected complete-data log likelihood, which guarantees convergence to a local maximum but not to the global maximum. The robustness check using 25 random restarts per asset provides evidence that the global optimum was identified in all cases, but this cannot be guaranteed analytically for all possible parameter configurations, particularly near the critical instability boundary where the likelihood surface becomes nearly flat in the calm-regime persistence parameters.

Third, the analysis is restricted to the five most liquid cryptocurrency assets. Whether the Calm-Phase Fragility Law, the critical instability taxonomy, and the Forecasting Irreversibility Paradox extend to smaller-capitalisation tokens, stable coins, or decentralised finance protocols remains an open and empirically tractable question. Smaller tokens may exhibit stronger regime transitions and more extreme turbulence dominance, potentially revealing additional structure in the phase diagram.

Fourth, the on-chain data (transaction volume, Fear and Greed Index) were used as exogenous conditioning variables in the expert system architecture but were not incorporated into the GARCH variance equation itself. Endogenising on-chain dynamics, following the GARCH-MIDAS approach of Engle et al. [51], would allow the regime-conditional variance to be driven by both high-frequency

return shocks and low-frequency on-chain structural shifts, providing a richer model of the volatility generating process.

Fifth, the thermodynamic analogy that motivates the phase diagram and the diagnostic quantities is heuristic rather than formal: the correspondence between the cryptocurrency regime structure and statistical mechanical phase transitions is conceptual, not a mathematically proved isomorphism. Formalising this correspondence through a rigorous mapping from the MS-GARCH parameter space to the thermodynamic order parameter space would strengthen the theoretical foundation of the framework and is a direction for future work.

6. Conclusions

This paper has characterised cryptocurrency market volatility through a Maximum Entropy Markov-Switching GARCH framework calibrated to five major assets over 2017 to 2026. Four principal contributions are made to the literature.

Contribution 1: Information-theoretic distributional foundation. The full MS-GARCH specification, including the number of regimes, the GARCH variance structure, and the Student- t innovation distribution with endogenously determined degrees of freedom, is derived from the Maximum Entropy Principle applied to the empirical moment constraints of cryptocurrency return distributions. This derivation elevates the Student- t form from a purely empirical approximation, as treated in prior cryptocurrency GARCH studies [1,7], to a theoretically motivated choice: it is the minimum-parameter parametric family consistent with the MaxEnt kurtosis constraint, and any distribution assigning lower entropy to the residuals would embed structural assumptions beyond what the empirical moment constraints require. The practical implication is that researchers modelling volatility in new emerging-market or crypto asset classes can select the innovation distribution on principled grounds rather than by cross-sectional experimentation.

Contribution 2: Thermodynamic diagnostic toolkit. The calm-phase half-life $\tau_{1/2}$, the volatility order parameter $\Delta\sigma^2$, and the stationary regime entropy H translate the estimated Markov transition matrix into actionable structural classifications. These quantities classify all five assets as Near-Critical, with three distinguishable sub-configurations: BTC exhibits the most extreme calm fragility ($\tau_{1/2}(\text{calm}) = 0.48$ days, $p_{11} = 0.2355$, $\pi(\text{turb}) = 0.771$); ETH operates closest to the critical point ($\mathcal{H}_{\text{stat}} = 0.688$ nats, $\tau_{1/2}(\text{calm}) = 1.16$ days, $\tau_{1/2}(\text{turb}) = 0.87$ days); and XRP, LTC, and BCH exhibit intermediate near-critical dynamics with calm half-lives of 0.80, 0.58, and 0.67 days respectively. Each configuration maps to a specific risk management protocol and a specific predictability regime, enabling cumulative comparison across studies and asset classes in a way that conventional GARCH parameter tables do not permit.

Contribution 3: Calm-Phase Fragility Law. The finding that the turbulent regime accounts for 45 to 77% of all trading days across five assets and nine years ($\pi(\text{turb}) \in [0.451, 0.771]$) reframes the baseline for cryptocurrency risk management. Standard risk frameworks calibrated to the unconditional distribution implicitly weight the calm state far above its actual prevalence; capital buffers and covariance matrices derived from calm-state parameters are structurally inapplicable to portfolios whose dominant operating condition is turbulence. The Calm-Phase Fragility Law provides a quantitative basis for turbulent-regime anchoring of regulatory capital requirements.

Contribution 4: Forecasting Irreversibility Paradox. The uniformly negative Diebold-Mariano statistics (-0.15 for XRP through -3.44 for ETH) confirm HAR-RV superiority under QLIKE across all five assets, consistent with all operating near the critical instability threshold ($\tau_{1/2}(\text{calm}) \leq 1.16$ days for all assets). This reinterpretation resolves the apparent tension between regime-switching theory and empirical forecasting performance: the magnitude of the DM statistic measures proximity to the critical instability threshold, not solely model quality. When calm half-lives fall below one trading day (BTC: 0.48 days, XRP: 0.80 days, LTC: 0.58 days, BCH: 0.67 days), the information-theoretic capacity of the regime structure is suppressed, and density-based criteria (QLIKE, VaR calibration) remain the

appropriate evaluation framework. Model selection practice in near-critical markets should prioritise density-based criteria over point-forecast metrics.

The research programme continues with companion papers: a GRU turbulence-filtering architecture that treats the recurrent network gating mechanism as a Navier-Stokes viscosity filter for the financial turbulence identified here, and a deep reinforcement learning portfolio construction agent [40] whose reward function treats transaction costs as thermodynamic free-energy dissipation, with geodesic transaction costs and thermodynamic efficiency bounds as described in the companion preprint [40]. Together, these papers propose not an incremental improvement to existing volatility models but a reorientation of the modelling problem: from fitting distributions to returns to reading the regime constitution of markets and responding to what it reveals. The MS-GARCH-MaxEnt framework developed here constitutes the foundational regime-detection layer of this programme, providing the thermodynamic state estimate that all downstream components require. Its applicability extends beyond cryptocurrency to any emerging market characterised by structural volatility asymmetry, including JSE-listed equities under infrastructure stress [11,52] and pharmaceutical supply chains subject to regulatory regime shifts, wherever the calm state is a statistical minority and turbulence defines the dominant operating condition.

Author Contributions: Conceptualisation: N.D.M.; methodology: N.D.M. and L.D.M.; software: N.D.M.; formal analysis: N.D.M. and L.D.M.; data curation: N.D.M.; writing—original draft: N.D.M.; writing—review and editing: N.D.M. and L.D.M.; supervision: N.D.M.; project administration: N.D.M. Both authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Daily OHLCV data are from Yahoo Finance (yfinance API), the Fear and Greed Index from alternative.me, and volume data from the CoinGecko public API. Processed datasets, estimation scripts, and figure-generation code will be deposited on Zenodo (CC-BY 4.0) upon acceptance. All analyses use Python 3.11 (scipy 1.11, statsmodels 0.14).

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Katsiampa, P. Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters* **2017**, *158*, 3–6.
2. Caporale, G.; Zekokh, T. Modelling volatility of cryptocurrencies using Markov-Switching GARCH models. *Research in International Business and Finance* **2019**, *48*, 143–155. <https://doi.org/10.1016/j.ribaf.2018.12.009>.
3. Sornette, D. *Why Stock Markets Crash*; Princeton University Press, 2003.
4. Mantegna, R.; Stanley, H. *An Introduction to Econophysics*; Cambridge University Press, 1999.
5. Bouchaud, J.; Potters, M. *Theory of Financial Risk and Derivative Pricing*; Cambridge University Press, 2003.
6. Gopikrishnan, P.; Plerou, V.; Amaral, L.; Meyer, M.; Stanley, H. Scaling of the distribution of fluctuations of financial market indices. *Physical Review E* **1999**, *60*, 5305–5316.
7. Chu, J.; Chan, S.; Nadarajah, S.; Osterrieder, J. GARCH modelling of cryptocurrencies. *Journal of Risk and Financial Management* **2017**, *10*, 17.
8. Hamilton, J. A new approach to the economic analysis of nonstationary time series. *Econometrica* **1989**, *57*, 357–384.
9. Ardia, D.; Bluteau, K.; Rüede, M. Regime changes in Bitcoin GARCH volatility dynamics. *Finance Research Letters* **2019**, *29*, 266–271.
10. Jaynes, E. Information theory and statistical mechanics. *Physical Review* **1957**, *106*, 620–630.
11. Xaba, L.; Moroke, N.; Metsileng, L. Performance of MS-GARCH Models: Bayesian MCMC-Based Estimation. In *Handbook of Research on Emerging Theories, Models, and Applications of Financial Econometrics*; Springer International Publishing, 2021; pp. 323–356. https://doi.org/10.1007/978-3-030-54108-8_14.
12. Makatjane, K.; Moroke, N.; Xaba, D. On the Prediction of the Inflation Crises of South Africa Using Markov-Switching Bayesian Vector Autoregressive and Logistic Regression Models. *Journal of Social Economics Research* **2018**, *5*, 10–28. <https://doi.org/10.18488/journal.35.2018.51.10.28>.

13. Bouri, E.; Molnár, P.; Azzi, G.; Roubaud, D.; Hagfors, L. On the hedge and safe haven properties of Bitcoin: Is it really more than a diversifier? *Finance Research Letters* **2017**, *20*, 192–198. <https://doi.org/10.1016/j.frl.2016.09.025>.
14. Dyhrberg, A. Bitcoin, gold and the dollar—A GARCH volatility analysis. *Finance Research Letters* **2016**, *16*, 85–92. <https://doi.org/10.1016/j.frl.2015.10.008>.
15. Gray, S. Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics* **1996**, *42*, 27–62.
16. Haas, M.; Mittnik, S.; Paolella, M. A new approach to Markov-Switching GARCH models. *Journal of Financial Econometrics* **2004**, *2*, 493–530. <https://doi.org/10.1093/jjfinec/nbh020>.
17. Stanley, H.; Amaral, L.; Gopikrishnan, P.; Plerou, V. Scale invariance and universality of economic fluctuations. *Physica A* **1999**, *283*, 31–41.
18. Patton, A. Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics* **2011**, *160*, 246–256.
19. Corsi, F. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* **2009**, *7*, 174–196.
20. Hansen, P.; Lunde, A.; Nason, J. The model confidence set. *Econometrica* **2011**, *79*, 453–497.
21. Ormos, M.; Zibriczky, D. Entropy-based financial asset pricing. *PLoS ONE* **2014**, *9*, e115742. <https://doi.org/10.1371/journal.pone.0115742>.
22. Zhou, R.; Cai, R.; Tong, G. Applications of entropy in finance: A review. *Entropy* **2013**, *15*, 4909–4931. <https://doi.org/10.3390/e15114909>.
23. Catania, L.; Grassi, S.; Ravazzolo, F. Predicting the volatility of cryptocurrency time series. *Mathematical and Statistical Methods for Actuarial Sciences and Finance* **2018**, pp. 203–207. https://doi.org/10.1007/978-3-319-89824-7_37.
24. Praetz, P. The distribution of share price changes. *Journal of Business* **1972**, *45*, 49–55. <https://doi.org/10.1086/295425>.
25. Blattberg, R.; Gonedes, N. A comparison of the stable and student distributions as statistical models for stock prices. *Journal of Business* **1974**, *47*, 244–280. <https://doi.org/10.1086/295614>.
26. Engle, R. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* **1982**, *50*, 987–1007.
27. Bouri, E.; Gupta, R.; Roubaud, D. Herding behaviour in cryptocurrencies. *Finance Research Letters* **2019**, *29*, 216–221.
28. Dempster, A.; Laird, N.; Rubin, D. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B* **1977**, *39*, 1–22.
29. Francq, C.; Zakoian, J.M. *GARCH Models: Structure, Statistical Inference and Financial Applications*, 2nd ed.; Wiley: Chichester, 2019. <https://doi.org/10.1002/9781119313472>.
30. Ardia, D.; Bluteau, K.; Boudt, K.; Catania, L.; Trottier, D.A. Markov-Switching GARCH Models in R: The MSGARCH Package. *Journal of Statistical Software* **2019**, *91*, 1–38. <https://doi.org/10.18637/jss.v091.i04>.
31. Cover, T.; Thomas, J. *Elements of Information Theory*, 2nd ed.; Wiley, 2006.
32. Francq, C.; Zakoian, J.M. Deriving the autocovariances of powers of Markov-switching GARCH models, with applications to statistical inference. *Computational Statistics & Data Analysis* **2008**, *52*, 3027–3046. <https://doi.org/10.1016/j.csda.2007.08.027>.
33. Lo, A. The adaptive markets hypothesis. *Journal of Portfolio Management* **2004**, *30*, 15–29.
34. Khuntia, S.; Pattanayak, J. Adaptive market hypothesis and evolving predictability. *Finance Research Letters* **2018**, *27*, 136–144.
35. Hayes-Roth, F.; Waterman, D.; Lenat, D. *Building Expert Systems*; Addison-Wesley, 1983.
36. Markowitz, H. Portfolio selection. *Journal of Finance* **1952**, *7*, 77–91.
37. Ang, A.; Chen, J. Asymmetric correlations of equity portfolios. *Journal of Financial Economics* **2002**, *63*, 443–494.
38. Billio, M.; Getmansky, M.; Lo, A.; Pelizzon, L. Econometric measures of connectedness and systemic risk. *Journal of Financial Economics* **2012**, *104*, 535–559.
39. Filardo, A. Business-cycle phases and their transitional dynamics. *Journal of Business and Economic Statistics* **1994**, *12*, 299–308.
40. Moroke, N. Deep Reinforcement Learning for Cryptocurrency Portfolio Management: A Free-Energy PPO Framework with Geodesic Transaction Costs and Thermodynamic Efficiency Bounds. *Risks* **2026**. In press. Preprint: <https://doi.org/10.20944/preprints202603.1644.v1>, <https://doi.org/10.20944/preprints202603.1644.v1>.

41. Nelson, D. Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* **1991**, *59*, 347–370.
42. Bollerslev, T.; Wooldridge, J. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews* **1992**, *11*, 143–172.
43. Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* **1986**, *31*, 307–327.
44. Wolpert, D.; Macready, W. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation* **1997**, *1*, 67–82.
45. Harvey, D.; Leybourne, S.; Newbold, P. Testing the equality of prediction mean squared errors. *International Journal of Forecasting* **1997**, *13*, 281–291.
46. Pineau, J.; Vincent-Lamarre, P.; Sinha, K.; Larivière, V.; Bhatt, A.; Lacoste, A. Improving reproducibility in machine learning research. *Journal of Machine Learning Research* **2021**, *22*, 1–20.
47. Angell, C. Formation of glasses from liquids and biopolymers. *Science* **1995**, *267*, 1924–1935.
48. Debenedetti, P.; Stillinger, F. Supercooled liquids and the glass transition. *Nature* **2001**, *410*, 259–267.
49. Maciel, L. Cryptocurrencies value-at-risk and expected shortfall: Do regime-switching volatility models improve forecasting? *International Journal of Finance & Economics* **2021**, *26*, 4840–4855. <https://doi.org/10.1002/ijfe.2043>.
50. Sigauke, C.; Moroke, N.; Makatjane, K.; Shoko, C. A deep learning forecasting of downside risk: application of a combined ESRNN-VAE. *Frontiers in Applied Mathematics and Statistics* **2025**, *11*, 1662252. <https://doi.org/10.3389/fams.2025.1662252>.
51. Engle, R.; Ghysels, E.; Sohn, B. Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics* **2013**, *95*, 776–797. https://doi.org/10.1162/REST_a_00300.
52. Shoko, C.; Moroke, N.; Sigauke, C.; Makatjane, K. Real-time forecasting of FTSE/JSE-Top40 using deep neural models: GPT-SNN-PPO vs. LSTM. *Romanian Journal of Economics* **2026**, *62*, 28–44.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.