computer sciences and mathematics forum

MDPI

Proceeding Paper

Detecting Financial Bubbles with Tail-Weighted Entropy †

Omid M. Ardakani

Department of Economics, Parker College of Business, Georgia Southern University, Savannah, GA 31419, USA; oardakani@georgiasouthern.edu

† Presented at the 11th International Conference on Time Series and Forecasting, Canaria, Spain, 16–18 July 2025.

Abstract

This paper develops a novel entropy-based framework to quantify tail risk and detect speculative bubbles in financial markets. By integrating extreme value theory with information theory, I introduce the Tail-Weighted Entropy (TWE) measure, which captures how information scales with extremeness in asset price distributions. I derive explicit bounds for TWE under heavy-tailed models and establish its connection to tail index parameters, revealing a phase transition in entropy decay rates during bubble formation. Empirically, I demonstrate that TWE-based signals detect crises in equities, commodities, and cryptocurrencies days earlier than traditional variance-ratio tests, with Bitcoin's 2021 collapse identified weeks prior to the peak. The results show that entropy decay—not volatility explosions—serves as the primary precursor to systemic risk, offering policymakers a robust tool for preemptive crisis management.

Keywords: Tail-Weighted Entropy; extreme value theory; financial bubbles; entropy concentration bounds; Kullback–Leibler divergence; systemic risk

1. Introduction

The quantification of tail risk constitutes a fundamental challenge in modern statistical science, with profound implications for financial stability and climate resilience [1,2]. While extreme value theory has matured significantly since its inception in the work of Fisher and Tippett [3], contemporary applications demand information-theoretic tools that transcend traditional moment-based analyses [4,5]. This imperative emerges from the limitations in existing risk measures, which were vividly exposed during the 2008 financial crisis [6], recent climate catastrophes [7], and the COVID-19 pandemic [8,9].

Traditional entropy metrics [10,11], though powerful for central tendency analysis, fail in tail regimes where standard distributional assumptions collapse. This failure concretely manifests in financial bubble detection latency [12–14] and climate risk underestimation [15,16]. This paper addresses this dichotomy through three pivotal advances: (1) the development of *Tail-Weighted Entropy (TWE)*, whose measures are compatible with peaks-over-threshold frameworks, (2) the derivation of entropy decay laws linking information dynamics to tail index parameters, and (3) the construction of entropic hypothesis tests with exact asymptotic properties. The financial implications are immediate: the proposed framework detects speculative bubbles earlier than current moment-based approaches.

Foundational work in extreme value theory by Gnedenko [17], Balkema and De Haan [18], Pickands III [19], and Leadbetter et al. [20] established mathematical frameworks for tail analysis, later operationalized through peaks-over-threshold (POT) methodologies. Parallel developments in information theory by Shannon [21], Kullback and Leibler [22], and

check for updates

Academic Editors: Olga Valenzuela, Fernando Rojas, Luis Javier Herrera, Hector Pomares and Ignacio Rojas

Published: 25 July 2025

Citation: Ardakani, O.M. Detecting Financial Bubbles with Tail-Weighted Entropy. *Comput. Sci. Math. Forum* **2025**, *11*, 3. https://doi.org/10.3390/cmsf2025011003

Jaynes [23] created powerful tools for uncertainty quantification, although their initial applications focused on central tendencies [10,24]. The limitations of Shannon entropy in tail regimes were first identified by Jaynes [25], who highlighted information loss in extreme value analysis. Modern risk management frameworks [26] subsequently revealed the concrete consequences of this limitation through a systematic underestimation of the financial tail risks.

Recent advances follow two main trajectories: Chen and Singh [27] developed climate-specific tail entropy measures, extended by Faranda et al. [28], who incorporated dynamical systems theory for extreme weather attribution, and Ragone and Bouchet [29] who applied information-theoretic rare event sampling in climate models. Concurrently, financial applications have seen innovations such as Billio et al. [30]'s entropy-based systemic risk indicators, augmented by Simsekli et al. [31] using neural networks to estimate tail entropy in high-dimensional markets and Ibragimov et al. [32] establishing geometric properties of divergence measures under heavy-tailed distributions. However, these approaches remain disconnected from POT frameworks—climate applications often rely on block maxima methods [15] while financial implementations use ad hoc thresholding [33]—resulting in estimators that are sensitive to distributional misspecification. Bridging this gap, Ardakani [8] recently proposed a POT-compatible entropy estimator but lacked the necessary theoretical links to extremal index estimation. This paper synthesizes these strands through measure-theoretic innovations that preserve asymptotic properties while enabling practical implementation.

The primary contributions of this paper are threefold. First, the TWE generalizes Shannon entropy to extreme regimes through conditional survival analysis. Second, the results establish asymptotic bounds for TWE under generalized Pareto distributions, revealing explicit dependence on the tail index. Third, an asymptotically normal test statistic is introduced to enable hypothesis testing for speculative bubbles. The empirical analysis demonstrates superior bubble detection by highlighting the average lead time compared to other methodologies. It also shows threshold stability with a higher consistency ratio compared to the Hill estimator. Applications to E-mini S&P 500, gold, crude oil, and Bitcoin futures markets establish entropy dynamics as fundamental market observables.

The remainder of this paper unfolds as follows. Section 2 introduces a tail-weighted entropy measure, establishing concentration bounds under generalized Pareto distributions and explicit entropy–decay relationships via digamma functions. Section 3 develops the dynamics of extremal entropy for explosive financial processes, deriving Kullback–Leibler (KL) divergence growth laws and entropy-based bubble detection thresholds. Section 4 characterizes information dissipation through entropy production laws and provides exact critical values for phase transition detection using catastrophe theory. Section 5 applies this framework across major asset classes, demonstrating improvements in early warning signals compared to moment-based approaches while validating threshold stability through Basel III-inspired robustness checks. Sections 6 and 7 discuss the implications for risk management and regulatory surveillance and provide concluding remarks.

2. Extremal Entropy Measures

Early efforts to quantify tail uncertainty focused on moment-based approaches [18], but these lack sensitivity to higher-order information structures. Jaynes [25] was the first to highlight the information-theoretic limitations in extreme value analysis, while Chen and Singh [27] developed tail-entropy measures for climate applications. However, existing formulations do not directly align with POT frameworks [19]. This section introduces the tail-weighted entropy measure, which integrates with the POT methodology and establishes fundamental concentration bounds for this information-theoretic measure under

generalized Pareto (GPD) assumptions. It then provides explicit connections between entropy decay rates and tail heaviness. The proposed TWE formulation addresses the compatibility issues in Chen and Singh [27] while generalizing the Shannon entropy framework to extreme regimes. The concentration bounds improve upon the moment inequalities in Balkema and De Haan [18] by incorporating information-theoretic constraints.

Definition 1. *For a random variable X with absolutely continuous distribution function (CDF) F and probability density function (PDF) f, the Tail-Weighted Entropy beyond threshold τ is*

$$\mathcal{H}(\tau) = -\int_{\tau}^{\infty} \frac{f(x)}{\bar{F}(\tau)} \log\left(\frac{f(x)}{\bar{F}(\tau)}\right) dx, \tag{1}$$

where $\bar{F}(\tau) = 1 - F(\tau)$ is the survival function. This quantifies the information content in the conditional distribution $X|X > \tau$.

The TWE in Definition 1 generalizes classical Shannon entropy [21], specifically focusing on tail regions and addressing the limitations identified by Jaynes [25] in extreme value analysis. Unlike the standard entropy $H(X) = -\int f(x) \log f(x) dx$, the TWE captures how information scales with extremeness, which is essential for risk quantification when tail dynamics dominate system behavior. Recently, Chen and Singh [27] demonstrate similar tail-focused measures in climate extremes, but the proposed formulation offers direct compatibility with peaks-over-threshold frameworks.

Lemma 1. *Let $X \sim \mathcal{GP}(\xi, \sigma)$ with shape parameter $\xi < 1$, scale $\sigma > 0$, and support $x \geq 0$. The TWE satisfies*

$$\mathcal{H}(\tau) = \log(\sigma + \xi\tau) + 1 + \xi + \varepsilon(\tau), \tag{2}$$

where $\varepsilon(\tau) = o(1)$ as $\tau \to \infty$. Furthermore, the entropy decay rate admits the asymptotic expansion

$$\mathcal{H}(\tau) = \log\sigma + (1 - \xi)\psi\left(\frac{1}{\xi}\right) + \frac{\xi}{1 - \xi} + o(1), \tag{3}$$

where $\psi(\cdot)$ is the digamma function. Equality holds in the exponential tail limit $\xi \to 0$.

Proof. Let $Y = X - \tau \mid X > \tau \sim \mathcal{GP}(\xi, \sigma_\tau)$ where $\sigma_\tau = \sigma + \xi\tau$. The conditional density is

$$f_Y(y) = \frac{1}{\sigma_\tau}\left(1 + \frac{\xi y}{\sigma_\tau}\right)^{-1/\xi - 1}.$$

The tail entropy becomes

$$\mathcal{H}(\tau) = \log\sigma_\tau + \left(1 + \frac{1}{\xi}\right) \underbrace{\mathbb{E}\left[\log\left(1 + \frac{\xi Y}{\sigma_\tau}\right)\right]}_{I}.$$

First, compute I. Let $Z = 1 + \frac{\xi Y}{\sigma_\tau}$. Then,

$$I = \int_1^{\infty} \log(z) \cdot \frac{1}{\xi} z^{-1/\xi - 1} dz.$$

Substitute $t = z^{-1/\xi} \Rightarrow dt = -\frac{1}{\xi} z^{-1/\xi - 1} dz$,

$$I = -\int_0^1 \log(t^{-\xi}) t dt = \xi \int_0^1 \log(t) \cdot t dt = \xi\left[\frac{t^2}{2}\log(t) - \frac{t^2}{4}\right]_0^1 = -\frac{\xi}{4} + \xi \lim_{t \to 0^+} \frac{t^2}{2}\log(t) = \xi.$$

Thus, $\mathcal{H}(\tau) = \log \sigma_\tau + 1 + \xi$. For $\tau \to \infty$,

$$\log \sigma_\tau = \log(\sigma + \xi\tau) = \log(\xi\tau) + \log\left(1 + \frac{\sigma}{\xi\tau}\right) = \log \tau + \log \xi + \frac{\sigma}{\xi\tau} + o\left(\tau^{-1}\right).$$

The survival function satisfies $\bar{F}(\tau) = (1 + \xi\tau/\sigma)^{-1/\xi}$, so

$$\log \tau = \log \sigma - \xi \log \bar{F}(\tau) - \log \xi + o(1).$$

Substitute into $\mathcal{H}(\tau)$,

$$\begin{aligned}\mathcal{H}(\tau) &= [\log \sigma - \xi \log \bar{F}(\tau) - \log \xi + o(1)] + \log \xi + 1 + \xi + o(1) \\ &= \log \sigma + 1 + \xi - \xi \log \bar{F}(\tau) + o(1).\end{aligned}$$

From the regular variation of $\bar{F}(\tau)$, $\log \bar{F}(\tau) = -\frac{1}{\xi} \log(1 + \xi\tau/\sigma)$. Using properties of the digamma function $\psi(z) = \frac{d}{dz} \log \Gamma(z)$,

$$-\xi \log \bar{F}(\tau) = \log(1 + \xi\tau/\sigma) = \psi\left(\frac{1}{\xi}\right) + \frac{\xi}{1 - \xi} + o(1),$$

where the equality follows from the asymptotic expansion of the incomplete gamma function. Combining terms yields the digamma-form expansion. The $\xi \to 0$ case follows from $\psi(1/\xi) \approx \log(1/\xi) - \gamma$ where γ is Euler–Mascheroni. □

Lemma 1 establishes limits on tail entropy for heavy-tailed models, complementing moment-based analyses in Balkema and De Haan [18]. The entropy decomposition reveals an explicit dependence on the tail index ξ through the digamma function $\psi(\cdot)$. Empirical studies of flood risks [34] confirm that $\xi > 0.3$ leads to entropy values exceeding $\log \sigma + 1$ by over 40%, aligning with the theoretical phase transition at $\psi(1/\xi)$. This results bridge information theory and extreme value analysis by showing how entropy production accelerates nonlinearly with tail heaviness via $\psi(1/\xi)$ asymptotics. The digamma relationship explains the threshold effects observed in financial bubbles [13] through entropy convexity in ξ, while matching the climate extremality patterns in Coles et al. [15]. The TWE framework demonstrates that heavier tails ($\xi \uparrow$) permit greater entropy production but reduce predictability—a critical insight for coherent risk measures [35–37]. Specifically, $\partial\mathcal{H}/\partial\xi > 0$ when $\xi > 0.3$, confirming Chen and Singh [27]'s empirical findings through POT theory.

3. Extreme Value Theory and Entropy

The intersection of extreme value theory and information theory has emerged as a frontier in financial econometrics, particularly for analyzing market instability. Bhaduri [38] introduces entropy methods for bubble detection, leveraging Shannon entropy to quantify tail risk information. Diks and Panchenko [39] also advance entropy-based causality testing, revealing nonlinear dependencies in financial markets. Traditional approaches relying on moment explosions [12–14] and variance ratios [40] often miss early warning signals due to their insensitivity to higher-order information structures. The recent literature in heavy-tailed time series provides theoretical foundations for modeling financial extremes but lacks explicit connections to information dynamics [41]. For instance, Chavez-Demoulin and Davison [42] formalize extremal dependence in multivariate settings, yet their framework does not incorporate entropy as a measure of tail uncertainty. Concurrently, Billio et al. [30] develop entropy-based crisis indicators, although their framework remains limited to stationary regimes.

To address nonstationarity, Dimpfl and Peter [43] and Ardakani [36] introduce transfer entropy to capture asymmetric information flows in time-varying markets, while Ren and Sun [44] propose wavelet entropy for detecting structural breaks in non-stationary financial series. Bridging these strands, Costa et al. [45] establish multiscale entropy frameworks for analyzing complexity in non-equilibrium systems, offering insights applicable to explosive financial processes. This section synthesizes these studies by establishing entropy–decay laws for explosive financial processes, unifying insights from extreme value analysis, non-stationary information dynamics, and entropy-based tail risk measurement.

This section (1) establishes exact asymptotic laws for tail entropy decay in explosive AR processes, (2) quantifies information divergence growth during bubble episodes via KL divergence bounds, and (3) provides explicit phase transition thresholds between stable and bubble regimes. The results generalize the entropy inequalities of Cover [10] to non-stationary heavy-tailed processes while improving the detection thresholds in Phillips et al. [13] by incorporating information-theoretic criteria.

Theorem 1. *Let $\{P_t\}$ follow an explosive autoregressive process $P_t = \delta P_{t-1} + \epsilon_t$ with $\delta > 1$ during bubble periods, where $\{\epsilon_t\}$ are i.i.d. innovations with regularly varying tails of index $-\alpha$. Then, the following occur:*

1. *The tail-weighted entropy grows logarithmically:*

$$\lim_{\tau \to \infty} \frac{\mathcal{H}(\tau)}{\log \tau} = 1 \quad a.s., \tag{4}$$

   where the convergence rate is determined by the tail index α.
2. *The KL divergence between consecutive tail distributions decays polynomially:*

$$\mathcal{D}_{KL}(f_{t+1}; f_t) \sim C(\delta, \alpha)\tau^{-\alpha} \quad as\ \tau \to \infty, \tag{5}$$

   where $C(\delta, \alpha) > 0$ depends on the process parameters.

Proof. Part 1. According to the Kesten–Goldie theorem [46], the stationary distribution of P_t under $\delta > 1$ has regularly varying tails with index $-\alpha$, where $\delta^\alpha \mathbb{E}[|\epsilon_1|^\alpha] = 1$. When exceeding threshold τ, the Pickands–Balkema–de Haan theorem [18,19] implies convergence to a generalized Pareto distribution $\mathcal{GP}(\xi, \sigma_\tau)$, where $\xi = 1/\alpha$.

From Lemma 1, the tail-weighted entropy is

$$\mathcal{H}(\tau) = \log(\sigma + \xi\tau) + 1 + \xi + o(1).$$

As $\tau \to \infty$, $\sigma + \xi\tau \sim \xi\tau$, so

$$\mathcal{H}(\tau) = \log \tau + \log \xi + 1 + \xi + o(1).$$

Normalizing by $\log \tau$,

$$\frac{\mathcal{H}(\tau)}{\log \tau} = 1 + \frac{\log \xi + 1 + \xi}{\log \tau} + o\left(\frac{1}{\log \tau}\right) \to 1 \quad \text{a.s.}$$

Part 2. For large τ, the explosive term dominates ($P_t \approx \delta P_{t-1}$). The densities satisfy $f_t(x) \approx \frac{1}{\delta} f_{t-1}\left(\frac{x}{\delta}\right)$ in the tail. The KL divergence becomes

$$\mathcal{D}_{KL}(f_t; f_{t-1}) \approx \int_\tau^\infty \frac{1}{\delta} f_{t-1}\left(\frac{x}{\delta}\right) \log\left(\frac{f_{t-1}(x/\delta)}{\delta f_{t-1}(x)}\right) dx.$$

Substituting $y = x/\delta$, we obtain

$$\mathcal{D}_{KL} \approx \int_{\tau/\delta}^{\infty} f_{t-1}(y) \left[-\log \delta + \log \frac{f_{t-1}(y)}{f_{t-1}(\delta y)} \right] dy.$$

For regularly varying $f_{t-1}(y) \sim Cy^{-\alpha-1}$,

$$\log \frac{f_{t-1}(y)}{f_{t-1}(\delta y)} \approx (\alpha+1) \log \delta.$$

Thus,

$$\mathcal{D}_{KL} \approx (\alpha+1) \log \delta \int_{\tau/\delta}^{\infty} f_{t-1}(y) dy = (\alpha+1) \log \delta \cdot \bar{F}_{t-1}(\tau/\delta).$$

Using $\bar{F}_{t-1}(\tau/\delta) \sim C(\tau/\delta)^{-\alpha}$,

$$\mathcal{D}_{KL} \sim C(\delta, \alpha) \tau^{-\alpha}.$$

□

Theorem 1 establishes an information-theoretic foundation for analyzing financial bubbles, complementing traditional moment-based approaches [13]. The entropy scaling law in Part 1 reveals how bubble dynamics fundamentally alter tail uncertainty quantification, while Part 2 characterizes the decay of distributional similarity in extreme regimes. The results advance beyond variance-ratio methods [40] by incorporating the higher-order information structures critical for detecting non-linear bubble mechanisms [35,47].

Example 1. *Consider a sequence of explosive AR(1) processes indexed by time step* Δt

$$P_{t+\Delta t} = (1 + \mu \Delta t) P_t + \epsilon_t(\Delta t), \quad \mu > 0,$$

where innovations $\epsilon_t(\Delta t)$ *are scaled to produce a diffusion limit. Let* $\epsilon_t(\Delta t) \sim \sigma P_t \sqrt{\Delta t}\, \eta_t$ *with* $\eta_t \sim$ *i.i.d.* $\mathcal{N}(0,1)$*, establishing the volatility scaling. Under the Kesten–Goldie regularity condition [46], the tail index* α *satisfies*

$$(1 + \mu \Delta t)^{\alpha} \mathbb{E}[|\eta_1|^{\alpha}] (\sigma \sqrt{\Delta t})^{\alpha} = 1.$$

As $\Delta t \to 0$*, Taylor expansion yields* $\alpha \sim \frac{2\mu}{\sigma^2} \Delta t^{-1} \to \infty$*, making tails exponentially decaying rather than heavy. Applying Theorem 1, tail entropy growth becomes dominated by Gaussian-like decay*

$$\lim_{\substack{\Delta t \to 0 \\ \tau \to \infty}} \frac{\mathcal{H}(\tau)}{\sqrt{\log \tau}} = \sqrt{\frac{\sigma^2}{2\mu}},$$

matching the entropy rate for geometric Brownian motion (GBM) $dP_t = \mu P_t dt + \sigma P_t dW_t$ *[48]. KL divergence decay transitions from polynomial to exponential*

$$\mathcal{D}_{KL}(f_{t+\Delta t}; f_t) \sim C \exp\left(-\frac{\mu}{\sigma^2} \log^2 \tau \right),$$

consistent with log-normal tail distinguishability [10].

This limit reveals three fundamental phase transitions. First, tail entropy scaling crosses over from $\log \tau$ *(heavy-tailed) to* $\sqrt{\log \tau}$ *(log-normal). Second, KL divergence decay accelerates from power law to Gaussian-type exponential rates. Finally, extreme value index* $\xi = 1/\alpha$ *collapses to 0, marking the disappearance of regular variation. This example demonstrates the framework's universality across tail regimes, with Theorem 1 recovering classical diffusion entropy laws through*

$$\lim_{\alpha \to \infty} \underbrace{C(\delta, \alpha) \tau^{-\alpha}}_{\text{Theorem}} = \underbrace{C \exp(-\lambda \log^2 \tau)}_{\text{GBM Limit}}.$$

The logarithmic entropy growth law establishes a universal scaling relationship for explosive processes. This can reveal that bubble formations fundamentally reconfigure tail information structures rather than merely amplifying volatility and explains why variance ratio tests [47] exhibit detection latency—they monitor second-moment explosions while entropy captures the underlying information geometry shift. The polynomial KL divergence decay quantifies the regulatory implications: distinguishability between successive tail distributions diminishes as τ increases, creating detection windows where α-estimates from Theorem 1 determine intervention deadlines. The example demonstrates the framework's universality, showing how TWE transitions between heavy-tailed and log-normal regimes—a phase change undetectable by extremal dependence measures [42]. By linking entropy production to the Kesten–Goldie tail index α, this result reconciles Diks and Panchenko [39]'s causality entropy with Embrechts et al. [41]'s extremal coupling models, providing microfoundations for multiscale crisis indicators [45].

Practically, these results enable the real-time monitoring of systemic risk through dual-entropy tracking: rising $\mathcal{H}(\tau)$ signals tail uncertainty proliferation, while decaying $\mathcal{D}_{KL}$ indicates reducing distinguishability between normal and bubble regimes. For regulators, the explicit phase transition thresholds permit the data-driven calibration of circuit-breakers based on α-estimates rather than ad hoc volatility limits [49]. Asset managers could leverage the TWE-KL divergence ratio to optimize tail hedge ratios, as $\partial\mathcal{H}/\partial\mathcal{D}_{KL}$ quantifies the marginal information cost of tail risk mitigation.

3.1. Optimal Threshold Selection

Threshold selection in extreme value analysis presents a dual challenge: maintaining fidelity to the generalized Pareto model while ensuring estimation stability. Traditional approaches emphasize moment-based criteria [50] or likelihood maximization [51], but these overlook the information-theoretic structure of tail dynamics. The framework proposed here addresses this by unifying distributional fidelity (via KL divergence) with entropy stability—an advancement in financial applications, where threshold sensitivity directly affects crisis prevention [13].

Proposition 1. *For price processes $\{P_t\}_{t=1}^n$ with empirical tail density $\hat{f}_\tau$ and GPD approximation $f_{\mathcal{GP}}(\cdot;\hat{\xi},\hat{\sigma})$, the entropy-optimal threshold τ^* solves*

$$\tau^* = \underset{\tau\in\mathcal{T}}{\operatorname{argmin}} \left[\underbrace{\mathcal{D}_{KL}(\hat{f}_\tau; f_{\mathcal{GP}})}_{\text{Model Fidelity}} + \lambda \underbrace{\mathbb{V}ar(\mathcal{H}(\tau))}_{\text{Estimation Stability}} \right], \tag{6}$$

where $\mathcal{T} = [\tau_{\min}, \tau_{\max}]$ is the admissible range and $\lambda > 0$ regulates variance penalization. Under a Lipschitz continuity of $\mathcal{H}(\tau)$, the solution exists and is unique.

Proof. Let $L(\tau) = \mathbb{E}[\mathcal{H}(\tau) - \mathcal{H}_{\mathcal{GP}}(\tau)]^2 + \lambda\mathbb{V}\mathrm{ar}(\mathcal{H}(\tau))$. Assuming $\mathcal{H}(\tau)$ is Fréchet differentiable over $\mathcal{T}$, the first variation is

$$\delta L(\tau; h) = 2\langle \mathcal{H} - \mathcal{H}_{\mathcal{GP}}, \delta\mathcal{H}\rangle_{L^2} + 2\lambda\langle \mathbb{C}\mathrm{ov}(\mathcal{H}, \delta\mathcal{H})\rangle.$$

From Lemma 1, $\delta\mathcal{H} = (\xi/(\sigma + \xi\tau) + \partial_\tau\varepsilon)h$. The Euler–Lagrange equation becomes

$$(\mathcal{H} - \mathcal{H}_{\mathcal{GP}})\left(\frac{\xi}{\sigma + \xi\tau} + \partial_\tau\varepsilon\right) + \lambda\mathbb{C}\mathrm{ov}\left(\mathcal{H}, \frac{\xi}{\sigma + \xi\tau}\right) = 0.$$

Existence and uniqueness follow from $L(\tau)$ convexity in τ (verified via positive-definite Hessian), $\mathcal{T}$ compactness (ensuring attainment), and $\mathcal{H}(\tau)$ monotonicity in τ (Lemma 1). □

Example 2. *Revisit Example 1 with* $dP_t = \mu P_t dt + \sigma P_t dW_t$. *The KL term becomes*

$$\mathcal{D}_{KL} \approx \frac{1}{2}\left(\frac{\mu}{\sigma^2}\log\tau - 1\right)^2,$$

while the entropy variance scales as $\mathbb{V}ar(\mathcal{H}(\tau)) \sim \sigma^4/(4\mu^2\log\tau)$. *Minimizing* $L(\tau)$ *yields optimal threshold*

$$\tau^* = \exp\left(\frac{\sigma^2}{\mu}\left(1 + \sqrt{1 + \lambda\sigma^2/\mu}\right)\right).$$

This explicitly links model parameters to threshold choice—heavier trading volumes ($\mu \uparrow$) *lower* τ^*, *while volatility* ($\sigma \uparrow$) *raises it. Compared to Zhang et al. [52]'s quantile-based rule* $\tau_Z = \exp(\sigma/\sqrt{\mu})$, *the proposed* τ^* *adapts to both drift and regulation intensity* λ.

Proposition 1 advances threshold selection by embedding information-theoretic stability directly into the optimization. The KL divergence enforces consistency with extreme value theory's POT framework [19], while the entropy variance term mitigates overfitting to local tail fluctuations—a key challenge in Thompson [53]'s MLE approach. Example 2 demonstrates how market microstructure parameters (μ, σ) directly govern threshold choice, enabling the scenario-specific calibration absent in heuristic methods.

Practically, implementing Proposition 1 requires (1) estimating $\mathcal{H}(\tau)$ via nonparametric entropy estimators [54], (2) computing $\mathcal{D}_{KL}$ using robust GPD fits [52], and (3) tuning λ via cross-validation on crisis periods. This operationalizes Basel III's "sound estimation practices" requirement [55] by providing auditable thresholds grounded in information theory. For bubble detection, setting λ inversely to market volatility (higher λ in turbulent periods) stabilizes threshold choices during crises—an improvement over static rules.

Remark 1. *The variance term* $\mathbb{V}ar(\mathcal{H}(\tau))$ *encodes regulatory risk appetite:* $\lambda \to 0$ *prioritizes model fit (aggressive bubble signals), while* $\lambda \to \infty$ *emphasizes stability (reduced false positives). This implements the "precautionary principle" advocated in Ardakani [35].*

By unifying information-theoretic fidelity with estimation stability, this optimal threshold framework addresses longstanding dilemmas in extreme value analysis. The explicit linkage between entropy variance and KL divergence (Proposition 1) provides justification for adaptive thresholding, while Example 2 demonstrates practical calibration to market conditions. This advances beyond Drees [56]'s bias–variance trade-off by incorporating tail information dynamics and subsumes Thompson [53]'s MLE approach as the $\lambda \to 0$ limit. For financial applications, the protocol enables dynamic threshold adjustment aligned with both market microstructure and regulatory constraints—a requirement for crisis prevention.

3.2. Entropic Hypothesis Testing for Tail Regimes

Modern financial surveillance requires hypothesis tests that detect regime shifts in tail behavior with both statistical power and economic interpretability. Traditional moment-based approaches [13,40] lack sensitivity to early-stage bubbles due to their reliance on variance explosions rather than information structure changes. While subsampling methods [57] improve detection latency, they fail in the non-stationary environments common to financial crises [2]. Recent advances in information-theoretic risk measures [35,58] motivate entropy-based testing, but require asymptotic justification for their regulatory adoption.

This section establishes an entropy-based hypothesis testing framework with three key advances: (a) asymptotic normality of tail entropy estimators under non-explosive regimes, (b) explicit influence function calculations for sensitivity analysis, and (c) composite hypothesis testing via the continuous mapping theorem. The proposed test generalizes moment-based approaches [13] while avoiding subsampling artifacts [57]. I derive closed-

form variance expressions under exponential and Pareto innovations, enabling exact power calculations. Practical applications span Basel III stress testing, cryptocurrency surveillance [59], and housing market risk monitoring [60].

Theorem 2. *Let $\{P_t\}$ follow $P_t = \delta P_{t-1} + \epsilon_t$ with $\epsilon_t \sim F_\epsilon$. For threshold $\tau_n \to \infty$ satisfying $n\bar{F}_\epsilon(\tau_n) \to \infty$, define $H_0 : \delta \leq 1$ with light-tailed F_ϵ (exponential decay) and $H_1 : \delta > 1$ with regularly varying F_ϵ (heavy tails). Under H_0, the normalized entropy ratio*

$$ER_n = \sqrt{n\bar{F}_\epsilon(\tau_n)}\left(\frac{\mathcal{H}_n(\tau_n) - \mathcal{H}_0}{\sigma_{\mathcal{H}_0}}\right) \xrightarrow{d} \mathcal{N}(0,1),$$

where $\mathcal{H}_0 = 1 + \log\lambda$ for $\epsilon_t \sim \mathcal{E}(\lambda)$. Reject H_0 at level α if $ER_n > z_{1-\alpha}$.

Proof. Under H_0, Theorem 3.1 of Embrechts et al. [26] implies $\bar{F}_{P_t}(x) \sim \bar{F}_\epsilon(x)$. For exponential $\epsilon_t \sim \mathcal{E}(\lambda)$,

$$\mathcal{H}_0(\tau) = -\int_\tau^\infty \frac{\lambda e^{-\lambda x}}{e^{-\lambda\tau}} \log\left(\frac{\lambda e^{-\lambda x}}{e^{-\lambda\tau}}\right) dx = 1 + \log\lambda.$$

The empirical entropy $\mathcal{H}_n(\tau_n)$ uses a kernel density estimator $\hat{f}_n$ with bandwidth $h_n = o(\bar{F}_\epsilon(\tau_n)^{-1/5})$. Define the entropy functional

$$\Phi(F) = -\int_{\tau_n}^\infty \frac{f(x)}{\bar{F}(\tau_n)} \log\left(\frac{f(x)}{\bar{F}(\tau_n)}\right) dx.$$

According to Van der Vaart [61]'s functional delta method, the influence function is

$$\phi_F(X) = \frac{\mathbb{I}_{\{X>\tau_n\}}}{\bar{F}(\tau_n)}\left(1 + \log\frac{f(X)}{\bar{F}(\tau_n)}\right) - \mathcal{H}_0.$$

Under H_0, $\sqrt{n\bar{F}_\epsilon(\tau_n)}(\mathcal{H}_n(\tau_n) - \mathcal{H}_0)$ converges to a normal distribution with variance

$$\sigma^2_{\mathcal{H}_0} = \mathbb{E}\Big[\phi_F(X)^2 | X > \tau_n\Big] = \mathbb{E}\left[\left(1 + \log\frac{f(X)}{\bar{F}(\tau_n)}\right)^2 | X > \tau_n\right] - \mathcal{H}_0^2.$$

For exponential tails,

$$\sigma^2_{\mathcal{H}_0} = \mathbb{E}\Big[(1 - \lambda(X - \tau_n))^2 | X > \tau_n\Big] - 1 = \mathbb{E}[Z^2] - 1 = 1,$$

where $Z \sim \mathcal{E}(1)$. Slutsky's theorem yields the asymptotic normality. □

Example 3. *For $H_1 : \delta > 1$ with $\epsilon_t \sim \mathcal{P}(\alpha)$, Theorem 1 obtains*

$$\mathcal{H}_n(\tau_n) \approx \log\tau_n + 1 + \frac{1}{\alpha} + o_P(1).$$

The test statistic becomes

$$ER_n \approx \sqrt{n\tau_n^{-\alpha}}\left(\log\tau_n + 1 + \frac{1}{\alpha} - (1 + \log\lambda)\right) \xrightarrow{p} \infty,$$

yielding consistent detection since $\tau_n^{-\alpha} \to 0$ slower than $\log\tau_n \to \infty$.

Theorem 2 addresses three limitations of existing bubble tests. The entropy ratio ER_n diverges under H_1 (Example 3) unlike variance-based statistics. Asymptotic normality under H_0 enables exact critical values. Rejection thresholds map to entropy units (nats/bits).

Remark 2. *The diverging threshold $\tau_n \to \infty$ must satisfy*

$$\tau_n = o\left(n^{1/\alpha}\right) \cap \omega(\log n)$$

to ensure consistency. Implement via the algorithm developed in Drees [56].

The entropic test supports (1) Basel III compliance by mapping ER_n exceedances to capital surcharges, (2) the detection of stablecoin collapses via hourly ER_n monitoring, and (3) the identification of regional bubbles through price entropy divergence. This enables regulators to phase out ad hoc volatility circuit breakers in favor of entropy-based triggers aligned with systemic risk accumulation [2]. By establishing entropic hypothesis testing on asymptotic foundations, this section provides regulators with a tool for real-time bubble detection. The test's dual interpretability—statistical significance and information-theoretic units—bridges the gap between econometric theory and prudent policy.

4. Information Dynamics

Recent advances in econophysics have revealed profound connections between information entropy and market instability, building on foundational work in nonequilibrium thermodynamics [62,63]. Traditional market efficiency theories [64] establish price discovery mechanisms but lack tools to quantify information transfer during extreme events. The emergence of stochastic thermodynamics [65] provided entropy production as a key measure of system irreversibility, later adapted to financial markets through thermodynamic analogies [66]. Despite progress, gaps remain in modeling non-stationary bubble regimes. Current approaches either rely on heuristic thresholds [67] or lack statistical rigor [68]. This section addresses these limitations through exact entropy production laws with time-varying thresholds and measure-theoretic phase transition detection with Benjamini–Hochberg error control.

4.1. Entropic Information Flow

I first establish exact entropy production laws for non-stationary financial processes and quantify information dissipation through KL divergence-regularized Brownian dynamics. A result also provides closed-form solutions for geometric Brownian motion benchmark cases. These results generalize the thermodynamic inequalities to adaptive threshold regimes while encompassing the entropy evolution laws.

Theorem 3. *For Itô process $dP_t = \mu_t dt + \sigma_t dW_t$ with $\tau_t = \tau_0 + \int_0^t \kappa_s ds + \int_0^t \nu_s dW_s$, the entropy production rate satisfies*

$$\frac{d\mathcal{H}_t}{dt} = \underbrace{-\gamma \mathcal{D}_{KL}(f_t; f_{t-\Delta})}_{\textit{Information Dissipation}} + \underbrace{\sigma_{\mathcal{H}} \dot{W}_t^{\mathcal{H}}}_{\textit{Exogenous Shock}} + \underbrace{\mathbb{E}\left[\frac{\partial \mathcal{H}}{\partial \tau} \dot{\tau}_t\right]}_{\textit{Threshold Adaptation}}, \tag{7}$$

where $\gamma = \mathbb{E}[\mu_t^2/\sigma_t^2 | P > \tau_t]$ is the Kyle-Lambda [69] and $\sigma_{\mathcal{H}}^2 = \mathbb{E}[(\partial \mathcal{H}/\partial \tau)^2 \nu_t^2]$.

Proof. From the threshold-adapted Fokker–Planck equation,

$$\frac{\partial f_t}{\partial t} = -\partial_p(\mu_t f_t) + \frac{1}{2}\partial_p^2(\sigma_t^2 f_t) - \dot{\tau}_t \delta(p - \tau_t) f_t.$$

Compute entropy production via Nelson's derivative [70],

$$\frac{d\mathcal{H}_t}{dt} = -\int_{\tau_t}^{\infty} \left[\frac{\partial_t f_t}{\bar{F}_t} - \frac{f_t \partial_t \bar{F}_t}{\bar{F}_t^2}\right] \left(1 + \log \frac{f_t}{\bar{F}_t}\right) dp.$$

Substituting $\partial_t \bar{F}_t = -f_t(\tau_t)\dot{\tau}_t + \int_{\tau_t}^{\infty} \partial_t f_t dp$ and applying Itô's lemma to τ_t-dynamics yields the three-component decomposition. The KL divergence emerges from the relative entropy rate, while threshold adaptation follows from the Leibniz rule. □

Example 4. *Consider high-frequency trading (HFT) strategies with* $\mu_t = \alpha sign(P_t - \tau_t)$ *and* $\sigma_t = \beta|P_t - \tau_t|^{1/2}$. *Theorem 3 shows the following:*

$$\frac{d\mathcal{H}_t}{dt} = -\frac{\alpha^2}{\beta^2}\mathbb{E}\left[\frac{1}{|P-\tau_t|}\middle| P > \tau_t\right] + \beta\sqrt{\mathbb{E}\left[\frac{1}{|P-\tau_t|}\right]}\dot{W}_t^{\mathcal{H}}.$$

Monitoring $d\mathcal{H}_t/dt$ *detects manipulative order flows through entropy dissipation anomalies [71].*

Theorem 3 establishes entropy production as a fundamental observable for analyzing market stability, synthesizing insights from stochastic thermodynamics [62,72] and information-based asset pricing [69,73]. The tripartite decomposition of entropy dynamics—information dissipation through KL divergence, exogenous shocks via Brownian drivers, and adaptive threshold effects—provides a thermodynamic lens for quantifying financial fragility. By linking the Kyle–Lambda (γ) to entropy dissipation rates, I formalize how informational efficiency governs market disorder during crises, addressing the microstructure indeterminacy in Fama [64]'s efficiency paradigm. Example 4 demonstrates this framework's practical power: the entropy production signature of high-frequency strategies reveals manipulative order flows through anomalous dissipation rates $\mathbb{E}[|P-\tau_t|^{-1}]$, empirically validating Cont [71]'s latency arbitrage models. For regulators, monitoring $d\mathcal{H}_t/dt$ enables the real-time detection of predatory trading patterns that elude variance-based surveillance [8,40], while adaptive thresholds τ_t implement dynamic risk measures consistent with recent robustness frameworks [74]. These results shift crisis prevention from static volatility monitoring to entropy flow regulation—a paradigm where markets stabilize by controlling information dissipation rather than merely constraining price swings. Future extensions could integrate this entropy production metric into central bank stress-testing frameworks, creating early warning systems that preempt systemic collapses triggered by informational cascades [75].

4.2. Phase Transition Detection

Detecting transitions in financial markets presents a fundamental challenge in systemic risk analysis, necessitating methods that are sensitive to nonlinear regime shifts. Catastrophe theory, first introduced to financial markets by Zeeman [76], offers mathematical foundations for modeling discontinuous transitions, while modern entropy methods provide mechanisms to quantify market disorder.

The recent literature developed heuristic bubble indicators using machine learning [67], but their black-box nature limits their regulatory adoption. Concurrently, the literature proposes physics-inspired early warning systems [68], although these lack statistical error control. The results in this section bridge these approaches through entropy-based phase transition theory, extending the thermodynamic framework to adaptive financial regimes. This section (1) develops an entropy-based phase transition indicator with exact critical values, (2) establishes links between catastrophe theory and multiple testing procedures, and (3) provides the measure-theoretic foundation for bubble collapse detection.

Theorem 4. *Let* $\{P_t\}$ *follow an explosive process with tail entropy* $\mathcal{H}_t(\tau)$ *and market efficiency parameter* γ *from Theorem 3. Define the Entropic Bubble Index (EBI) as*

$$EBI_t = \frac{\mathcal{H}_t(\tau)}{\mathcal{H}_t^{eq}}, \quad \mathcal{H}_t^{eq} = \log\sigma_t + (1-\xi_t)\psi\left(\frac{1-\xi_t}{\xi_t}\right) + \frac{\xi_t}{1-\xi_t}, \tag{8}$$

where ψ is the digamma function. A systemic phase transition occurs when

$$\exists \tau : \frac{dEBI_t}{dt}\Big|_{\tau^-} < -\eta \text{ and } \frac{d^2\mathcal{H}_t}{dt^2}\Big|_{\tau^+} > \nu. \tag{9}$$

Critical values (η, ν) satisfy

$$\eta = \sqrt{2\log(1/\alpha_{BH})}, \quad \nu = \frac{\pi^2}{6\gamma^2}, \tag{10}$$

where α_{BH} is the Benjamini–Hochberg false discovery rate and γ is the Kyle–Lambda from Theorem 3.

Proof. Construct the entropy potential $V(\mathcal{H}) = \frac{\gamma^2}{2}(\delta - 1)\mathcal{H}^2 + \frac{\pi^2}{24}\mathcal{H}^4$, where δ is the autoregressive parameter from Theorem 1. Critical points satisfy

$$\frac{\partial V}{\partial \mathcal{H}} = \gamma^2(\delta - 1)\mathcal{H} + \frac{\pi^2}{6}\mathcal{H}^3 = 0 \Rightarrow \mathcal{H}^* = \sqrt{\frac{6\gamma^2(1-\delta)}{\pi^2}}.$$

The Hessian condition for saddle-node bifurcation

$$\det D^2 V = \gamma^2(\delta - 1) + \frac{\pi^2}{2}\mathcal{H}^2 = 0 \Rightarrow \delta^* = 1 - \frac{\pi^2\mathcal{H}^2}{2\gamma^2}.$$

Taylor expands EBI_t near $\mathcal{H}^*$,

$$\frac{dEBI_t}{dt} \approx -\sqrt{\frac{2\pi^2}{3\gamma^2}(1-\delta)} + O((\mathcal{H} - \mathcal{H}^*)^2).$$

Set $1 - \delta = \sqrt{2\log(1/\alpha_{BH})}$ for False Discovery Rate (FDR) control [77], yielding η. For convexity,

$$\frac{d^2\mathcal{H}_t}{dt^2} > \frac{\pi^2}{6\gamma^2} \Rightarrow \nu = \frac{\pi^2}{6\gamma^2}.$$

□

Theorem 4 establishes the first entropy-based phase transition detection framework with exact critical values, addressing the heuristic calibration issues in Zhao and Sornette [67]. The detection threshold $\eta = \sqrt{2\log(1/\alpha_{BH})}$ directly embeds Benjamini–Hochberg false discovery rate control [77], while the convexity requirement $\nu = \pi^2/(6\gamma^2)$ ties market efficiency (γ) to crisis predictability through the Kyle–Lambda [69].

This framework addresses three systemic challenges: (1) replacing arbitrary volatility triggers with entropy-based critical values, (2) formalizing Zeeman [76]' catastrophe theory through measure-theoretic bifurcation analysis, and (3) addressing Astill et al. [57]'s critique via α_{BH}-controlled detection latency. Practically, Federal Reserve systemic risk dashboards can integrate EBI thresholds for real-time monitoring. Furthermore, circuit breaker calibration transitions from fixed percentage rules to η-ν phase diagrams dynamically adjusted via market efficiency parameter γ.

In summary, this section establishes a framework for analyzing information dynamics in non-stationary financial markets using entropy methods. Theorem 3 quantifies entropy production as the sum of three components: information dissipation through KL divergence, exogenous shocks via Brownian drivers, and adaptive threshold effects. This decomposition generalizes stochastic thermodynamics to time-varying regimes while incorporating traditional market efficiency metrics. Theorem 4 further connects catastrophe theory and multiple testing procedures by deriving exact critical values for phase transition

detection with Benjamini–Hochberg error control. The Entropic Bubble Index transforms heuristic indicators into statistically robust early warning signals. This framework advances several research strands. First, it extends thermodynamic inequalities to non-equilibrium financial systems via Theorem 3. Second, it formalizes entropy-based bubble detection through the critical values of Theorem 4 and quantifies microstructure fragility through HFT entropy signatures. Lastly, it incorporates adaptive thresholds into risk management.

5. Application in Financial Markets

This empirical study implements the proposed framework across major asset classes. While the existing literature establishes entropy's theoretical utility in extreme value analysis [27,37], few studies operationalize these concepts for real-time market surveillance across heterogeneous instruments. This section addresses this gap by evaluating the TWE framework's capacity to detect speculative bubbles and systemic risks through (1) an adaptive threshold optimization that jointly minimizes KL divergence and entropy variance, (2) entropy ratio statistics with exact asymptotic distributions, and (3) phase transition detection via thermodynamic critical values.

By applying these tools to equity, commodity, and cryptocurrency markets, I quantify how tail entropy decay mechanisms govern crisis formation. The results demonstrate that entropy-driven signals provide statistically and economically significant early warnings compared to traditional variance-ratio tests, while robustness checks validate the framework's stability across threshold selection methods and finite-sample regimes. This empirical validation confirms the theoretical entropy scaling laws from Section 3 and provides practitioners with a replicable protocol for integrating information-theoretic risk measures into prudential oversight systems [35,78].

This empirical study uses daily data on the E-mini S&P 500, gold, crude oil (WTI), and Bitcoin futures over the period from 1 January 2015, to 24 March 2025, with $N = 1812$. The data are obtained as continuous front-month contracts, and Bitcoin spot prices serve as a proxy for cryptocurrency market behavior. All data are sourced from Bloomberg. Daily log returns are computed using the continuously compounded transformation to ensure time additivity and statistical tractability. Days with missing data across any series are excluded from the analysis to maintain a balanced panel.

Table 1 presents summary statistics for the daily log returns. For each asset, annualized mean returns are calculated by multiplying the average daily return by 252. Annualized standard deviations are derived from the daily standard deviation multiplied by the square root of 252. The annualized Sharpe ratio is defined as the ratio of the annualized mean to the annualized standard deviation. The S&P 500 exhibits the highest Sharpe ratio, indicating the most favorable risk-adjusted performance among the assets. Bitcoin and oil demonstrate the highest volatility, with extreme minimum and maximum returns. Bitcoin also shows strong negative skewness and excess kurtosis, which are consistent with its heavy-tailed return distribution. In contrast, gold returns are relatively stable, with lower volatility and skewness. These distributional features highlight the heterogeneity in risk–return profiles across traditional and alternative assets.

Table 1. Summary statistics of daily log returns (2015–2025).

	Mean (Annual)	SD (Annual)	Sharpe Ratio	Min	Max	Skewness	Kurtosis
S&P 500	0.14	0.17	0.83	−0.10	0.09	−0.19	11.58
Gold	0.11	0.15	0.75	−0.05	0.06	−0.11	3.48
Oil	0.31	0.45	0.70	−0.28	0.32	1.14	21.50
Bitcoin	0.30	0.62	0.48	−0.46	0.23	−1.02	13.88

Then, the daily log returns are preprocessed to remove volatility clustering and potential seasonality. First, a GARCH(1,1) model is applied to each return series individually. The model is specified as

$$r_t = \mu + \epsilon_t, \quad \epsilon_t = \sigma_t z_t, \quad \sigma_t^2 = \omega + \alpha \epsilon_{t-1}^2 + \beta \sigma_{t-1}^2,$$

where σ_t^2 captures conditional heteroskedasticity and $z_t \sim \mathcal{N}(0,1)$. The standardized residuals $z_t = \epsilon_t / \sigma_t$ are used from the model as filtered returns to ensure homoskedasticity. This is important for inference in models sensitive to variance dynamics. To account for potential intra-year seasonality in volatility and return magnitudes, the filtered returns are standardized using monthly z-score normalization $z_t^m = (z_t - \bar{z}_m)/s_m$, where $\bar{z}_m$ and s_m are the mean and standard deviation of returns within each month m. The resulting series is used in subsequent analyses to ensure comparability across assets and time.

5.1. Tail Entropy Dynamics

Table 2 reports the optimal thresholds τ^*, tail-weighted entropy estimates $\mathcal{H}(\tau^*)$, and generalized Pareto shape parameters $\hat{\xi}$ for each asset in the preprocessed return series. Entropy is computed over exceedances $X_t > \tau$, where τ^* is selected via the regularized optimization procedure described in Proposition 1. $\mathcal{H}(\tau)$ quantifies the disorder among tail events, and its decay rate reflects the informational dynamics governing extremes.

The results suggest distinct entropy behaviors across assets. Gold exhibits the highest positive entropy ($\mathcal{H}(\tau^*) = 17.93$), suggesting relatively uniform tail dispersion. In contrast, the S&P 500 and Bitcoin display strongly negative entropy values, indicating sharper tail concentration and greater predictability in extreme episodes. Despite their different volatility regimes, all assets exhibit negative shape parameters consistent with bounded tails and subexponential decay. These findings support the entropy decay implications of Theorem 1 and underscore the heterogeneity in tail informativeness across asset classes.

Table 2. Tail-weighted entropy estimates and shape parameters.

Asset	Optimal Threshold τ^*	Entropy $\mathcal{H}(\tau^*)$	Shape Index $\hat{\xi}$
S&P 500	1.61	−19.45	−0.35
Gold	1.52	17.93	−0.32
Oil	1.02	2.40	−0.23
Bitcoin	1.48	−12.53	−0.34

5.2. Bubble Detection Analysis

The detection process is formalized through Algorithm 1. This integrates (i) adaptive threshold optimization (Proposition 1), (ii) entropy ratio monitoring (Theorem 2), and (iii) phase transition criteria (Theorem 4) into a computationally replicable workflow. The algorithm decouples model specification (Step 1) from real-time monitoring (Step 2) and regime classification (Step 3).

The entropic detection framework was implemented on the four major financial assets. The three-stage methodology operationalizes the theoretical results. Table 3 reveals consistent early warnings from the entropic test versus the Phillips et al. [13] (PSY) benchmark. S&P 500 detected the COVID crash days earlier (19 February vs. 23 March 2020), with stronger signal ($ER = -3.21$ vs. -2.87). Oil's March 2020 collapse and gold's Ukraine crisis response were detected days faster. Bitcoin flagged the 2021 speculative peak days prior (20 October vs. 10 November), showing extreme ER (-3.89) from heavy-tailed entropy ($\mathcal{H} = -12.53$).

Algorithm 1 Entropic bubble detection framework

Input: Return series $\{r_t\}_{t=1}^T$, regulation intensity λ, FDR level α_{BH}
Output: Bubble signals $\{B_t\}_{t=1}^T$, Entropic Bubble Index $\{EBI_t\}_{t=1}^T$

1. **Threshold Optimization** (Proposition 1):
   1: For each asset, solve
$$\tau^* = \underset{\tau \in \mathcal{T}}{\operatorname{argmin}} \Big[\mathcal{D}_{KL}(\hat{f}_\tau; f_{\mathcal{GP}}(\cdot; \hat{\xi}, \hat{\sigma})) + \lambda \mathbb{V}\mathrm{ar}(\mathcal{H}(\tau)) \Big].$$
   2: Estimate $\hat{\xi}, \hat{\sigma}$ via Probability Weighted Moments [79].
2. **Entropy Monitoring** (Theorem 2):
   1: **for** each window $t \in [252, T]$ **do**
   2: Compute tail entropy
$$\mathcal{H}_t(\tau^*) = -\int_{\tau^*}^{\infty} \frac{\hat{f}_t(x)}{\bar{F}_t(\tau^*)} \log \frac{\hat{f}_t(x)}{\bar{F}_t(\tau^*)} dx.$$
   3: Calculate entropy ratio
$$ER_t = \sqrt{n} \frac{\mathcal{H}_t(\tau^*) - (1 + \log \hat{\lambda}_t)}{\hat{\sigma}_{\mathcal{H}_0}}.$$
   4: **end for**
3. **Phase Transition Detection** (Theorem 4):
   1: Compute EBI
$$EBI_t = \frac{\mathcal{H}_t(\tau^*)}{\log \hat{\sigma}_t + (1 - \hat{\xi}_t)\psi(\frac{1-\hat{\xi}_t}{\hat{\xi}_t}) + \frac{\hat{\xi}_t}{1-\hat{\xi}_t}}.$$
   2: Detect transitions where
$$\frac{dEBI_t}{dt}\Big|_{\tau^-} < -\sqrt{2\log(1/\alpha_{BH})} \quad \text{and} \quad \frac{d^2\mathcal{H}_t}{dt^2}\Big|_{\tau^+} > \frac{\pi^2}{6\hat{\gamma}_t^2}.$$
   3: Return $B_t = \mathbb{I}_{\{ER_t > \Phi^{-1}(1-\alpha_{BH})\}}$.

Table 3. Bubble detection performance: entropic test vs. Phillips et al. [13].

Asset	Detection Method	First Signal Date	ER Statistic	*p*-Value
S&P 500	Entropic Test	19 February 2020	−3.21	0.0013
	PSY Test	23 March 2020	−2.87	0.0021
Gold	Entropic Test	24 February 2022	−3.12	0.0009
	PSY Test	7 March 2022	−2.68	0.0037
Oil	Entropic Test	9 March 2020	−3.45	0.0006
	PSY Test	20 April 2020	−2.95	0.0016
Bitcoin	Entropic Test	20 October 2021	−3.89	0.0001
	PSY Test	10 November 2021	−3.01	0.0013

The superiority of the entropic framework arises from the entropy decay mechanism of Theorem 1. For Bitcoin ($\xi = -0.34$), we observe

$$\frac{d\mathcal{H}}{d\tau} = -0.34 \log \tau + o(1)$$

compared to PSY's variance-based detection. Bubble formation is demonstrated through dual criteria: EBI crossing $\eta = 2.33$ (99% quantile) and entropy convexity $\frac{d^2\mathcal{H}}{dt^2} > \nu = 1.47$.

During Bitcoin's 2021 bubble, EBI breached η on October 20 when prices reached USD 64,382, while the PSY test triggered only after the November 10 peak (USD 68,982). The entropic signal provided a 21-day lead time with rising convexity ($\frac{d^2\mathcal{H}}{dt^2} = 1.83 > \nu$), enabling preventive action before the 48.6% collapse to USD 35,104 by January 2022. The proposed framework's average 23-day early warning capability across assets (vs. 9 days for PSY) translates to critical risk management advantages. For a USD 10B portfolio, USD 1.4B with PSY. Basel III countercyclical buffers could be activated before momentum traps form. High-frequency surveillance systems can incorporate entropy decay rates for circuit-breakers.

This empirical study confirms that tail entropy dynamics provide an informationally efficient channel for bubble detection, overcoming the latency and distributional assumptions of moment-based approaches. The methodology's foundations in extreme value theory make it suited for modern heavy-tailed markets.

5.3. Robustness Checks

To validate the reliability of the entropic bubble detection framework, three critical robustness analyses were conducted: (i) alternative threshold selection methodologies, (ii) finite-sample properties through subsampling, and (iii) comparison with traditional moment-based approaches.

The regularized entropy threshold (Proposition 1) was compared against two established methods: MLE thresholding and the Hill estimator. MLE Thresholding maximizes the GPD likelihood [51]

$$\tau_{\text{MLE}} = \underset{\tau}{\text{argmax}} \sum_{x_i > \tau} \log f_{\mathcal{GP}}(x_i; \hat{\xi}, \hat{\sigma}),$$

while the Hill estimator focuses on tail index stability [80]

$$\tau_{\text{Hill}} = \underset{\tau}{\text{argmin}} \left| \hat{\xi}_k(\tau) - \hat{\xi}_{k+1}(\tau) \right|.$$

Table 4 compares the proposed entropy-regularized threshold selection against traditional methods across four key metrics. The Average Detection Lag measures the temporal delay in identifying bubble onsets, where the entropy method shows 34% faster detection than MLE approaches. This acceleration stems from Proposition 1's dual optimization of KL divergence and entropy variance, enabling the earlier recognition of tail distribution shifts. The Threshold Stability Ratio quantifies threshold consistency over time, calculated as $\mathbb{P}(|\tau_t - \tau_{t+1}| < 0.5\sigma)$. The 0.87 ratio versus 0.52 for the Hill estimator demonstrates superior robustness to market noise. This stability directly results from the entropy variance regularization term $\lambda \mathbb{V}\text{ar}(\mathcal{H}(\tau))$ in our loss function.

KL divergence values confirm the entropy method's distributional fidelity, with 33% lower divergence than MLE thresholds ($p < 0.05$, paired t-test). This aligns with Theorem 2's requirement for accurate GPD approximation in entropy ratio calculation. The Variance Ratio Test p-values further validate our thresholds' homoskedastic properties ($p = 0.43$), satisfying the i.i.d. assumptions underlying extreme value theory.

Table 4. Threshold selection method comparison. Stability ratio measures $\mathbb{P}(|\tau_t - \tau_{t+1}| < 0.5\sigma)$.

Metric	Entropy-Regularized	MLE	Hill
Average Detection Lag (days)	23.1	34.7	41.2
Threshold Stability Ratio	0.87	0.65	0.52
KL Divergence	0.12	0.18	0.27
Variance Ratio Test (p-value)	0.43	0.12	0.08

The proposed method demonstrates superior performance in threshold stability ($p < 0.01$, Wilcoxon test) while preserving distributional fidelity, especially for heavy-tailed assets like Bitcoin, where Hill estimator thresholds differ by 2.1σ compared to 0.7σ.

Through block bootstrap subsampling (B = 10,000 replications), we can assess

$$\mathbb{P}(\widehat{ER}_n < -z_\alpha | H_0) \leq \alpha + \frac{C}{\sqrt{n}}, \tag{11}$$

where C depends on the tail entropy properties from Theorem 2. The results show that (a) test size remains below 1.5α for $n \geq 100$, (b) power exceeds 90% for $n \geq 200$ across assets, and (c) entropic test dominates PSY in subsample agreement ($\kappa = 0.82$ vs. 0.63).

To evaluate operational effectiveness in crisis prevention, Table 5 benchmarks the proposed framework against classical approaches: the Variance Ratio Test [40], which detects heteroskedasticity bursts, and Tail Index Monitoring [80], which tracks changes in the shape parameter. The Average Lead Time advantage of 23.1 days over PSY's 8.9 days translates into critical risk management windows. For a USD 10B portfolio, this represents USD 2.1B compared to \$0.7B in preventable losses based on the dynamics of the 2021 Bitcoin crash. Crisis Recall rates reflect each method's ability to detect true bubble events. The 92% recall rate versus 71% for PSY results from the sensitivity of entropy to tail information decay (Theorem 1), in contrast to variance-based methods that overlook early log-scale entropy changes. The differential in False Positive Rates (2.1% vs. 6.2% for PSY) stems from the method's dual thresholding: statistical significance via $ER_n < -z_\alpha$ and economic significance through phase transition criteria (η, ν). This differs from moment-based methods that tend to overfit volatility clusters instead of identifying genuine bubble formations.

Table 5. Performance comparison across bubble detection methodologies.

Method	Avg. Lead Time (Days)	Crisis Recall	False Positive Rate
Entropic Test	23.1	92%	2.1%
Variance Ratio	9.3	68%	4.8%
Tail Index	14.7	79%	3.9%
PSY Test	8.9	71%	6.2%

The proposed framework achieves a $2.5\times$ higher early warning efficiency ($p < 0.001$, McNemar test) while maintaining low false positives. The entropy advantage stems from simultaneously capturing distributional shape (through $\mathcal{H}(\tau)$ dynamics), information decay (via ξ-dependence), and regime transitions (using KL divergence). These robustness checks confirm the proposed method's superiority in modern markets characterized by heavy tails and volatile extremes [41]. The entropy framework's mathematical coherence (Lemma 1) and adaptive thresholding provide unique advantages over both moment-based and alternative threshold selection approaches.

6. Discussion and Implications

For financial practitioners, the results offer three valuable insights. First, the entropy decay law provides a real-time diagnostic for tail risk accumulation, enabling portfolio managers to adjust hedging strategies before volatility spikes. The early warning window during Bitcoin's 2021 crash could prevent significant losses on a portfolio. Second, the entropic phase transition criteria allow central banks to activate countercyclical buffers preemptively, addressing Basel III's mandate for dynamic risk management. Third, the convex optimization framework for threshold selection addresses backtesting overfitting in expected shortfall models, enhancing Basel III compliance.

The findings challenge three conventions in financial econometrics. First, variance-based bubble tests are suboptimal in heavy-tailed markets—entropy's sensitivity to distributional shape accounts for its higher crisis recall. Second, regulatory stress tests can incorporate entropy production metrics to quantify information dissipation during crises, advancing beyond static scenario analysis. Third, high-frequency trading surveillance can utilize the entropy–Brownian noise decomposition to distinguish speculative exuberance from fundamental shocks. The TWE measure bridges two recent research subjects: information-theoretic risk measures [1] and non-stationary EVT [41]. While Zhang et al. [58] develop coherent entropy measures, the GPD concentration bounds developed in this paper address extremal compatibility. The entropy decay rate extends Phillips et al. [13]'s moment explosion criteria to information dynamics, formalizing Billio et al. [30]'s intuition about entropy as a crisis precursor.

Methodologically, the phase transition detection enhances Zhao and Sornette [67]'s machine learning approach by providing exact critical values. The threshold optimization framework incorporates Zhang et al. [52]'s data-driven method through entropy regularization, reducing overfitting in quantile estimation. These contributions align with four emerging trends: (1) thermodynamic finance [66], where the entropy production laws quantify market irreversibility; (2) tail risk network analysis [81], where TWE can measure systemic information contagion; (3) crypto-market surveillance [59], where Bitcoin's entropy dynamics reveal speculative cycles; and (4) climate–finance integration [34], where the proposed GPD–entropy linkage enables compound risk assessment. Future research can extend TWE to multivariate extremes, exploring broader applications in statistical modeling and potentially unifying EVT and deep learning for high-dimensional risk management.

7. Conclusions

This paper integrates extreme value theory with information theory by introducing tail-weighted entropy, a measure that quantifies the information contents in extremal regions while remaining compatible with peaks-over-threshold frameworks. Theoretically, I establish concentration bounds for TWE under generalized Pareto assumptions and derive explicit entropy decay laws related to tail indices through the digamma function. These results address the incompatibility between Shannon entropy and peaks-over-threshold frameworks while generalizing recent entropy measures to financial extremes. Empirically, I demonstrate that TWE provides superior early detection of financial bubbles compared to moment-based approaches. The entropy ratio test identifies crises earlier than the moment-based methods. The phase transition framework highlights thresholds that formalize catastrophe theory in financial markets, offering regulators an alternative to heuristic indicators. By optimizing thresholds using KL divergence and entropy variance regularization, the proposed method achieves faster detection than MLE-based approaches while preserving distributional fidelity.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request.

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

References

1. Ardakani, O.M. Capturing information in extreme events. *Econ. Lett.* **2023**, *231*, 111301. [CrossRef]
2. Ardakani, O.M.; Ajina, R. Tail risks in household finance. *Financ. Res. Lett.* **2024**, *69*, 106065. [CrossRef]

3. Fisher, R.A.; Tippett, L.H.C. Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Math. Proc. Camb. Philos. Soc.* **1928**, *24*, 180–190. [CrossRef]
4. Nystrom, P.C.; Soofi, E.S. Rare, outlier and extreme: Beyond the Gaussian model and measures. *Int. J. Complex. Leadersh. Manag.* **2012**, *2*, 6–38. [CrossRef]
5. Soofi, E.S. A generalizable formulation of conditional logit with diagnostics. *J. Am. Stat. Assoc.* **1992**, *87*, 812–816. [CrossRef]
6. Blyth, M.; Matthijs, M. Black swans, lame ducks, and the mystery of IPE's missing macroeconomy. *Rev. Int. Political Econ.* **2017**, *24*, 203–231. [CrossRef]
7. Orlowsky, B.; Seneviratne, S.I. Global changes in extreme events: Regional and seasonal dimension. *Clim. Change* **2012**, *110*, 669–696. [CrossRef]
8. Ardakani, O.M. Information content of inflation expectations: A copula-based model. *Stud. Nonlinear Dyn. Econom.* **2025**, *29*, 71–93. [CrossRef]
9. Goodell, J.W. COVID-19 and finance: Agendas for future research. *Financ. Res. Lett.* **2020**, *35*, 101512. [CrossRef] [PubMed]
10. Cover, T.M. *Elements of Information Theory*; John Wiley & Sons: Hoboken, NJ, USA, 1999.
11. Maasoumi, E. A compendium to information theory in economics and econometrics. *Econom. Rev.* **1993**, *12*, 137–181. [CrossRef]
12. Phillips, P.C.; Yu, J. Dating the timeline of financial bubbles during the subprime crisis. *Quant. Econ.* **2011**, *2*, 455–491. [CrossRef]
13. Phillips, P.C.; Shi, S.; Yu, J. Testing for multiple bubbles: Limit theory of real-time detectors. *Int. Econ. Rev.* **2015**, *56*, 1079–1134. [CrossRef]
14. Phillips, P.C.; Shi, S. Real time monitoring of asset markets: Bubbles and crises. In *Handbook of Statistics*; Elsevier: Amsterdam, The Netherlands, 2020; Volume 42, pp. 61–80.
15. Coles, S.; Bawa, J.; Trenner, L.; Dorazio, P. *An Introduction to Statistical Modeling of Extreme Values*; Springer: Berlin/Heidelberg, Germany, 2001; Volume 208.
16. Cooley, D. Extreme value analysis and the study of climate change: A commentary on Wigley 1988. *Clim. Change* **2009**, *97*, 77–83. [CrossRef]
17. Gnedenko, B. Sur la distribution limite du terme maximum d'une serie aleatoire. *Ann. Math.* **1943**, *44*, 423–453. [CrossRef]
18. Balkema, A.A.; De Haan, L. Residual life time at great age. *Ann. Probab.* **1974**, *2*, 792–804. [CrossRef]
19. Pickands, J., III. Statistical inference using extreme order statistics. *Ann. Stat.* **1975**, *3*, 119–131. [CrossRef]
20. Leadbetter, M.R.; Lindgren, G.; Rootzén, H. Maxima and minima and extremal theory for dependent processes. In *Extremes and Related Properties of Random Sequences and Processes*; Springer: New York, NY, USA, 1983; pp. 205–215.
21. Shannon, C.E. A mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [CrossRef]
22. Kullback, S.; Leibler, R.A. On information and sufficiency. *Ann. Math. Stat.* **1951**, *22*, 79–86. [CrossRef]
23. Jaynes, E.T. Information theory and statistical mechanics. *Phys. Rev.* **1957**, *106*, 620. [CrossRef]
24. Csiszár, I.; Körner, J. *Information Theory: Coding Theorems for Discrete Memoryless Systems*; Cambridge University Press: Cambridge, UK, 2011.
25. Jaynes, E.T. On the rationale of maximum-entropy methods. *Proc. IEEE* **1982**, *70*, 939–952. [CrossRef]
26. Embrechts, P.; Klüppelberg, C.; Mikosch, T. *Modelling Extremal Events: For Insurance and Finance*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2013; Volume 33.
27. Chen, L.; Singh, V.P. Entropy-based derivation of generalized distributions for hydrometeorological frequency analysis. *J. Hydrol.* **2018**, *557*, 699–712. [CrossRef]
28. Faranda, D.; Bourdin, S.; Ginesta, M.; Krouma, M.; Messori, G.; Noyelle, R.; Pons, F.; Yiou, P. A climate-change attribution retrospective of some impactful weather extremes of 2021. *Weather Clim. Dyn. Discuss.* **2022**, *2022*, 1–37. [CrossRef]
29. Ragone, F.; Bouchet, F. Rare event algorithm study of extreme warm summers and heatwaves over Europe. *Geophys. Res. Lett.* **2021**, *48*, e2020GL091197. [CrossRef]
30. Billio, M.; Casarin, R.; Costola, M.; Pasqualini, A. An entropy-based early warning indicator for systemic risk. *J. Int. Financ. Mark. Inst. Money* **2016**, *45*, 42–59. [CrossRef]
31. Simsekli, U.; Sagun, L.; Gurbuzbalaban, M. A tail-index analysis of stochastic gradient noise in deep neural networks. In Proceedings of the International Conference on Machine Learning. PMLR, Long Beach, CA, USA, 9–15 June 2019; pp. 5827–5837.
32. Ibragimov, M.; Ibragimov, R.; Walden, J. *Heavy-Tailed Distributions and Robustness in Economics and Finance*; Springer: Berlin/Heidelberg, Germany, 2015; Volume 214.
33. Adrian, T.; Boyarchenko, N. Liquidity policies and systemic risk. *J. Financ. Intermediat.* **2018**, *35*, 45–60. [CrossRef]
34. Sakthivel, K.; Nandhini, V. An entropy-based validation of threshold selection Technique for Extreme Value Analysis and Risk Assessment. *Lobachevskii J. Math.* **2024**, *45*, 1633–1651. [CrossRef]
35. Ardakani, O.M. Coherent measure of portfolio risk. *Financ. Res. Lett.* **2023**, *57*, 104222. [CrossRef]
36. Ardakani, O.M. Portfolio optimization with transfer entropy constraints. *Int. Rev. Financ. Anal.* **2024**, *96*, 103644. [CrossRef]
37. Rockafellar, R.T.; Uryasev, S. Optimization of conditional value-at-risk. *J. Risk* **2000**, *2*, 21–42. [CrossRef]

38. Bhaduri, S.N. Applying approximate entropy (ApEn) to speculative bubble in the stock market. *J. Emerg. Mark. Financ.* **2014**, *13*, 43–68. [CrossRef]
39. Diks, C.; Panchenko, V. Rank-based entropy tests for serial independence. *Stud. Nonlinear Dyn. Econom.* **2008**, *12*, 1–19. [CrossRef]
40. Homm, U.; Breitung, J. Testing for speculative bubbles in stock markets: A comparison of alternative methods. *J. Financ. Econom.* **2012**, *10*, 198–231. [CrossRef]
41. Embrechts, P.; Klüppelberg, C.; Mikosch, T. Modern extreme value theory at the interface of risk management, Bayesian networks and heavy-tailed time series. In *Mathematics Going Forward: Collected Mathematical Brushstrokes*; Springer: Berlin/Heidelberg, Germany, 2022; pp. 115–139.
42. Chavez-Demoulin, V.; Davison, A.C. Generalized additive modelling of sample extremes. *J. R. Stat. Soc. Ser. C Appl. Stat.* **2005**, *54*, 207–222. [CrossRef]
43. Dimpfl, T.; Peter, F.J. Using transfer entropy to measure information flows between financial markets. *Stud. Nonlinear Dyn. Econom.* **2013**, *17*, 85–102. [CrossRef]
44. Ren, W.X.; Sun, Z.S. Structural damage identification by using wavelet entropy. *Eng. Struct.* **2008**, *30*, 2840–2849. [CrossRef]
45. Costa, M.; Goldberger, A.L.; Peng, C.K. Multiscale entropy analysis of biological signals. *Phys. Rev. E—Stat. Nonlinear Soft Matter Phys.* **2005**, *71*, 021906. [CrossRef]
46. Goldie, C.M. Implicit renewal theory and tails of solutions of random equations. *Ann. Appl. Probab.* **1991**, *1*, 126–166. [CrossRef]
47. Balcilar, M.; Ozdemir, H.; Agan, B. Effects of COVID-19 on cryptocurrency and emerging market connectedness: Empirical evidence from quantile, frequency, and lasso networks. *Phys. A Stat. Mech. Its Appl.* **2022**, *604*, 127885. [CrossRef]
48. Hida, T.; Hida, T. *Brownian Motion*; Springer: Berlin/Heidelberg, Germany, 1980.
49. Ardakani, O.M. Bayesian extreme value models: Asymptotic behavior, hierarchical convergence, and predictive robustness. *Econom. Stat.* **2024**, 1–21. [CrossRef]
50. Davison, A.C.; Smith, R.L. Models for exceedances over high thresholds. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **1990**, *52*, 393–425. [CrossRef]
51. Smith, R.L. Estimating tails of probability distributions. *Ann. Stat.* **1987**, *15*, 1174–1207. [CrossRef]
52. Zhang, Y.; Wang, H.J.; Zhu, Z. Single-index thresholding in quantile regression. *J. Am. Stat. Assoc.* **2022**, *117*, 2222–2237. [CrossRef]
53. Thompson, W.C. Shifting decision thresholds can undermine the probative value and legal utility of forensic pattern-matching evidence. *Proc. Natl. Acad. Sci. USA* **2023**, *120*, e2301844120. [CrossRef] [PubMed]
54. Gadjiev, B.; Mohammad-Djafari, A.; Bercher, J.F.o.; Bessière, P. Generalized Maximum Entropy Principle, Superstatistics and Problem of Networks Classification. In Proceedings of the AIP Conference Proceedings-American Institute of Physics, Belfast, Ireland, 19–22 July 2011; Volume 1305, p. 407.
55. Embrechts, P.; Wüthrich, M.V. Recent challenges in actuarial science. *Annu. Rev. Stat. Its Appl.* **2022**, *9*, 119–140. [CrossRef]
56. Drees, H. Extreme quantile estimation for dependent data, with applications to finance. *Bernoulli* **2003**, *9*, 617–657. [CrossRef]
57. Astill, S.; Harvey, D.I.; Leybourne, S.J.; Sollis, R.; Robert Taylor, A. Real-time monitoring for explosive financial bubbles. *J. Time Ser. Anal.* **2018**, *39*, 863–891. [CrossRef]
58. Zhang, J.; Liu, T.; Tao, D. Going deeper, generalizing better: An information-theoretic view for deep learning. *IEEE Trans. Neural Netw. Learn. Syst.* **2024**, *35*, 16683–16695. [CrossRef] [PubMed]
59. Fry, J.; Ibiloye, O. Towards a taxonomy for crypto assets. *Cogent Econ. Financ.* **2023**, *11*, 2207266. [CrossRef]
60. Case, K.E.; Shiller, R.J. Is there a bubble in the housing market? *Brook. Pap. Econ. Act.* **2003**, *2003*, 299–362. [CrossRef]
61. Van der Vaart, A.W. *Asymptotic Statistics*; Cambridge University Press: Cambridge, UK, 2000; Volume 3.
62. Seifert, U. Stochastic thermodynamics: Principles and perspectives. *Eur. Phys. J. B* **2008**, *64*, 423–431. [CrossRef]
63. Seifert, U. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.* **2012**, *75*, 126001. [CrossRef] [PubMed]
64. Fama, E.F. Efficient capital markets. *J. Financ.* **1970**, *25*, 383–417. [CrossRef]
65. Prokopenko, M.; Boschetti, F.; Ryan, A.J. An information-theoretic primer on complexity, self-organization, and emergence. *Complexity* **2009**, *15*, 11–28. [CrossRef]
66. Aguilar, J.P.; Korbel, J.; Pesci, N. On the quantitative properties of some market models involving fractional derivatives. *Mathematics* **2021**, *9*, 3198. [CrossRef]
67. Zhao, D.; Sornette, D. Bubbles for Fama from Sornette. Available online: https://ssrn.com/abstract=3995526 (accessed on 20 July 2025).
68. Sornette, D.; Wei, R. Multiple Outlier Detection in Samples with Exponential & Pareto Tails; Swiss Finance Institute Research Paper No. 21-94; 2024, pp. 1–34. Available online: https://ssrn.com/abstract=4957557 (accessed on 20 July 2025).
69. Kyle, A.S.; Todorov, K. *The Cumulant Risk Premium*; Bank for International Settlements, Monetary and Economic Department: Basel, Switzerland, 2023.
70. Nelson, N.J.; Strauss, E. Dynamical chaos in a simple model of a knuckleball. *Appl. Math. Comput.* **2021**, *391*, 125661. [CrossRef]

71. Cont, R. Empirical properties of asset returns: Stylized facts and statistical issues. *Quant. Financ.* **2001**, *1*, 223. [CrossRef]
72. Peliti, L.; Pigolotti, S. *Stochastic Thermodynamics: An Introduction*; Princeton University Press: Princeton, NJ, USA, 2021.
73. Ardakani, O.M. Option pricing with maximum entropy densities: The inclusion of higher-order moments. *J. Futures Mark.* **2022**, *42*, 1821–1836. [CrossRef]
74. Embrechts, P.; Schied, A.; Wang, R. Robustness in the optimization of risk measures. *Oper. Res.* **2022**, *70*, 95–110. [CrossRef]
75. Bookstaber, R. Agent-based models for financial crises. *Annu. Rev. Financ. Econ.* **2017**, *9*, 85–100. [CrossRef]
76. Zeeman, E. Levels of structure in catastrophe theory illustrated by applications in the social and biological sciences. In Proceedings of the International Congress of Mathematicians, Vancouver, BC, Canada, 21–29 August 1974; Volume 2, pp. 533–548.
77. Benjamini, Y.; Heller, R.; Yekutieli, D. Selective inference in complex research. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **2009**, *367*, 4255–4271. [CrossRef] [PubMed]
78. Basten, C. Higher bank capital requirements and mortgage pricing: Evidence from the counter-cyclical capital buffer. *Rev. Financ.* **2020**, *24*, 453–495.
79. Hosking, J.R.; Wallis, J.R. Parameter and quantile estimation for the generalized Pareto distribution. *Technometrics* **1987**, *29*, 339–349. [CrossRef]
80. Hill, B.M. A simple general approach to inference about the tail of a distribution. *Ann. Stat.* **1975**, *3*, 1163–1174. [CrossRef]
81. Barigozzi, M.; Cho, H.; Owens, D. FNETS: Factor-adjusted network estimation and forecasting for high-dimensional time series. *J. Bus. Econ. Stat.* **2024**, *42*, 890–902. [CrossRef]