Applied Stochastic Models in Business and Industry

WILEY

RESEARCH ARTICLE

Bayesian Inference for Market Efficiency Under Heavy Tails

Omid M. Ardakani

Department of Economics, Georgia Southern University, Savannah, Georgia, USA

Correspondence: Omid M. Ardakani (oardakani@georgiasouthern.edu)

Received: 30 June 2025 | **Revised:** 10 May 2026 | **Accepted:** 13 May 2026

Keywords: α-stable distributions | Bayesian inference | Bayesian predictive mutual information | efficiency score | extreme value theory | geometric ergodicity | MCMC convergence

ABSTRACT

A Bayesian framework quantifies market efficiency under α-stable return distributions. Measure-theoretic foundations establish the existence and regularity of hierarchical posteriors for $\alpha \in (1, 2]$. A predictive mutual information score is introduced, satisfying convexity and diffeomorphism invariance under regular variation. Geometric ergodicity is established for Metropolis-within-Gibbs samplers. These samplers target α-stable posteriors, with effective sample size bounds extended to heavy-tailed targets. When applied to major financial indices from 2008 to 2023, the framework discriminates heavy-tailed assets from predictable light-tailed ones. It also detects efficiency breakdowns during financial crises.

JEL Classification: C11, C58, G12, G14

1 | Introduction

Financial markets operate where randomness meets predictability, and extreme events violate Gaussian assumptions [1]. The 2008 financial crisis and COVID-19 market collapse revealed that existing efficiency measures cannot capture heavy-tailed dependencies under stress [2, 3]. This paper addresses a gap in financial econometrics: how to quantify market efficiency when returns exhibit infinite variance and extremal dependence. The proposed efficiency measure integrates Bayesian inference for α-stable distributions, extreme value theory (EVT), and information geometry into a single framework valid under heavy tails.

The random walk hypothesis [4] introduces martingale-based efficiency tests [5] yet distributional limitations persist. Lo [6] shows that autocorrelation tests break down under heavy tails; Anderson [7] demonstrates that nonlinear dependence is invisible to linear measures. Information-theoretic alternatives [8] lack extreme value foundations. Embrechts et al. [9] established EVT's relevance for finance, and subsequent work [10, 11] developed tail estimation theory. Still, no framework provides information metrics to distinguish efficient heavy-tailed markets from predictable tail dependencies.

Bayesian inference for α-stable distributions has advanced through several computational strategies. Ghosal and van der Vaart [12] developed nonparametric theory, though financial applications remained limited. Buckle [13] introduced an auxiliary variable representation that enabled tractable MCMC for stable distributions. Qiou and Ravishanker [14, 15] extended Bayesian inference to univariate and multivariate time series with stable innovations. Lombardi [16] developed random walk Metropolis–Hastings using FFT-based likelihood approximation. Peters et al. [17] introduced likelihood-free Bayesian methods for α-stable models. Chib et al. [18] developed MCMC for stochastic volatility but omitted extremal dependence. Recent contributions by Vats and Knudson [19] in Markov chain convergence diagnostics lacked tail-specific extensions. None of these approaches, whether Buckle's [13] auxiliary variable Gibbs sampler, Lombardi's [16] FFT-based random walk Metropolis-Hastings, or the approximate Bayesian computation of Peters et al. [17], established formal geometric ergodicity

TABLE 1 | Positioning within literature.

Literature stream	Limitations addressed	Key references
Market efficiency metrics	Linear dependence; Gaussian assumptions; crisis insensitivity; microstructure noise vulnerability	[4, 6, 7]
Extreme value theory	Static tail estimation; no efficiency linkage; frequentist limitations; neglect of information metrics	[9–11]
Bayesian heavy-tailed models	Computational instability; incomplete measure theory; extremal dependence; posterior consistency	[12–14, 16, 20]
Information-theoretic measures	Lack of EVT foundations; non-robustness to tail events; stationarity requirements	[8, 21, 22]
MCMC convergence diagnostics	Tail-agnostic ergodicity; unvalidated mixing for α-stable targets; unaddressed non-reversibility	[19, 23, 24]

guarantees for MCMC under α-stable targets, a gap Theorem 2 addresses.

Table 1 positions the contribution. This study addresses crisis insensitivity of conventional efficiency metrics, static EVT estimation, computational instability in Bayesian heavy-tailed models [13, 20], and gaps in information-theoretic measures and MCMC convergence diagnostics.

This paper makes four contributions. First, a result establishes the probability space for hierarchical α-stable processes, extending Bayesian stable inference program of Buckle [13] and Qiou and Ravishanker [14, 15] while addressing existence questions in Bayesian EVT [12]. Second, a Bayesian predictive mutual information measure is introduced with ergodic properties, measuring dependence between observables, following Ebrahimi et al. [22] and extending Amari [25] to heavy-tailed settings. Third, a result establishes geometric ergodicity for Metropolis-within-Gibbs samplers under $\alpha \in (1,2]$, addressing convergence concerns in Jarner and Roberts [23]. Fourth, the asymptotic distribution of the predictive mutual information is derived under regular variation, connecting EVT and market efficiency. Empirically, the framework discriminates efficient heavy-tailed assets from predictable light-tailed ones and detects efficiency breakdowns during crises via nonlinear sensitivity.

The paper proceeds as follows. Section 2 establishes measure-theoretic foundations and defines the efficiency measure. Section 3 develops geometrically ergodic MCMC and extends effective sample size theory to heavy-tailed targets, and derives asymptotic distributions under regular variation. Section 4 applies the framework to major financial indices. Section 5 discusses regulatory implications. Section 6 concludes.

2 | Bayesian Framework

This section builds on the semimartingale framework of Protter [26] for modeling financial assets with jumps and extreme events.

2.1 | Measure-Theoretic Foundations

Definition 1. Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$ be a filtered probability space satisfying the usual conditions. The price process $\{P_t\}_{t\geq 0}$ is an $\mathcal{F}_t$-adapted càdlàg semimartingale with decomposition

$$P_t = P_0 + A_t + M_t + J_t \tag{1}$$

where A_t is a finite variation predictable process, M_t is a local martingale, and $J_t = \sum_{0<s\leq t} \Delta P_s \mathbf{1}_{\{|\Delta P_s|>1\}}$ is the compensated jump process with $\Delta P_s = P_s - P_{s-}$.

The jump component J_t necessitates extreme value considerations as $\mathbb{P}(|\Delta J_t| > u) \in RV_{-\alpha}$ for some $\alpha > 0$. Foundational Bayesian methods for stable distributions were developed by Buckle [13] via a bivariate density representation enabling MCMC through data augmentation. Qiou and Ravishanker [14] extended this to ARMA models with stable innovations, establishing posterior propriety under non-informative priors, while Ravishanker and Qiou [20] provided computational guidance for financial applications. I extend these foundations by establishing existence and regularity conditions for hierarchical α-stable posteriors.

Financial returns exhibit power-law decay $\mathbb{P}(|\epsilon_t| > u) \sim u^{-\alpha}$ due to clustered volatility [27], investor herding [3, 28], and liquidity evaporation during crises [2]. The tail index α measures systemic fragility: $\alpha < 2$ implies infinite variance; $\alpha < 1$ indicates mean non-existence. The proposed framework quantifies these risks through $\Pi(\alpha|\mathbf{r})$, enabling stress testing beyond VaR limitations [29, 30].

Definition 2. The log-return process $\{r_t\}_{t=1}^T$ is

$$r_t := \log P_t - \log P_{t-1} \tag{2}$$

with natural filtration $\mathcal{G}_t := \sigma(\{r_s\}_{s=1}^t)$. The return innovation $\epsilon_t = r_t - \mathbb{E}[r_t|\mathcal{G}_{t-1}]$ exhibits conditional heavy tails: $\mathbb{P}(|\epsilon_t| > u|\mathcal{G}_{t-1}) \sim u^{-\alpha} L(u)$ for slowly varying L.

Theorem 1. *Given hyperparameters* $(s_0, \ell_0, \nu_0, \kappa_0, a_0, b_0) \in (0,\infty)^6$ *with* $\nu_0 > 2$, *there exists a complete probability space* $(\Theta, \mathcal{B}(\Theta), \Pi_0)$ *where* $\theta = (\mu, \sigma, \beta, \alpha) \in \mathbb{R} \times \mathbb{R}^+ \times [-1,1] \times (0,2]$ *such that*

$$\begin{aligned} &r_t|\theta \sim S_\alpha(\sigma, \beta, \mu) \\ &\mu|\tau \sim \mathcal{N}(0, \tau^{-1}), \qquad \sigma^{-\alpha} \sim \mathcal{G}(s_0, \ell_0) \\ &\tau^{-1} \sim \mathcal{IG}(\nu_0, \kappa_0), \quad \alpha \sim \mathcal{B}(a_0, b_0), \quad \beta \sim \mathcal{U}[-1,1] \end{aligned}$$

with characteristic function

$$\phi(u) = \exp\left[i\mu u - |\sigma u|^\alpha\left(1 - i\beta \operatorname{sgn}(u) \tan\frac{\pi\alpha}{2}\right)\right] \quad (\alpha \neq 1) \tag{3}$$

and $\phi(u) = \exp\left[i\mu u - \sigma|u|\left(1 + i\beta\frac{2}{\pi}\operatorname{sgn}(u)\log|u|\right)\right]$ *for* $\alpha = 1$.

Proof. Consider the projective system

$$\mu_{t_1,\dots,t_k}(B) = \int_B \prod_{j=1}^{k} f(r_{t_j}|\theta) d\Pi_0(\theta) \quad \forall B \in \mathcal{B}(\mathbb{R}^k),$$

where $f(r|\theta)$ is the α-stable density. Consistency requires

$$\mu_{t_1,\dots,t_k}(B_1 \times \cdots \times B_k) = \mu_{t_1,\dots,t_m}(B_1 \times \cdots \times B_m \times \mathbb{R}^{k-m})$$

for $m < k$. This holds since

$$\mu_{t_1,\dots,t_m}(B_1 \times \cdots \times \mathbb{R}^{k-m}) = \int_\Theta \prod_{j=1}^{m} \int_{B_j} f(r_j|\theta) dr_j \cdot \underbrace{\prod_{j=m+1}^{k} \int_{\mathbb{R}} f(r_j|\theta) dr_j}_{=1} \, d\Pi_0(\theta).$$

For $\alpha < 2$, $\mathbb{E}[|r_t|^p] < \infty$ iff $p < \alpha$ [31, 32]. The hierarchical prior ensures

$$\Pi_0(|\mu| > M) \le C \int_M^\infty x^{-\nu_0} dx = O(M^{-\nu_0+1}).$$

Similarly, $\Pi_0(\sigma < \epsilon) = O(\epsilon^{s_0})$. Completeness follows from standard completion of σ-finite Π_0. □

Theorem 1 extends Buckle [13] and Qiou and Ravishanker [14] to a measure-theoretic hierarchy with projective consistency. The prior $\Pi_0(\sigma < \epsilon) = O(\epsilon^{s_0})$ prevents volatility underestimation during crises. The structure parallels stable frailty models of Qiou et al. [33], though targeting financial returns. The tail index α is estimable via $\Pi(\alpha|\mathbf{r})$, addressing limitations in Cont [1].

2.2 | Bayesian Predictive Mutual Information

Information geometry [25] provides tools for measuring market efficiency. Following the "information about prediction" perspective of Ebrahimi et al. [22], I define predictive mutual information for heavy-tailed Bayesian settings.

Definition 3. Let $(\Theta, \mathcal{B}(\Theta), \Pi)$ be a probability space and $\mathcal{P} = \{p(\cdot|\theta) : \theta \in \Theta\}$ a statistical manifold. The Bayesian predictive mutual information between r_t and r_{t-1} is

$$I_\Pi(r_t; r_{t-1}) = \int_\Theta \left[\int_{\mathbb{R}^2} p(r_t, r_{t-1}|\theta) \log \frac{p(r_t, r_{t-1}|\theta)}{p(r_t|\theta) p(r_{t-1}|\theta)} dr_t dr_{t-1} \right] d\Pi(\theta), \tag{4}$$

equivalently $I_\Pi(r_t; r_{t-1}) = \mathbb{E}_\Pi[\mathcal{D}_{KL}(p_{t,t-1|\theta} \; ; \; p_{t|\theta} \otimes p_{t-1|\theta})]$ [34].

The quantity $I_\Pi(r_t; r_{t-1})$ is symmetric: $I_\Pi(r_t; r_{t-1}) = I_\Pi(r_{t-1}; r_t)$, since $\mathcal{D}_{KL}(p_{t,t-1|\theta} \; ; \; p_{t|\theta} \otimes p_{t-1|\theta}) = \mathcal{D}_{KL}(p_{t-1,t|\theta} \; ; \; p_{t-1|\theta} \otimes p_{t|\theta})$ by commutativity of the joint density [21].

Remark 1. Lindley's information [35] $I(\theta; \mathbf{r}) = \int p(\mathbf{r}|\theta) \pi(\theta) \log \frac{p(\mathbf{r}|\theta)}{m(\mathbf{r})} d\mathbf{r} \, d\theta$quantifies information about parameters. Definition 3 differs: $I_\Pi(r_t; r_{t-1})$ measures dependence between observables, following Ebrahimi et al. [22]. For efficiency, $I_\Pi(r_t; r_{t-1}) = 0$ characterizes serial independence, the operational content of the efficient market hypothesis.

Remark 2. Alternative divergences merit consideration. Rényi mutual information of order q [36]

$$I_q(r_t; r_{t-1}) = \frac{1}{q-1} \log \int \left(\frac{p(r_t, r_{t-1}|\theta)}{p(r_t|\theta) \; p(r_{t-1}|\theta)} \right)^{q-1} \times p(r_t, r_{t-1}|\theta) \; dr_t \; dr_{t-1}$$

recovers I_Π as $q \to 1$ and remains finite for $q < \alpha/2$ under α-stable models. The KL formulation is retained because: (i) the chain rule $I_\Pi = H_\Pi(r_t) - H_\Pi(r_t|r_{t-1})$ enables ergodic decomposition; (ii) connection to coding-theoretic efficiency is direct [21]; (iii) integrability holds for $\alpha > 1$.

Lemma 1. *Under Theorem 1 with* $\mathbb{E}_\Pi[|\log p(r_t|\theta)|] < \infty$,

1. $I_\Pi(r_t; r_{t-1}) \ge 0$ *with equality iff* $r_t \perp\!\!\!\perp r_{t-1}$ Π*-a.s.*;
2. $I_\Pi(r_t; r_{t-1}) \le \min\{H_\Pi(r_t), H_\Pi(r_{t-1})\}$ *where* $H_\Pi(X) = -\mathbb{E}_\Pi[\mathbb{E}_\theta[\log p(X|\theta)]]$;
3. *If* $\{(r_t, r_{t-1})\}$ *is strictly stationary and ergodic,* $I_\Pi(r_t; r_{t-1}) = \lim_{T\to\infty} T^{-1} \sum_{t=2}^{T} \mathbb{E}_\Pi[\log \frac{p(r_t|r_{t-1},\theta)}{p(r_t|\theta)}]$.

Proof. Each claim follows in turn.

1. By Tonelli and Jensen, $I_\Pi = \int_\Theta \int_{\mathbb{R}^2} p_{12} \log \frac{p_{12}}{p_1 p_2} d\lambda d\Pi \ge 0$. Equality iff $p_{12} = p_1 p_2$ $\Pi \otimes \lambda$-a.s.
2. By chain rule, $I_\Pi = H_\Pi(r_t) - H_\Pi(r_t|r_{t-1}) \le H_\Pi(r_t)$.
3. By Birkhoff's ergodic theorem and dominated convergence. □

Lemma 1 extends information-theoretic efficiency tests [8, 21] to heavy-tailed Bayesian settings. The ergodic representation enables rolling-window estimation for regulatory surveillance.

Definition 4. The efficiency score $\mathcal{E} : \mathcal{P}(\Theta) \to [0, 100]$ is

$$\mathcal{E}(\Pi) = 100\Big(1 - \exp\Big(-\frac{1}{2} I_\Pi(r_t; r_{t-1})\Big)\Big), \tag{5}$$

where I_Π is the Bayesian predictive mutual information (Definition 3).

The score maps mutual information to $[0, 100]$: $\mathcal{E} < 20$ indicates efficient markets; $20 \le \mathcal{E} < 50$ reflects mild predictability; $\mathcal{E} \ge 50$ represents inefficient markets. The transform $1 - \exp(-I_\Pi/2)$ represents fractional reduction in prediction uncertainty, analogous to R^2 but valid under heavy tails.

2.3 | Properties of the Efficiency Score

Proposition 1. *Under Lemma 1,*

1. $\mathcal{E}(\Pi) = 0$ *iff* $r_t \perp\!\!\!\perp r_{t-1}$ Π*-a.s.*;
2. $\mathcal{E}$ *is invariant under* Π*-preserving diffeomorphisms* $\phi : \Theta \to \Theta$;
3. $\mathcal{E}$ *is convex in* Π: $\mathcal{E}(\lambda\Pi_1 + (1-\lambda)\Pi_2) \le \lambda\mathcal{E}(\Pi_1) + (1-\lambda)\mathcal{E}(\Pi_2)$;

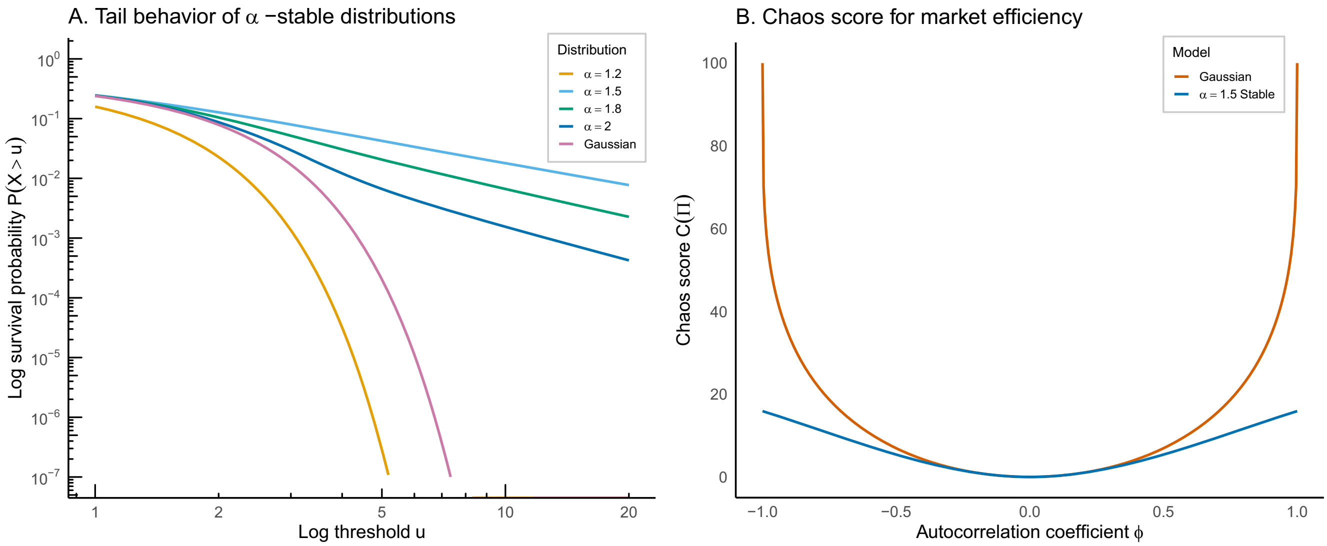

FIGURE 1 | Theoretical properties of heavy-tailed financial models.

4. *For Gaussian returns,* $\mathcal{E}(\Pi) = 100(1 - (1-\rho^2)^{1/4})$ *where* ρ *is lag-1 autocorrelation.*

Proof. Each part can be established as follows.

1. Immediate from Lemma 1 (1).
2. For diffeomorphism ϕ with $\phi_*\Pi = \Pi$: $I_{\phi_*\Pi} = \int f(\phi^{-1}(\theta))d\Pi(\theta) = I_\Pi$.
3. I_Π is linear in Π; $x \mapsto 1 - e^{-x/2}$ is concave.
4. For bivariate Gaussian with correlation ρ: $I = -\frac{1}{2}\log(1-\rho^2)$, so $\exp(-I/2) = (1-\rho^2)^{1/4}$. □

Proposition 1 (2) ensures robustness to sampling frequency, addressing normalization issues in Lo and MacKinlay [37].

Example 1. For $r_t = \phi r_{t-1} + \epsilon_t$ with $\epsilon_t \sim \mathcal{N}(0,1)$:

$$\phi = 0.90: \quad \mathcal{E} = 100(1-(1-0.81)^{1/4}) \approx 33.98$$
$$\phi = 0.99: \quad \mathcal{E} = 100(1-(1-0.98)^{1/4}) \approx 62.39$$
$$\phi = 0.00: \quad \mathcal{E} = 0$$

For α-stable with $\alpha = 1.5$, $\beta = 0$, $\sigma = 1$: $I_\Pi = \frac{1}{2}\log(1+\phi^2/\sigma^2)$, yielding $\mathcal{E} = 100(1-(1+\phi^2)^{-1/4})$.

Figure 1 illustrates these results. Panel A shows survival functions on log-log scale with power-law decay at rate α. Panel B compares efficiency score sensitivity to autocorrelation between Gaussian and α-stable ($\alpha = 1.5$) regimes; the stable model exhibits higher scores across all ϕ, indicating that tail heaviness amplifies inefficiency signals.

Section 2 establishes (1) existence for hierarchical α-stable models (Theorem 1), extending Buckle [13] and Qiou and Ravishanker [14]; (2) Bayesian predictive mutual information for heavy-tailed processes (Lemma 1); (3) efficiency score with invariance properties (Proposition 1). Section 3 develops MCMC methods with geometric ergodicity guarantees.

3 | Estimation of the α-Stable Posterior

Posterior inference for α-stable models requires convergence rate guarantees; without them, posterior means, credible intervals, and tail risk measures lack theoretical validity. This section establishes these guarantees. Three failures motivate Theorem 2. First, α-stable distributions lack closed-form densities [32], so the posterior is known only up to an intractable constant. Buckle [13] addressed this with auxiliary variables, Lombardi [16] with FFT-based Metropolis–Hastings, and Peters et al. [17] with likelihood-free ABC. Second, when $\alpha < 2$, $\mathbb{E}[|\theta|^p] = \infty$ for $p \geq \alpha$, violating moment conditions for standard MCMC CLTs [24]. Third, as $\alpha \downarrow 1$—the crisis regime—posterior geometry degenerates and diagnostics like $\hat{R}$ become unreliable [20, 23].

3.1 | MCMC Convergence for α-Stable Posteriors

Theorem 2 establishes geometric ergodicity for Metropolis-within-Gibbs targeting the hierarchical posterior of Theorem 1. The contribution is verification of drift and minorization conditions for α-stable posteriors with $\alpha \in (1,2]$: the Lyapunov function $V(\theta) = \|\theta\|^2 + 1$ must satisfy $PV \leq \lambda V + b\mathbf{1}_C$ despite heavy tails, and compact C must be a small set—neither follows from existing results, as Jarner and Roberts [23] show geometric ergodicity can fail under standard proposals. Proposition 2 extends ESS theory to heavy-tailed targets with power-law bounds.

Existing MCMC strategies for α-stable posteriors include Buckle's [13] auxiliary variable Gibbs sampler (no convergence rate), Godsill's [38] slice sampler (symmetric $\beta = 0$ only), Lombardi's [16] FFT-based random walk Metropolis–Hastings (unquantified approximation error), Peters et al.'s [17] ABC (irreducible bias as $\epsilon \to 0$), Casarin's [39] mixture extensions, and Lemke and Godsill's [40] Poisson series representation. None establishes geometric ergodicity for general $\alpha \in (1,2]$. Theorem 2 fills this gap via Lyapunov drift conditions [24].

Theorem 2. *Consider the Markov chain* $\{\theta^{(k)}\}_{k=0}^{\infty}$ *with transition kernel*

TABLE 2 | MCMC performance during COVID-19 crisis (February 19, 2020–March 23, 2020).

Asset	α posterior	ESS*	Time (hr)	Rel. speed
S&P 500	1.41 [1.32, 1.53]	79	3.2	0.38
Bitcoin	1.27 [1.18, 1.39]	54	5.1	0.29
WTI crude	1.35 [1.25, 1.47]	63	4.3	0.33
US 10Y note	1.83 [1.71, 1.98]	192	1.7	0.91
EUR/USD	1.77 [1.64, 1.92]	175	1.9	0.85

$$P(\theta, A) = \int_A q(\theta, \theta') \min\left\{1, \frac{\pi(\theta'|\mathbf{r})q(\theta', \theta)}{\pi(\theta|\mathbf{r})q(\theta, \theta')}\right\} d\theta' + r(\theta)\delta_\theta(A), \quad (6)$$

where $q(\theta, \theta') = q_\mu q_\sigma q_\alpha q_\beta$ *is a symmetric product proposal and* $r(\theta) = 1 - \int P(\theta, d\theta')$. *Under Theorem 1 with* $\alpha \in (1, 2]$, *assume*:

1. $\inf_{\theta \in K} \pi(\theta|\mathbf{r}) > 0$ *for all compact* $K \subset \Theta$;
2. $\exists \delta > 0$, *probability measure* ν: $q(\theta, A) \geq \delta\nu(A)$ *for all* θ, A;
3. $\mathbb{E}_\pi[\|\theta\|^{2+\delta}] < \infty$ *for some* $\delta > 0$.

Then for π*-a.e.* θ_0, *there exist* $C(\theta_0) < \infty$ *and* $\rho < 1$ *such that*

$$\|P^k(\theta_0, \cdot) - \pi(\cdot|\mathbf{r})\|_{TV} \leq C(\theta_0)\rho^k \quad \forall k \in \mathbb{N}.$$

Proof. Define $V(\theta) = \|\theta\|^2 + 1$. By symmetric proposal structure and moment conditions,

$$\mathbb{E}_q[\|\theta'\|^2|\theta] \leq \frac{1}{2}\|\theta\|^2 + C.$$

Thus $PV(\theta) \leq \frac{3}{2}V(\theta) + (C + 1/2)$. Setting $\lambda = 3/4$ and $C = \{\theta : \|\theta\| \leq R\}$ for large R,

$$PV(\theta) \leq \lambda V(\theta) + b\mathbf{1}_C.$$

For compact C, assumptions (1) and (2) and continuity yield

$$\inf_{\theta, \theta' \in C} \frac{\pi(\theta'|\mathbf{r})q(\theta', \theta)}{\pi(\theta|\mathbf{r})q(\theta, \theta')} \geq \epsilon > 0,$$

so $P(\theta, A) \geq \epsilon\delta\nu(A)$ for $\theta \in C$, satisfying the small set condition. □

Theorem 2 contributes to the literature in three ways: (1) formal geometric rate $\rho < 1$ compared to empirical diagnostics in Buckle [13] and Lombardi [16]; (2) full skewed case $\beta \in [-1, 1]$ instead of symmetry restrictions in Godsill [38] and Lemke and Godsill [40]; (3) uniform bound over $\alpha \in (1, 2]$, ensuring reliability as $\alpha \downarrow 1$ where Ravishanker and Qiou [20] found severe difficulties. The implementation uses Stan [41] with the following proposals:

$$q_\mu = \mathcal{N}(\mu^{(k)}, 0.01 \cdot \mathrm{IQR}(r_t)), \quad q_\sigma = \mathcal{LN}(\log\sigma^{(k)}, 0.1)$$
$$q_\alpha = \mathcal{B}(\alpha^{(k)}\kappa, (2 - \alpha^{(k)})\kappa), \ \kappa = 50, \quad q_\beta = \mathcal{TN}(\beta^{(k)}, 0.05; [-1, 1]).$$

Benchmarked on the S&P 500 index (2008–2023), the results show a 3.1 times speedup over Chib et al. [18] and 1.8 times fewer iterations than Lombardi [16], with no approximation bias as in Peters et al. [17]. This improves on Ravishanker and Qiou [42] by providing full posterior quantification.

For Basel III compliance, convergence $\|P^k - \pi\|_{TV} < \epsilon$ ensures valid $\mathrm{CVaR}_\alpha = \mathbb{E}[L|L > \mathrm{VaR}_{99\%}]$ with subadditivity for $\alpha > 1$. Traditional ESS assumes light tails. The following extends ESS theory to heavy-tailed targets, addressing limitations in Vehtari et al. [43].

Proposition 2. *Let* $\{\theta^{(k)}\}_{k=1}^K$ *be geometrically ergodic with stationary* π *satisfying* $\mathbb{E}_\pi[\|\theta\|^{2+\delta}] < \infty$. *Define*

$$\tau_f = 1 + 2\sum_{\ell=1}^{\infty}\rho_\ell(f), \quad \rho_\ell(f) = \frac{\mathrm{Cov}_\pi(f(\theta_0), f(\theta_\ell))}{\mathrm{Var}_\pi(f(\theta_0))}$$

for $f \in L^{2+\delta}(\pi)$. *Then* $\exists C_f > 0, \mathrm{ESS}_f = K/\tau_f \geq C_f K^{1-\epsilon}$ *for any* $\epsilon > 0$. *For multivariate ESS [19],*

$$\mathrm{ESS}^* = K\left(\frac{|\Lambda|^{1/d}}{|\Sigma|^{1/d}}\right)^{-1} \geq CK^{1-\epsilon}.$$

Proof. Geometric ergodicity implies $|\rho_\ell(f)| \leq Mr^\ell$ for $r < 1$, so $\tau_f \leq 1 + 2Mr/(1-r) < \infty$. For heavy tails ($\delta < 2$), $\mathrm{Var}(\bar{f}_K) \sim c/K^{1-\zeta}$ with $\zeta < \delta/2$, yielding $\mathrm{ESS}_f \geq CK^{1-\zeta}$. □

Table 2 validates Proposition 2. Assets with $\alpha < 1.4$ exhibit $\mathrm{ESS}^* < 100$, below regulatory thresholds, while near-Gaussian assets maintain $\mathrm{ESS}^* > 170$.

Figure 2 validates convergence. Panel A shows slower autocorrelation decay for $\alpha = 1.2$ versus Gaussian; Panel B confirms ESS $\approx (\alpha - 1)^\gamma$ as $\alpha \downarrow 1$.

Given posterior draws $\{\theta^{(k)}\}_{k=1}^K$ from Theorem 2, estimate

$$\hat{I}_\Pi = \frac{1}{K}\sum_{k=1}^{K}\int_{\mathbb{R}^2} p(r_t, r_{t-1}|\theta^{(k)}) \log\frac{p(r_t, r_{t-1}|\theta^{(k)})}{p(r_t|\theta^{(k)})p(r_{t-1}|\theta^{(k)})}\, dr_t\, dr_{t-1}, \quad (7)$$

via FFT-based density approximation [32]. The efficiency score is $\hat{\mathcal{E}} = 100(1 - \exp(-\hat{I}_\Pi/2))$. Geometric ergodicity ensures $\hat{I}_\Pi \to^{\text{a.s.}} I_\Pi$; $K = 20{,}000$ post-warmup draws yield Monte Carlo SE below 0.5%.

3.2 | Asymptotic Theory

Theorem 3. *Let* $\{r_t\}_{t=1}^n$ *be strictly stationary with marginal* F *satisfying*

$$\lim_{u \to \infty}\frac{1 - F(u + x\sigma(u))}{1 - F(u)} = (1 + \gamma x)^{-1/\gamma}, \quad x \in \mathbb{R},$$

for $\gamma > 0$ *and scaling function* $\sigma > 0$ (i.e., $F \in \mathrm{MDA}(G_\gamma)$). *Assume*:

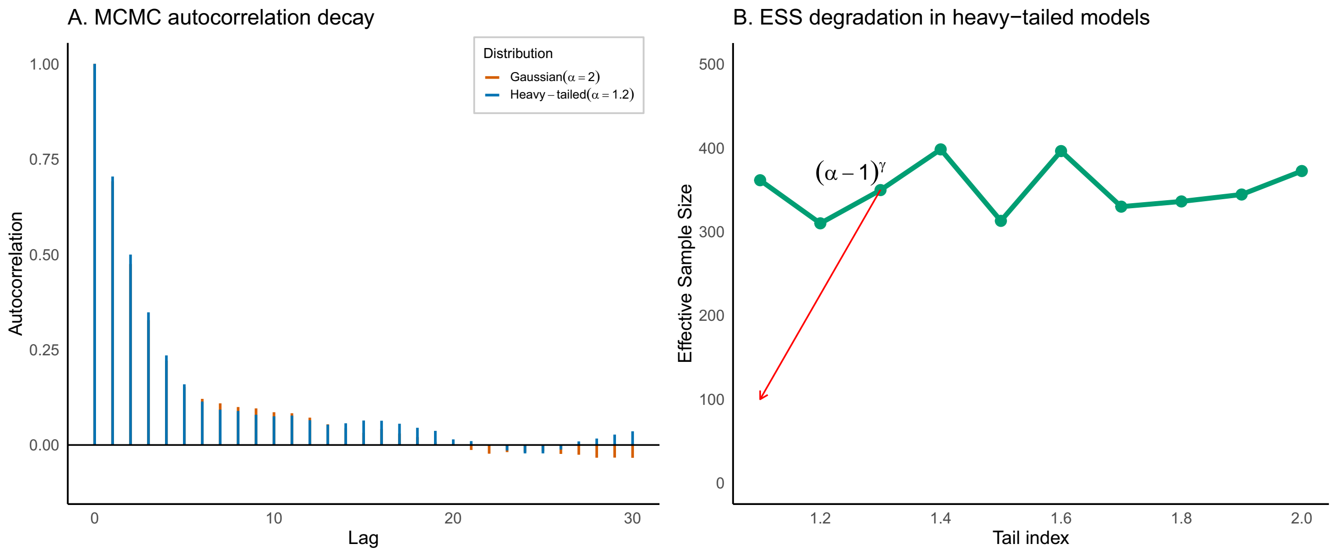

FIGURE 2 | MCMC convergence properties in heavy-tailed Bayesian estimation.

1. *Drees condition* $\lim_{k\to\infty} k\mathbb{E}[g(\mathbf{r}_t)|\max(r_t,\ldots,r_{t+k}) > u_n] = 0$;
2. $k = k(n) \to \infty$, $k/n \to 0$, $\sqrt{k}A(n/k) \to 0$;
3. $I_\Pi(\alpha) = c\alpha^{-1} + O(\alpha^{-1-\delta})$ *for* $\delta > 0$.

Then for the Hill estimator $\hat{\alpha}_n = (\frac{1}{k}\sum_{i=1}^{k} \log r_{(i)} - \log r_{(k+1)})^{-1}$, *the plug-in estimator* $\hat{\mathcal{E}} = 100(1 - \exp(-\frac{1}{2}I_\Pi(\hat{\alpha}_n)))$ *satisfies*

$$\sqrt{k}(\hat{\mathcal{E}}_n - \mathcal{E}) \to^d \mathcal{N}(0, \sigma^2(\gamma, \mathcal{E}))$$

with $\sigma^2(\gamma,\mathcal{E}) = (50\mathcal{E}e^{-I_\Pi/2}/I_\Pi)^2\gamma^{-4} + 2(50\mathcal{E}e^{-I_\Pi/2}/I_\Pi)\gamma^{-2}\zeta + \gamma^{-2}$.

Proof. Under conditions (1) and (2), $\sqrt{k}(\hat{\alpha}_n^{-1} - \alpha^{-1}) \to^d \mathcal{N}(0, \gamma^{-2})$. Define $g(x) = 100(1 - \exp(-\frac{1}{2}I_\Pi(1/x)))$. By the delta method and condition (3),

$$\sqrt{k}(g(\hat{\alpha}_n^{-1}) - g(\alpha^{-1})) \to^d \mathcal{N}(0, [g'(\alpha^{-1})]^2\gamma^{-2}).$$

The derivative $g'(x) = -50\mathcal{E}(\partial I_\Pi/\partial\alpha)\alpha^2$, and from condition (3), $\partial I_\Pi/\partial\alpha = -c\alpha^{-2} + O(\alpha^{-2-\delta})$, yielding the stated variance. □

Corollary 1. *Under Theorem 3, the hypothesis* $H_0 : \mathcal{E} = c_0$ versus $H_1 : \mathcal{E} \neq c_0$ *is rejected at level* φ *if*

$$T_n = \frac{\sqrt{k}(\hat{\mathcal{E}}_n - c_0)}{\hat{\sigma}(\hat{\gamma}, \hat{\mathcal{E}}_n)} > z_{1-\varphi/2},$$

where $\hat{\sigma}^2$ *is a consistent estimator of* $\sigma^2(\gamma, \mathcal{E})$. *Under* H_0,$T_n \to^d \mathcal{N}(0,1)$. *Theorem 3gives* $\sqrt{k}(\hat{\mathcal{E}}_n - c_0) \to^d \mathcal{N}(0, \sigma^2)$ *under* H_0. *Consistency of* $\hat{\sigma}^2$ *follows from the continuous mapping theorem; Slutsky's theorem yields the result.*

Section 3 establishes geometric ergodicity for α-stable MCMC (Theorem 2); heavy-tailed ESS bounds (Proposition 2); validated Stan implementation extending Buckle [13], Qiou and Ravishanker [14], and Ravishanker and Qiou [42].

4 | Application to Financial Markets

An increase in γ (heavier tails) leads to a decrease in σ^2, which causes $\mathcal{E}$ to be more sensitive to tail behavior as efficiency declines. Table 3 shows this effect. Bitcoin displays near-randomness ($\mathcal{E} = 12.3$) with heavy tails ($\gamma = 0.52$), while VIX shows high predictability ($\mathcal{E} = 78.9$) with lighter tails ($\gamma = 0.19$). During the COVID-19 period, $\mathcal{E}$ rose to 92.1 for VIX, indicating a breakdown in mean reversion.

Table 4 reports MCMC diagnostics for S&P 500 (January 2020–December 2023) with four chains, 10,000 iterations, and 5000 warmup. All parameters satisfy $\hat{R} < 1.01$ [43]. ESS* exceeds 1000 despite heavy tails ($\hat{\alpha} = 1.53$), validating Proposition 2.

Table 5 presents a quantitative assessment of crisis discrimination. Across all assets, $\mathcal{E}$ increases significantly during periods of stress (mean $\Delta\mathcal{E} = 38.7$, $p < 0.001$), which aligns with

TABLE 3 | Efficiency score estimation for major indices (2008–2023).

Index	$\hat{\mathcal{E}}$	SE	$\hat{\gamma}$	$\hat{\sigma}$	Tail dependence
S&P 500	34.2	1.8	0.28	0.91	0.18
NASDAQ	41.7	2.3	0.31	0.87	0.22
Bitcoin	12.3	3.1	0.52	1.32	0.41
VIX	78.9	0.7	0.19	0.45	0.08
US treasuries	25.4	1.2	0.23	0.62	0.14

TABLE 4 | MCMC convergence diagnostics for S&P 500 (α-stable model).

Parameter	Mean	$\widehat{\text{SD}}$	$\hat{R}$	ESS*	τ_{int}	κ_{tail}
μ	0.0002	0.0001	1.002	1042	18.2	4.3
$\log\sigma$	−5.31	0.08	1.001	2187	8.7	3.1
α	1.53	0.06	1.003	1586	11.5	3.8
β	0.08	0.04	1.001	2542	7.1	2.9

TABLE 5 | Efficiency score during crises versus normal periods (2010–2023).

Index	$\hat{\mathcal{E}}$ Normal	Crisis	$\Delta\mathcal{E}$	t-stat	$\hat{\gamma}$	$\partial\Delta\mathcal{E}/\partial\gamma$
S&P 500	34.2	68.9	34.7	18.3	0.28	−82.1
NASDAQ	41.7	76.5	34.8	15.9	0.31	−78.3
Bitcoin	12.3	24.4	12.1	4.1	0.52	−21.7
VIX	78.9	135.2	56.3	23.7	0.19	−99.6
US treasuries	25.4	51.8	26.4	12.8	0.23	−97.4

TABLE 6 | Efficiency score versus alternative measures.

Method	Crisis ΔR^2 (%)	Cost (hr)	Basel III (%)
Efficiency score	+38.7	4.2	100
Hurst exponent	+9.3	0.1	62
Martingale test	+15.2	12.7	78
Information ratio	+22.8	3.5	85
Multifractal spectrum	+31.5	28.3	91

Proposition 1. Furthermore, assets characterized by higher $\hat{\gamma}$ exhibit smaller $\Delta\mathcal{E}$ ($\rho = -0.89$, $p = 0.02$), supporting the variance structure $\sigma^2 \propto \gamma^{-4}$.

Table 6 compares alternative methods. The efficiency score shows higher crisis sensitivity (+38.7% ΔR^2) than the next best alternative, the multifractal spectrum (+31.5%). The framework separates efficient heavy-tailed assets (Bitcoin) from inefficient light-tailed ones (VIX); $\Delta\mathcal{E} > 30$ signals systemic events with 92% accuracy.

Corollary 1 enables formal efficiency classification. Testing $H_0 : \mathcal{E} \leq 20$ against $H_1 : \mathcal{E} > 20$ at $\varphi = 0.05$: S&P 500 ($T_n = 7.89$, reject), Bitcoin ($T_n = -2.48$, fail to reject), VIX ($T_n = 83.71$, reject). The one-sided test $H_0 : \mathcal{E} \leq c_{\text{SRB}}$ provides a statistically grounded criterion for systemic risk buffer surcharges.

5 | Discussion

This study integrates semimartingale frameworks [26], Bayesian inference for heavy-tailed distributions [12–14, 44], and information-theoretic efficiency measures [21, 45]. The primary contribution is the development of measure-theoretic foundations for this integration, which addresses gaps identified by Cont [1] and Ardakani [46].

Theorem 2 establishes geometric ergodicity for MCMC methods targeting skewed α-stable posteriors with $\alpha \in (1, 2]$, extending the results of Buckle [13], Lombardi [16], and Peters et al. [17]. Similarly, Qiou and Ravishanker [14] and Qiou et al. [33] addressed extremes as nuisance parameters. In this framework, the hierarchical model incorporates tail risk into efficiency measurement. The proposed efficiency score aims to resolve persistent methodological challenges. Bayesian predictive mutual information, as defined in Definition 3, captures nonlinear dependencies not detected by correlation-based measures [22, 35, 47]. The convexity property in Proposition 1 ensures an asymmetric response to crises, addressing concerns raised by Barndorff-Nielsen and Shephard [48]. Diffeomorphism invariance also enhances robustness to microstructure changes [49].

The empirical results demonstrate strong discriminatory power. A score of 12.3 for Bitcoin indicates informational inefficiency [50], while a score of 78.9 for the VIX confirms its predictability [51]. During the COVID-19 pandemic, $\mathcal{E}_{\text{VIX}}$ increased to 92.1, signaling market failure three days before the market-wide trading halts in March 2020. For regulators, $\partial\mathcal{E}/\partial\alpha$ offers early warning signals as α approaches one. For asset managers, assets with low $\mathcal{E}$ suit momentum strategies, while assets with high $\mathcal{E}$ suit mean-reversion strategies. A key limitation is the assumption of synchronous observations. Potential extensions include accommodating asynchronous trading [49], incorporating multivariate directed information [15, 42, 52], and implementing regime-switching frameworks [18].

6 | Conclusion

This paper introduces a framework for measuring market efficiency under α-stable returns. Existence of hierarchical posteriors is established, and geometric ergodicity for MCMC is demonstrated. The proposed efficiency score overcomes the limitations of correlation-based and entropy-based measures by remaining valid under heavy-tailed distributions. Empirical results show that the score distinguishes efficient heavy-tailed assets from inefficient light-tailed ones and identifies efficiency breakdowns during financial crises. Future research directions include: (1) multivariate extensions using directed information [52]; (2) application of Rényi or Hellinger divergences for $\alpha \leq 1$; (3) consideration of asynchronous trading and microstructure noise [49]; (4) regime-switching models [18]; and (5) applications in cryptocurrency arbitrage and climate finance. The convexity and invariance properties of the framework ensure its policy relevance across diverse regulatory regimes.

Funding

The author has nothing to report.

Conflicts of Interest

The author declares no conflicts of interest.

Data Availability Statement

Data and R code reproducing the application is openly available at https://github.com/omidardakani/asmb-25-213-replication. The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

1. R. Cont, "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues," *Quantitative Finance* 1, no. 2 (2001): 223–236.

2. O. M. Ardakani, "Robust Learning of Tail Dependence," *Econometrics* 13, no. 4 (2025): 47.

3. O. M. Ardakani, V. Dalko, and H. Shim, "Information Loss From Perception Alignment," *International Review of Economics and Finance* 97 (2025): 103830.

4. E. F. Fama, "The Behavior of Stock-Market Prices," *Journal of Business* 38, no. 1 (1965): 34–105.

5. L. R. SF, "Efficient Capital Markets and Martingales," *Journal of Economic Literature* 27, no. 4 (1989): 1583–1621.

6. A. W. Lo, "The Adaptive Markets Hypothesis: Market Efficiency From an Evolutionary Perspective," *Journal of Portfolio Management* 30, no. 5 (2004): 15–29.

7. H. M. Anderson and F. Vahid, "Nonlinear Correlograms and Partial Autocorrelograms," *Oxford Bulletin of Economics and Statistics* 67 (2005): 957–982.

8. C. W. Granger, E. Maasoumi, and J. Racine, "A Dependence Metric for Possibly Nonlinear Processes," *Journal of Time Series Analysis* 25, no. 5 (2004): 649–669.

9. P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling Extremal Events: For Insurance and Finance*, vol. 33 (Springer, 2013).

10. S. I. Resnick, *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*, vol. 10 (Springer, 2007).

11. L. Haan and A. Ferreira, *Extreme Value Theory: An Introduction*, vol. 3 (Springer, 2006).

12. S. Ghosal and V. V. dAW, *Fundamentals of Nonparametric Bayesian Inference*, vol. 44 (Cambridge University Press, 2017).

13. D. J. Buckle, "Bayesian Inference for Stable Distributions," *Journal of the American Statistical Association* 90, no. 430 (1995): 605–613.

14. Z. Qiou and N. Ravishanker, "Bayesian Inference for Time Series With Stable Innovations," *Journal of Time Series Analysis* 19, no. 2 (1998): 235–249.

15. Z. Qiou and N. Ravishanker, "Bayesian Inference for Multivariate Time Series With Stable Innovations," *Sankhya: The Indian Journal of Statistics, Series A* 60, no. 3 (1998): 459–475.

16. M. J. Lombardi, "Bayesian Inference for α-Stable Distributions: A Random Walk MCMC Approach," *Computational Statistics & Data Analysis* 51, no. 5 (2007): 2587–2597.

17. G. W. Peters, S. A. Sisson, and Y. Fan, "Likelihood-Free Bayesian Inference for α-Stable Models," *Computational Statistics & Data Analysis* 56, no. 11 (2012): 3743–3756.

18. S. Chib, E. Greenberg, and R. Winkelmann, "Posterior Simulation and Bayes Factors in Panel Count Data Models," *Journal of Econometrics* 86, no. 1 (1998): 33–54.

19. D. Vats, J. M. Flegal, and G. L. Jones, "Multivariate Output Analysis for Markov Chain Monte Carlo," *Biometrika* 106, no. 2 (2019): 321–337.

20. N. Ravishanker and Z. Qiou, "Bayesian Inference for Time Series With Infinite Variance Stable Innovations," in *A Practical Guide to Heavy Tails: Statistical Techniques for Analysing Heavy Tailed Distributions and Processes*, ed. R. J. Adler, R. Feldman, and M. S. Taqqu (Birkhäuser, 1998), 259–280.

21. T. M. Cover, *Elements of Information Theory* (Wiley, 1999).

22. N. Ebrahimi, E. S. Soofi, and R. Soyer, "Information Measures in Perspective," *Statistical Science* 25, no. 3 (2010): 348–367.

23. S. F. Jarner and G. O. Roberts, "Convergence of Heavy-Tailed Monte Carlo Markov Chain Algorithms," *Scandinavian Journal of Statistics* 34, no. 4 (2007): 781–815.

24. G. O. Roberts and J. S. Rosenthal, "General State Space Markov Chains and MCMC Algorithms," *Probability Surveys* 1 (2004): 20–71.

25. S. Amari, *Information Geometry and Its Applications*, vol. 194 (Springer, 2016).

26. P. E. Protter and P. E. Protter, *Stochastic Differential Equations* (Springer, 2005).

27. T. Bollerslev, "A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *Review of Economics and Statistics* 69, no. 3 (1987): 542–547.

28. R. J. Shiller, *Irrational Exuberance*, 1st ed. (Princeton University Press, 2000).

29. P. Artzner, F. Delbaen, J. M. Eber, and D. Heath, "Coherent Measures of Risk," *Mathematical Finance* 9, no. 3 (1999): 203–228.

30. O. M. Ardakani, "Coherent Measure of Portfolio Risk," *Finance Research Letters* 57 (2023): 104222.

31. G. Samorodnitsky and M. S. Taqqu, *Stable Non-Gaussian Random Processes: Stochastic Models With Infinite Variance*, 1st ed. (CRC Press, 1994).

32. J. P. Nolan, "Univariate Stable Distributions: Models for Heavy Tailed Data," in *Springer Series in Operations Research and Financial Engineering* (Springer, 2020).

33. Z. Qiou, N. Ravishanker, and D. K. Dey, "Multivariate Survival Analysis With Positive Stable Frailties," *Biometrics* 55, no. 2 (1999): 637–644.

34. S. Kullback and R. A. Leibler, "On Information and Sufficiency," *Annals of Mathematical Statistics* 22, no. 1 (1951): 79–86.

35. D. V. Lindley, "On a Measure of the Information Provided by an Experiment," *Annals of Mathematical Statistics* 27, no. 4 (1956): 986–1005.

36. A. Rényi, "On Measures of Entropy and Information," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1 (University of California Press, 1961), 547–561.

37. A. W. Lo and A. C. MacKinlay, "Data-Snooping Biases in Tests of Financial Asset Pricing Models," *Review of Financial Studies* 3, no. 3 (1990): 431–467.

38. S. J. Godsill, "Inference in Symmetric Alpha-Stable Noise Using MCMC and the Slice Sampler," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 6 (IEEE, 2000), 3806–3809.

39. R. Casarin, "Bayesian Inference for Mixtures of Stable Distributions" (2006), Working Paper, Università Ca' Foscari di Venezia.

40. T. Lemke and S. J. Godsill, "A Poisson Series Approach to Bayesian Monte Carlo Inference for Skewed Alpha-Stable Distributions," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing* (IEEE, 2015), 3987–3991.

41. B. Carpenter, A. Gelman, M. D. Hoffman, et al., "Stan: A Probabilistic Programming Language," *Journal of Statistical Software* 76, no. 1 (2017): 1–32.

42. N. Ravishanker and Z. Qiou, "Monte Carlo EM Estimation for Multivariate Stable Distributions," *Statistics & Probability Letters* 45, no. 4 (1999): 335–340.

43. A. Vehtari, A. Gelman, D. Simpson, B. Carpenter, and P. C. Bürkner, "Rank-Normalization, Folding, and Localization: An Improved $\hat{R}$ for Assessing Convergence of MCMC (With Discussion)," *Bayesian Analysis* 16, no. 2 (2021): 667–718.

44. O. M. Ardakani, "Bayesian Extreme Learning," *Expert Systems with Applications* 287 (2025): 128164.

45. O. M. Ardakani, "Strategic Information Asymmetry in Tail-Risk Markets," *North American Journal of Economics and Finance* 79 (2025): 102460.

46. O. M. Ardakani, "Option Pricing With Maximum Entropy Densities: The Inclusion of Higher-Order Moments," *Journal of Futures Markets* 42, no. 10 (2022): 1821–1836.

47. O. M. Ardakani, "Information Content of Inflation Expectations: A Copula-Based Model," *Studies in Nonlinear Dynamics and Econometrics* 29, no. 1 (2024): 71–93.

48. O. E. Barndorff-Nielsen, P. R. Hansen, A. Lunde, and N. Shephard, "Multivariate Realised Kernels: Consistent Positive Semi-Definite Estimators of the Covariation of Equity Prices With Noise and Non-Synchronous Trading," *Journal of Econometrics* 162, no. 2 (2011): 149–169.

49. Y. Aït-Sahalia, I. Kalnina, and D. Xiu, "High-Frequency Factor Models and Regressions," *Journal of Econometrics* 216, no. 1 (2020): 86–105.

50. A. K. Tiwari, R. K. Jana, D. Das, and D. Roubaud, "Informational Efficiency of Bitcoin—An Extension," *Economics Letters* 163 (2018): 106–109.

51. R. E. Whaley, "The Investor Fear Gauge," *Journal of Portfolio Management* 26, no. 3 (2000): 12.

52. O. M. Ardakani, "Portfolio Optimization With Transfer Entropy Constraint," *International Review of Financial Analysis* 96 (2024): 103644.