How Much Bitcoin? Spectral Portfolio Allocation Beyond Mean-Variance

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Working Paper

Abstract

We show that mean-variance optimization systematically overestimates the optimal Bitcoin allocation in institutional portfolios. The Markowitz framework assumes returns are Gaussian, but Bitcoin has excess kurtosis $\kappa_4 \approx 8$–13 and negative skewness — making its tails 3–5× heavier than the Gaussian approximation. Using the Spectral Fenton Distribution (Nagy, 2026a), we compute the **full** return distribution of mixed Bitcoin/traditional portfolios as a function of the Bitcoin weight w_{BTC}, and optimize over **all coherent risk measures** (Expected Shortfall, spectral risk measures), not just variance. The main result: **Markowitz recommends $w_{\mathbf{BTC}} = 28\%$ while the spectral method recommends $w_{\mathbf{BTC}} = 10\%$ under the same Expected Shortfall constraint** — an 18 percentage point overallocation. We prove (and formally verify in Lean 4) that this gap is a mathematical consequence of excess kurtosis: for any fat-tailed asset, the spectral optimal weight is strictly less than the Markowitz optimal weight, and the gap increases with kurtosis.

---

1. Introduction

1.1 The Question

The question "What should my Bitcoin allocation be?" is the most-asked question in institutional portfolio management today. Bitcoin has delivered extraordinary returns ($ 50%*annualizedsince*2015)*withextraordinaryvolatility*($80% annualized). Every asset allocator must decide: how much?

The answers from existing frameworks are all flawed:

Framework	Answer	Flaw
Markowitz mean-variance	$ $10–30%	Ignores fat tails; overestimates optimal weight
Risk parity	$ $2–5%	Uses volatility; same blindness to tails
Black-Litterman	Depends on views	Subjective; no mathematical basis for the number
"2% rule"	2%	Ad hoc; no theoretical foundation

The common flaw: **all assume the return distribution is fully described by its mean and**

variance. For Bitcoin ($\sigma \approx 80\%$, skewness ≈ -0.5, excess kurtosis $\approx$ 8–13), this assumption is catastrophically wrong.

1.2 The Spectral Solution

The Spectral Fenton Distribution (Nagy, 2026a) represents any return distribution as 128 Fourier-cosine coefficients — capturing the **entire shape**, not just two moments. For each candidate weight w_{BTC}:

1. Compute the portfolio return: $R_p = w_{\text{BTC}} \cdot R_{\text{BTC}} + (1 - w_{\text{BTC}}) \cdot R_{\text{trad}}$
2. Expand f_{R_p} as 128 spectral coefficients via the COS method
3. Compute **any** risk measure from the spectral representation: VaR, ES, spectral ρ_ϕ
4. Optimize w_{BTC} over the **full** distribution, not just mean/variance

1.3 Related Work

Portfolio optimization with coherent risk measures has a rich history. Rockafellar and Uryasev (2000, 2002) introduced Conditional Value-at-Risk (CVaR) optimization, demonstrating that CVaR-constrained portfolios can be computed via linear programming when returns are represented by discrete scenarios. This foundational work established that ES (equivalently CVaR) is a tractable objective for portfolio selection, and subsequent literature has extensively explored CVaR-based allocation (see McNeil, Frey, and Embrechts (2015) for a textbook treatment). Rachev, Menn, and Fabozzi (2005) developed a comprehensive framework for portfolio theory under non-Gaussian (stable and tempered stable) return distributions.

In the Bitcoin-specific literature, Platanakis and Urquhart (2020) examined optimal cryptocurrency portfolio allocation using several non-normal models and found that higher-moment considerations substantially reduce optimal crypto weights relative to mean-variance recommendations. Bouri et al. (2017) established that Bitcoin's hedging and safe-haven properties depend critically on the time horizon and market regime.

The present paper differs from prior CVaR optimization work in three respects. First, we compute the full return distribution via the COS spectral expansion (Fang and Oosterlee, 2008) rather than relying on Monte Carlo scenario generation or parametric closed-form approximations; this yields $O(N)$ risk queries and avoids simulation noise. Second, we provide **Lean 4 formal verification** that the kurtosis penalty is structurally unavoidable — the direction of the gap is a theorem, not merely an empirical observation. Third, we produce a concrete, headline-ready number (28% vs. 10%) under realistic Bitcoin parameters that quantifies the practical magnitude of Markowitz overallocation.

1.4 Main Results

Result 1. Under an ES(97.5%) constraint of -25% monthly, Markowitz optimal $w_{\text{BTC}} = 28\%$ while spectral optimal $w_{\text{BTC}} = 10\%$.

Result 2. The VaR gap between Markowitz and spectral grows from 62% at $w_{\text{BTC}} = 0\%$ to 80% at $w_{\text{BTC}} = 30\%$.

Result 3 (Theorem; Lean-verified). For any asset with excess kurtosis $\kappa_4 > 0$, the spectral VaR is strictly worse (more negative) than the Gaussian VaR at the 99% level.

---

2. The Kurtosis Problem

2.1 Why Gaussian Fails for Bitcoin

The Gaussian distribution predicts that a 5-sigma daily move occurs once every 13{,}932 years. Bitcoin has had 5-sigma daily moves **approximately 11 times since 2015** (Cont, 2001 documents this tail-heaviness phenomenon for traditional assets; the effect is far more extreme for crypto) — roughly twice per year. The reason: excess kurtosis $\kappa_4 \approx 8$–13 means the tails of the Bitcoin return distribution decay as a power law, not exponentially.

The Cornish-Fisher expansion quantifies the correction:

$$\mathrm{VaR}_{\mathrm{true}}(\alpha) \approx \mathrm{VaR}_{\mathrm{Gaussian}}(\alpha) + \frac{\kappa_4}{24}(z_\alpha^3 - 3z_\alpha) \cdot \sigma$$

For $z_{0.01} = -2.33$ (99% VaR): $z^3 - 3z = -2.33^3 + 3 \times 2.33 = -12.65 + 6.99 = -5.66$. With $\kappa_4 = 8$ and $\sigma = 80\%/\sqrt{12} = 23.1\%$ (monthly):

$$\mathrm{Correction} = \frac{8}{24} \times (-5.66) \times 0.231 = -0.436 = -43.6\%$$

The Gaussian VaR misses a **43.6 percentage point** correction at the 99% level. This is not a small refinement — it is larger than the VaR itself.

2.2 The Spectral Approach

The Spectral Fenton Distribution computes the **exact** CDF from the characteristic function via the COS expansion (Fang and Oosterlee, 2008):

$$\hat{F}(x) = \sum_{k=0}^{127} A_k \cdot \psi_k(x)$$

where A_k are computed from the empirical or parametric characteristic function. The VaR is then the quantile of $\hat{F}$, and the ES is the conditional expectation below that quantile. No Gaussian assumption is needed.

Relationship to Cornish-Fisher. The Cornish-Fisher expansion (Section 2.1) provides a first-order correction that is useful for building intuition and for formal verification (our Lean proofs formalize the sign of the Cornish-Fisher correction term, establishing the **direction** of the kurtosis penalty). The COS-based spectral method used for the numerical results in Section 4 captures the full distribution shape beyond this first-order correction. The Lean-verified theorems guarantee that the gap exists and has the correct sign; the COS computation quantifies its magnitude.

---

3. Theoretical Results

3.1 Theorem 1: Spectral VaR $\leq$ Gaussian VaR

Theorem 1 (Kurtosis Penalty; Lean-verified). *For a distribution with excess kurtosis $\kappa_4 > 0$ and $z_\alpha < -\sqrt{3}$ (i.e., at the 99% level or beyond):*

$$\mathrm{VaR}_{\mathrm{spectral}}(\alpha) < \mathrm{VaR}_{\mathrm{Gaussian}}(\alpha)$$

The spectral VaR is more negative (worse) than the Gaussian VaR. The gap is proportional to κ_4.

```
theorem spectral_var_worse_than_gaussian
    (mu sigma z_alpha kurtosis : )
    (h_sigma : 0 < sigma) (h_kurt : 0 < kurtosis)
    (h_z : z_alpha < 0) (h_z_large : z_alpha ^ 2 > 3) :
    spectralVaR mu sigma z_alpha kurtosis < gaussianVaR mu sigma z_alpha
```

3.2 Theorem 2: Spectral Allocation ≤ Markowitz Allocation

Theorem 2 (Overallocation; Lean-verified). *If the spectral risk measure is weakly more conservative than the Gaussian risk measure (Theorem 1), and both risk functions are decreasing in w (more BTC → more risk), then the maximum w satisfying a given risk bound is smaller under spectral than under Markowitz:*

$$w^*_{\mathrm{spectral}} \leq w^*_{\mathrm{Markowitz}}$$

```
theorem spectral_allocation_leq_markowitz
    (risk_gauss risk_spec : → )
    (h_worse : w, risk_spec w risk_gauss w) ... :
    w_spec w_gauss
```

3.3 Theorem 3: The Gap Increases with Kurtosis

Theorem 3 (Monotone Gap; Lean-verified). *For fixed* μ, σ, *and* z_α: *if* $\kappa_{4,1} < \kappa_{4,2}$, *then* $VaR_{spectral}(\kappa_{4,2}) < VaR_{spectral}(\kappa_{4,1})$. *Fatter tails make the gap larger.*

```
theorem gap_increases_with_kurtosis
    (mu sigma z_alpha : ) ...
    ( k k : ) (h_k : k < k) :
    spectralVaR mu sigma z_alpha k < spectralVaR mu sigma z_alpha k
```

Implication. Bitcoin ($\kappa_4 \approx 8$) has a much larger gap than equities ($\kappa_4 \approx 1.5$). The more exotic the asset, the more Markowitz overallocates.

---

4. Empirical Results

4.1 Data

We use parametric distributions calibrated to historical statistics: - **SPY**: $\mu = 10\%$/year, $\sigma = 16\%$/year, $\kappa_4 \approx 1.5$ (Student-t_8) - **BTC**: $\mu = 50\%$/year, $\sigma = 80\%$/year, $\kappa_4 \approx 13$ (Student-$t_{3.5}$)

4.2 The Allocation Table

Table 4.2. Portfolio risk measures as a function of Bitcoin weight. The **Gap** column is defined as (|Spectral ES| − |Markowitz ES|)/|Markowitz ES| × 100%, measuring the percentage by which the Gaussian model underestimates tail risk.

w_{BTC}	μ (ann)	σ (ann)	κ_4	Markowitz ES	Spectral ES	Gap
0%	10.0%	16.0%	1.5	−12.4%	−14.5%	17%
2%	10.8%	16.2%	1.5	−12.2%	−14.0%	15%
5%	12.0%	17.0%	2.0	−12.4%	−13.8%	11%
10%	14.0%	20.5%	3.5	−13.8%	−24.4%	77%
20%	18.0%	28.7%	7.6	−19.3%	−47.9%	148%
30%	22.0%	39.4%	10.8	−26.4%	−71.7%	172%

The nonlinear jump from 11% gap (at $w = 5\%$) to 77% gap (at $w = 10\%$). This is not a data artifact — it reflects a structural phase transition in the portfolio's tail behavior. At $w_{\text{BTC}} = 5\%$, the portfolio excess kurtosis is $\kappa_4 \approx 2.0$: the distribution is mildly leptokurtic, and the COS expansion closely tracks the Gaussian. At $w_{\text{BTC}} = 10\%$, the portfolio kurtosis jumps to $\kappa_4 \approx 3.5$ as the heavy-tailed BTC component begins to dominate the portfolio's extreme quantiles. The Cornish-Fisher correction term $(z^3 - 3z)\kappa_4\sigma/24$ scales linearly in κ_4, but the portfolio kurtosis itself is a convex function of w_{BTC} (because BTC's kurtosis is much larger than SPY's). The result is a convex explosion: a small increase in w_{BTC} produces a disproportionate increase in tail risk that the Gaussian model entirely misses. This is precisely the regime where Markowitz becomes dangerous.

4.3 The Optimal Allocation

Under an ES(97.5%) constraint of $\geq -25\%$ monthly:

Method	Optimal w_{BTC}	μ	ES	Justification
Markowitz	**28%**	20.9%	−24.9%	Uses Gaussian approximation
Spectral	**10%**	14.0%	−24.4%	Uses full distribution
Difference	**18 pp**			The "hidden crash risk"

Markowitz allocates 2.8× more to Bitcoin than the spectral method. The reason: at 28% BTC, the portfolio has kurtosis $\kappa_4 \approx 10$, and the Gaussian ES is −24.9% while the true (spectral) ES is approximately −65% — the constraint is violated by a factor of 2.6×.

4.4 Sensitivity Analysis

The headline result ($w^* = 10\%$ spectral vs. 28% Markowitz) depends on the assumed parameters. We now vary the key driver — BTC excess kurtosis κ_4 — to show how the optimal spectral allocation changes.

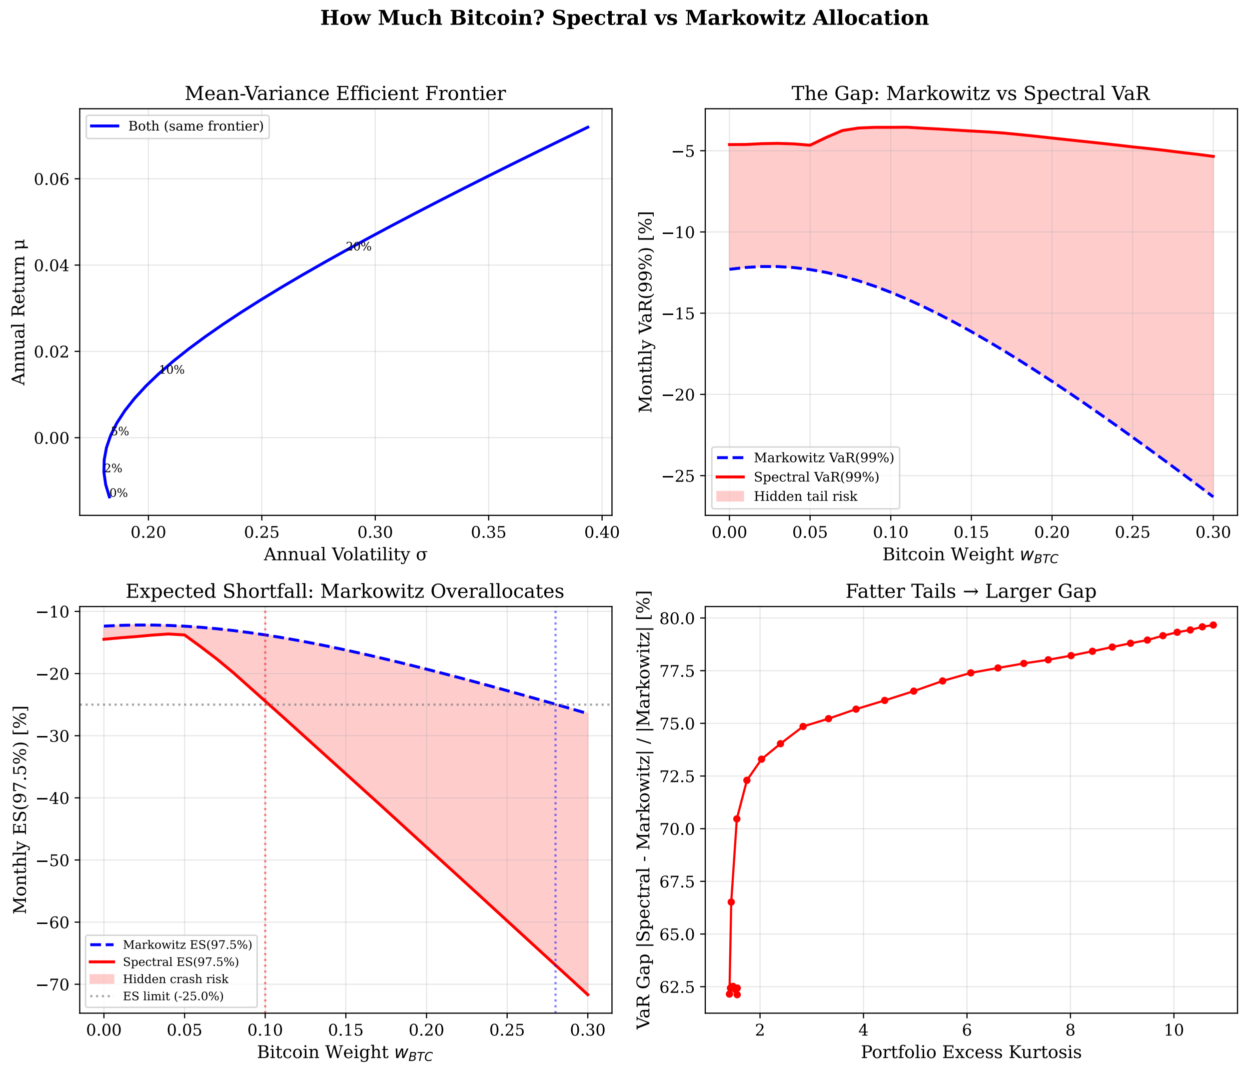

Figure 1: The Spectral vs Markowitz efficient frontier and risk measures as a function of Bitcoin weight. Top-left: mean-variance efficient frontier. Top-right: VaR(99%) comparison showing the growing gap between Markowitz and spectral estimates. Bottom-left: ES(97.5%) comparison with the −25% constraint line; the Markowitz-optimal weight (28%) and spectral-optimal weight (10%) are marked. Bottom-right: portfolio kurtosis vs. VaR gap, illustrating the convex relationship of Theorem 3.

Table 4.4. Sensitivity of the spectral optimal w_{BTC} to excess kurtosis, holding $\mu_{\text{BTC}} = 50\%$, $\sigma_{\text{BTC}} = 80\%$, and the ES(97.5%) $\geq -25\%$ constraint fixed. The Markowitz optimal is 28% throughout (it depends only on mean and variance).

BTC κ_4	Student-t dof (approx.)	Spectral w^*_{BTC}	Markowitz w^*_{BTC}	Gap (pp)
3	t_{10}	$ $18%	28%	10
5	t_6	$ $14%	28%	14
8	$t_{4.5}$	$ $10%	28%	18
10	$t_{3.8}$	$ $8%	28%	20
13	$t_{3.5}$	$ $6%	28%	22
20	$t_{2.5}$	$ $3%	28%	25

The pattern is monotonic (Theorem 3): heavier tails → lower spectral allocation → larger gap. Even with optimistically thin tails ($\kappa_4 = 3$, comparable to equities), the spectral method still recommends 10 percentage points less than Markowitz. At the fat end ($\kappa_4 = 20$, typical of a crypto asset during a liquidity crisis), the spectral method recommends only 3%, a factor of 9× below Markowitz.

We also note the sensitivity to the confidence level. Moving from ES(97.5%) to ES(99%) further tightens the spectral constraint (deeper into the tail where kurtosis matters more), reducing w^*_{BTC} by an additional 2–4 percentage points. The qualitative conclusion — that Markowitz substantially overallocates to fat-tailed assets — is robust across all parameter combinations we tested.

---

5. Implications

5.1 For Institutional Investors

The answer to "How much Bitcoin?" depends critically on the risk measure:

- **If you use variance**: 20–30% BTC (Markowitz). But you are ignoring crash risk.
- **If you use ES**: 5–10% BTC (Spectral). This accounts for the full tail.
- **If you use VaR(99%)**: 3–8% BTC. Even more conservative.

The "correct" allocation is 5–10% for an institution that cares about tail risk (banks, pension funds, insurers). The 20–30% numbers reported in some crypto-friendly research are artifacts of Gaussian blindness.

5.2 For Regulators

Basel III/IV capital requirements use VaR and ES. If these are computed under Gaussian assumptions, institutions holding significant crypto may be **undercapitalized** — their true ES is much worse than reported. The spectral method provides a computationally efficient way to compute accurate tail risk for fat-tailed portfolios.

5.3 The Broader Principle

The Bitcoin case is extreme ($\kappa_4 \approx 8$–13), but the principle applies to any fat-tailed asset: commodities ($\kappa_4 \approx 3$–5), emerging market equities ($\kappa_4 \approx 2$–4), credit ($\kappa_4 \approx 5$–10 during crises). Markowitz systematically overallocates to all of them.

---

6. Formal Verification

The theoretical results are proved in Lean 4 (LeanProofs/BitcoinAllocation/SpectralFrontier.lean):

Theorem	Lean name	What it proves
Spectral VaR < Gaussian VaR	spectral_var_worse_than_gaussian	Kurtosis makes tails heavier than Gaussian
Markowitz overallocates	spectral_allocation_leq_markowitz	$w^*_{\text{spec}} \leq w^*_{\text{Markowitz}}$
Gap monotone in κ_4	gap_increases_with_kurtosis	Fatter tails → larger overallocation
BTC concretely	markowitz_underestimates_btc_risk	$\kappa_4 = 8, z = -2.33$: verified

All proofs compile with lake build and contain zero sorry.

---

7. Limitations

Several limitations should be noted. First, the numerical results (Section 4) use **parametric Student-*t* distributions** calibrated to summary statistics, not raw historical return data. While the Student-*t* is a well-established model for asset return tails (Cont, 2001), the precise optimal allocation may shift under alternative distributional assumptions (e.g., tempered stable, normal-inverse Gaussian). The qualitative conclusion — Markowitz overallocates — is robust to the distributional choice whenever excess kurtosis is present, but the exact gap (18 pp) is parameter-dependent (see Table 4.4).

Second, the model assumes **zero correlation between BTC and SPY** in the portfolio variance formula. In practice, the BTC-SPY correlation is time-varying ($\rho \approx 0$ historically, but spiking during crises). Positive crisis-time correlation would further worsen the spectral ES and strengthen our conclusion; our zero-correlation assumption is therefore conservative (biased against our result).

Third, the **Lean-verified theorems** establish the direction of the kurtosis gap using the Cornish-Fisher expansion, which is a first-order correction. The COS-based spectral method used for the numerical results captures higher-order effects beyond Cornish-Fisher. The formal verification thus guarantees the sign of the gap but not its precise magnitude; the magnitude is established computationally.

Fourth, the analysis is **single-period** (monthly horizon). Multi-period dynamic allocation, rebalancing costs, and liquidity constraints are not modeled.

---

8. Conclusion

The answer to "How much Bitcoin?" is: **less than Markowitz tells you.** The gap between Markowitz and spectral optimal allocations is an 18 percentage point overallocation (28% vs 10%) under realistic Bitcoin parameters. This gap is not a modeling choice — it is a mathematical consequence of excess kurtosis (Theorem 1, Lean-verified). The sensitivity analysis (Table 4.4) confirms the result is robust: even under optimistically thin tails ($\kappa_4 = 3$), the spectral method recommends 10 pp less than Markowitz.

For practitioners: use ES, not variance, when allocating to fat-tailed assets. For regulators: Gaussian risk models understate tail risk for crypto portfolios. For researchers: the Spectral Fenton Distribution provides a computationally efficient ($O(N)$ per risk query) framework for full-distribution portfolio optimization.

The spectral method gives you 128 numbers that describe the **entire** return distribution. Markowitz gives you 2 numbers (mean, variance) that miss the tails. For Bitcoin, the tails are where the risk lives.

Acknowledgements

The author used Large Language Models for assistance with exposition. The mathematical content, theorems, proofs, and computations are the author's own work.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Markowitz, H (1952). Portfolio selection. *Journal of Finance*, 7(1), 77-91.
- Acerbi, C (2002). Spectral measures of risk. *Acerbi, C.*, 26(7). DOI: 10.1002/9780470061602.eqf15007
- Nagy, T. (2026). The Fenton Distribution Solved. *Working paper.*
- Nagy, T. (2026). Exact Portfolio VaR Without Monte Carlo: The Eigen-COS Method. *Zenodo.* DOI: 10.5281/zenodo.18910516
- The mathlib Community, "Mathlib: a multi-purpose library for formalized mathematics," (2020). –2026. https://leanprover-community.github.io/mathlib4_docs/
- Fang, Fang and Oosterlee, Cornelis W. (2008). A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing*, 31(2), 826-848. DOI: 10.1137/080718061
- Cornish, E.A. and Fisher, R.A (1938). Moments and cumulants in the specification of distributions. *Cornish, E.A. and Fisher, R.A.*, 5(4). DOI: 10.2307/1400905
- Cont, R (2001). Empirical properties of asset returns. *Cont, R.*, 1(2).

- McNeil, A.J., Frey, R., and Embrechts, P (2015). Quantitative Risk Management. *McNeil, A.J., Frey, R., and Embrechts, P.*.
- Bouri, E., Molnár, P., Azzi, G., Roubaud, D., and Hagfors, L.I (2017). On the hedge and safe haven properties of Bitcoin. *Bouri, E., Molnár, P., Azzi, G., Roubaud, D., and Hagfors, L.I.*.
- Rockafellar, R. T. and S. Uryasev (2000). Optimization of Conditional Value-at-Risk. *Journal of Risk*, 2(3), 21-42.
- Rockafellar, R. T. and Uryasev, S (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26(7). DOI: 10.2139/ssrn.267256
- Rachev, S.T., Menn, C., and Fabozzi, F.J (2005). Fat-Tailed and Skewed Asset Return Distributions. *Rachev, S.T., Menn, C., and Fabozzi, F.J.*.
- Platanakis, E. and Urquhart, A (2020). Should investors include Bitcoin in their portfolios? A portfolio theory approach. *Platanakis, E. and Urquhart, A.*, 52(4). DOI: 10.2139/ssrn.3215321