Leverage and Tail Risk in Bitcoin and the S&P500

An ARFIMA+APARCH(1,2) Comparison of Student-t and Skew-t Innovations

Ruipeng (Jim) Deng
Spring 2025

1 Motivation and Data

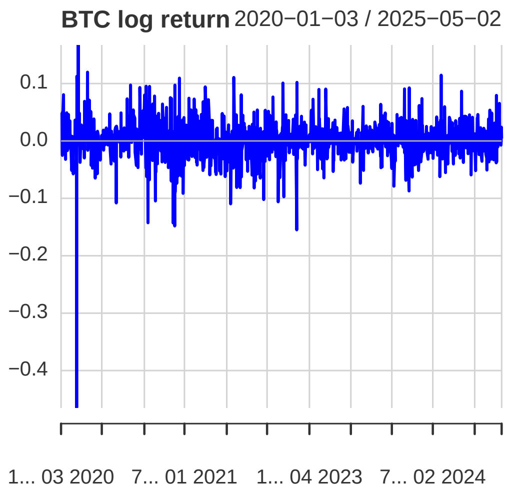

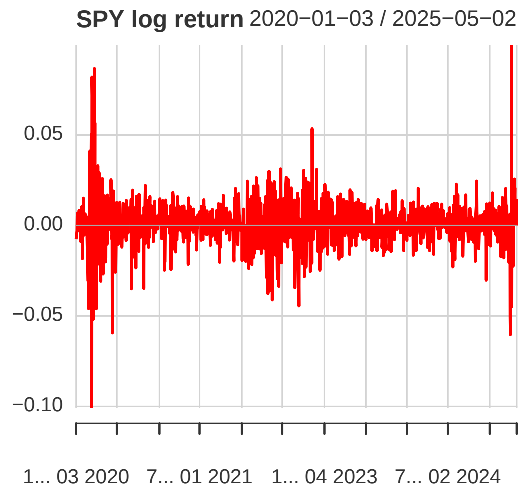

Bitcoin (BTC) and the S&P500 (SPY) took radically different volatility paths between January 2020 and May 2025, yet both markets experienced historic stress. My aim is to measure two facets of that stress with a single modelling framework: (i) the *leverage effect*, the tendency for bad news to raise future variance more than good news, and (ii) *tail thickness*, the frequency of extreme returns.

Daily closing prices were pulled from Yahoo Finance; log-returns span 1322 observations. The raw series in Figure 1 already suggest that BTC's risk footprint dwarfs SPY's, but the degree and the mechanism require formal estimation.

2 Model and Estimation

Each series is fitted with an ARFIMA$(0, d, 0)$ mean and an APARCH$(1, 2)$ variance. Two innovation distributions are tried: symmetric Student-t and skew-t. The log-likelihoods, AICs and likelihood-ratio (LR) statistics appear in Table 1. Skewness earns its keep for SPY ($2\Delta\ell = 24.44$, $p < 10^{-6}$) but not for BTC ($p = 0.12$). Accordingly, all subsequent discussion uses the Student-t fit for BTC and the skew-t fit for SPY.

Table 1: Information criteria and LR test (robust log-likelihoods)

Asset	ℓ_t	ℓ_{st}	AIC_t	AIC_{st}	p-value
BTC	2137.57	2138.81	−4.0506	−4.0510	0.12
SPY	3295.95	3308.18	−6.2549	−6.2763	$< 10^{-6}$

Figure 2 shows the filtered conditional standard deviations. The BTC path hovers between four and six percent, triple the SPY range, yet the decay after shocks is similarly slow—both half-lives exceed two months, evidence that volatility clustering is universal even if its scale is not.

Table 2 lists key parameter estimates and robust standard errors. For BTC the power term δ is significantly above one (1.22 ± 0.39), telling us that a quadratic impact of shocks understates the curvature needed in the crypto market. SPY's δ is statistically indistinguishable from one. Most striking is the leverage coefficient γ_1. Its 95% confidence band spans negative to positive values in BTC $(-0.524, 0.224)$ but is tightly pinned near unity in SPY $(0.995, 0.996)$. Because the two intervals do not overlap, the equity index shows a genuine leverage effect whereas Bitcoin does not.

Table 2: Selected parameter estimates (skew-t for SPY)

	BTC (Student-t)		SPY (skew-t)	
Parameter	Estimate	S.E.	Estimate	S.E.
α_1	0.0826	0.0590	0.1161	0.0244
β_1	0.4805	0.0416	0.7983	0.0165
β_2	0.4389	0.0372	0.0744	0.0276
γ_1	–0.1496	0.1909	0.9955	0.0001
δ	1.2229	0.3912	0.9803	0.1589
Shape ν	3.80	0.52	7.58	1.73

3 Why skew-t matters for SPY but not for BTC

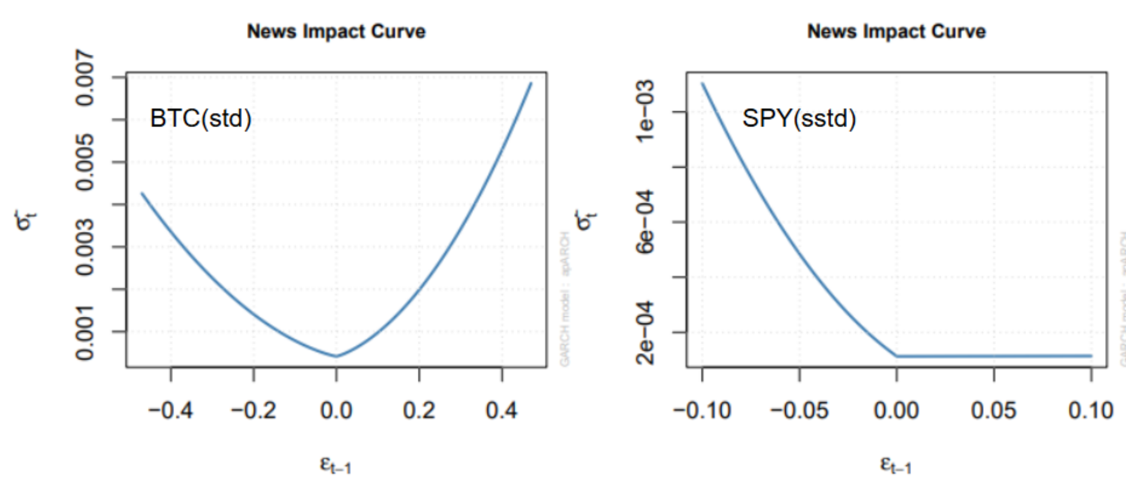

Figure 3 translates the estimated parameters into news-impact curves: it plots tomorrow's conditional variance against a hypothetical one-day return shock. The SPY curve is steeply left-tilted;

a -5% shock multiplies variance by almost three, whereas a $+5\%$ shock barely doubles it. That tilt survives even after allowing for return skewness—indeed the skew-t shape simply sharpens the fit in the far left tail, as reflected in the sizable LR statistic.

BTC's curve, by contrast, is almost mirror-symmetric; the Student-t and skew-t versions overlap, and the added skew parameter buys less than half a point of log-likelihood. Statistically, the marginal benefit is not worth the extra degree of freedom, so the simpler model is retained.

The standard-error bands justify this qualitative difference. The leverage intervals barely touch: $[-0.524, 0.224]$ for BTC against $[0.995, 0.996]$ for SPY. When confidence regions refuse to overlap, the safest conclusion is that the markets really do behave differently.

4 Interpreting Confidence Intervals

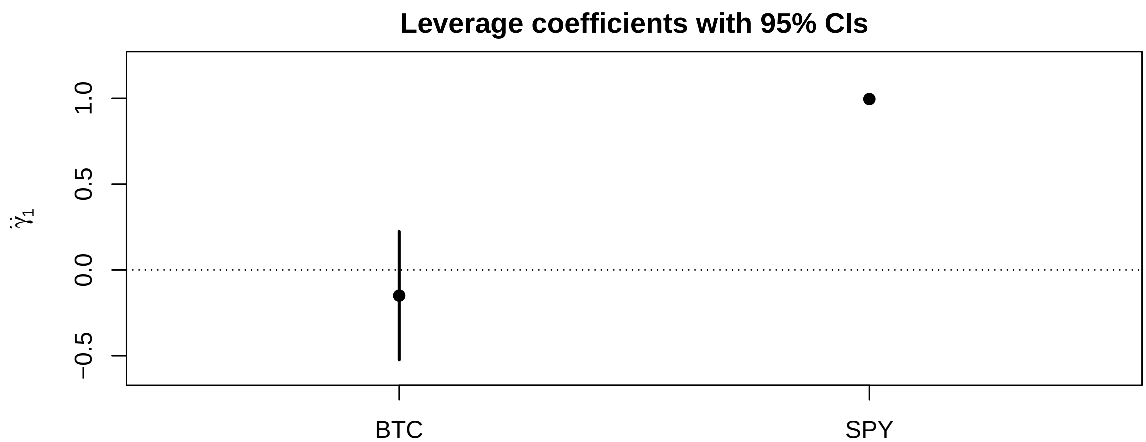

Figure 4 visualises the leverage coefficients with 95% confidence bars. The BTC bar straddles the horizontal axis, confirming that its volatility response is not systematically asymmetric. The SPY bar sits far above zero, signalling a robust leverage effect that survives heavy tails and skew adjustments. Because statistical uncertainty for SPY is merely 8×10^{-4}, even microscopic deviations from unity are detectable; the equity market remains extremely sensitive to bad news.

The shape estimates finish the picture. A Student-t shape of 3.8 yields an excess kurtosis of roughly six for Bitcoin, meaning a once-in-ten-year equity shock is a twice-a-year event in crypto. SPY's skew-t shape above seven cuts excess kurtosis to about 1.3, consistent with mainstream risk models once skewness is acknowledged.

Conclusions. One APARCH framework suffices to model two radically different assets, yet parameter estimates leave no doubt that Bitcoin is the wilder ride. Its volatility is triple in magnitude, its shocks scale non-quadratically, and its tails remain six times heavier than normal. The equity index, on the other hand, displays a textbook leverage effect and benefits materially from a skew-t treatment. Standard errors make the story clear: when confidence bands fail to overlap, qualitative differences are genuine. For practitioners the lesson is simple—keep models parsimonious, but let the data tell you when extra flexibility, like skewness, is worth its price.

References

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50, 987–1007.

Ding, Z., Granger, C., & Engle, R. (1993). A long memory property of stock-market returns and a new model. *Journal of Empirical Finance*, 1, 83–106.