

Statistics of Financial Markets

Chapter 3: α-stable distributions



Daniel Traian PELE

Bucharest University of Economic Studies
IDA Institute Digital Assets
Blockchain Research Center
AI4EFin Artificial Intelligence for Energy Finance
Romanian Academy, Institute for Economic Forecasting
MSCA Digital Finance

Acronyms used in this material

- **AI**: Artificial Intelligence
- **AIC**: Akaike Information Criterion
- **BET**: Bucharest Exchange Trading (index)
- **CLT**: Central Limit Theorem
- **CMS**: Chambers–Mallows–Stuck (simulation method)
- **EODHD**: EOD Historical Data (market-data provider)
- **ES**: Expected Shortfall
- **FFT**: Fast Fourier Transform
- **GCLT**: Generalised Central Limit Theorem
- **IBM**: International Business Machines (company)
- **KS**: Kolmogorov–Smirnov (test)
- **ML**: Maximum Likelihood
- **QQ**: Quantile–Quantile (plot)
- **SAS**: Statistical Analysis System (software)
- **SE**: Standard Error

Today's Question, Route and Learning Outcomes

- **Question**: which probability laws can describe daily returns and also their weekly and monthly sums?
  - Chapter 2: daily returns have heavier tails than the Normal distribution
  - log returns add up over time (Chapter 1): the law of a sum matters
- **Route and outcomes**: after this chapter you can
  - state the stability property and compute the scale of a sum of stable variables
  - explain the generalised central limit theorem and when sums do not approach the Normal distribution
  - read the four parameters $(\alpha, \beta, \gamma, \delta)$ and convert between the S0 and S1 parameterisations
  - say which moments exist for a given α and compare tail probabilities with the Normal distribution
  - simulate stable variables and estimate their parameters by quantiles and by maximum likelihood
  - judge, on BET, S&P 500, DAX and Bitcoin returns, where the stable model works and where it fails

Reading and Tools

- Main reference: Nolan (2020), Ch. 1 (definitions, parameterisations, tails) and Ch. 4 (estimation)
  - short survey with code: Borak, Härdle & Weron (2005); updated in Borak, Misiorek & Weron (2011)
- Textbook: Franke, Härdle & Hafner (2019), Ch. 3 (probability) and Ch. 18 (heavy tails and extreme risks)
- The original papers: Mandelbrot (1963), Fama (1965)
- Python Quantlets of this chapter: Quantlets/Ch_03
  - ported from the Quantlet sim_stable; each chart links to the Quantlet that draws it
- Lecture notebook: open in Google Colab
- Video course: Stable Distribution

Cotton: an Old Market with Long Price Records

- In the 19th century cotton was one of the most traded commodities in the world
  - exchanges in New Orleans, New York and Liverpool quoted prices every day
  - the result: some of the longest daily and monthly price series in economics
- In the early 1960s these series reached Benoit Mandelbrot, then at IBM
  - he looked at the **changes** of the logarithm of the price, i.e. log returns
- His paper Mandelbrot (1963) started the study of heavy tails in finance



Edgar Degas, A Cotton Office in New Orleans (1873)

Image: Edgar Degas (1873); public domain; Wikimedia Commons

Mandelbrot's Observations on Cotton Prices

- **Too many large changes**: far more big moves than the Normal distribution allows
  - the sample variance did not settle as the sample grew
- **The same shape at every horizon**: daily and monthly changes looked alike after rescaling
  - a sum of many daily changes kept the shape of one daily change
- **His proposal**: stable Paretian laws with $\alpha \approx 1.7$ Mandelbrot (1963)
  - "Paretian": the tails decay like a power of x, as in Pareto's law of incomes
  - consequence: the variance of price changes would be infinite



Benoit Mandelbrot (1924–2010)

Photo: Rama (2007); CC BY-SA 2.0 fr; Wikimedia Commons

Fama's Tests on Stock Returns

- Fama (1965): daily returns of the 30 stocks of the Dow Jones Industrial Average
  - more large returns than the Normal distribution predicts, for every stock
  - estimates of α below 2: support for Mandelbrot's hypothesis
- Estimation methods for symmetric stable laws followed
  - Fama & Roll (1968), Fama & Roll (1971): quantile-based estimates and tables
- Within ten years the evidence started to turn: see the critique at the end of the chapter



Eugene Fama at the Nobel Prize ceremony, 2013

Photo: Bengt Nyman (2013); CC BY 2.0; Wikimedia Commons

The Importance of the Law of a Sum

- A monthly log return is the sum of about 21 daily log returns
  - a model for daily returns implies a model for weekly, monthly and yearly returns
- The Normal distribution keeps its shape under summation (Chapter 2)
  - but it has thin tails: it misses the large daily moves
- **Question for the room**: the Student-t distribution has heavy tails; is a sum of 21 independent Student-t variables again Student-t?
- **Answer**
  - no: the Student-t family is not closed under summation
  - if its variance is finite, the sum moves towards the Normal distribution (the central limit theorem)
- **Goal**: find every law whose sums keep the same shape

Definition: Stable Laws

- X is **stable** if, for X_1, X_2 independent copies of X and any $a, b > 0$, there are $c > 0$ and d with
  - $aX_1 + bX_2 \overset{d}{=} cX + d$, where $\overset{d}{=}$ means "has the same distribution as"
  - a, b: the weights of the two copies; c: the scale and d: the shift of the result
  - **strictly stable** if $d = 0$ for all a, b
- In words: a weighted sum of two copies has the same shape as one copy
  - only the scale (c) and the location (d) change
- By induction, for n independent copies: $X_1 + \cdots + X_n \overset{d}{=} c_n X + d_n$
  - the only possible scale factor is $c_n = n^{1/\alpha}$ with $0 < \alpha \leq 2$ Nolan (2020)
- α is called the **index of stability** (also: characteristic exponent, tail index)

Two Examples: Normal and Cauchy

- **Normal**: $X_1, X_2 \sim N(0, \sigma^2)$ independent
  - $aX_1 + bX_2 \sim N(0, (a^2 + b^2)\sigma^2)$, so $c = \sqrt{a^2 + b^2}$
  - variances add: $c^2 = a^2 + b^2$, the case $\alpha = 2$
- **Cauchy**: density $f(x) = \dfrac{1}{\pi(1 + x^2)}$
  - $aX_1 + bX_2$ is Cauchy with scale $c = a + b$
  - scales add linearly: $c^1 = a^1 + b^1$, the case $\alpha = 1$
  - the average of n Cauchy draws is again standard Cauchy: averaging does not help



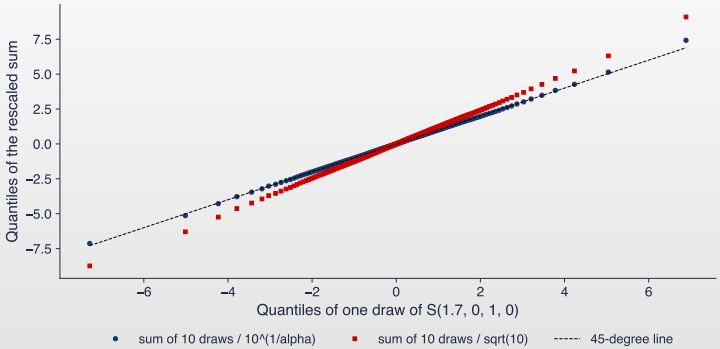
Augustin-Louis Cauchy (1789–1857)

Lithograph: J. Boilly (1821), Wellcome Collection; CC BY 2.0; Wikimedia Commons

The General Rule and a Worked Example

- For X_1, X_2 independent copies of a strictly stable X with index α:
  - $c^\alpha = a^\alpha + b^\alpha$, i.e. $c = (a^\alpha + b^\alpha)^{1/\alpha}$
- **Example**: $a = 3$, $b = 4$
  - $\alpha = 2$: $c = \sqrt{9 + 16} = 5.00$
  - $\alpha = 1.5$: $c = (3^{1.5} + 4^{1.5})^{1/1.5} = 5.58$
  - $\alpha = 1$: $c = 3 + 4 = 7.00$
- The smaller α, the larger the scale of the sum: one large term can dominate the sum
- **What do you think?** For $\alpha = 1.5$, what is the scale factor of a sum of two copies?
- **Answer**
  - $c_2 = 2^{1/1.5} = 2^{2/3} = 1.59$, larger than $\sqrt{2} \approx 1.41$ for the Normal distribution

Stability in a Simulation



- Notation (Section 4): $S(\alpha, \beta, \gamma, \delta)$ is the stable law with index α, skewness β, scale γ and location δ
- Blue: sums of 10 draws of $S(1.7, 0, 1, 0)$ divided by $10^{1/1.7} = 3.87$ lie on the 45-degree line: same law as one draw
- Red: dividing by $\sqrt{10} = 3.16$ (the Normal rule) leaves the tails too wide: 99% quantile 6.31 vs 5.04

SFM_ch3_stability_gclt

Consequences: Horizons and Diversification

- **Horizon**: if daily returns are i.i.d. (independent, identically distributed) $S(\alpha, 0, \gamma, 0)$, the 20-day sum has scale $20^{1/\alpha}\gamma$
  - $\alpha = 2$: $4.47\,\gamma$ (the square-root-of-time rule); $\alpha = 1.7$: $5.83\,\gamma$; $\alpha = 1.5$: $7.37\,\gamma$
  - with $\alpha < 2$, the risk of long horizons grows faster than $\sqrt{n}$
- **Diversification**: an equally weighted portfolio of N i.i.d. stable assets has scale $N^{1/\alpha - 1}\gamma$
  - $N = 10$: $\alpha = 2$ gives $0.32\,\gamma$; $\alpha = 1.5$ gives $0.46\,\gamma$; $\alpha = 1$ gives $1.00\,\gamma$
  - with $\alpha = 1$ diversification does not reduce the scale at all; Fama (1965) discussed this point

Recap: Stability

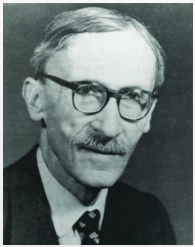
- Stable law: a sum of independent copies has the same shape, only rescaled and shifted
- Scale rule: $c^\alpha = a^\alpha + b^\alpha$; n copies: factor $n^{1/\alpha}$, $0 < \alpha \leq 2$
- Normal: $\alpha = 2$; Cauchy: $\alpha = 1$
- Small α: risk grows faster with the horizon and diversification helps less

From the Classical to the Generalised CLT

- **Classical CLT** (central limit theorem, Chapter 2): X_i i.i.d. with mean μ and finite variance σ^2
  - $\dfrac{X_1 + \cdots + X_n - n\mu}{\sigma\sqrt{n}} \to N(0,1)$ in distribution
- What if the variance is infinite? The normalisation $\sqrt{n}$ is then wrong
- **GCLT** (generalised central limit theorem) Gnedenko & Kolmogorov (1954): if $(X_1 + \cdots + X_n - b_n)/a_n$ converges to a non-degenerate law, that law is stable
  - $a_n > 0$, b_n: normalising constants that rescale and recentre the sum (for the classical CLT, $a_n = \sigma\sqrt{n}$, $b_n = n\mu$)
  - stable laws are the **only** possible limits of normalised sums of i.i.d. variables
  - the Normal distribution is the special case $\alpha = 2$

Domains of Attraction

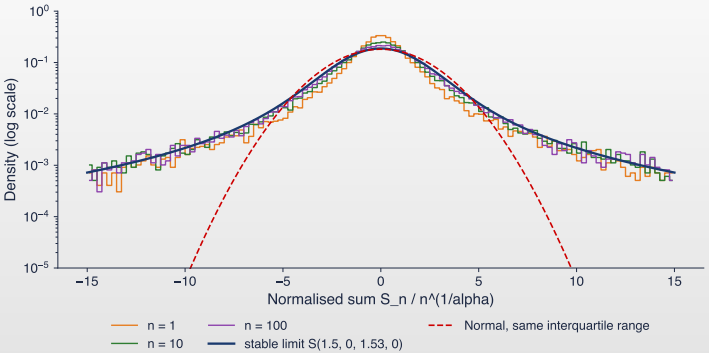
- The **domain of attraction** of a stable law: all laws whose normalised sums converge to it
  - finite variance: domain of attraction of the Normal distribution, $a_n \propto \sqrt{n}$
  - power-law tails $P(|X| > x) \approx Cx^{-\alpha}$ with $0 < \alpha < 2$: the stable law with the same α, $a_n \propto n^{1/\alpha}$
  - $C > 0$: a constant; $\propto$: proportional to
- So the **tail** of one return decides the law of long sums
- Paul Lévy characterised the stable laws in the 1920s; the full theory is in Gnedenko & Kolmogorov (1954)



Paul Lévy (1886–1971)

Photo: Konrad Jacobs, Oberwolfach Photo Collection; CC BY-SA 2.0 de; Wikimedia Commons

The GCLT in a Simulation



- Student-t with $\nu = 1.5$ degrees of freedom: infinite variance, tails $\propto x^{-1.5}$; sums of n draws divided by $n^{1/1.5}$
- $n = 1, 10, 100$: the histograms approach the stable limit $S(1.5, 0, 1.53, 0)$; $P(|\cdot| > 10)$: 2.73% for $n = 100$, 2.64% in the limit
- A Normal law with the same interquartile range gives 0.0005%: it misses the tails completely

SFM_ch3_stability_gclt

What the GCLT Means for Returns

- Two possible worlds for daily returns
  - **tail index below 2** (infinite variance): monthly returns are stable with the **same** α
  - **tail index above 2** (finite variance): monthly returns move towards the Normal distribution
- The two worlds give different testable predictions about returns at longer horizons
- **Question for the room**: in which world do the S&P 500 returns live?
- **Answer**
  - we test it with real data in the last sections: estimates of α at daily, weekly and monthly horizons

Recap: The Generalised CLT

- Classical CLT: finite variance, normalisation $\sqrt{n}$, Normal limit
- GCLT: stable laws are the only limits of normalised sums
- Power-law tails with index $\alpha < 2$: normalisation $n^{1/\alpha}$, stable limit with the same α
- The tail of a single return decides the law of long sums

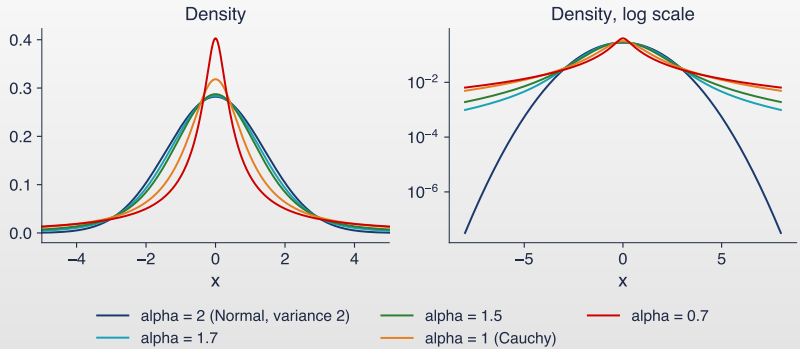
The Characteristic Function

- **Characteristic function** of X: $\varphi_X(t) = E[e^{itX}] = E[\cos(tX)] + i\,E[\sin(tX)]$, $t \in \mathbb{R}$
  - $i = \sqrt{-1}$: the imaginary unit; t: a real argument, not time
  - it always exists, because $|e^{itX}| = 1$, even when no moment exists
  - it determines the distribution uniquely
- Two properties used in this chapter
  - independent X, Y: $\varphi_{X+Y}(t) = \varphi_X(t)\,\varphi_Y(t)$: sums become products
  - $\varphi_{aX+b}(t) = e^{ibt}\varphi_X(at)$
- Normal $N(\mu, \sigma^2)$: $\varphi(t) = \exp(i\mu t - \sigma^2 t^2/2)$
- Why we need it: apart from three cases, stable laws have **no closed-form density**; they are defined by φ

The Stable Characteristic Function (S1)

- $X \sim S(\alpha, \beta, \gamma, \delta; 1)$ if, for $\alpha \neq 1$, Samorodnitsky & Taqqu (1994), Nolan (2020):
  - $\varphi(t) = \exp\left\{ -\gamma^\alpha |t|^\alpha \left[1 - i\beta \operatorname{sign}(t) \tan \frac{\pi\alpha}{2}\right] + i\delta t \right\}$
  - for $\alpha = 1$: $\varphi(t) = \exp\left\{ -\gamma|t|\left[1 + i\beta \frac{2}{\pi} \operatorname{sign}(t) \ln|t|\right] + i\delta t \right\}$
- The four parameters
  - $\alpha \in (0, 2]$: **index of stability**; smaller α, heavier tails
  - $\beta \in [-1, 1]$: **skewness**; $\beta > 0$ heavier right tail, $\beta < 0$ heavier left tail
  - $\gamma > 0$: **scale** (it plays the role of σ); $\delta \in \mathbb{R}$: **location**
- $\operatorname{sign}(t)$: $+1$ for $t > 0$, -1 for $t < 0$; the term with β makes the law asymmetric
- Stability is visible: $\varphi(t)^n$ has the same form, with γ^α replaced by $n\gamma^\alpha$

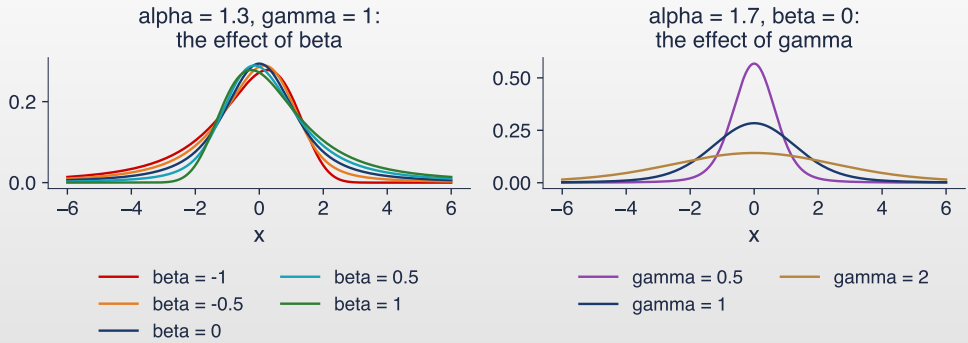
The Effect of α



- Symmetric laws $S(\alpha, 0, 1, 0)$: smaller α gives a taller, narrower peak and much heavier tails
- Log scale: the Normal tail ($\alpha = 2$) falls like $e^{-x^2/4}$; the others fall like a power of x

SFM_ch3_densities

The Effect of β and γ



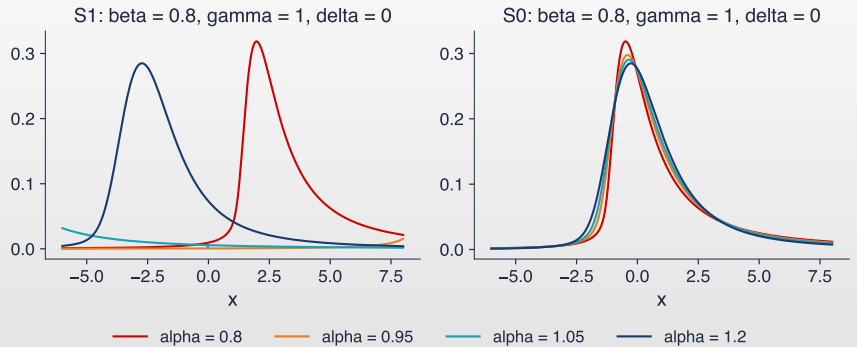
- β moves weight between the tails: $\beta = 1$ (green) has the heavier right tail, $\beta = -1$ (red) the heavier left tail
- γ stretches the density: doubling γ doubles every quantile around δ

SFM_ch3_densities

Two Parameterisations: S0 and S1

- **S0** (Nolan's recommendation for data analysis), $\alpha \neq 1$:
  - $\varphi(t) = \exp\Big\{ -\gamma^\alpha |t|^\alpha \big[1 + i\beta\,\mathrm{sign}(t) \tan\tfrac{\pi\alpha}{2}\big(|\gamma t|^{1-\alpha} - 1\big)\big] + i\delta_0 t \Big\}$
- Only the location differs: $\delta_0 = \delta_1 + \beta\gamma \tan\tfrac{\pi\alpha}{2}$ $(\alpha \neq 1)$
  - δ_0, δ_1: the location parameter in S0 and in S1; α, β and γ are the same in both
  - for $\beta = 0$ the two coincide
- Why two?
  - S1: simpler algebra for sums; used in most theory books Samorodnitsky & Taqqu (1994)
  - S0: the density changes continuously with α and β; better for estimation and plots Nolan (2020)

S1 vs S0 near $\alpha = 1$



- S1 (left): with $\beta = 0.8$, the density runs away to $+\infty$ as $\alpha \uparrow 1$ and comes back from $-\infty$ as $\alpha \downarrow 1$, because $\tan \frac{\pi\alpha}{2}$ explodes
- S0 (right): the same four laws, shifted; the mode stays near 0 and changes smoothly

SFM_ch3_densities

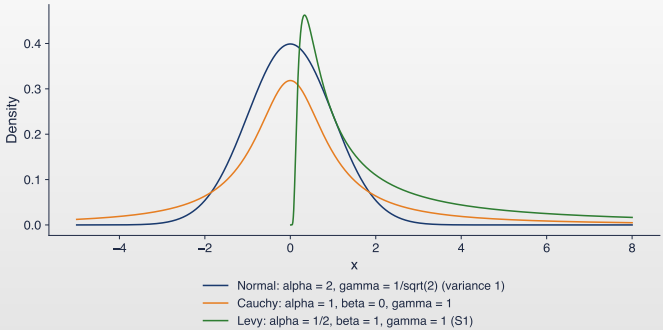
Parameterisations of Stable Laws in Software

- `scipy.stats.levy_stable` uses **S1** by default
  - switch with `levy_stable.parameterization = 'S0'`
  - arguments: `levy_stable.pdf(x, alpha, beta, loc=delta, scale=gamma)`
- Papers and programs report δ in S0, in S1 or in other parameterisations: check which one before comparing
- Scale trap: $S(2, 0, \gamma, \delta)$ is $N(\delta, 2\gamma^2)$, not $N(\delta, \gamma^2)$
  - so $\gamma = \sigma/\sqrt{2}$ for the Normal distribution
- In this course: estimates in S0, with δ_1 also reported

Three Laws with a Closed-Form Density

- **Normal**: $\alpha = 2$ (β plays no role)
  - $\varphi(t) = \exp(-\gamma^2 t^2 + i\delta t)$, i.e. $N(\delta, 2\gamma^2)$
- **Cauchy**: $\alpha = 1$, $\beta = 0$
  - $f(x) = \dfrac{\gamma}{\pi\left(\gamma^2 + (x-\delta)^2\right)}$; $\varphi(t) = \exp(-\gamma|t| + i\delta t)$
- **Lévy**: $\alpha = 1/2$, $\beta = 1$ (S1); totally skewed to the right
  - $f(x) = \sqrt{\dfrac{\gamma}{2\pi}}\,(x-\delta)^{-3/2} \exp\left(-\dfrac{\gamma}{2(x-\delta)}\right)$ for $x > \delta$
  - the first time a Brownian motion hits a level has a Lévy law
- All other stable densities are computed numerically from φ Nolan (1997), Mittnik, Doganoglu & Chenyao (1999)

Normal, Cauchy and Lévy



- Cauchy: lower peak, much heavier tails than the Normal distribution with variance 1
- Lévy: only positive values, a long right tail; its mean is infinite

SFM_ch3_densities

Worked Example: S1 vs S0 at $t = 1$

- X with $\alpha = 1.5$, $\beta = 0.3$, $\gamma = 1$, $\delta = 0$; note $\tan \frac{3\pi}{4} = -1$
- **S1**: $\ln \varphi(1) = -1 \cdot [1 - i \cdot 0.3 \cdot (-1)] = -1 - 0.3i$
  - $\varphi(1) = e^{-1}(\cos 0.3 - i \sin 0.3) = 0.351 - 0.109i$; $|\varphi(1)| = e^{-1} = 0.368$
- **S0**: the factor $|t|^{1-\alpha} - 1$ is 0 at $t = 1$, so $\varphi(1) = e^{-1} = 0.368$, a real number
  - same modulus: the two laws differ only by a shift
  - shift $= \beta\gamma \tan \frac{\pi\alpha}{2} = -0.3$: the S1 law with $\delta_1 = 0$ is the S0 law with $\delta_0 = -0.3$

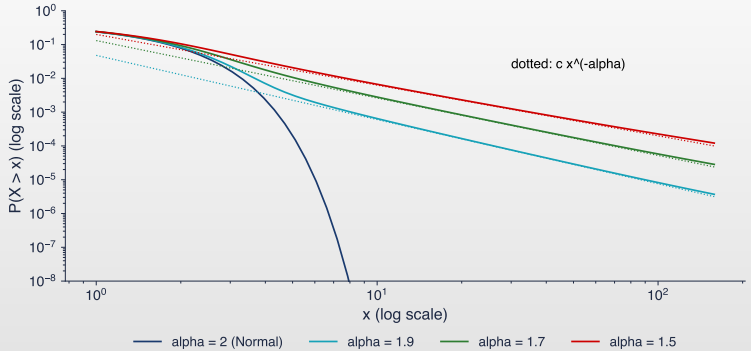
Recap: Parameters and the Characteristic Function

- Stable laws are defined by $\varphi(t)$; the density is computed numerically
- α tails, β skewness, γ scale, δ location
- S0 vs S1: only δ differs, $\delta_0 = \delta_1 + \beta\gamma \tan \frac{\pi\alpha}{2}$; scipy uses S1 by default
- Closed forms: Normal ($\alpha = 2$, variance $2\gamma^2$), Cauchy, Lévy

Power-Law Tails

- For $\alpha < 2$, as $x \to \infty$ Nolan (2020):
  - $P(X > x) \sim \gamma^\alpha c_\alpha (1+\beta)\, x^{-\alpha}$ and $P(X < -x) \sim \gamma^\alpha c_\alpha (1-\beta)\, x^{-\alpha}$
  - $c_\alpha = \Gamma(\alpha)\sin(\pi\alpha/2)/\pi$, with Γ the gamma function; $f \sim g$ means $f/g \to 1$
- On log-log axes the tail is a straight line with slope $-\alpha$
  - doubling the threshold divides the tail probability by 2^α: $\times 0.31$ for $\alpha = 1.7$
- The Normal tail falls like $e^{-x^2/(4\gamma^2)}$: doubling the threshold makes it vanish
- $\beta = \pm 1$ with $\alpha < 1$: one tail is missing (the law lives on a half-line)

Tails on Log-Log Axes



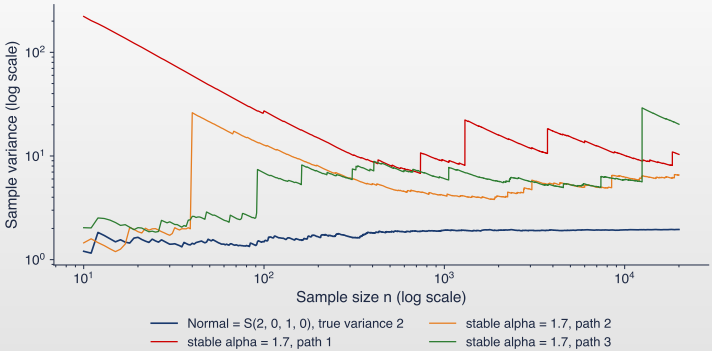
- Solid: $P(X > x)$ of $S(\alpha, 0, 1, 0)$; dotted: the power law $c_\alpha x^{-\alpha}$, reached already near $x = 5$
- $P(X > 10)$: 0.279% for $\alpha = 1.7$, 0.667% for $\alpha = 1.5$, about 8×10^{-11}% for the Normal distribution

SFM_ch3_tails_moments

Existence of Moments

- For $0 < \alpha < 2$: $E|X|^p < \infty$ if and only if $p < \alpha$
  - $E|X|^p$: the absolute moment of order $p > 0$; $p = 1$ gives the mean, $p = 2$ the variance
  - reason: $|x|^p$ times a density of order $|x|^{-\alpha-1}$ is integrable only if $p < \alpha$
- Consequences
  - $\alpha < 2$: **infinite variance**; skewness and kurtosis are not defined
  - $1 < \alpha < 2$: the mean exists; in S1 it equals δ_1
  - $\alpha \leq 1$: not even the mean exists (Cauchy, Lévy)
  - $\alpha = 2$: all moments exist
- **What do you think?** Mandelbrot found $\alpha \approx 1.7$ for cotton. Does the mean daily change exist?
- **Answer**
  - yes, because $1 < 1.7$; but the variance does not, because $2 > 1.7$

Instability of the Sample Variance



- Normal (blue): the sample variance converges to the true value 2
- Stable $\alpha = 1.7$: each large draw makes the variance jump; after 20 000 draws the three paths end at 10.4, 6.6, 20.2
- Mandelbrot saw exactly this in cotton prices

SFM_ch3_tails_moments

Tail Probabilities: Stable vs Normal

k	$P(\lvert X\rvert > k)$, Normal %	$P(\lvert X\rvert > k)$, $\alpha = 1.7$ %	ratio
2	15.73	18.59	1.2
3	3.39	7.25	2.1
5	0.041	2.13	52.4
10	1.5×10^{-10}	0.56	3.6×10^{9}

- Both laws with $\beta = 0$, $\gamma = 1$, $\delta = 0$: the Normal law is $N(0, 2)$
- Near the centre the two laws are close; far in the tails they differ by orders of magnitude
- A risk model is judged in the tails, where the choice of law matters most

SFM_ch3_tails_moments

Infinite Variance and Risk Measures

- Tools that need a finite variance
  - volatility, the Sharpe ratio and its standard error (Chapter 1); mean–variance portfolios; confidence intervals based on the CLT
  - with $\alpha < 2$ these numbers exist in every sample but estimate nothing
- Tools that survive
  - quantiles always exist: $\mathrm{VaR}_{1\%} = -q_{1\%}$, the loss exceeded with probability 1%
  - $\mathrm{ES}_{2.5\%}$ (expected shortfall, the mean loss beyond the VaR) exists only if $\alpha > 1$
  - the scale γ and the interquartile range replace the standard deviation
- VaR and ES are developed in Chapter 10

Recap: Tails and Moments

- Tails: $P(|X| > x) \propto x^{-\alpha}$; a straight line on log-log axes
- Moments: $E|X|^p < \infty$ iff $p < \alpha$; $\alpha < 2$ means infinite variance
- The sample variance of a stable sample jumps and does not settle
- Quantiles, VaR and the scale γ remain meaningful

The Chambers–Mallows–Stuck Method (Symmetric Case)

- No closed-form density, yet stable variables are easy to simulate Chambers, Mallows & Stuck (1976)
- **CMS** (Chambers–Mallows–Stuck) for $S(\alpha, 0, 1, 0)$:
  - draw $V \sim U(-\pi/2, \pi/2)$, a uniform angle, and $W \sim \mathrm{Exp}(1)$, an exponential variable with mean 1, independent
  - $X = \dfrac{\sin(\alpha V)}{(\cos V)^{1/\alpha}} \left(\dfrac{\cos((1-\alpha)V)}{W} \right)^{(1-\alpha)/\alpha}$
  - then $\gamma X + \delta \sim S(\alpha, 0, \gamma, \delta)$
- Checks on special cases
  - $\alpha = 1$: $X = \tan V$, a Cauchy variable
  - $\alpha = 2$: $X = 2 \sin V \sqrt{W}$, a $N(0, 2)$ variable

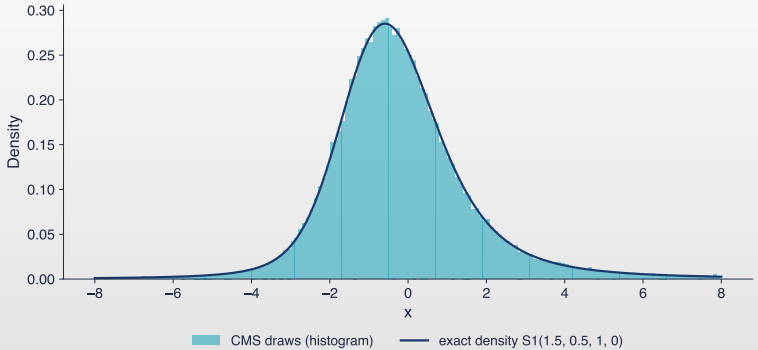
The General Case ($\beta \neq 0$)

- For $\alpha \neq 1$, S1 parameterisation Weron (1996):
  - $B = \dfrac{\arctan(\beta \tan \frac{\pi\alpha}{2})}{\alpha}, \quad S = \left(1 + \beta^2 \tan^2 \frac{\pi\alpha}{2}\right)^{1/(2\alpha)}$
  - $X = S \dfrac{\sin(\alpha(V+B))}{(\cos V)^{1/\alpha}} \left(\dfrac{\cos(V - \alpha(V+B))}{W}\right)^{(1-\alpha)/\alpha} \sim S(\alpha, \beta, 1, 0; 1)$
- B, S: auxiliary constants that depend only on α and β; for $\beta = 0$, $B = 0$ and $S = 1$
- From the standard draw to any parameters
  - S1: $\gamma X + \delta_1$; S0: $\gamma X + \delta_0 - \beta\gamma \tan \frac{\pi\alpha}{2}$
- Weron (1996) corrected a formula of the original paper for $\beta \neq 0$; `scipy.stats.levy_stable.rvs` uses the same method

Worked Example: One Draw Step by Step

- $\alpha = 1.5$, $\beta = 0$; a uniform draw $U = 0.8$ and an exponential draw $W = 1.2$
- Steps
  - $V = \pi(U - 1/2) = 0.942$: turns U, uniform on $(0, 1)$, into a uniform angle on $(-\pi/2, \pi/2)$
  - $\sin(\alpha V) = 0.988$; $(\cos V)^{1/\alpha} = 0.702$
  - $\big(\cos((1-\alpha)V)/W\big)^{(1-\alpha)/\alpha} = 1.104$
  - $X = 0.988/0.702 \times 1.104 = 1.554$
- Repeat with new (U, W) pairs to obtain a sample of any size

CMS Draws against the Exact Density



- 100 000 CMS draws of $S(1.5, 0.5, 1, 0; 1)$: the histogram follows the density computed from φ
- Quantiles 1% and 99%: sample -5.47 and 9.52, exact (scipy `levy_stable`) -5.39 and 9.79
- Two-sample KS (Kolmogorov–Smirnov) test against `levy_stable.rvs`: statistic 0.0044, p-value 0.30, no difference

SFM_ch3_sim_stable

Recap: Simulation

- CMS: a uniform angle V and an exponential W give one stable draw
- Special cases check the formula: $\tan V$ (Cauchy), $2 \sin V \sqrt{W}$ (Normal)
- Skewed case: the formula of Weron (1996); shift by $\beta\gamma \tan \frac{\pi\alpha}{2}$ for S0

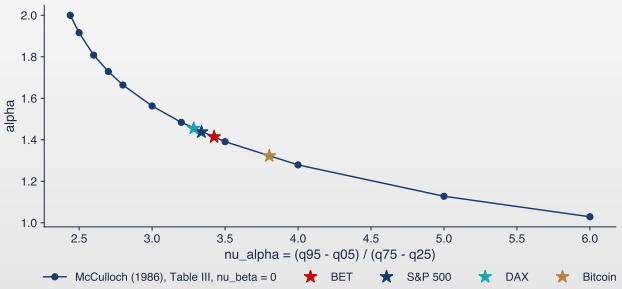
Three Families of Estimators

- **Quantile methods**: match sample quantiles to those of the stable law
  - Fama & Roll (1971) for symmetric laws; McCulloch (1986) for all four parameters
- **Characteristic-function methods**: regress the empirical $\ln(-\ln|\hat{\varphi}(t)|^2)$ on $\ln|t|$
  - $\hat{\varphi}(t) = \frac{1}{n}\sum_{j=1}^{n} e^{itr_j}$: the empirical characteristic function; $|\varphi(t)|^2 = e^{-2\gamma^\alpha |t|^\alpha}$, so the log-log relation is linear in $\ln|t|$
  - Koutrouvelis (1980); the slope estimates α
- **Maximum likelihood** (ML): the most precise, once the density can be computed
  - Nolan (2001), Mittnik et al. (1999)
- A SAS implementation of these estimators is described in Pele (2014)

McCulloch's Quantile Method

- Two ratios of sample quantiles $\hat{q}_p$, free of location and scale McCulloch (1986):
  - $\nu_\alpha = \dfrac{\hat{q}_{0.95} - \hat{q}_{0.05}}{\hat{q}_{0.75} - \hat{q}_{0.25}}$ (tail width relative to the body)
  - $\nu_\beta = \dfrac{\hat{q}_{0.95} + \hat{q}_{0.05} - 2\hat{q}_{0.5}}{\hat{q}_{0.95} - \hat{q}_{0.05}}$ (asymmetry)
- $\hat{q}_p$: the sample quantile of order p, the value below which a share p of the returns lies
- Tables III and IV of the paper map (ν_α, ν_β) to $(\hat{\alpha}, \hat{\beta})$
  - Normal distribution: $\nu_\alpha = 2.439$; larger ν_α means smaller α
- Then $\hat{\gamma} = (\hat{q}_{0.75} - \hat{q}_{0.25})/\phi_3(\hat{\alpha}, \hat{\beta})$ (Table V) and $\hat{\delta}_0$ from Table VII; ϕ_3: the interquartile range of $S(\alpha, \beta, 1, 0)$
- Fast, consistent, valid for $0.6 \leq \alpha \leq 2$; less precise than ML; a good starting point for ML

McCulloch's Map from ν_α to α



- Line: Table III for a symmetric sample ($\nu_\beta = 0$)
- Stars: daily log returns since 2000 (Bitcoin since 2014)
- ν_α: BET 3.425, S&P 500 3.338, DAX 3.285, Bitcoin 3.803
- All well above 2.439: tails much wider than the Normal distribution allows

SFM_ch3_estimation

Worked Example: McCulloch on the BET

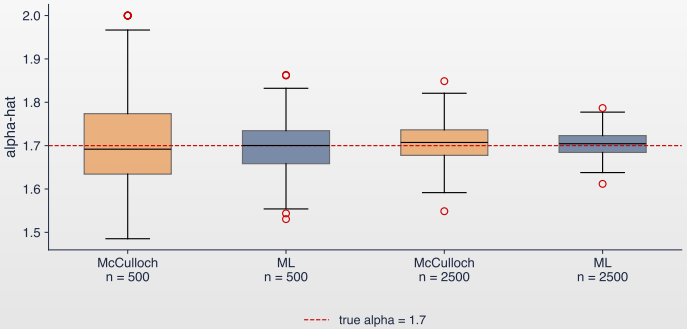
- BET daily log returns (%), 2000–2026, $n = 6\,670$
- Step 1: quantiles
  - $\hat{q}_{0.05} = -1.9345$, $\hat{q}_{0.25} = -0.5025$, $\hat{q}_{0.5} = 0.0771$, $\hat{q}_{0.75} = 0.6649$, $\hat{q}_{0.95} = 2.0641$
- Step 2: ratios
  - $\nu_\alpha = 3.9986/1.1674 = 3.425$; $\nu_\beta = -0.006$, practically symmetric
- Step 3: Table III, column $\nu_\beta = 0$: $\nu_\alpha = 3.2 \mapsto 1.484$, $\nu_\alpha = 3.5 \mapsto 1.391$
  - linear interpolation: $\hat{\alpha} \approx 1.414$
- Step 4: $\hat{\gamma} = 0.60$ (% a day), $\hat{\beta} = -0.011$, $\hat{\delta}_0 = 0.079$

SFM_ch3_estimation

Maximum Likelihood

- Log-likelihood of a sample $r_1, \ldots, r_n$: $\ell(\theta) = \sum_{t=1}^{n} \ln f(r_t; \theta)$, $\theta = (\alpha, \beta, \gamma, \delta_0)$
  - θ: the vector of the four parameters; $f(r_t; \theta)$: the stable density at r_t; $\hat{\theta}$ maximises ℓ numerically, starting from the McCulloch estimate
  - standard errors from the curvature of ℓ: the inverse of the Hessian, the matrix of second derivatives of ℓ
- The density f is computed from φ
  - $f(x) = \frac{1}{2\pi} \int e^{-itx} \varphi(t)\, dt$, evaluated on a grid with the FFT (fast Fourier transform) Mittnik, Doganoglu & Chenyao (1999)
  - or by direct integration Nolan (1997); `levy_stable.fit` does this and needs minutes for a few thousand returns
- ML in S0 is regular for $\alpha \in (0, 2)$: standard errors shrink like $1/\sqrt{n}$ Nolan (2001)

McCulloch vs ML in a Monte Carlo Study



- 200 samples of $S(1.7, 0, 1, 0)$ for each size; both estimators are centred on 1.7
- Standard deviation of $\hat{\alpha}$: $n = 500$: McCulloch 0.1065, ML 0.0611; $n = 2500$: 0.0462 vs 0.0298
- ML is more precise for the same data: McCulloch's standard deviation is 1.74 times that of ML at $n = 500$, 1.55 times at $n = 2500$

SFM_ch3_estimation

Recap: Estimation

- McCulloch: five quantiles, two ratios, four tables; fast and robust
- ML: numerical density from φ, numerical maximisation; the most precise
- Report the parameterisation (S0 or S1) and the units of γ and δ

Stable Fits of Daily Log Returns

	n	$\hat{\alpha}$ McC.	$\hat{\alpha}$ ML (SE)	$\hat{\beta}$ ML (SE)	$\hat{\gamma}$	$\hat{\delta}_0$	$\hat{\delta}_1$
BET	6 670	1.414	1.486 (0.020)	−0.016 (0.034)	0.62	0.088	0.078
S&P 500	6 714	1.438	1.540 (0.021)	−0.181 (0.036)	0.58	0.093	0.000
DAX	6 786	1.455	1.606 (0.021)	−0.179 (0.040)	0.74	0.090	−0.004
Bitcoin	4 384	1.322	1.410 (0.026)	0.019 (0.036)	1.54	0.132	0.171

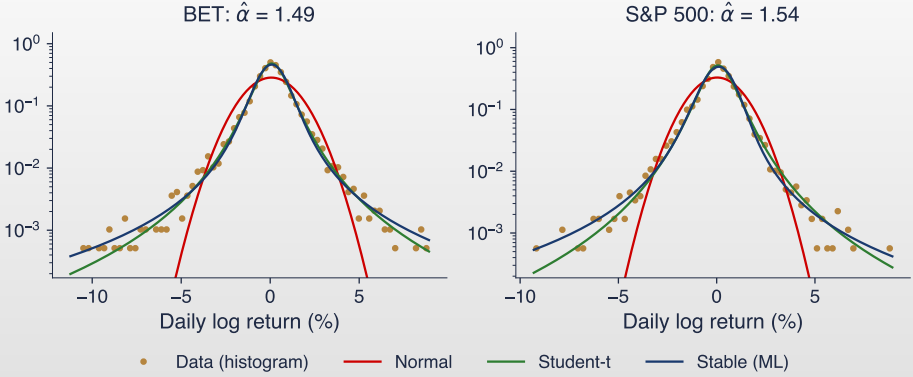
- Daily log returns in %; indices 2000–2026, Bitcoin 2014–2026 (7 days a week); data: EODHD
- McC.: McCulloch; ML in S0 with standard errors; $\hat{\gamma}$, $\hat{\delta}_0$, $\hat{\delta}_1$ in % a day

SFM_ch3_fit_returns

Reading the Estimates

- $\hat{\alpha}$ well below 2 for all four series
  - S&P 500: $(2 - 1.540)/0.021 \approx 22$ standard errors from the Normal case
  - the heaviest tails: Bitcoin ($\hat{\alpha} = 1.410$)
- $\hat{\beta} < 0$ for the S&P 500 and the DAX: the left tail is heavier (crashes)
  - BET and Bitcoin: $\hat{\beta}$ within two standard errors of 0
- $\hat{\gamma}$: Bitcoin 1.54% vs S&P 500 0.58% a day; compare the sample standard deviations 3.49% and 1.21%
- McCulloch gives lower $\hat{\alpha}$ than ML for every series: the two methods weigh the tails differently

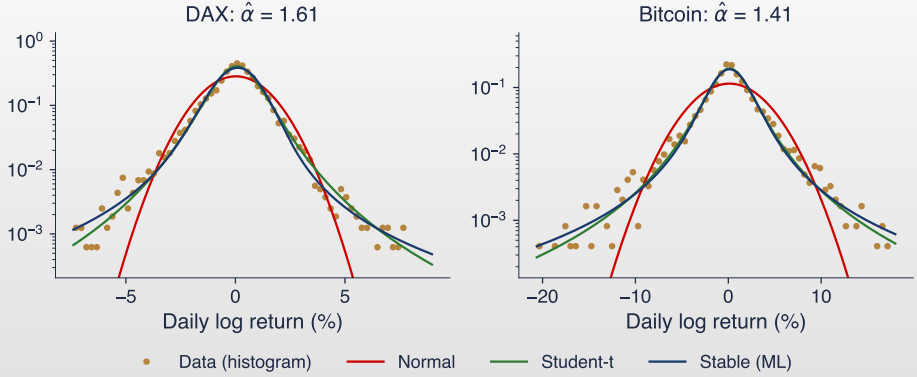
Fitted Densities on a Log Scale (1/2)



- Normal (red): far too few large returns; Student-t (green) and stable (blue) both follow the body

SFM_ch3_fit_returns

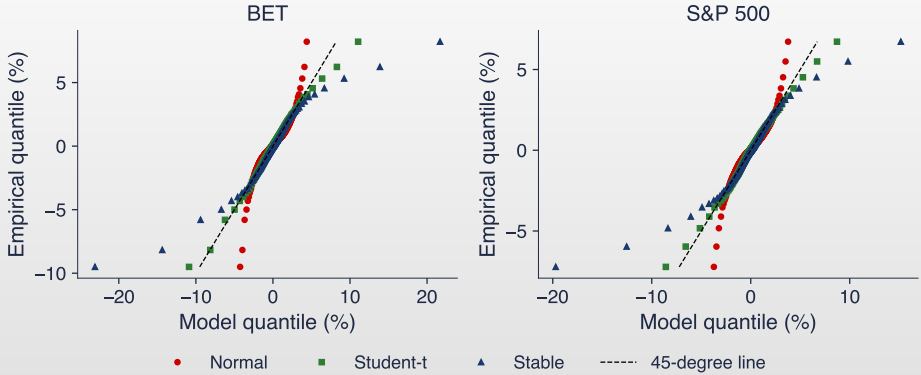
Fitted Densities on a Log Scale (2/2)



- Far tails: the stable density stays above the data points; the Student-t is closer

SFM_ch3_fit_returns

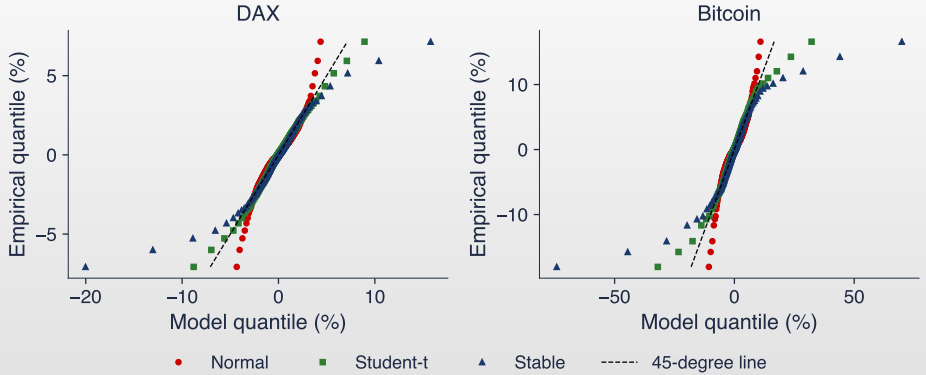
QQ Plots against the Three Models (1/2)



- A QQ (quantile–quantile) plot puts the empirical quantiles against the model quantiles; a good model lies on the 45-degree line

SFM_ch3_fit_returns

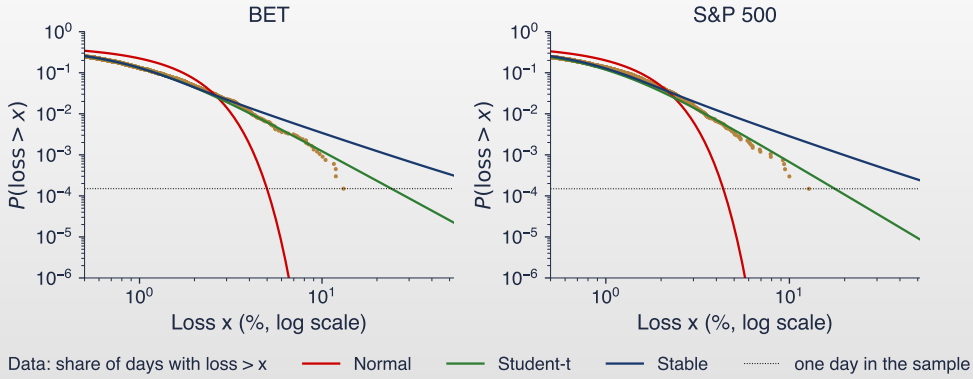
QQ Plots against the Three Models (2/2)



- Normal: too-short tails (steep ends); stable: too-long tails (flat ends, model quantiles far beyond the data); Student-t: the closest

SFM_ch3_fit_returns

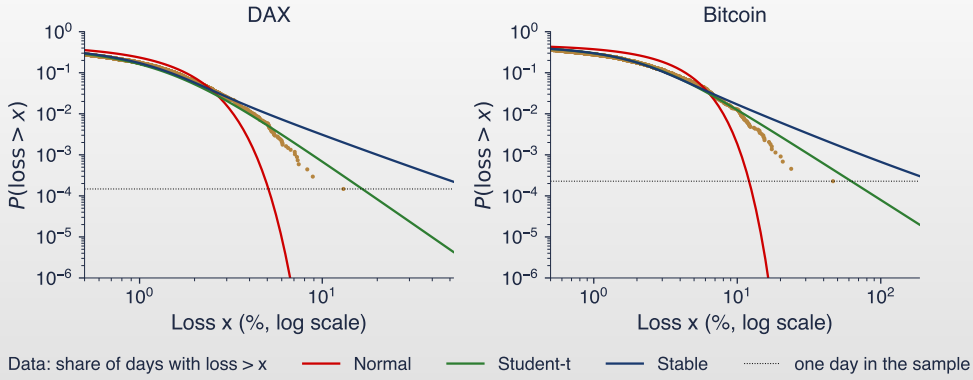
The Left Tail on Log-Log Axes (1/2)



- Points: share of days with a loss above x; lines: the same probability under each fitted model

SFM_ch3_fit_returns

The Left Tail on Log-Log Axes (2/2)



- The empirical tail bends down faster than the stable line: the tail index of the data is larger than $\hat{\alpha}$

SFM_ch3_fit_returns

Counting Large Losses

	loss above	observed	Normal	Student-t	stable
S&P 500	5%	23	0.1	27.0	57.3
S&P 500	10%	1	≈ 0	4.5	19.1
S&P 500	15%	0	≈ 0	1.5	10.2
BET	10%	6	≈ 0	8.2	22.8
Bitcoin	10%	56	8.1	54.2	73.9
Bitcoin	20%	4	≈ 0	12.3	26.9

- Expected number of days $= n \times P(r < -x)$ under each fitted model
- Normal: almost no large losses; stable: too many, and the gap grows with the threshold; Student-t: closest to the counts
- S&P 500: 1 day below -10% in 6 714 days; the stable fit expects 19.1

SFM_ch3_fit_returns

VaR 1% and VaR 0.1% under the Three Models

	VaR 1% (%)				VaR 0.1% (%)			
	data	Normal	t	stable	data	Normal	t	stable
BET	4.1	3.2	4.1	4.9	9.5	4.3	10.9	23.1
S&P 500	3.5	2.8	3.5	4.5	7.2	3.7	8.6	19.7
DAX	4.2	3.3	4.0	5.0	7.1	4.3	8.8	20.0
Bitcoin	10.5	8.0	11.1	14.3	18.1	10.7	31.9	74.2

- VaR at level p: $\mathrm{VaR}_p = -q_p$, the daily loss exceeded with probability p (Chapter 10); here p, because α is the stable index
- VaR 1%: the Normal model is too low, the stable model too high
- VaR 0.1%: the stable model gives roughly two to four times the empirical quantile

SFM_ch3_fit_returns

Comparison of the Three Models

	AIC Normal	AIC Student-t	AIC stable	$\hat{\nu}$ (Student-t)
BET	23 407	20 996	21 088	2.43
S&P 500	21 650	19 702	19 794	2.67
DAX	23 901	22 595	22 697	3.09
Bitcoin	23 391	22 056	22 159	2.23

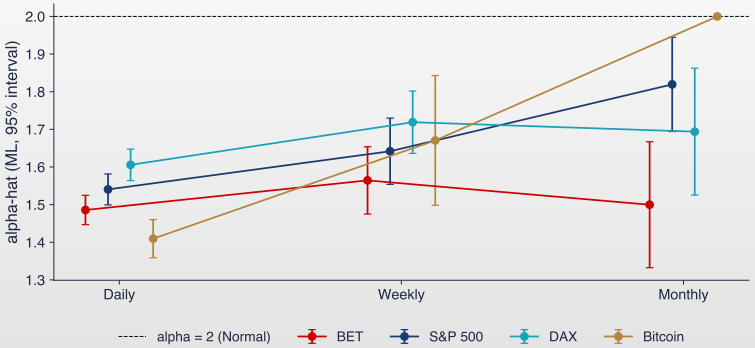
- AIC (Akaike information criterion) $= 2k - 2\ell(\hat{\theta})$, with k the number of parameters; lower is better
- Both heavy-tailed models beat the Normal distribution: Student-t by 1 306 (DAX) to 2 411 (BET) AIC points, stable by 1 204 to 2 319
- Student-t wins for every series, with fewer parameters; S&P 500: by 92 points
- $\hat{\nu}$ between 2 and 3.1: finite variance, but few finite moments beyond it

SFM_ch3_fit_returns

Recap: Fitting Real Returns

- All four series: $\hat{\alpha}$ between 1.410 and 1.606, far from 2
- The stable law fits the body much better than the Normal distribution
- It overstates the far tails: too many extreme losses, VaR 0.1% too high
- By AIC, the Student-t distribution fits better

Test 1: Stability of α under Aggregation



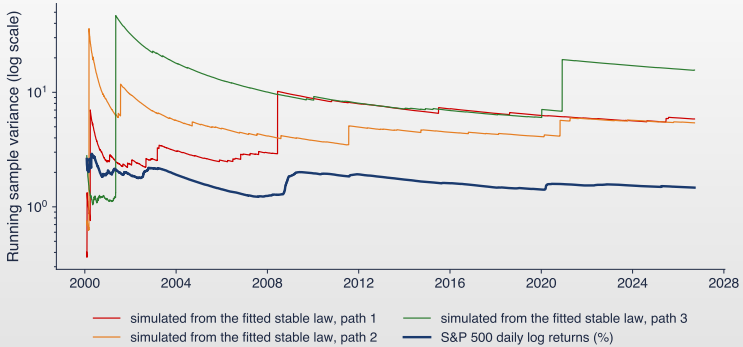
- If daily returns were i.i.d. stable, weekly and monthly sums would have the **same** α (stability)
- ML estimates: S&P 500 1.54 $\rightarrow$ 1.64 $\rightarrow$ 1.82; Bitcoin 1.41 $\rightarrow$ 1.67 $\rightarrow$ 2.00 (at the bound 2)

SFM_ch3_aggregation

Reading the Aggregation Test

- $\hat{\alpha}$ rises with the horizon for the S&P 500, DAX and Bitcoin
  - DAX: $1.61 \to 1.72 \to 1.69$; the daily and monthly intervals do not overlap for the S&P 500
- BET: $1.49 \to 1.56 \to 1.50$, no clear rise
  - monthly samples are short (321 months): wide intervals
- The same pattern in the classic studies: Officer (1972), Akgiray & Booth (1988)
- A rising α is what the GCLT predicts for sums of finite-variance returns that converge slowly to the Normal distribution

Test 2: Convergence of the Sample Variance



- S&P 500 (blue): after the jumps of 2008 and 2020 the running variance settles near 1.47
- Paths simulated from the fitted stable law, same length: final variances between 5.4 and 15.6, far above the data

SFM_ch3_aggregation

Test 3: Thickness of the Far Tail

- The stable α is driven by the whole distribution, mostly by the body
  - the tail index of the far tail alone is estimated with the Hill estimator Hill (1975) (Chapter 5)
- For stock indices the far tail decays roughly like x^{-3}, the "inverse cubic law" Gopikrishnan et al. (1999)
  - a tail index near 3 means finite variance, and no stable law has it
- Our counts agree: the stable fits expect many more losses beyond 10% than observed
- Lau, Lau & Wingender (1990): further evidence against the stable model from the behaviour of sample moments

Origin of the Apparent Stability of Daily Returns

- **Volatility clustering** (Chapters 2 and 9): calm and turbulent periods alternate
  - a mixture of Normal laws with changing variance has heavy tails and a finite variance Blattberg & Gonedes (1974)
  - fitted to one stable law, the mixture gives $\hat{\alpha} < 2$
- **Truncated and tempered stable laws**: stable at moderate scales, lighter tails far out
  - Mantegna & Stanley (1995) fitted a truncated Lévy flight to the S&P 500
  - Grabchak & Samorodnitsky (2010): such laws explain why returns look stable at short horizons and Normal at long ones

Conclusion of the Critique

- **Mandelbrot was right** about the facts
  - daily returns have heavy tails; the Normal distribution fails badly
- **The infinite variance is not supported**
  - $\hat{\alpha}$ rises with aggregation; the sample variance settles; the far tail is lighter than any stable tail
- Where stable laws remain useful
  - a benchmark for heavy tails; simulation; the scaling idea; data with truly infinite variance (some insurance losses, some crypto tokens)
  - the GCLT: a reminder that $\sqrt{n}$-rules need a finite variance
- Next: Chapter 5 estimates the tail index directly with extreme value theory

Recap: The Critique

- Stability predicts the same α at every horizon: the data show a rising α
- The sample variance of real returns settles; that of stable samples does not
- Volatility clustering and tempered tails explain the apparent $\alpha < 2$
- Heavy tails: yes; infinite variance: no

An Open Research Question

- **Have the tails of Bitcoin become lighter as the market matured?**
  - more participants, futures (since 2017) and spot ETFs (exchange-traded funds, since 2024)
  - if the tails got lighter, risk models fitted on old data overstate today's risk
- Why it is open: $\hat{\alpha}$ moves with volatility regimes, and yearly samples are short
- Full sample: $\hat{\alpha} = 1.410$ (SE 0.026), the lowest of our four series
- AI tools can speed up such a study; they do not replace checking it Wang et al. (2023)

Possible Uses of AI Tools

- **Literature**: list papers on tail estimates for crypto assets and summarise their methods
- **Code**: draft a rolling-window McCulloch estimator with bootstrap intervals
- **Robustness**: propose other windows, ML instead of quantiles, Student-t ν instead of α
- **Explanation**: a first draft of the interpretation of a table of yearly estimates
- Example prompt
  - `Write a Python function that estimates the stable alpha of daily log returns on rolling 365-day windows with McCulloch (1986), and adds 95% bootstrap intervals.`

What to Check

- Parameterisation: S0 or S1; scipy uses S1 by default; the location differs by $\beta\gamma \tan \frac{\pi\alpha}{2}$
- Units: returns in % or in decimals; γ and δ change, α and β do not
- Overlapping windows are not independent: do not read the rolling estimates as separate tests
- Numbers: recompute one window with a second method (ML or the quantile tables)
- References: every cited paper must exist; check the DOI

Project Idea

- **Question**: is the tail of Bitcoin returns lighter now than in 2014–2018?
  - data: Bitcoin daily closes from the course data (EODHD), 2014–2026
- Steps
  - estimate α by McCulloch and by ML for each calendar year, with bootstrap intervals
  - repeat on returns divided by a rolling volatility: does the trend survive?
  - compare with the Student-t $\hat{\nu}$ of each year
- Deliverable: one table, one chart, and a paragraph on what the data can and cannot show
- Declare any AI use, and list the errors of the AI that you corrected

Key Formulas

Quantity	Formula
Stability	$aX_1 + bX_2 \stackrel{d}{=} cX + d, \quad c^\alpha = a^\alpha + b^\alpha$
Sum of n copies	$X_1 + \cdots + X_n \stackrel{d}{=} n^{1/\alpha} X + d_n$
Characteristic function (S1)	$\exp\{-\gamma^\alpha \lvert t\rvert^\alpha [1 - i\beta \operatorname{sign}(t) \tan \frac{\pi\alpha}{2}] + i\delta_1 t\}$
S0 vs S1	$\delta_0 = \delta_1 + \beta\gamma \tan \frac{\pi\alpha}{2}$
Normal case	$S(2, \beta, \gamma, \delta) = N(\delta, 2\gamma^2)$
Tails	$P(X > x) \sim \gamma^\alpha c_\alpha (1+\beta) x^{-\alpha}$
Moments	$E\lvert X\rvert^p < \infty \iff p < \alpha \quad (\alpha < 2)$
CMS ($\beta = 0$)	$\frac{\sin(\alpha V)}{(\cos V)^{1/\alpha}} \left(\frac{\cos((1-\alpha)V)}{W}\right)^{(1-\alpha)/\alpha}$
McCulloch	$\nu_\alpha = \frac{q_{0.95} - q_{0.05}}{q_{0.75} - q_{0.25}}, \quad \nu_\beta = \frac{q_{0.95} + q_{0.05} - 2q_{0.5}}{q_{0.95} - q_{0.05}}$

Self-Assessment

- **Question**: daily returns are i.i.d. $S(1.5, 0, \gamma, 0)$; by what factor does the scale grow from 1 day to 20 days?
  - **Answer**: $20^{1/1.5} = 7.37$, i.e. 1.65 times the square-root-of-time factor $\sqrt{20}$
- **Question**: a fitted law has $\alpha = 1.2$; do the mean and the variance exist?
  - **Answer**: the mean exists ($1 < 1.2$), the variance does not ($2 > 1.2$)
- **Question**: scipy reports $\delta = 0.05$ with $\alpha = 1.7$, $\beta = -0.2$, $\gamma = 0.6$; what is δ_0?
  - **Answer**: scipy uses S1, so $\delta_0 = 0.05 + (-0.2)(0.6)\tan(0.85\pi) = 0.05 + 0.061 = 0.111$
- Next: Chapter 4, probability for finance

References (1/2)

- Akgiray, V., Booth, G.G. (1988). The stable-law model of stock returns. *Journal of Business & Economic Statistics*, 6(1), 51–57.
- Blattberg, R.C., Gonedes, N.J. (1974). A comparison of the stable and Student distributions as statistical models for stock prices. *Journal of Business*, 47(2), 244–280.
- Borak, S., Härdle, W., Weron, R. (2005). Stable distributions. In Čížek, P., Härdle, W., Weron, R. (eds.), *Statistical Tools for Finance and Insurance*, 21–44. Springer.
- Borak, S., Misiorek, A., Weron, R. (2011). Models for heavy-tailed asset returns. In Čížek, P., Härdle, W.K., Weron, R. (eds.), *Statistical Tools for Finance and Insurance*, 2nd ed., 21–55. Springer.
- Chambers, J.M., Mallows, C.L., Stuck, B.W. (1976). A method for simulating stable random variables. *Journal of the American Statistical Association*, 71(354), 340–344.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236.
- Fama, E.F. (1965). The behavior of stock-market prices. *Journal of Business*, 38(1), 34–105.
- Fama, E.F., Roll, R. (1968). Some properties of symmetric stable distributions. *Journal of the American Statistical Association*, 63(323), 817–836.
- Fama, E.F., Roll, R. (1971). Parameter estimates for symmetric stable distributions. *Journal of the American Statistical Association*, 66(334), 331–338.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). *Statistics of Financial Markets: An Introduction*, 5th ed. Springer.
- Gnedenko, B.V., Kolmogorov, A.N. (1954). *Limit Distributions for Sums of Independent Random Variables*. Addison-Wesley.
- Gopikrishnan, P., Plerou, V., Amaral, L.A.N., Meyer, M., Stanley, H.E. (1999). Scaling of the distribution of fluctuations of financial market indices. *Physical Review E*, 60(5), 5305–5316.
- Grabchak, M., Samorodnitsky, G. (2010). Do financial returns have finite or infinite variance? A paradox and an explanation. *Quantitative Finance*, 10(8), 883–893.
- Hill, B.M. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics*, 3(5), 1163–1174.
- Koutrouvelis, I.A. (1980). Regression-type estimation of the parameters of stable laws. *Journal of the American Statistical Association*, 75(372), 918–928.

References (2/2)

- Lau, A.H.-L., Lau, H.-S., Wingender, J.R. (1990). The distribution of stock returns: new evidence against the stable model. *Journal of Business & Economic Statistics*, 8(2), 217–223.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36(4), 394–419.
- Mantegna, R.N., Stanley, H.E. (1995). Scaling behaviour in the dynamics of an economic index. *Nature*, 376(6535), 46–49.
- McCulloch, J.H. (1986). Simple consistent estimators of stable distribution parameters. *Communications in Statistics – Simulation and Computation*, 15(4), 1109–1136.
- Mittnik, S., Doganoglu, T., Chenyao, D. (1999). Computing the probability density function of the stable Paretian distribution. *Mathematical and Computer Modelling*, 29(10–12), 235–240.
- Mittnik, S., Rachev, S.T., Doganoglu, T., Chenyao, D. (1999). Maximum likelihood estimation of stable Paretian models. *Mathematical and Computer Modelling*, 29(10–12), 275–293.
- Nolan, J.P. (1997). Numerical calculation of stable densities and distribution functions. *Communications in Statistics. Stochastic Models*, 13(4), 759–774.
- Nolan, J.P. (2001). Maximum likelihood estimation and diagnostics for stable distributions. In Barndorff-Nielsen, O.E., Mikosch, T., Resnick, S.I. (eds.), *Lévy Processes*, 379–400. Birkhäuser.
- Nolan, J.P. (2020). *Univariate Stable Distributions: Models for Heavy Tailed Data*. Springer.
- Officer, R.R. (1972). The distribution of stock returns. *Journal of the American Statistical Association*, 67(340), 807–812.
- Pele, D.T. (2014). A SAS approach for estimating the parameters of an alpha-stable distribution. *Procedia Economics and Finance*, 10, 68–77.
- Rachev, S.T., Mittnik, S. (2000). *Stable Paretian Models in Finance*. Wiley.
- Samorodnitsky, G., Taqqu, M.S. (1994). *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. Chapman & Hall.
- Wang, H., Fu, T., Du, Y., Gao, W., et al. (2023). Scientific discovery in the age of artificial intelligence. *Nature*, 620, 47–60.
- Weron, R. (1996). On the Chambers–Mallows–Stuck method for simulating skewed stable random variables. *Statistics & Probability Letters*, 28(2), 165–171.