






Statistics of Financial Markets

Chapter 5: Heavy tails and extreme value theory



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
- **BET**: Bucharest Exchange Trading (index)
- **BET-TR**: BET Total Return (index)
- **BVB**: Bursa de Valori București (the Bucharest Stock Exchange)
- **CDF**: Cumulative Distribution Function
- **CI**: Confidence Interval
- **CLT**: Central Limit Theorem
- **COVID-19**: COronaVIrus Disease 2019
- **EODHD**: EOD Historical Data (market-data provider)
- **ES**: Expected Shortfall
- **EVT**: Extreme Value Theory
- **GARCH**: Generalised AutoRegressive Conditional Heteroskedasticity
- **GEV**: Generalised Extreme Value (distribution)
- **GPD**: Generalised Pareto Distribution
- **OMV**: Österreichische Mineralölverwaltung — Austrian Mineral Oil Administration (OMV group)
- **POT**: Peaks Over Threshold
- **PP**: Probability–Probability (plot)
- **QQ**: Quantile–Quantile (plot)
- **SE**: Standard Error
- **SFE**: Statistics of Financial Markets, Exercises (Quantlet collection of the textbook)

Today's Question and Route

- **Question**: how large can a daily loss be, and how often should we expect it?
  - Chapters 2 and 3: returns have heavier tails than the Normal distribution
  - risk managers and regulators care about the tail, not about the centre
- **Route** of the chapter
  - three crashes: 1987, 2008, 2020
  - heavy and light tails; regular variation and the tail index α
  - the Hill estimator and the mean excess function
  - block maxima and the GEV distribution; peaks over threshold and the GPD
  - VaR 1%, ES 2.5% and VaR 0.1% by EVT, compared with historical simulation and the Normal distribution

Learning Outcomes

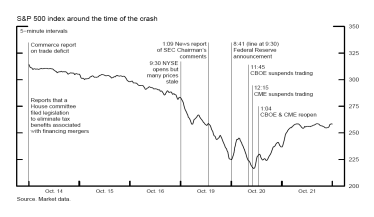
- Tell a heavy tail from a light tail and read a tail on log-log axes
- Estimate the tail index with the Hill estimator, with a standard error and a Hill plot
- State the Fisher–Tippett–Gnedenko and the Pickands–Balkema–de Haan theorems in words and in formulas
- Fit a GEV to block maxima and a GPD to the losses above a threshold; check the fit with PP and QQ plots
- Compute return levels, VaR and ES from an EVT fit, step by step
- Judge, on real data, when EVT improves on historical simulation and on the Normal distribution

Reading and Tools

- Textbook: Franke, Härdle & Hafner (2019), Ch. 18 (statistics of extreme risks)
  - exercises with solutions: Borak, Härdle & López-Cabrera (2013), Ch. 18
- Risk management view: McNeil, Frey & Embrechts (2015), Ch. 5 (extreme value theory)
- Reference books: Embrechts, Klüppelberg & Mikosch (1997); Coles (2001)
- Python Quantlets of this chapter: Quantlets/Ch_05
  - ported from the SFE Quantlets SFElshill, SFEhillquantile, SFEMeanExcessFun, SFEevt1, SFEtailGEV, SFEtailGPareto, block_max and var_pot
- Lecture notebook: open in Google Colab
- Video course: Measuring Statistical Risk

19 October 1987: Black Monday

- On 19 October 1987 the S&P 500 fell about 20% in one day Carlson (2007)
  - as a log return: $\ln(1 - 0.20) = -22.3\%$
  - larger than any daily loss in our data since 1990 (the worst: -12.8% on 16 March 2020)
- Measure it with the Normal distribution fitted to the S&P 500, 1990–2026
  - daily standard deviation $\hat{\sigma} = 1.13\%$; $z = (r - \hat{\mu})/\hat{\sigma} = -19.7$ standard deviations
  - $P(Z \leq -19.7) \approx 10^{-86}$, with $Z \sim N(0, 1)$: impossible, yet it happened
- The Normal distribution says nothing useful about such days: we need a model of the tail



The S&P 500, 14–21 October 1987, 5-minute data (Federal Reserve)

Chart: M. Carlson, Federal Reserve Board (2006); public domain; Wikimedia Commons

September 2008: Lehman Brothers

- Lehman Brothers filed for bankruptcy on 15 September 2008; the global financial crisis followed
  - S&P 500: -9.5% on 15 October 2008 and -9.4% on 1 December 2008
  - BET: -13.1% on 7 January 2009
- Large losses came in clusters, not one at a time
  - volatility clustering (Chapter 2) concentrates the extremes in a few months
- A VaR model based on the Normal distribution puts almost zero probability on such days



Lehman Brothers headquarters, New York, 2007

Photo: David Shankbone (2007); CC BY-SA 3.0; Wikimedia Commons

March 2020: the COVID-19 Crash

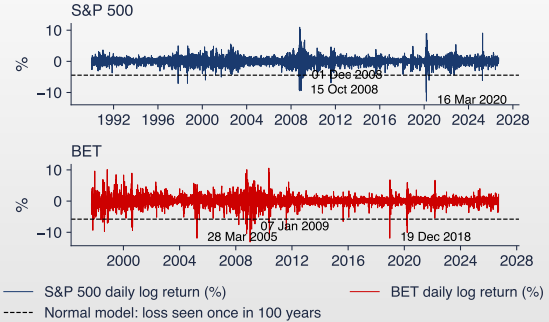
- In one week of March 2020 three of our series had their largest daily loss
  - S&P 500: -12.8% on 16 March 2020
  - DAX: -13.1% on 12 March 2020
  - Bitcoin: -46.5% on 12 March 2020
- The data of this chapter
  - daily closes from EODHD; S&P 500 and DAX since 1990, BET since 1997, Bitcoin since 2014



The empty Financial District of New York, 29 March 2020

Photo: Billie Grace Ward (2020); CC BY 2.0; Wikimedia Commons

Crash Days against the Normal Distribution



- Dashed line: the daily loss that a Normal law fitted to each series exceeds once in 100 years (S&P 500: -4.4%; BET: -5.8%)
- Observed: 35 such days for the S&P 500 since 1990 and 34 for the BET since 1997

SFM_ch5_crashes

The Largest Daily Losses under the Normal Distribution

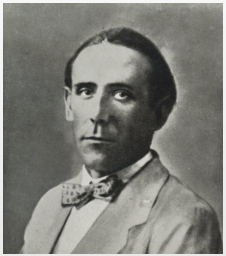
	worst day	loss (%)	z	Normal: once in (years)	days below -4σ	expected
S&P 500	16 March 2020	-12.8	-11.3	10^{27}	35	0.29
DAX	12 March 2020	-13.1	-9.5	10^{19}	27	0.29
BET	7 January 2009	-13.1	-8.9	10^{16}	33	0.23
Bitcoin	12 March 2020	-46.5	-13.4	10^{38}	18	0.14

- $z = (r - \hat{\mu})/\hat{\sigma}$ with the mean and the standard deviation of the whole sample
  - $p = \Phi(z)$: the Normal probability of a day at least as bad, with Φ the $N(0,1)$ CDF
  - once in $1/(p \times 252)$ years (365 days a year for Bitcoin)
- The Normal waiting times are absurd: far longer than the history of any market
- Losses beyond 4σ: the Normal law expects fewer than one day per series; we see between 18 and 35

SFM_ch5_crashes

Origins of Extreme Value Theory

- **EVT** (extreme value theory): the statistics of the largest values of a sample
  - first used for floods, storms and the strength of materials
  - Gumbel (1958): the yearly maximum of a river decides the height of a dam
- In finance the same questions appear
  - how large is the loss exceeded once in 10 years?
  - how much capital covers the worst 1% of days?
- EVT models only the tail and lets the data speak there; early financial uses: Jansen & de Vries (1991), Longin (1996)



Emil Julius Gumbel (1891–1966)

Photo: unknown author (c. 1931); public domain; Wikimedia Commons

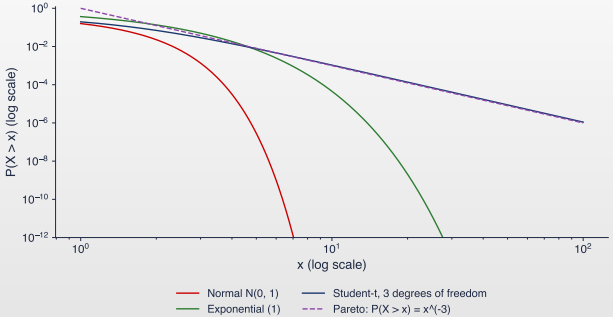
The Role of the Tail in Risk Measurement

- **VaR** (Value at Risk) at level p: the loss exceeded with probability p
  - $\mathrm{VaR}_p = -q_p(r)$, with q_p the p-quantile of the return; for example VaR 1%
- **ES** (expected shortfall) at level p: the average loss on the worst p share of days
  - $\mathrm{ES}_p = -E[r \mid r \leq q_p(r)]$; banks report ES 2.5% under the Basel Committee (MAR33)
  - ES looks inside the tail and is a coherent risk measure Acerbi & Tasche (2002)
- Both live in the tail: a model that misses the tail misses the risk
- **Question for the room**: with 250 trading days a year, how many years of data contain about 10 days below VaR 0.1%?
- **Answer**
  - $10/(0.001 \times 250) = 40$ years: historical data alone are too thin; EVT extrapolates the tail with a model

Light and Heavy Tails

- Work with losses $L = -r$ (in %): a large positive L is a large loss
  - the **tail** of L is described by the survival function $\bar{F}(x) = P(L > x)$
- **Light tail**: $\bar{F}(x)$ falls at least as fast as an exponential, e^{-cx} with some $c > 0$
  - examples: Normal ($\bar{F}(x) \approx e^{-x^2/2}$), Exponential ($e^{-x}$); all moments are finite
- **Heavy (power-law) tail**: $\bar{F}(x)$ falls like a power, $x^{-\alpha}$
  - examples: Pareto, Student-t, the α-stable laws of Chapter 3
  - $\alpha > 0$ is the **tail index**: the smaller α, the heavier the tail

Power-Law Tails and Light Tails



- On log-log axes a power law is a straight line with slope $-\alpha$; light tails bend down and vanish
- $P(X > 10)$: Normal 7.6×10^{-24}, Exponential 4.5×10^{-5}, Student-t with 3 degrees of freedom 1.1×10^{-3}

SFM_ch5_tails_loglog

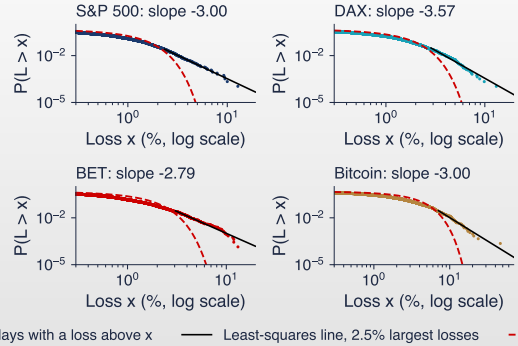
Regular Variation and the Tail Index

- L has a **regularly varying** tail with index $\alpha > 0$ if $\bar{F}(x) = x^{-\alpha}\ell(x)$
  - in words: far in the tail, $\bar{F}$ behaves like the power $x^{-\alpha}$, up to a factor $\ell(x)$ that changes very slowly
  - ℓ is **slowly varying**: $\ell(tx)/\ell(x) \to 1$ as $x \to \infty$, for every $t > 0$ (e.g. a constant, $\ln x$)
- Equivalent: $\dfrac{P(L > tx)}{P(L > x)} \to t^{-\alpha}$ for large x
  - doubling the loss divides its probability by 2^{α}, whatever the level
  - Student-t with 3 degrees of freedom: $P(X > 20)/P(X > 10) = 0.128$, close to $2^{-3} = 0.125$
- Tail indices of common laws
  - Pareto with $\bar{F}(x) = (x/x_0)^{-\alpha}$: index α; Student-t with ν degrees of freedom: $\alpha = \nu$
  - α-stable with $\alpha < 2$ (Chapter 3): the same α; Normal: no power tail ($\alpha = \infty$)

The Tail Index and the Existence of Moments

- For a regularly varying tail: $E|L|^m < \infty$ if $m < \alpha$ and $E|L|^m = \infty$ if $m > \alpha$
  - $E|L|^m$: the moment of order m; it is finite only for orders below the tail index
  - $\alpha \leq 1$: no mean; $\alpha \leq 2$: no variance; $\alpha \leq 4$: no kurtosis
- The "inverse cubic law" of stock returns: $\alpha \approx 3$ Gopikrishnan et al. (1999)
  - variance finite, kurtosis infinite: the sample kurtosis of Chapter 2 does not settle
- Link to Chapter 3
  - the stable α (≈ 1.5) is fitted to the whole distribution; the tail index here uses only the tail
  - a tail index near 3 is the evidence against infinite variance announced in Chapter 3 Mandelbrot (1963)

Real Losses on Log-Log Axes



- Dots: the share of days with a loss above x; black: least-squares line through the 2.5% largest losses; red: Normal
- Slopes: S&P 500 -3.00, DAX -3.57, BET -2.79, Bitcoin -3.00: straight lines, power tails

SFM_ch5_tails_loglog

Reading a Log-Log Tail Plot

- If $\bar{F}(x) \approx Cx^{-\alpha}$, then $\ln \bar{F}(x) \approx \ln C - \alpha \ln x$: a line with slope $-\alpha$
  - **least-squares estimator**: regress $\ln(i/n)$ on $\ln L_{(i)}$ for the k largest losses $L_{(1)} \geq \cdots \geq L_{(k)}$
  - $L_{(i)}$ is the i-th largest loss, an **order statistic**
- The body of the distribution (small losses) bends away from the line: the power law holds only in the tail
- The Normal tail drops below the data already at 3–4% for the indices
- The least-squares slope is simple but its points are not independent; the Hill estimator is the standard tool

SFM_ch5_tails_loglog

Worked Example: a Pareto Tail

- Daily losses have a Pareto tail with $\alpha = 3$ and $P(L > 3\%) = 1\%$
  - $P(L > x) = 1\% \times (x/3)^{-3}$ for $x \geq 3\%$
- Loss above 6%: $1\% \times 2^{-3} = 0.125\%$; above 9%: $1\% \times 3^{-3} = 0.037\%$
  - doubling the loss divides the probability by $2^3 = 8$
- **What do you think?** Which loss is exceeded with probability 0.1% (VaR 0.1%)?
- **Answer**
  - solve $1\% \times (x/3)^{-3} = 0.1\%$: $x = 3 \times 10^{1/3} = 3 \times 2.15 = 6.46\%$
  - a probability ten times smaller only moves the loss by a factor $10^{1/\alpha}$

The Hill Estimator

- Sort the losses: $L_{(1)} \geq L_{(2)} \geq \cdots \geq L_{(n)}$; keep the k largest
  - **Hill estimator** Hill (1975): $\hat{\alpha}_k = \left[\frac{1}{k} \sum_{i=1}^{k} \ln \frac{L_{(i)}}{L_{(k+1)}} \right]^{-1}$
  - in words: the inverse of the average log ratio between the k largest losses and the threshold
  - the threshold is the $(k+1)$-th largest loss, $u = L_{(k+1)}$
- Why it works: if $P(L > x \mid L > u) = (x/u)^{-\alpha}$ (Pareto above u), then $\ln(L/u)$ given $L > u$ is Exponential with mean $1/\alpha$
  - the mean of the k log excesses estimates $1/\alpha$; its inverse estimates α
  - Hill = the maximum likelihood estimator of a Pareto tail above u
- Ported from the SFE Quantlet SFElshill (least-squares and Hill estimators of the tail index) Franke, Härdle & Hafner (2019)

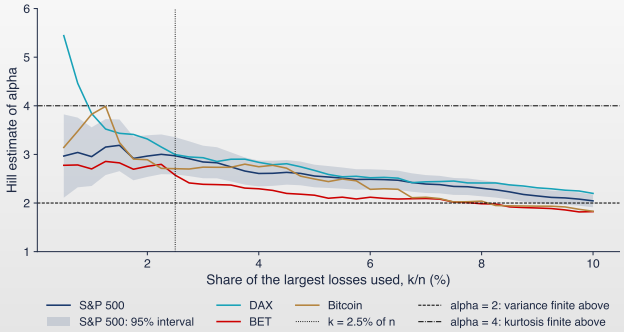
Worked Example: Hill on Six Losses

- The six largest daily losses (%): $9.5, 7.2, 6.1, 5.4, 4.8, 4.0$; take $k = 5$, so $L_{(6)} = 4.0$
- Log excesses $\ln(L_{(i)}/4.0)$
  - 0.865, 0.588, 0.422, 0.300, 0.182; mean 0.471
- $\hat{\alpha}_5 = 1/0.471 = 2.12$
- Standard error under independence: $\mathrm{SE}(\hat{\alpha}) \approx \hat{\alpha}/\sqrt{k} = 0.95$
  - with $k = 5$ the estimate is very imprecise: real applications use hundreds of losses

Choosing k: the Hill Plot

- **Small** k: only the far tail, few points: high variance
- **Large** k: the body of the distribution enters, where the power law fails: bias
- The **Hill plot**: $\hat{\alpha}_k$ against k; look for a region where it is flat Drees, Resnick & de Haan (2000)
  - here k is shown as a share of the sample, k/n between 0.5% and 10%
- Fix the rule before looking at the result: here $k = 2.5\%$ of n
  - and report the estimates at $k = 1\%$ and 5% of n as a check

Hill Plots of Daily Losses



- Between 1% and 5% of n the curves are fairly flat, mostly between 2 and 3.5; band: 95% interval for the S&P 500
- Beyond 5% the estimates drift down: the body of the distribution biases $\hat{\alpha}$

SFM_ch5_hill

Inference for the Hill Estimator

- For i.i.d. data with a Pareto-type tail: $\sqrt{k}\,(\hat{\alpha}_k - \alpha) \to N(0, \alpha^2)$, for $k \to \infty$, $k/n \to 0$
  - i.i.d.: independent and identically distributed; for large k, $\hat{\alpha}_k$ is approximately Normal around α
  - $k/n \to 0$: the number of tail losses grows, but is a vanishing share of the sample
  - $\mathrm{SE} = \hat{\alpha}/\sqrt{k}$; 95% CI (confidence interval): $\hat{\alpha}(1 \pm 1.96/\sqrt{k})$
- Returns are not independent: large losses come in clusters (volatility clustering)
  - **moving-block bootstrap**: resample blocks of 20 consecutive days, recompute $\hat{\alpha}$, take the standard deviation
  - the blocks keep the clusters together, so the standard error is honest about the dependence
- Every series: 500 bootstrap samples, with the k fixed on the Hill-plot slide

Hill Estimates for Four Markets

	n	k	$\hat{\alpha}$	95% CI	SE i.i.d.	SE block	$k = 1\%$ / 5%	gains
S&P 500	9 240	231	2.97	[2.59, 3.35]	0.20	0.30	2.95 / 2.56	2.96
DAX	9 288	232	3.00	[2.61, 3.38]	0.20	0.20	3.84 / 2.67	3.22
BET	7 243	181	2.57	[2.20, 2.95]	0.19	0.23	2.70 / 2.16	2.82
Bitcoin	4 384	109	2.71	[2.20, 3.21]	0.26	0.23	3.83 / 2.49	3.37

- Losses: $\hat{\alpha}$ between 2.57 and 3.00
  - the intervals exclude 2 (finite variance) and lie below 4 (no finite kurtosis)
- The block-bootstrap SE is close to the i.i.d. one for the DAX, BET and Bitcoin, but larger for the S&P 500 (0.30 vs 0.20); the estimates move with k by more than one SE
- Last column: the Hill index of gains ($r_t > 0$), same k

SFM_ch5_hill

The Hill (Weissman) Quantile Estimator

- **Hill (Weissman) quantile** Weissman (1978): $\hat{x}_p = L_{(k+1)} \left(\frac{k}{np} \right)^{1/\hat{\alpha}}$
  - $\hat{x}_p$: the loss exceeded with probability p; k/n: the share of losses above the threshold $L_{(k+1)}$
  - extrapolated from the threshold with the power law; SFE Quantlet SFEhillquantile
- **Worked example**, S&P 500, VaR 0.1%: $n = 9\,240$, $k = 231$, $L_{(k+1)} = 2.34\%$, $\hat{\alpha} = 2.97$
  - $k/(np) = 231/9.24 = 25.0$; $25.0^{1/2.97} = 2.95$
  - $\hat{x}_{0.001} = 2.34 \times 2.95 = 6.92\%$; empirical 0.1% quantile: 6.94%
- The extrapolation is only as good as $\hat{\alpha}$: an error of one SE in $\hat{\alpha}$ moves $\hat{x}_p$ noticeably

SFM_ch5_hill

Four Stocks of the Bucharest Stock Exchange

	n	$\hat{\alpha}$	95% CI	SE block	gains
OMV Petrom (SNP)	3 815	2.96	[2.36, 3.56]	0.39	3.71
BRD	3 771	3.22	[2.57, 3.87]	0.33	3.77
Banca Transilvania (TLV)	3 788	2.49	[1.98, 2.99]	0.28	3.55
Transgaz (TGN)	3 725	3.02	[2.41, 3.64]	0.31	3.56

- Daily log returns of the adjusted close since 2010, $k = 2.5\%$ of n
- TLV: the two days 30–31 May 2016 are removed, an adjustment applied one day late (Chapter 2)
- All four: $\hat{\alpha}$ between 2.49 and 3.22, as for the indices; the losses have heavier tails than the gains

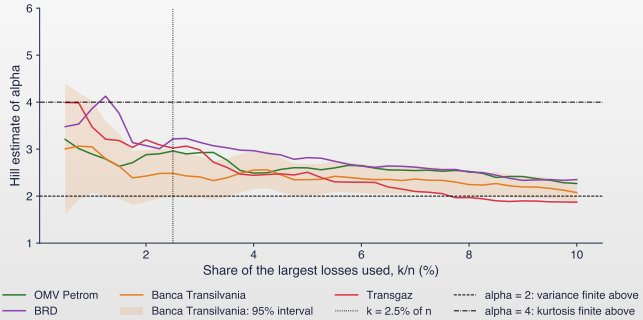


The Stock Exchange Palace, Bucharest

Photo: Neoclassicism Enthusiast (2019); CC BY-SA 4.0; Wikimedia Commons

SFM_ch5_hill

Hill Plots of the BVB Stocks



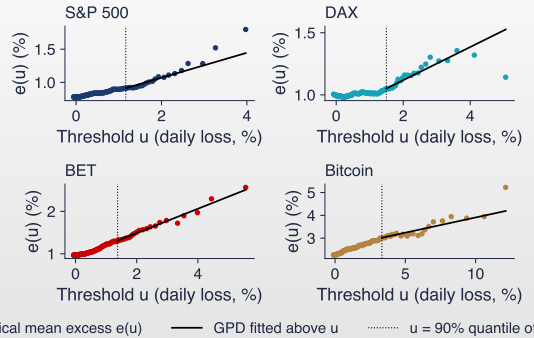
- Shorter samples ($n \approx 3\,800$): wider bands (TLV shown) and more wiggles at small k
- No stock has $\hat{\alpha}$ below 2 in the flat region: variances are finite

SFM_ch5_hill

The Mean Excess Function

- **Mean excess function**: $e(u) = E[L - u \mid L > u]$, the average loss beyond u, given that u is exceeded
  - empirical version: the average of $L_i - u$ over the losses with $L_i > u$
- Its shape tells the type of tail
  - Exponential: $e(u)$ constant (no memory)
  - Pareto with $\alpha > 1$: $e(u) = u/(\alpha - 1)$, a rising straight line
  - Normal: $e(u)$ falls towards 0 (roughly $1/u$ for a standard Normal)
- **Example**: Pareto tail with $\alpha = 3$: $e(u) = u/2$
  - beyond a 2% loss the average extra loss is 1.0%; beyond 4% it is 2.0%: the deeper, the worse

Empirical Mean Excess of Daily Losses



- All four functions rise above the 90% quantile u (dotted): a heavy, Pareto-type tail; black: the straight line of the tail model fitted above u (Section 6)
- The last points rest on very few losses and jump around: do not read them literally

SFM_ch5_mean_excess

Reading the Mean Excess Plot

- Look for the threshold above which $e(u)$ becomes roughly a straight line
  - a rising line means a heavy tail; its slope gives a first guess of the shape of the tail
- At the 90% quantile: $e(u)$ is 0.92% for the S&P 500 ($u = 1.16\%$) and 1.30% for the BET ($u = 1.38\%$)
  - on a day with a loss above u, the loss beyond u is about as large as u itself
- The mean excess function is one of the tools for choosing the threshold of the peaks over threshold method (Section 6)
- Ported from the SFE Quantlet SFEMeanExcessFun

SFM_ch5_mean_excess

Block Maxima

- Split the losses into blocks of equal length and keep the largest loss of each block
  - $M_n = \max(L_1, \ldots, L_n)$, here the largest daily loss of each calendar month ($n \approx 21$)
- Question: what is the law of M_n for large n?
  - the analogue of the central limit theorem, for maxima instead of sums
- Fisher & Tippett (1928) found the three possible limit laws, studying the largest of a sample



Ronald A. Fisher (1890–1962), in 1913

Photo: unknown author (1913); public domain; Wikimedia Commons

The Fisher–Tippett–Gnedenko Theorem

- $L_1, L_2, \ldots$ i.i.d. If there are constants $a_n > 0$, b_n such that $P\left(\frac{M_n - b_n}{a_n} \leq x\right) \to H(x)$, non-degenerate,
  - then H is a **GEV** (generalised extreme value) distribution Fisher & Tippett (1928), Gnedenko (1943)
- Notation
  - M_n: the maximum of n losses; b_n centres it, a_n rescales it, as $\bar{X}_n$ is standardised in the CLT
  - non-degenerate: H is a true distribution, not a mass concentrated in a single point
- Whatever the law of the losses, the normalised maximum has only one family of possible limits

The GEV Distribution

- **GEV** with shape ξ, location μ, scale $\sigma > 0$ (the unified form of Jenkinson (1955))
  - $H(x) = \exp\left\{ -\left[1 + \xi \frac{x-\mu}{\sigma}\right]^{-1/\xi} \right\}$ for $1 + \xi(x-\mu)/\sigma > 0$
  - $\xi = 0$: $H(x) = \exp\{-e^{-(x-\mu)/\sigma}\}$, the limit as $\xi \to 0$
- What each parameter does
  - μ: shifts the distribution (the typical size of the maximum)
  - σ: stretches it (the spread of the maxima)
  - ξ: the shape of the right tail; the larger ξ, the heavier the tail

The Three Types of the GEV Distribution

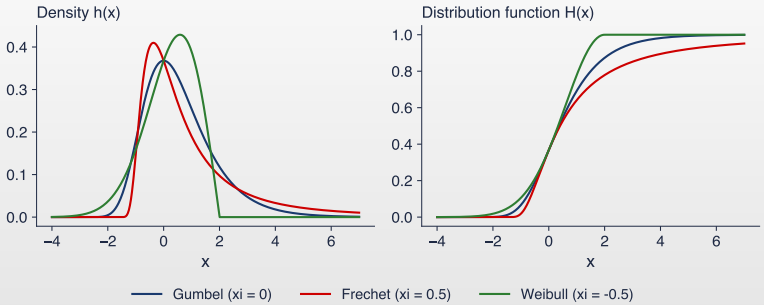
- **Fréchet** ($\xi > 0$): heavy tail, $1 - H(x) \sim x^{-1/\xi}$
  - tail index $\alpha = 1/\xi$; the case of financial losses
- **Gumbel** ($\xi = 0$): light, exponential tail
  - maxima of Normal, lognormal and Exponential variables
- **Weibull** ($\xi < 0$): a finite upper end point $\mu - \sigma/\xi$
  - maxima of bounded variables (Uniform, Beta)
- The set of laws whose maxima converge to a given H is its **maximum domain of attraction**



Maurice Fréchet (1878–1973)

Photo: unknown author; public domain; Wikimedia Commons

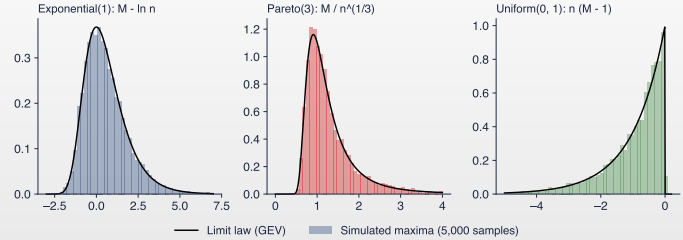
The Three Extreme Value Laws



- Standard location and scale ($\mu = 0$, $\sigma = 1$): Fréchet has the longest right tail, Weibull stops at $x = 2$
- Ported from the SFE Quantlet SFEevt1

SFM_ch5_gev_block_maxima

The Theorem in a Simulation



- 5,000 samples of $n = 250$; Exponential maxima minus $\ln n \to$ Gumbel; Pareto(3) maxima over $n^{1/3} \to$ Fréchet with $\xi = 1/3$; Uniform: $n(M_n - 1) \to$ Weibull
- **Question for the room**: daily losses have a tail index near 3; which type do their monthly maxima follow?
- **Answer**: Fréchet, with $\xi = 1/\alpha \approx 0.3$

SFM_ch5_gev_block_maxima

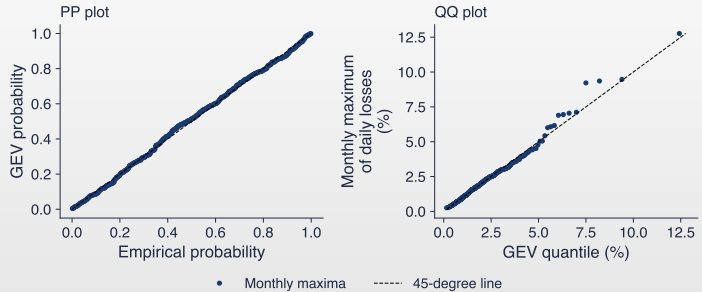
GEV Fits to Monthly Maxima

	months	$\hat{\xi}$	95% CI	$\hat{\mu}$ (%)	$\hat{\sigma}$ (%)	$1/\hat{\xi}$
S&P 500	440	0.21	[0.13, 0.29]	1.31	0.74	4.7
DAX	440	0.17	[0.09, 0.25]	1.69	0.91	5.9
BET	342	0.36	[0.26, 0.46]	1.47	0.92	2.7
Bitcoin	144	0.23	[0.08, 0.38]	5.19	2.85	4.3

- Maximum likelihood on the monthly maxima of daily losses; SE from the curvature of the log-likelihood
- $\hat{\xi} > 0$ for all four series: Fréchet type, heavy tails; the intervals exclude 0 (Gumbel)
- $1/\hat{\xi}$ is a second estimate of the tail index; with only one loss per month it is less precise than Hill

SFM_ch5_gev_block_maxima

Goodness of Fit of the GEV: PP and QQ Plots



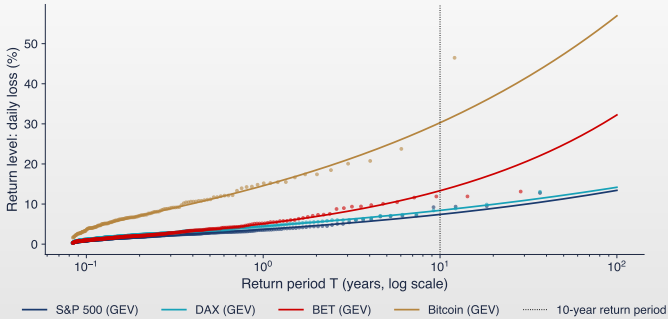
- **PP plot**: fitted probability $H(m_{(i)})$ against the empirical $(i - 0.5)/n$; **QQ plot**: observed maxima against fitted quantiles; S&P 500
- Both close to the 45-degree line; the QQ plot shows the few largest months (2008, 2020) a little above: the hardest part to fit

SFM_ch5_gev_block_maxima

Return Levels

- **Return level** for a return period of T years: the level that the monthly maximum exceeds on average once every T years
  - $z_T = H^{-1}(1 - 1/m)$ with $m = 12T$ months: $z_T = \mu + \frac{\sigma}{\xi}\left[(-\ln(1 - 1/m))^{-\xi} - 1\right]$
- **Worked example**, S&P 500, $T = 10$ years: $\hat{\xi} = 0.211$, $\hat{\mu} = 1.311$, $\hat{\sigma} = 0.737$
  - $m = 120$; $-\ln(1 - 1/120) = 0.00837$; raised to $-\hat{\xi}$: 2.747
  - $z_{10} = 1.311 + (0.737/0.211)(2.747 - 1) = 7.4\%$
- A daily loss of about 7.4% in the S&P 500 is a "10-year event"; it was exceeded in 4 months of 36.7 years

Return Level Plot



- Lines: GEV return levels; dots: observed monthly maxima at their empirical return periods
- Beyond the length of the sample (36.7 years for the S&P 500, 12 for Bitcoin) the curves are extrapolation

SFM_ch5_gev_block_maxima

Return Levels of Daily Losses: 1, 10 and 50 Years

	1 year	10 years	95% interval	50 years	largest observed	years of data
S&P 500	3.7	7.4	[6.2, 8.9]	11.3	12.8	36.7
DAX	4.5	8.4	[7.1, 9.9]	12.2	13.1	36.7
BET	5.1	13.3	[10.3, 17.5]	24.8	13.1	28.5
Bitcoin	14.6	30.3	[21.2, 41.1]	47.4	46.5	12

- Daily loss (%) exceeded on average once in T years; interval: bootstrap of the monthly maxima (200 samples)
- BET: $\hat{\xi} = 0.36$ makes the 50-year level 24.8%, almost twice the largest loss seen
  - a small change in $\hat{\xi}$ gives a large change far in the tail

SFM_ch5_gev_block_maxima

Peaks over Threshold

- **POT** (peaks over threshold): fix a high threshold u and keep every loss above it
  - the **excess** over u: $Y = L - u$, given $L > u$
  - excess distribution: $F_u(y) = P(L - u \leq y \mid L > u)$
- Uses all large losses, not one per block: more data in the tail
- The method behind most EVT estimates of VaR and ES McNeil & Frey (2000), McNeil, Frey & Embrechts (2015)



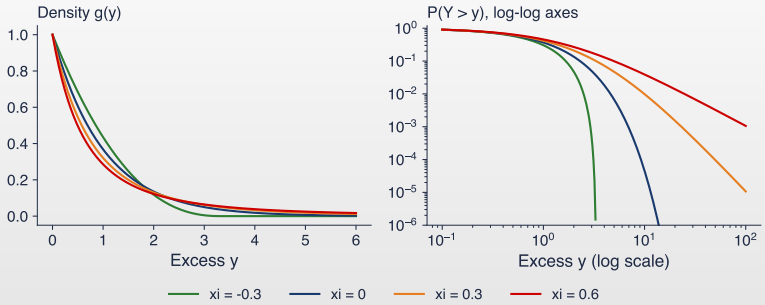
Laurens de Haan, Oberwolfach, 1987

Photo: Konrad Jacobs (1987), Oberwolfach Photo Collection; CC BY-SA 2.0 de; Wikimedia Commons

The Pickands–Balkema–de Haan Theorem

- If the maxima of L converge to a GEV with shape ξ, then for high thresholds
  - $F_u(y) \approx G_{\xi,\beta(u)}(y)$, the **GPD** (generalised Pareto distribution) Pickands (1975), Balkema & de Haan (1974)
- $G_{\xi,\beta}(y) = 1 - \left(1 + \xi \frac{y}{\beta}\right)^{-1/\xi}$, $y \geq 0$ (for $\xi < 0$: $y \leq -\beta/\xi$); $\xi = 0$: $1 - e^{-y/\beta}$
  - y: the excess over u; shape ξ: the **same** as in the GEV; scale $\beta > 0$ depends on u
  - $\xi > 0$: Pareto-type tail with $\alpha = 1/\xi$; $\xi = 0$: Exponential; $\xi < 0$: bounded
- Moments of the GPD: mean $\beta/(1-\xi)$ if $\xi < 1$; variance finite if $\xi < 1/2$
  - mean excess of a GPD: $e(v) = \frac{\beta + \xi(v-u)}{1-\xi}$, a straight line in $v > u$, as seen in Section 4

GPD Densities and Tails



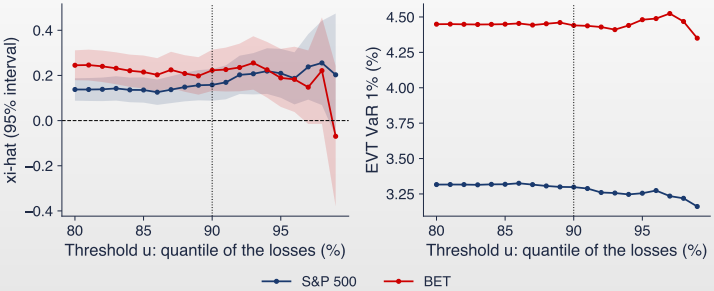
- $\beta = 1$; the densities look alike near 0, the tails do not: on log-log axes only $\xi > 0$ gives straight lines
- $\xi = 0.3$ corresponds to a tail index $\alpha \approx 3.3$, close to stock returns

SFM_ch5_pot_gpd

Choosing the Threshold

- The same trade-off as k in the Hill estimator
  - u too low: the GPD approximation is poor (bias); u too high: few excesses (variance)
- Graphical tools
  - mean excess plot: roughly linear above u
  - parameter stability plot: $\hat{\xi}$ should not change much when u is raised further
- Rule used here: $u =$ the 90% quantile of the losses, i.e. the 10% largest losses
  - McNeil & Frey (2000) use $k = 100$ excesses in windows of $n = 1000$ days, the same share

Parameter Stability



- Left: $\hat{\xi}$ with 95% interval for u from the 80% to the 99% quantile; right: the EVT VaR 1% for the same thresholds
- Between the 85% and 95% quantiles: S&P 500 $\hat{\xi} \in [0.13, 0.22]$, VaR 1% $\in [3.25, 3.33]$%; the VaR barely moves

SFM_ch5_pot_gpd

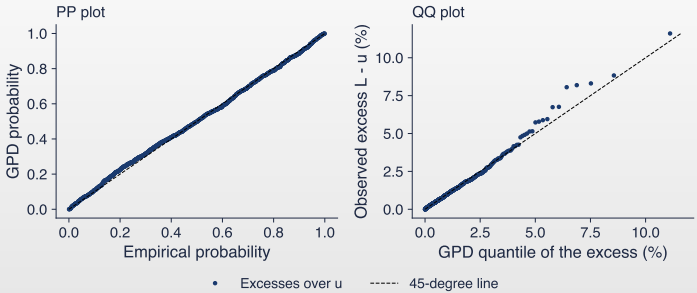
GPD Fits above the 90% Quantile

	n	u (%)	N_u	$\hat{\xi}$	SE	$\hat{\beta}$	$1/\hat{\xi}$
S&P 500	9 240	1.16	924	0.158	0.04	0.769	6.3
DAX	9 288	1.50	929	0.118	0.04	0.927	8.5
BET	7 243	1.38	725	0.223	0.05	1.018	4.5
Bitcoin	4 384	3.33	439	0.118	0.05	2.663	8.5

- Maximum likelihood on the excesses $L - u$ Davison & Smith (1990); SE from the curvature of the log-likelihood
- $\hat{\xi}$ between 0.12 and 0.22: heavy tails with a finite variance ($\xi < 0.5$) for all four
- $1/\hat{\xi}$ is larger than the Hill $\hat{\alpha}$ (2.5% of the losses)
  - above the 90% quantile, a GPD with a free scale β needs a smaller ξ
  - the tail index is hard to pin down: report more than one estimator

SFM_ch5_pot_gpd

Goodness of Fit of the GPD: PP and QQ Plots



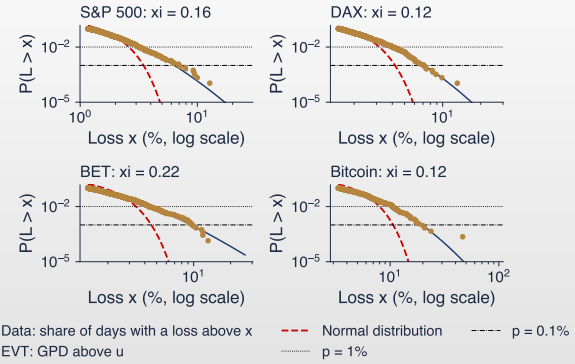
- S&P 500 excesses over u: the PP plot is on the line; the QQ plot is on the line up to excesses of about 4%
- The largest excess (11.6%, March 2020) against a fitted quantile of 11.1%; ported from the SFE Quantlets SFEtailGPareto_pp and _qq

SFM_ch5_pot_gpd

The GPD Tail Estimator of Losses

- For $x > u$: $P(L > x) = P(L > u)\, P(L > x \mid L > u)$
  - estimate $P(L > u)$ by N_u/n and $P(L > x \mid L > u)$ by the GPD
- **Tail estimator**: $\hat{\bar{F}}(x) = \dfrac{N_u}{n}\left(1 + \hat{\xi}\,\dfrac{x-u}{\hat{\beta}}\right)^{-1/\hat{\xi}}$, $x > u$
  - N_u: the number of losses above u; n: the sample size; the bracket is the GPD probability of exceeding x
- Historical data inside the sample, a fitted power law beyond it
  - the model is used only where it is justified: above u
- Inverting $\hat{\bar{F}}(x) = p$ gives VaR; integrating beyond it gives ES (Section 7)

The Fitted Tails



- Dots: share of days with a loss above x (above u); blue: EVT tail estimator; red: Normal; dotted lines at $p = 1\%$ and 0.1%
- The GPD follows the data down to 10^{-4}; the Normal tail leaves them before $p = 1\%$

SFM_ch5_pot_gpd

Waiting Times of the Worst Days: Normal Law and EVT

	day	loss (%)	once in (years) Normal law	EVT
S&P 500	16 March 2020	−12.8	10^{27}	89
DAX	12 March 2020	−13.1	10^{19}	85
BET	7 January 2009	−13.1	10^{16}	12
Bitcoin	12 March 2020	−46.5	10^{38}	237
S&P 500, Black Monday	19 October 1987	−22.3	10^{84}	1 602

- EVT waiting time: $1/(\hat{\bar{F}}(x) \times 252)$ years (365 for Bitcoin), with the POT fit of the whole sample
- EVT makes the COVID-19 days rare but plausible events of a lifetime; the Normal law makes them impossible
- 1987 is outside our data: EVT extrapolates it to a waiting time of 1 602 years, with a wide uncertainty

SFM_ch5_pot_gpd

Estimating VaR and ES

- VaR and ES of Section 1, written for the losses $L = -r$, with distribution function F_L
  - $\mathrm{VaR}_p = F_L^{-1}(1-p) = -q_p(r)$
  - $\mathrm{ES}_p = \dfrac{1}{p}\displaystyle\int_0^p \mathrm{VaR}_s \, ds = E[L \mid L \geq \mathrm{VaR}_p]$ for continuous laws
- Three ways to estimate them from daily losses
  - **historical simulation**: the empirical quantile and the mean of the losses beyond it
  - **Normal**: $\mathrm{VaR}_p = \hat{\mu}_L + \hat{\sigma} z_{1-p}$, $\mathrm{ES}_p = \hat{\mu}_L + \hat{\sigma}\, \varphi(z_{1-p})/p$
  - $\hat{\mu}_L$, $\hat{\sigma}$: mean and standard deviation of the losses; z_{1-p}: the $(1-p)$ quantile of $N(0,1)$; φ: the standard Normal density
  - **EVT (POT)**: invert the tail estimator
- Levels used in practice: VaR 1% (daily, Basel backtesting), ES 2.5% Basel Committee (MAR33), VaR 0.1% (stress)

EVT Formulas for VaR and ES

- Solve $\frac{N_u}{n}\left(1+\xi\frac{x-u}{\beta}\right)^{-1/\xi} = p$ for x
  - $\mathrm{VaR}_p = u + \frac{\beta}{\xi}\left[\left(\frac{np}{N_u}\right)^{-\xi} - 1\right]$, valid for $p < N_u/n$
  - the threshold plus a GPD quantile of the excess; np/N_u: the target probability relative to the share of excesses
- The excess over VaR_p is again GPD with the same ξ: its mean is $(\beta + \xi(\mathrm{VaR}_p - u))/(1-\xi)$
  - $\mathrm{ES}_p = \frac{\mathrm{VaR}_p}{1-\xi} + \frac{\beta - \xi u}{1-\xi}$, valid for $\xi < 1$ McNeil & Frey (2000)
  - ES = VaR plus the mean excess beyond VaR
- With $\xi > 0$, ES is a fixed multiple of VaR far in the tail: $\mathrm{ES}_p/\mathrm{VaR}_p \to 1/(1-\xi)$
- Ported from the SFE Quantlet var_pot (VaR estimates with the POT model)

Worked Example: S&P 500, Step by Step

- Fit: $n = 9\,240$, $u = 1.163\%$, $N_u = 924$ (10.0%), $\hat{\xi} = 0.158$, $\hat{\beta} = 0.769$, $\hat{\beta}/\hat{\xi} = 4.860$
- **VaR 1%**: $np/N_u = 0.100$; $0.100^{-\hat{\xi}} = 1.439$
  - $\mathrm{VaR}_{1\%} = 1.163 + 4.860\,(1.439 - 1) = 3.30\%$
- **ES 2.5%**: $np/N_u = 0.250$, $\mathrm{VaR}_{2.5\%} = 1.163 + 4.860\,(1.245 - 1) = 2.35\%$
  - $\mathrm{ES}_{2.5\%} = \dfrac{2.35 + 0.769 - 0.158 \times 1.163}{1 - 0.158} = 3.49\%$
- **VaR 0.1%**: $np/N_u = 0.010$, $0.010^{-\hat{\xi}} = 2.072$, $\mathrm{VaR}_{0.1\%} = 6.37\%$

VaR 1% and ES 2.5%: Three Methods

	VaR 1% (%)			ES 2.5% (%)		
	EVT	historical	Normal	EVT	historical	Normal
S&P 500	3.30	3.15	2.60	3.49	3.48	2.62
DAX	3.95	4.02	3.16	4.13	4.15	3.18
BET	4.44	4.33	3.39	4.81	4.80	3.41
Bitcoin	10.38	10.51	7.99	10.90	10.84	8.03

- At these levels EVT and historical simulation differ by at most 0.14 percentage points: there are enough data at 1% and 2.5%
- The Normal law is too low
  - the EVT VaR 1% is 27% higher for the S&P 500 and 31% higher for the BET
  - ES 2.5% by a similar margin

SFM_ch5_evt_var

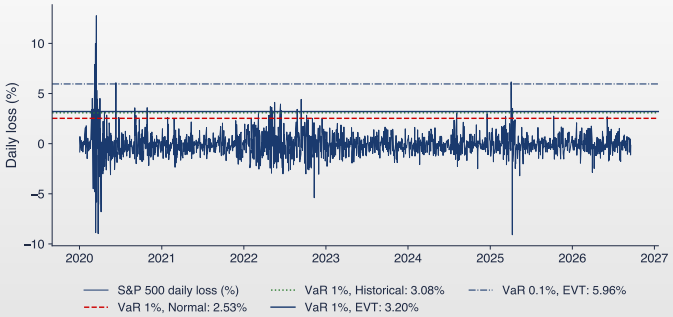
VaR at the 0.1% Level

	EVT (POT)	Hill	historical	Normal	largest loss
S&P 500	6.37	6.92	6.94	3.47	12.8
DAX	7.17	8.37	6.97	4.21	13.1
BET	9.56	10.69	9.64	4.52	13.1
Bitcoin	19.61	23.30	18.05	10.65	46.5
OMV Petrom (SNP)	8.89	9.29	8.32	4.94	15.8
BRD	8.39	8.93	7.85	5.07	18.2
Banca Transilvania (TLV)	10.68	12.10	10.50	5.35	22.2
Transgaz (TGN)	8.29	9.18	8.23	4.68	10.1

- VaR 0.1% (%): the loss exceeded on one day in 1 000; BVB stocks since 2010
- EVT is about 1.8 times the Normal value for the S&P 500 and 2.1 times for the BET
- Historical VaR 0.1% rests on the 3–9 largest losses only; EVT and Hill smooth them with a model of the tail

SFM_ch5_evt_var

Out of Sample: Estimate until 2019, Test 2020–2026



- S&P 500 daily losses after 2019 with the VaR 1% of the three methods, all estimated on 1990–2019
- A good VaR 1% is exceeded on about 1% of the test days, here 16.9 days

SFM_ch5_evt_var

Counting the Exceedances

	test days	VaR 1% expected	VaR 1% EVT	VaR 1% hist.	VaR 1% Normal	VaR 0.1% expected	VaR 0.1% EVT	VaR 0.1% hist.	VaR 0.1% Normal
S&P 500	1 687	16.9	25	26	44	1.7	5	3	23
DAX	1 711	17.1	15	12	31	1.7	2	2	8
BET	1 675	16.8	4	5	13	1.7	0	1	4
Bitcoin	2 453	24.5	8	10	23	2.5	1	1	8

- The Normal VaR is exceeded far too often, especially at 0.1%: 23 days instead of 1.7 for the S&P 500
- EVT and historical (hist.) VaR are close to each other; S&P 500 above the target (the 2020 cluster), BET below it (calm years)
- A VaR estimated once ignores the changing volatility: conditional EVT (GARCH + POT) McNeil & Frey (2000), Chapters 9 and 10

SFM_ch5_evt_var

Strengths and Limitations of EVT

- **Can**: give a principled model of the tail and extrapolate beyond the largest observed loss
  - stable VaR and ES at 0.1% where historical simulation has only a handful of points
- **Cannot**: remove the uncertainty of the far tail
  - $\hat{\xi}$ has a standard error of 0.04–0.06 here; return levels beyond the sample have wide intervals
- **Assumes**: i.i.d. losses, or at least a stable tail over time
  - apply EVT to returns divided by a GARCH volatility (Chapter 9) for daily risk; see also entropy-based ES Pele, Lazar & Mazurencu-Marinescu-Pele (2019)
- Surveys of EVT in finance: Rocco (2014); McNeil, Frey & Embrechts (2015), Ch. 5

An Open Research Question

- **Has the tail of BET losses changed as the Bucharest market grew?**
  - Hill index, $k = 2.5\%$ of n, 1997–2009: $\hat{\alpha} = 2.90$ (block SE 0.33)
  - 2010–2026: $\hat{\alpha} = 2.16$ (block SE 0.27)
  - $z = (\hat{\alpha}_1 - \hat{\alpha}_2)/\sqrt{\mathrm{SE}_1^2 + \mathrm{SE}_2^2} = 1.73$: a heavier tail after 2010, but not significant at 5% ($|z| < 1.96$)
- Why it is open: the threshold fell from 3.95% to 2.01% with the volatility
  - is the change in the tail or in the volatility?
- AI tools can speed up such a study; they do not replace checking it Wang et al. (2023)

SFM_ch5_hill

Possible Uses of AI Tools

- **Literature**: list studies of tail indices in emerging markets and summarise their methods
- **Code**: a first draft of a Hill and POT analysis on rolling windows, with block-bootstrap intervals
- **Robustness**: other thresholds, GARCH-standardised returns, other BVB indices
- Example prompt
  - `Write a Python function that computes the Hill tail index of daily BET losses with k = 2.5% of n on rolling 5-year windows, with moving-block bootstrap standard errors (blocks of 20 days).`

What to Check

- Sign: losses are $-r$; a Hill estimate on returns instead of losses measures the gains
- Parameterisation: scipy's `genextreme` uses $c = -\xi$; `genpareto` uses $c = \xi$
- The rule for k or u is fixed before seeing the result, not tuned to get a significant change
- Rolling windows overlap: their estimates are not independent tests
- References: every cited paper must exist; check the DOI

Project Idea

- **Question**: is the tail of BET losses lighter or heavier today than before 2010, once volatility is removed?
  - data: BET daily closes since 1997 (EODHD); optionally BET-TR since 2014
- Steps
  - Hill and POT (u at the 90% quantile) for the two periods, with block-bootstrap intervals
  - repeat on returns divided by a rolling 60-day volatility: does the difference survive?
  - compare with the S&P 500 over the same periods
- 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 Takeaways

- Daily losses have power-law tails with index $\alpha \approx 3$: finite variance, infinite kurtosis
- Hill: the tail index from the k largest losses; choose k on the Hill plot; block-bootstrap SE
- Block maxima $\to$ GEV; excesses over a threshold $\to$ GPD; the same shape $\xi = 1/\alpha$
- Check every fit with mean excess, parameter stability, PP and QQ plots
- EVT VaR and ES: close to historical at 1% and 2.5%, far above the Normal values at 0.1%
- Extrapolation far beyond the data remains uncertain: report intervals

Key Formulas

Quantity	Formula
Regular variation	$P(L > x) = x^{-\alpha}\ell(x), \quad P(L > tx)/P(L > x) \to t^{-\alpha}$
Hill	$\hat{\alpha}_k = [\frac{1}{k}\sum_{i=1}^{k} \ln(L_{(i)}/L_{(k+1)})]^{-1}, \quad \mathrm{SE} \approx \hat{\alpha}/\sqrt{k}$
Hill quantile	$\hat{x}_p = L_{(k+1)}(k/(np))^{1/\hat{\alpha}}$
Mean excess	$e(u) = E[L - u \mid L > u]; \quad$ GPD: $(\beta + \xi(v-u))/(1-\xi)$
GEV	$H(x) = \exp\{-[1 + \xi(x-\mu)/\sigma]^{-1/\xi}\}$
Return level	$z_T = \mu + \frac{\sigma}{\xi}[(-\ln(1 - \frac{1}{127}))^{-\xi} - 1]$
GPD	$G(y) = 1 - (1 + \xi y/\beta)^{-1/\xi}$
VaR (POT)	$u + \frac{\beta}{\xi}[(np/N_u)^{-\xi} - 1]$
ES (POT)	$(\mathrm{VaR}_p + \beta - \xi u)/(1-\xi)$

Self-Assessment

- **Question**: a GPD fit gives $\hat{\xi} = 0.25$; which moments of the losses exist?
  - **Answer**: $\alpha = 1/\xi = 4$: mean, variance and skewness exist; the kurtosis does not ($m < 4$ only)
- **Question**: why use the 2.5% largest losses and not all losses for the Hill estimator?
  - **Answer**: the power law holds only in the tail; with the body included $\hat{\alpha}$ is biased
- **Question**: is VaR 0.1% by historical simulation reliable with 10 years of data?
  - **Answer**: no: about 2 500 days leave only 2–3 losses beyond it; EVT uses the whole tail
- Next: Chapter 6, model selection and risk management

References (1/2)

- Acerbi, C., Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7), 1487–1503.
- Balkema, A.A., de Haan, L. (1974). Residual life time at great age. *Annals of Probability*, 2(5), 792–804.
- Basel Committee on Banking Supervision. MAR33: Internal models approach – capital requirements calculation. *The Basel Framework*. Bank for International Settlements.
- Borak, S., Härdle, W.K., López-Cabrera, B. (2013). *Statistics of Financial Markets: Exercises and Solutions*, 2nd ed. Springer.
- Carlson, M. (2007). A brief history of the 1987 stock market crash with a discussion of the Federal Reserve response. *Finance and Economics Discussion Series* 2007-13, Board of Governors of the Federal Reserve System.
- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236.
- Daníelsson, J., de Vries, C.G. (1997). Tail index and quantile estimation with very high frequency data. *Journal of Empirical Finance*, 4(2–3), 241–257.
- Davison, A.C., Smith, R.L. (1990). Models for exceedances over high thresholds. *Journal of the Royal Statistical Society, Series B*, 52(3), 393–425.
- Drees, H., Resnick, S., de Haan, L. (2000). How to make a Hill plot. *Annals of Statistics*, 28(1), 254–274.
- Embrechts, P., Klüppelberg, C., Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer.
- Fisher, R.A., Tippett, L.H.C. (1928). Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society*, 24(2), 180–190.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). *Statistics of Financial Markets: An Introduction*, 5th ed. Springer.
- Gnedenko, B. (1943). Sur la distribution limite du terme maximum d'une série aléatoire. *Annals of Mathematics*, 44(3), 423–453.
- 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.

References (2/2)

- Gumbel, E.J. (1958). *Statistics of Extremes*. Columbia University Press.
- Hill, B.M. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics*, 3(5), 1163–1174.
- Jansen, D.W., de Vries, C.G. (1991). On the frequency of large stock returns: putting booms and busts into perspective. *Review of Economics and Statistics*, 73(1), 18–24.
- Jenkinson, A.F. (1955). The frequency distribution of the annual maximum (or minimum) values of meteorological elements. *Quarterly Journal of the Royal Meteorological Society*, 81(348), 158–171.
- Longin, F.M. (1996). The asymptotic distribution of extreme stock market returns. *Journal of Business*, 69(3), 383–408.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36(4), 394–419.
- McNeil, A.J., Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7(3–4), 271–300.
- McNeil, A.J., Frey, R., Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools*, revised ed. Princeton University Press.
- Pele, D.T., Lazar, E., Dufour, A. (2017). Information entropy and measures of market risk. *Entropy*, 19(5), 226.
- Pele, D.T., Lazar, E., Mazurencu-Marinescu-Pele, M. (2019). Modeling expected shortfall using tail entropy. *Entropy*, 21(12), 1204.
- Pickands, J. (1975). Statistical inference using extreme order statistics. *Annals of Statistics*, 3(1), 119–131.
- Resnick, S.I. (2007). *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer.
- Rocco, M. (2014). Extreme value theory in finance: a survey. *Journal of Economic Surveys*, 28(1), 82–108.
- Wang, H., Fu, T., Du, Y., Gao, W., et al. (2023). Scientific discovery in the age of artificial intelligence. *Nature*, 620, 47–60.
- Weissman, I. (1978). Estimation of parameters and large quantiles based on the k largest observations. *Journal of the American Statistical Association*, 73(364), 812–815.