Asset Allocation Workshop

Malcolm Kemp

Managing Director, Nematrian Limited and Adjunct Professor, Imperial College Business School, London

Global Markets Training, 29 – 30 November 2012

Underlying thesis

- Taking due account of extreme events is a key issue for (many) market professionals
  - They are a fact of life, can severely disrupt but also provide opportunities for gain
- No amount of clever data analysis can circumvent the need for expert judgement and subjective practitioner input
  - But ignoring mathematical and market insights isn't the right answer
- The limiting case where fat-tailed behaviour is insignificant covers most strands of traditional portfolio construction theory
  - A 2-for-1 opportunity: cover all main strands of basic mean-variance investing, risk budgeting, etc. *as well as* refinements needed to handle extreme events

Extreme Events: Robust Portfolio Construction in the Presence of Fat Tails

- Chapters:

  1. Introduction
  2. Fat tails – in single (i.e. univariate) return series
  3. Fat tails – in joint (i.e. multivariate) return series
  4. Identifying factors that significantly influence markets
  5. Traditional portfolio construction techniques
  6. Robust mean-variance portfolio construction
  7. Regime switching and time-varying risk and return parameters
  8. Stress testing
  9. Really extreme events

  Plus Principles (Chapter 10) and Exercises (Appendix). Each chapter also includes specific sections covering *practitioner perspectives* and *implementation challenges*.

- Book provides practitioners and students with all main recipes (plus author views on them) with a toolkit provided through www.nematrian.com

Selected highlights from "Extreme Events"

Modelling fat tails for *individual* risks

- ‘Fat-tailed’ means probability of extreme-sized outcomes seems to be higher than if coming from (usually) a (log) Normal distribution
- There are various ways of visualising fat tails in a *single* return distribution. Easiest to see in format (c) below, i.e. QQ-plots
- Note: portfolio construction usually involves *multiple* assets / risk exposures

(a) probability density function

(b) cumulative distribution function

(c) quantile-quantile (QQ) plot

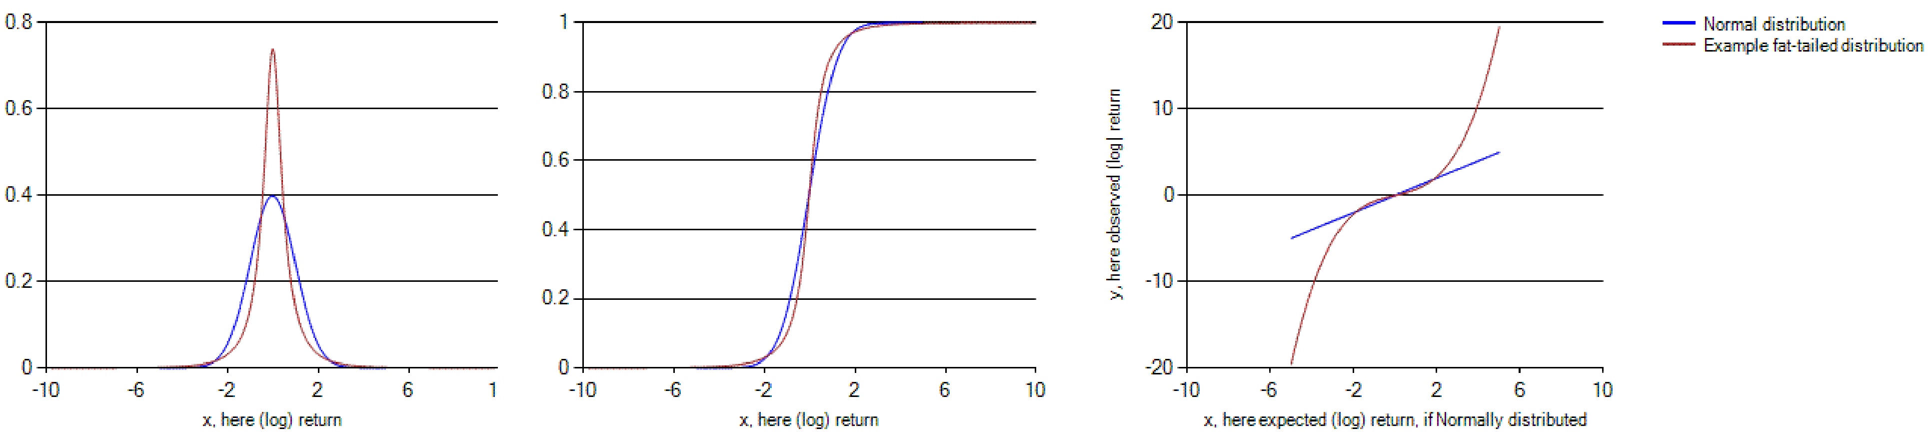

Nematrian

Source: www.nematrian.com

Many (most?) investment return series are 'fat-tailed'

- Some instrument types intrinsically skewed (e.g. high-grade bonds, options)
- Others (e.g. equities) still exhibit fat-tails, particularly higher frequency data

(1) Monthly returns

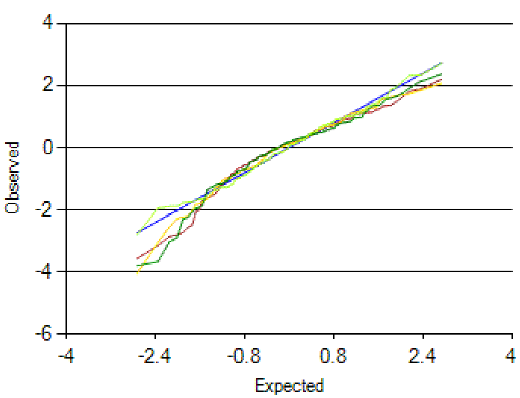

(2) Weekly returns

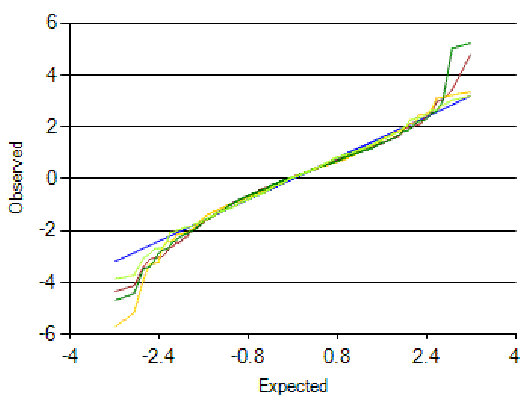

(3) Daily returns

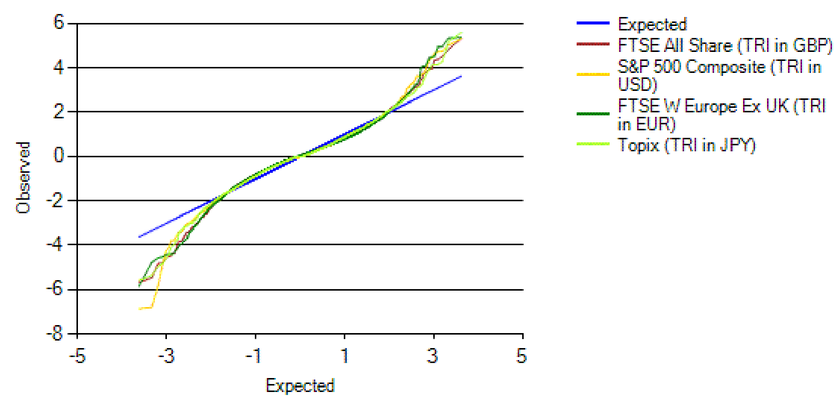

Source: www.nematrian.com, Threadneedle, S&P, FTSE, Thomson Datastream

Returns from end June 1994 to end Dec 2007, charts show standardised logged returns

Nematrian

Why are return series often fat-tailed?

- Time-varying nature of the world in which we live
  - Market / sector / instrument volatility (and maybe other distributional characteristics) change through time
- Crowded trades and leverage
- Selection effects, e.g. manager behaviour may (consciously or unconsciously) bias towards fat-tailed behaviour
- As well as intrinsically skewed behaviour such as for individual bonds

Explains some equity index fat fails, particularly upside

Raw Data

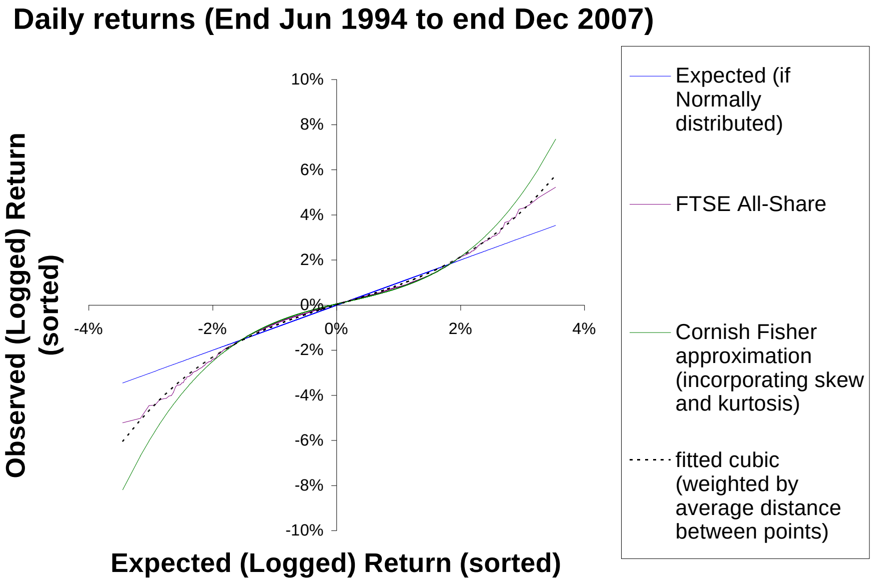

With Short-term Volatility Adjustment

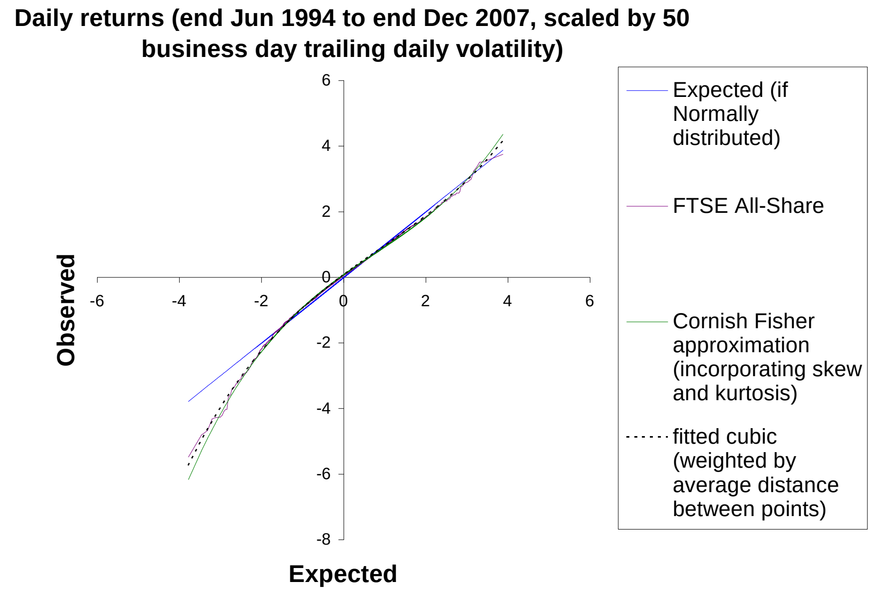

Average extent to which tail exceeds expected level (average of 6 most extreme outcomes)				
	Downside (%)		Upside (%)	
	Unadj	Adj for vol	Unadj	Adj for vol
FTSE All-Share (in GBP)	54	41	42	3
S&P 500 (in USD)	68	70	50	7
FTSE Eur ex UK (in EUR)	48	53	54	-3
Topix (in JPY)	54	72	42	39

Source:

Threadneedle, FTSE

Thomson Datastream

And over longer time periods

Raw Data

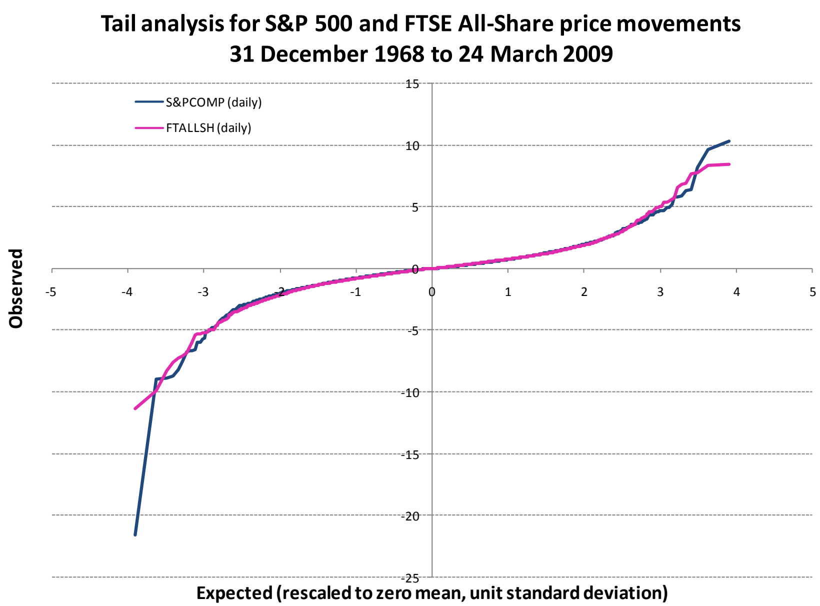

With Short-term Volatility Adjustment

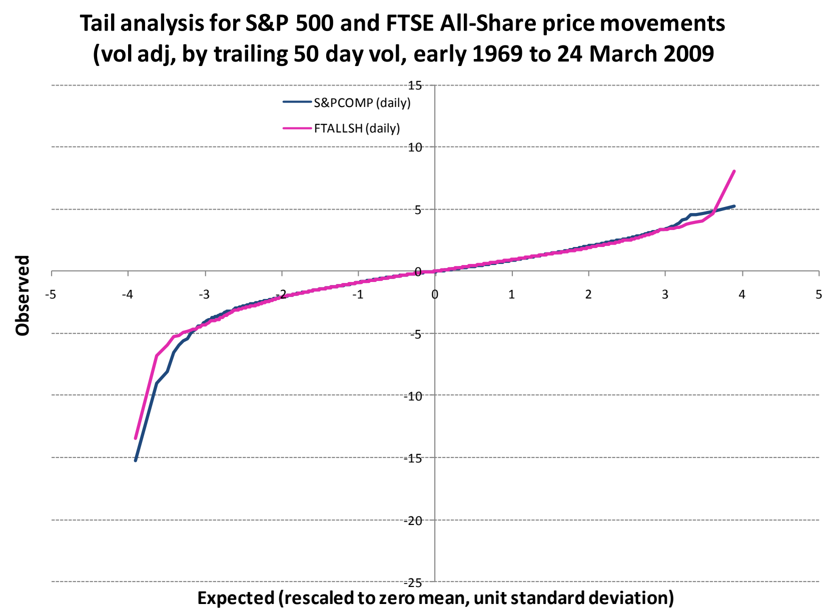

Source: Threadneedle, S&P, FTSE, Thomson Datastream

Crowded trades and leverage

- Some fat tails still seem to come "out of the blue"
  - E.g. Quant funds in August 2007
  - Too many investors in the same *crowded trades*? Behavioural finance implies potentially unstable
  - For less liquid investments, impact may be via an apparent shift in price basis
- System-wide equivalents via leverage?
  - Leverage introduces/magnifies *liquidity* risk, *forced unwind* risk and *variable borrow cost* risk
  - Like selection, involves behavioural finance effects

Extreme Value Theory (EVT)

- Potentially relevant to risk management (and pricing)
  - Capital adequacy seeks to protect against (we hope) relatively rare events
  - Pricing often dominated by potential magnitude and likelihood of large losses, which are (we hope) rare
  - EVT appears to offer a convenient way of identifying shape of the ‘tail’ distribution, which should be very valuable for such purposes
- But bear in mind
  - Possibility (indeed probability) that the world is not time stationary
  - Inherent unreliability of extrapolation – including extrapolation into the tails of a probability distribution

Potential weaknesses

- EVT seems very helpful
  - Characterises limiting distributions very succinctly
  - But required regularity conditions are potentially strong
  - Relies on existence of a limiting distribution but this is not guaranteed
- At issue is potential unreliability of extrapolation

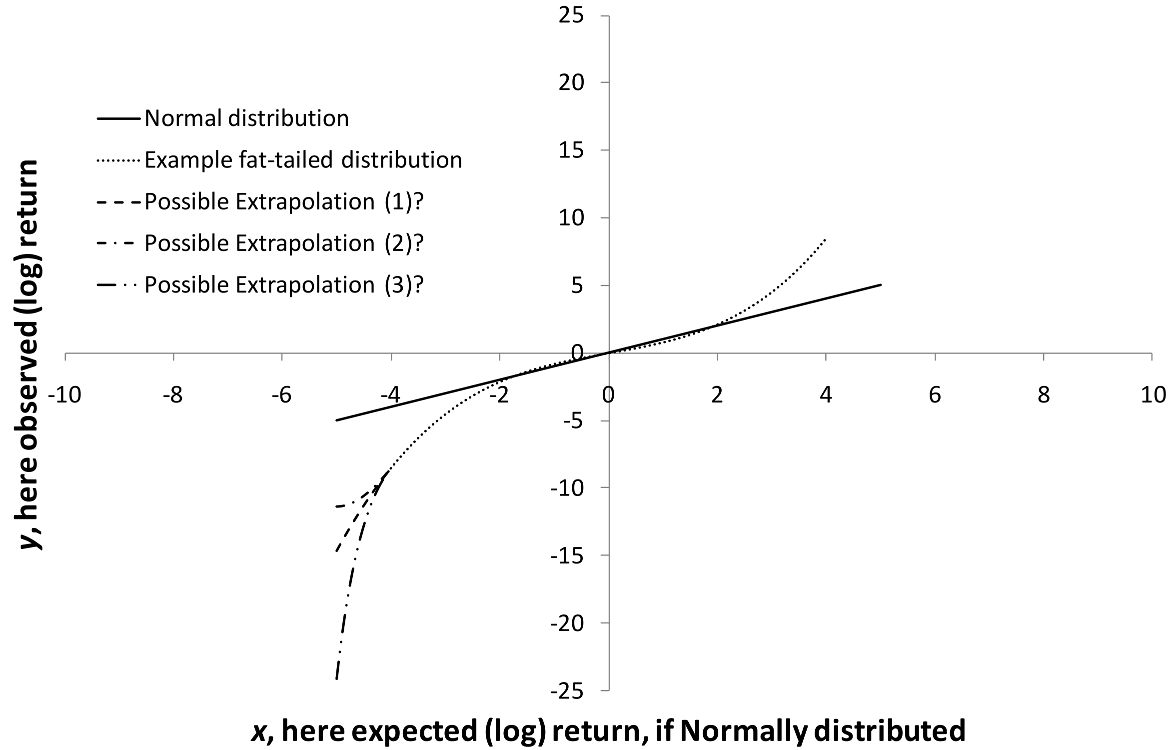

Source: Nematrian

Selected highlights from “Extreme Events”

Joint fat-tailed behaviour

- Crucial to the portfolio construction problem
- Can split the probability distribution into two components:
  a) Marginals (i.e. distributions of each individual risk in isolation); and
  b) Copula (i.e. the remainder, the ‘co-dependency’ between risks)
- However
  - Fat-tailed characteristics then difficult to visualise
  - Copulas are akin to (indeed are) cumulative distribution functions
  - Many problems depend on a) and b) in tandem

Copulas: a well trodden (mathematical) path

probability density function

copula density

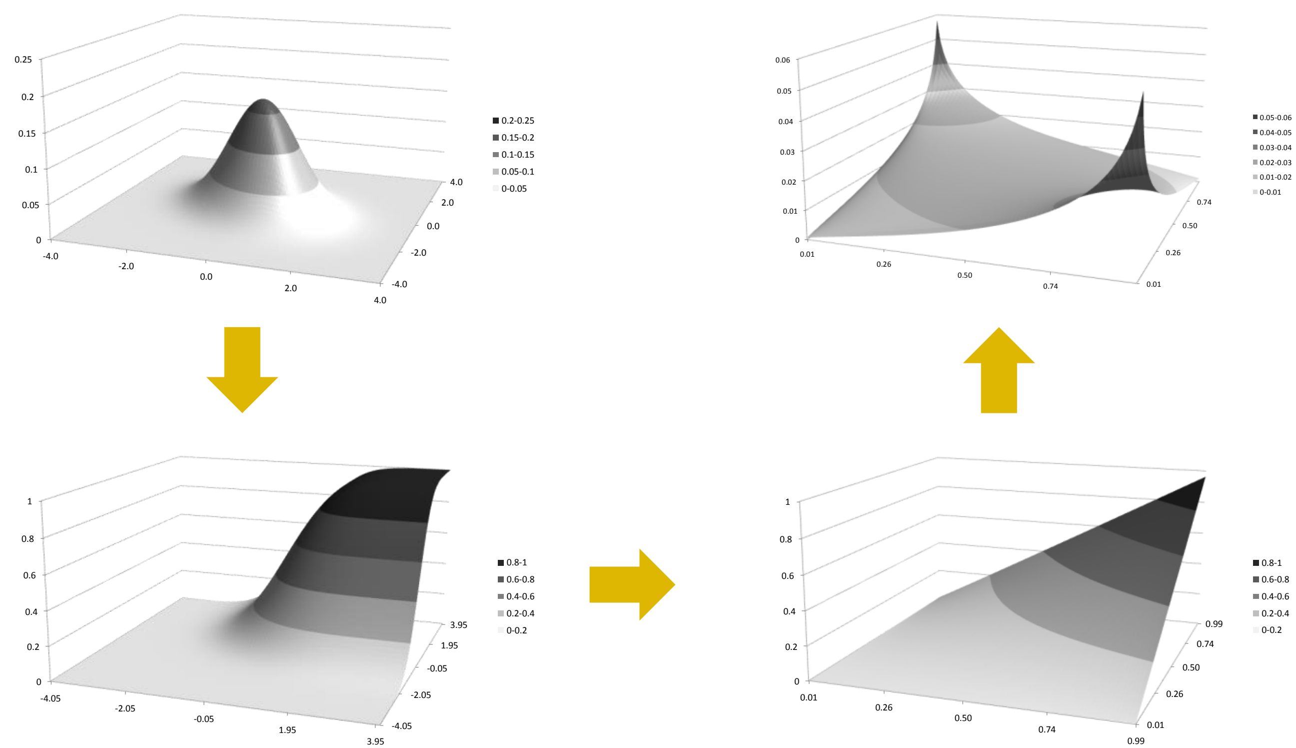

cumulative distribution function

copula

- The copula involves rescaling (stretching/squashing) each axis so that the distribution is uniform between 0 and 1 along each axis
  - Allows models to exhibit non-zero tail dependency (i.e. ‘correlation’ in tail)

But extreme behaviour shows up better in QQ plots

- Book suggests how to refine quantile-quantile plots to show joint (here 2 risk) extreme behaviour
  - Uses 'upwards' QQ plots (right half corresponds to $X_{(i)}$, left half to $-X_{(i)}$)
  - Then create surface plot that encapsulates upwards QQ plots for all (linear) combinations of X and Y, relevant combination given by angle of rotation around centre
  - Encapsulates in a single chart fat-tailed behaviour arising from co-dependency characteristics *and* marginal distributions
  - Like a one-dimensional QQ plot, places greater visual emphasis on extreme events

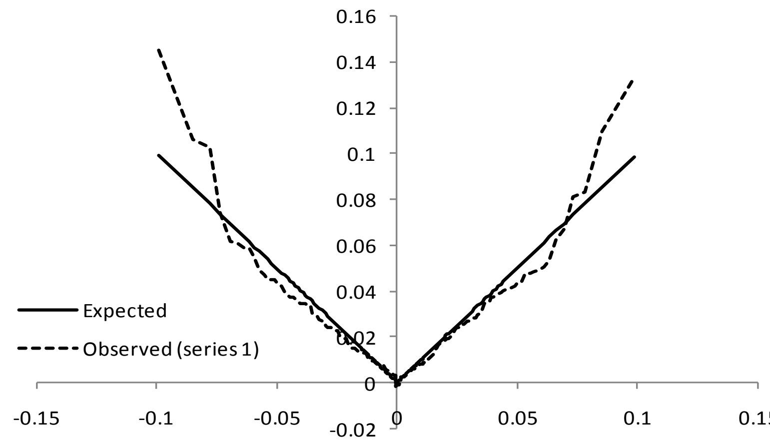

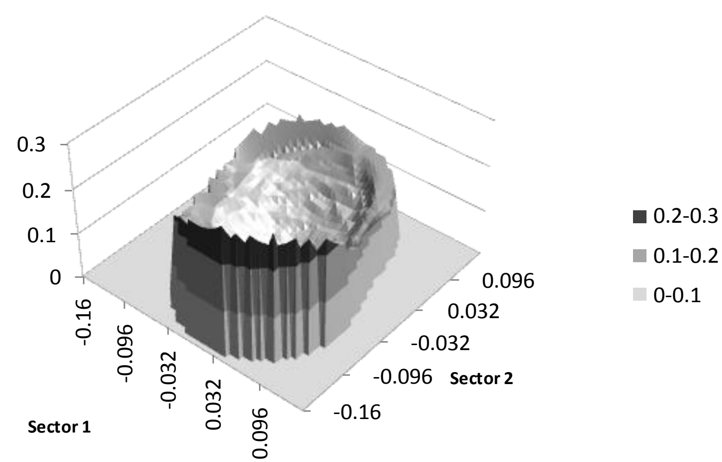

Nematrian

Selected highlights from "Extreme Events"

Factor identification

- Copulas are rather complicated mathematically. Typically simpler correlation based aggregation techniques are used instead
- In a portfolio construction context involves a factor-based model of the world
  - Vastly reduces number of parameters that need estimating (if large universe)
- An entire risk model vendor industry focuses on how to create such models, involving one or more of the following:
  - Fundamental risk models
  - Econometric risk models
  - Statistical risk models

Fundamental data limitations

- Suppose we have *N* instruments and estimate the factor structure from *T* observations per instrument where *T* much less than *N* (e.g. as would normally be the case for a whole market model)
  - Then at most *T*-1 non-zero factors and random matrix theory (RMT) suggests most of the smaller ones often indistinguishable from ones that would arise randomly
  - Places fundamental limits on reliability of factor analysis (or any other risk modelling derived from historic return series)
- Means that fine structure of an optimised portfolio inherently depends on practitioner's (or model creator's) subjective views

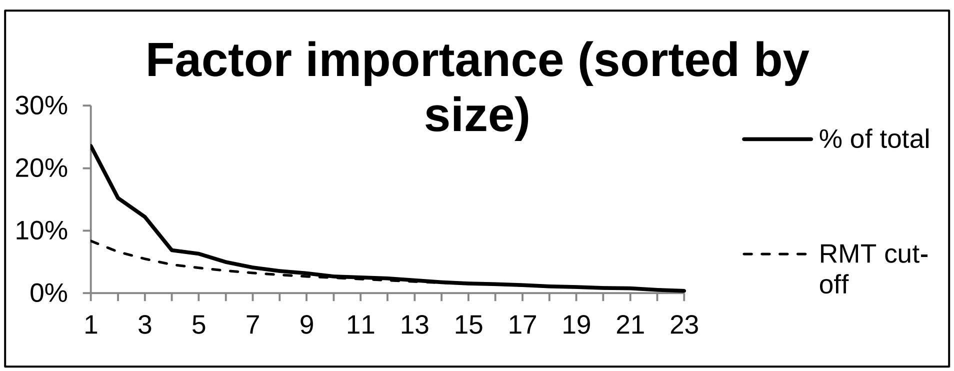

Selection effects

- 'Selection' effects are common in finance, e.g. annuitants typically have higher than average life expectancies
- Can also apply to portfolios being analysed by risk models
  - Many risk models assume behaviour that is (approximately) Gaussian, i.e. multivariate (log) Normal, akin to lots of different sources of random noise
  - Can decompose multiple series return data into *principal components*, the most important of which contribute the most to the aggregate variability exhibited by securities in the relevant universe
- But what if portfolios are structured to seek 'meaning' (e.g. if actively managed!) and 'meaning' is (partly) associated with non-Normality?
  - Both *meaning* and *magnitude* are important

PCA vs. ICA

- Both are examples of ‘blind source separation’, aiming to identify ‘signals’ (i.e. sources / factors) that explain (observed) market behaviour
- Principal Components Analysis (PCA)
  - Seeks to identify the largest contributors to variance, i.e. magnitude of impact
  - ‘Signals’ maximise sum of variances of returns of each security within universe
- Independent Components Analysis (ICA)
  - Seeks to identify contributors to market behaviour that are meaningful
  - ‘Signals’ maximise independence, non-Normality and/or complexity

Selection effects are potentially very important

	PCA, only StDev			*Blend* (*c* = 0.39)			*ICA*, Only Kurtosis		
Component	StDev	Kurt	CF4 est	StDev	Kurt	CF4 est	StDev	Kurt	CF4 est
1st	10.6%	3.1	23%	8.3%	14.9	57%	4.5%	24.2	47%
Av (top 6)	5.9%	1.6	**10%**	5.3%	16.8	**40%**	4.7%	18.5	**40%**

- Assume Cornish-Fisher (CF4) is a good estimator of '1 in 200' risk level
  - Risk then proportional to *std* . (1 + *c* . *kurt*) where *c* = 0.39 at '1 in 200' level
  - *Principal Components Analysis* (PCA) focuses just on standard deviation, *Independent Components Analysis* (ICA) just on kurtosis, blend on both
  - Analysis based on monthly MSCI World sector relative returns Jun 1996 to Feb 2009
- Sizes of '1 in 200' events potentially underestimated several-fold by PCA, if portfolio built on the basis of 'meaning' (e.g. if actively managed)

Selection effects: do they occur in practice?

- Banks that failed during 2007-09 Credit Crisis were disproportionately biased towards strategies that depended on continuing favourable liquidity conditions
- Liquidity risk is highly skewed, i.e. highly fat-tailed
- I.e. these banks were (consciously or unconsciously) biasing their business strategies towards ones that had fat-tailed characteristics
- No wonder traditional risk models appear to have underestimated potential magnitudes of adverse outcomes!

Selected highlights from “Extreme Events”

Portfolio construction

- Traditional (quantitative) approach involves portfolio optimisation
  - Identify expected return ('alpha') from each position
  - Maximise expected return for a given level of risk (subject to constraints, e.g. weights sum to unity)
  - Typically focus on mean-variance optimisation

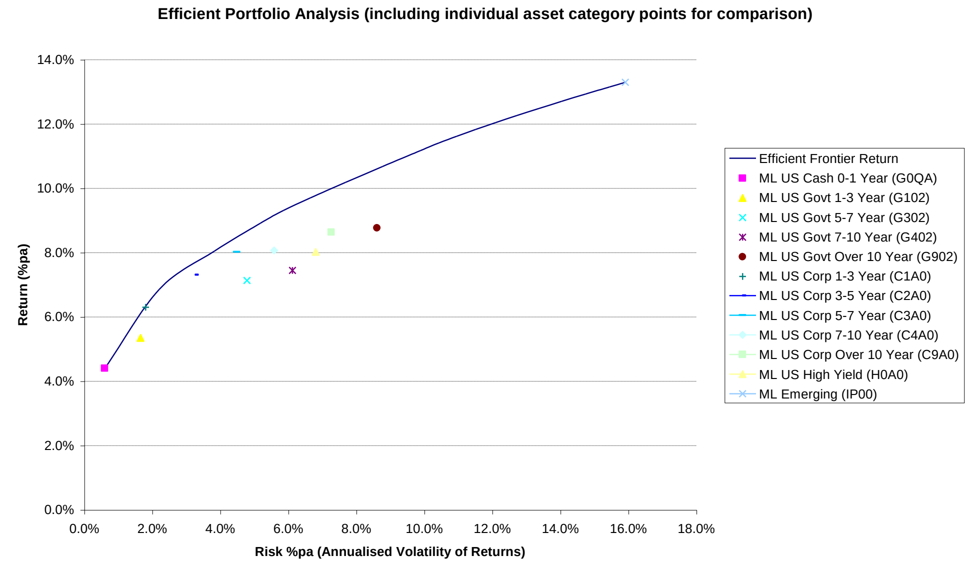

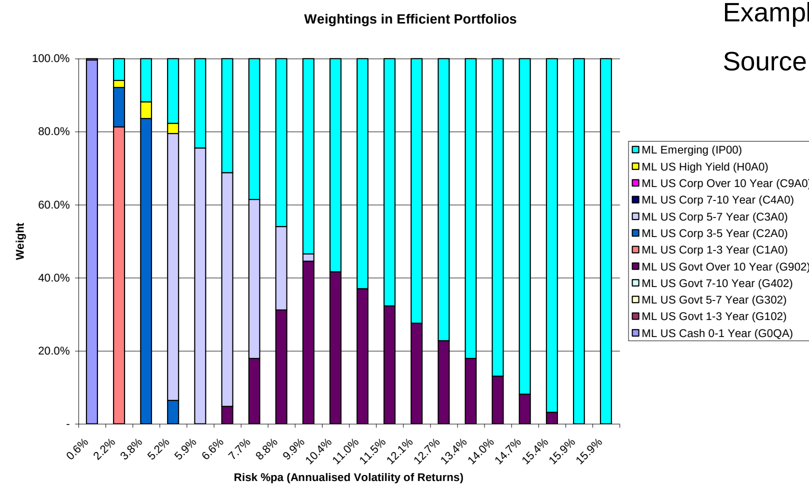

Example for illustration only

Source: Threadneedle

Portfolio construction: sensitivities

- Output results are notoriously sensitive to input assumptions. Possible responses include:
  - Treat quant models with scepticism (the fundamental manager's approach?)
  - Focus on reverse optimisation
- Book covers all the main quantitative refinements, including:
  - Robust approaches and Bayesian priors/anchors, e.g. Black-Litterman
  - Shrinkage
  - Resampled optimisation
- And ties them back to earlier chapters
  - E.g. how resampled optimisation doesn't avoid 'fine structure' problem, instead it just inherits it from the dataset being used for bootstrapping purposes

Impact of fat tails (1)

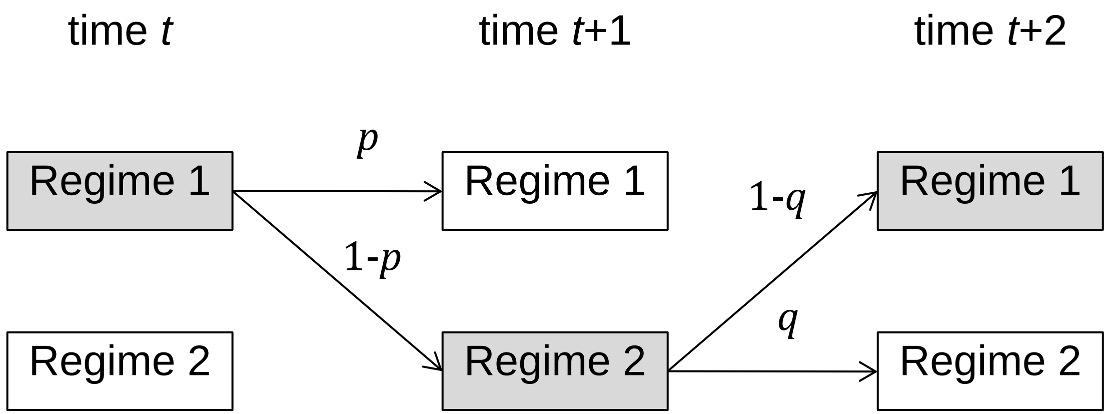

- Regime-switching covered in some detail in book
- Builds on premise that a high proportion of fat-tailed behaviour observed in practice derives from time-varying nature of the world in which we live

Impact of fat tails (2)

- Most important (predictable) single contributor to fat tails seems to be time-varying volatility. So:
  - Calculate covariance matrix between return series after stripping out effect of time-varying volatility?
  - Optimise as you think fit (standard, “robust”, Bayesian, BL, ...), using adjusted covariance matrix
  - Adjust risk aversion/risk budget appropriately
  - Then unravel time-varying volatility adjustment
  - Or derive implied alphas using same adjusted covariance matrix
- Implicitly assumes all adjusted return series ‘equally’ fat-tailed

Other approaches

- Model with a (time-stationary?) mixture of multivariate Normal distributions, see e.g. Scherer (2007)
  - If time-varying then involves regime switching
  - Even more difficult to estimate reliably
- Or use lower partial moments
  - But also challenging to estimate reliably

Selected highlights from “Extreme Events”

Stress testing

- Even a model with good back-test properties may fail to model future risks effectively
  - Because the past is not necessarily a good guide to the future. Markets are not like precisely defined external physical systems whose actions are perfectly predictable.
  - **Magnitude** of impact of a given scenario usually (relatively) easy to calculate
  - The challenge is how to identify the scenario's **likelihood**
- Stress testing focuses **more on magnitude**, and what makes the scenario adverse, and pays **less attention to likelihood**

Observations

- There are several types of stress test each with different nuances:
  - Analysis of impact on portfolio (or firm) of movements in specific market drivers, within an envelope of the plausible range of outcomes
  - Specific industry-wide stresses mandated by the regulator (e.g. used in capital computations)
  - A focus on configurations of events that might lead to large losses (e.g. reverse stress testing)
- Some commentators encourage us to revert to a more statistical emphasis when choosing stress tests
- Challenge: **portfolio construction must ultimately trade off risk versus reward, so needs to include some link back to likelihood of occurrence**

Really extreme events

- We need to promote the right **mind-set** as there is insufficient data to allow rigorous mathematical analysis
  - Think outside the box (c.f. reverse stress testing)
  - Accept that (Knightian) uncertainty is a fact of life
    - But bear in mind that market behaviour can inform us about how markets react to intrinsic uncertainty and that there is typically a premium for flexibility
  - Be particularly aware of exposures that are sensitive to aggregate ‘market risk appetite’, including liquidity risk and other selection-sensitive risk types
  - Don’t undervalue good governance and operational management (including ERM principles)
  - Remember markets are driven in part by (hard to quantify) behavioural factors

Extreme Events: Robust Portfolio Construction in the Presence of Fat Tails

- Chapters:

  1. Introduction
  2. Fat tails – in single (i.e. univariate) return series
  3. Fat tails – in joint (i.e. multivariate) return series
  4. Identifying factors that significantly influence markets
  5. Traditional portfolio construction techniques
  6. Robust mean-variance portfolio construction
  7. Regime switching and time-varying risk and return parameters
  8. Stress testing
  9. Really extreme events

  Plus Principles (Chapter 10) and Exercises (Appendix). Each chapter also includes specific sections covering *practitioner perspectives* and *implementation challenges*.

- Book provides practitioners and students with all main recipes (plus author views on them) with a toolkit provided through www.nematrian.com

Important Information