## Coping with Extreme Events

Malcolm Kemp
Head of Quantitative Research
Threadneedle Asset Management

15-17 JUNE 2008
HILTON DEANSGATE, MANCHESTER



## Agenda

- Introduction
  - What are fat tails and why are we interested in them?
  - What impact ought they to have on portfolio construction?
- Fat-tailed behaviour in individual return series
  - Behaviour of individual (equity) index series and of high conviction portfolios
  - Modelling fat tails – why focusing on skew and kurtosis may not be the best approach
  - Possible causes of fat tails
- Fat-tailed behaviour in multiple (joint) return series
  - Copulas and other co-movement analysis tools
  - Analysing fat-tailed behaviour more directly
- Portfolio construction
  - Traditional tools
  - Refinements designed to cater better for fat-tailed behaviour



## What are fat tails and why are we interested in them?

- Extreme events / outcomes seem to occur more often than would be the case if the return series were coming from a (log) Normal distribution
  - Typical academic theory revolves around Normality (law of large numbers, mean-variance optimisation etc.) but this doesn't match observed behaviour
  - Extreme events can materially disrupt (or, on the upside, materially benefit) portfolio progress
  - Natural for risk managers to consider (e.g. for stress testing). Also important for portfolio managers.
- There are various ways of characterising fat tails in a return distribution
  - Easiest to see in (c) below, so this is the format we concentrate on in the rest of the presentation

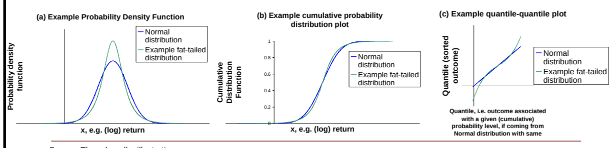


Source: Threadneedle, illustrative

## What impact ought they to have on portfolio construction?

- Clients want *good performance at an acceptable level of risk*
  - i.e. efficient use of the available risk budget
- Conceptually involves:
  - *Choosing the right level of risk to run* (i.e. the risk budget, and
  - *Constructing a portfolio* to deliver versus this budget
- If *all* opportunities (and combinations) '*equally*' (jointly) fat-tailed
  - Answers the same as using traditional mean-variance optimisation approaches
  - With the risk budget adjusted accordingly
- If *different* combinations exhibit *differential* fat-tailed behaviour
  - Portfolio construction ought in principle to change
  - If you can reliably estimate these differentials
  - And if investors do not have solely quadratic utility functions

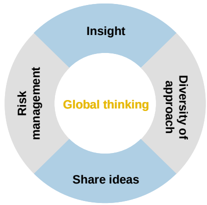


out-think out-perform

The Actuarial Profession
making financial sense of the future

## Fat-tailed behaviour in individual return series

- For some markets, fat tails are intrinsically to be expected
  - e.g. high grade bonds should default infrequently, but when they do, their price movement is typically large
- Even when not intrinsically to be expected, they seem to appear anyway!
  - Although their extent may vary according to timescale
  - E.g. Monthly, weekly and daily returns for major equity market indices (end June 1994 to end Dec 2007

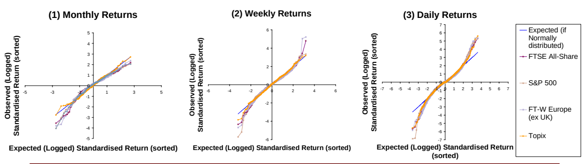


Source: Threadneedle, S&P, FTSE, Thomson Datastream

The Actuarial Profession
making financial sense of the future

## And in high conviction management

- Fat tails are potentially particularly important in the hedge fund arena
- High conviction typically leads to
  - Concentrated portfolios, and/or
  - Portfolios expressing strong thematic exposures
  - Either can make portfolios more susceptible to fat tails

| Moments of daily returns | | |
|---|---|---|
| | FTSE All-Share | Threadneedle Crescendo UK |
| Mean | 0.02% | 0.06% |
| Standard deviation | 1.0% | 0.4% |
| Skew | -0.3 | 0.1 |
| (Excess) Kurtosis | 3.5 | 6.3 |

FTSE All-Share & Threadneedle Crescendo

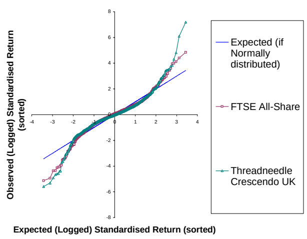


Source: Threadneedle, Thomson Datastream, daily returns June 2001 (inception of fund) to Dec 2007

The Actuarial Profession
making financial sense of the future

## Skew(ness) and kurtosis

- Fat tails involve deviation from Normality
- Hence at least some of the higher moments of the distribution must deviate from Normality
- So one might think that a natural way of assessing fat tailed behaviour is by reference to the skew and (excess) kurtosis of the return distribution. Both are zero for a Normal distribution.

$$mean = \mu = E(x)$$

$$standard\ deviation = \sigma = E\left((x-\mu)^2\right)$$

$$skew = \gamma_1 = E\left(\left(\frac{x-\mu}{\sigma}\right)^3\right)$$

$$(excess)\,kurtosis = \gamma_2 = E\left(\left(\frac{x-\mu}{\sigma}\right)^4\right) - 3$$

- Whilst focusing on skew and kurtosis is the most common way of analysing and coping with fat tails, it may not necessarily be the best approach



## Skew and kurtosis have a natural interpretation via the so-called Cornish-Fisher asymptotic expansion

- Cornish-Fisher (4th moment version) is a way of estimating the shape of the distribution from merely the first 4 moments, i.e. mean, standard deviation, skew and (excess) kurtosis
- Regularly appears in risk management academic literature
- In effect, for standardised returns (zero mean, unit standard deviation), it involves estimating the shape of the quantile plot using the following cubic equation

$$y_{CF4}(x) = x + \frac{\gamma_1(x^2-1)}{6} + \frac{3\gamma_2(x^3-3x) - 2\gamma_1^2(2x^3-5x)}{72}$$

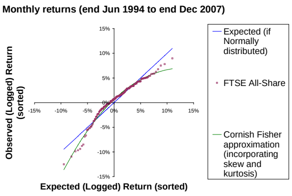


Source: Threadneedle, FTSE, Thomson Datastream



## Flaws in Cornish Fisher (and hence in skew/kurtosis) – (1)

- Cornish-Fisher doesn't actually seem to model index return distributions particularly well
  - Particularly for parts of the distribution in which Risk Managers might be most interested, i.e. the downside tails
- Because computation of skew and kurtosis may not give the right weight to parts of the distribution we are interested in
  - Gives less weight to observations in the tails, because most observations are towards the middle of the distribution

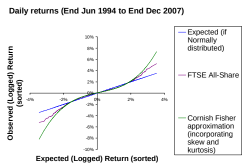


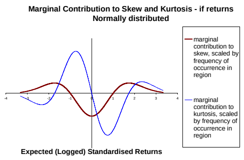


Source: Threadneedle, FTSE, Thomson Datastream

## Flaws in Cornish Fisher (and hence in skew/kurtosis) – (2)

- Cornish-Fisher seems to get worse the more fat-tailed the distribution is
  - Despite its use in the hedge fund community
- Lacks an intrinsically desirable stability property

| Impact of applying CF adjustment to a distribution already derived from the CF approach | | | |
|---|---|---|---|
| Before | | After | |
| skew | kurtosis | skew | kurtosis |
| 0 | 0 | 0 | 0 |
| 0 | 0.3 | 0 | 0.3 |
| 0 | 0.6 | 0 | 0.7 |
| 0 | 1 | 0 | 1.3 |
| 0 | 3 | 0 | 7 |
| 0 | 5 | 0 | 16 |
| -0.4 | 2.0 | -0.6 | 3.4 |

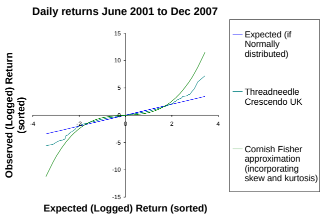


Source: Threadneedle



## A better approach to modelling tail behaviour?

- How about fitting the quantile plot directly, using normal curve-fitting techniques?
  - E.g. with a cubic curve
- Calculation is somewhat more complicated
- Maybe this is why skew and kurtosis are so widely used
- Skew/kurtosis:
  - Do not need data to be ordered
  - Come as pre-canned functions in Microsoft Excel, i.e. SKEW() and KURT()

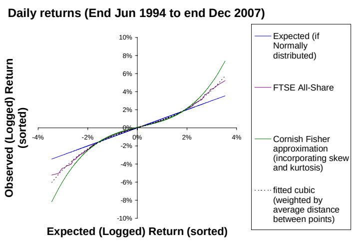


Source: Threadneedle, Thomson Datastream



## One possible source of fat tails is *time-varying* volatility

- Because a *mixture* of Normal distributions with different standard deviations can have fat tails
- Does appear to explain a significant proportion of (upside) volatility for some major market indices
- And potentially some portfolio fat-tailed behaviour

| 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 |

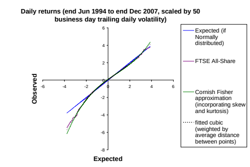


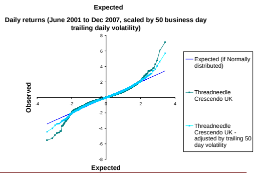


Source: Threadneedle, FTSE, Thomson Datastream

## Another possible source: crowded trades

- Some fat tails appear to be due to crowded trades
- E.g. Quant funds in Aug 2007
- Are the views being adopted similar to lots of other people's views?
  - If they are and there is a stress, will the portfolio be able to ride out the storm
  - Or will it be one of the portfolios that has to unwind at the "wrong" time?
- Is it possible to work out how others are positioned and position accordingly?



## Implications for risk managers and portfolio managers

- For risk managers
  - Skew and kurtosis may be relatively straightforward tools to use to analyse fat tails, but do not seem to be ideal, viz Cornish-Fisher approach
  - Better seems to be to estimate the distributional form directly from observed (sorted) values, although the calculations are more complicated
  - For major western markets, an important source of fat tails (particularly on the upside) seem to be time-varying volatility
  - Conversely, a significant proportion not explained by this effect. Hence merits of stress tests etc.
- For portfolio managers
  - Changing volatility levels can be assessed using implied volatility
  - E.g. from VIX, VDAX, variance swaps, or option prices more generally
  - Or from credit spreads and Merton-style firm models
  - But some fat-tails still seem to come out of the blue. Importance of understanding "crowded trades"?



## Fat-tailed behaviour in joint return series

- One obvious approach is to consider:
  - How fat-tailed each series is in isolation, i.e. each *marginal* distribution, and
  - How they might co-move together, i.e. their (joint) *copula* function
- Any multivariate distribution can be expressed this way via *Sklar's theorem*:
  - Suppose that $X_1$, $X_2$, ..., $X_N$ are random variables
  - With marginal distribution functions, i.e. individual cumulative probability distribution functions, say, $F_1(x_1)$, $F_2(x_2)$, ..., $F_N(x_N)$
  - And a joint distribution function $F(x_1, x_2, ..., x_N)$
  - Then $F$ can be characterised by the $N$ marginal distributions and an $N$-dimensional copula, $C$, i.e. a function that maps a vector of $N$ numbers each between 0 and 1 onto some value in the range 0 to 1, using:

$$F(x_1, x_2, ..., x_N) = C(x_1, x_2, ..., x_N) \times F_1(x_1) \times F(x_2) \times ... \times F(x_N)$$

# Visualising multivariate distributions (1)

- Unfortunately, as in the univariate case, traditional ways of visualising multivariate distributions aren't always easy to interpret, e.g. Bivariate Normal / Gaussian pdf and cdf, see below
- Even harder to visualise if number of dimensions is greater than 2!

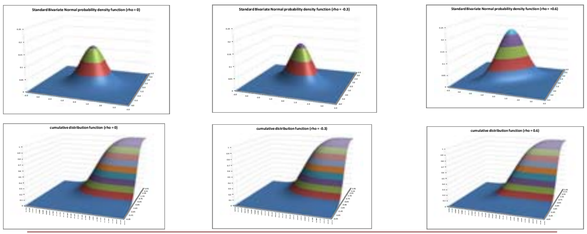

Source: Threadneedle



# Visualising multivariate distributions (2)

- Copulas aren't much easier to visualise
- Although easier to visualise is the difference between two copulas

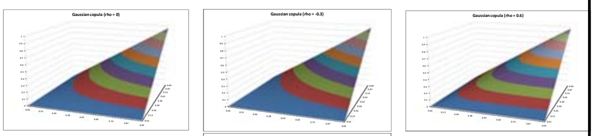

Gaussian copula with rho = 0 also called the "product" or "independence" copula

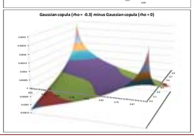

Source: Threadneedle



# Fractile analysis

- Monthly pair-wise differences between (log) sector relatives (31/12/1998 to 31/12/2007) on average have peaks in all four corners
  - But this partly due to different pairs exhibiting different correlations
- Principal components (orthogonal by construction) also on average show 4 peaks
  - Less marked, although still statistically significant
- Size of peaks reduce materially if adjust for time-varying volatility

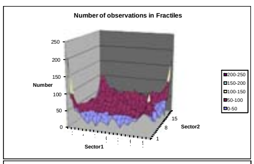


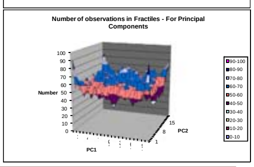


Source: Threadneedle, Thomson Datastream

## Fat-tailed behaviour of sector relative returns

- Weekly sector relative returns (industry group, June 2002 to March 2008) appear to exhibit fat tails
- Both on upside and on downside
- Extent seems to depend on sector

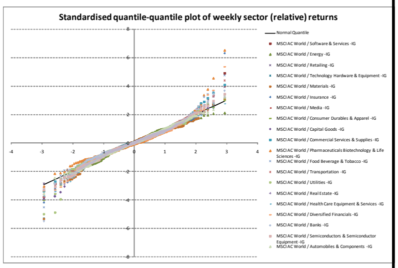


Source: Threadneedle, MSCI, Factset, Thomson Datastream



## Cross-sectional dispersion of quantiles

- Displaying fat tails exhibited by pair-wise combinations of series presents challenges
  - *n* individual series have *n*(*n*-1)/2 pair-wise combinations for any given weight combination
  - So we need a simpler way of presenting spread of fat tails
- E.g. measure cross-sectional dispersion using quantiles
  - i.e. quantiles of quantiles
  - Say 95th, 75th, 50th, 25th, 5th cross-sectional and average for a given expected quantile level (i.e. *x* coordinate)

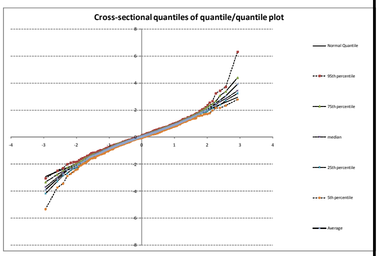


Source: Threadneedle, MSCI, Factset, Thomson Datastream



## Cross-sectional dispersion of pair-wise combinations (1)

- We might combine two sector positions in any linear combination, i.e. combine relative returns ($r_1$ and $r_2$) as $a_1 \times r_1 + a_2 \times r_2$ using any values for $a_1$ and $a_2$
  - We will concentrate on "X + Y", i.e. ($a_1$, $a_2$) = (1,1), (-1,-1) ,... (i.e. positions in same direction)
  - And on "X-Y", i.e. ($a_1$, $a_2$) = (1,1), (-1,-1) .... (i.e. positions in opposite directions)
  - Combinations based on standardised return series, i.e. combinations are volatility weighted
- Spreads pretty similar, maybe a little lower than for individual return series

Individual Sector Relatives

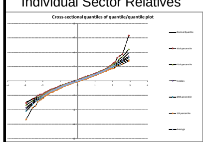


All "X+Y" combinations

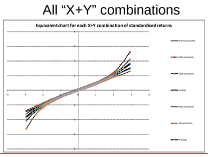


All "X-Y" combinations

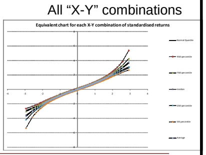


Source: Threadneedle, MSCI, Factset, Thomson Datastream

## Cross-sectional dispersion of pair-wise combinations (2)

- Time varying volatility again seems to explain a significant proportion of the fat-tailed behaviour
  - Although better now to adjust by the recent average dispersion between sectors
  - E.g. preceding 10 weeks worth (measured by cross-sectional standard deviation, weighted by entire period standard deviation of returns)
- Magnitude of residual fat-tailed behaviour similar to that of market indices

Individual Sector Relatives

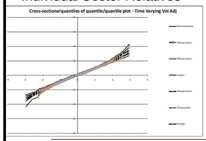

All "X+Y" combinations

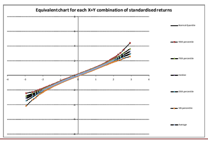

All "X-Y" combinations

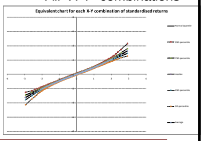

Source: Threadneedle, MSCI, Factset, Thomson Datastream

The Actuarial Profession
making financial sense of the future

## Fat-tailed behaviour in joint return series

- Behaviour of multiple (joint) return series can be decomposed into :
  - Behaviour of each series in isolation, via marginal distribution
  - Dependency between series, via copula
- We are interested in linear combinations, which can also be analysed directly
- Both main market index and sector relative returns appear to be fat tailed
  - As do combinations of sectors, with some bunching into corners of the copula
- Again, time-varying volatility seems to explain some fraction of fat tails
  - Better modelled by time-varying *cross-sectional* variability rather than each sector in isolation
  - Hence focus on implied correlation as well as implied volatility
- But some fat-tailed behaviour is not explained by time-varying volatility
  - Crowded trades?

The Actuarial Profession
making financial sense of the future

## Portfolio construction

- Traditional (quantitative) approach to portfolio construction involves *portfolio optimisation*
  - Typically *mean-variance* optimisation
  - Identify alpha (return) expected from each position
  - Maximise expected return for a given level of risk subject to constraints (e.g. weights sum to unity)
  - Or equivalently maximise $a.r - lambda.a^{T}Va$
- Results notoriously sensitive to input assumptions
  - Treat quant models with scepticism (the fundamental manager's approach)?
  - Or otherwise use "robust" approaches, or Bayesian priors/anchors, e.g. Black-Litterman?
  - Or reverse optimisation, i.e. *implied alphas*
- Mean-variance can be shown to be optimal
  - *either* if return distribution is multivariate (log) Normal (as then distribution characterised merely by mean and covariance matrix)
  - *or* if investors have quadratic utility, i.e. indifferent to deviations from non-Normality
- The risk aversion parameter, lambda, not specified a priori – it is in effect defined by the "risk budget"

The Actuarial Profession
making financial sense of the future

## Portfolio construction – impact of fat-tails

- We reiterate: if *all* return opportunities (and combinations of them) are '*equally*' (jointly) fat-tailed then optimal portfolios are the same as those arising if you use traditional mean-variance optimisation approaches (and adjust the risk budget accordingly)
- Most important (predictable) single contributor to fat-tails seems to be time-varying volatility
  - Both of individual return series in isolation
  - And, more importantly for multiple return series, changing cross-sectional volatility
  - Market implied equivalents are implied volatility and implied correlation
- Suggested prescription (if your investment process requires an optimiser)?
  - Calculate co-dependency between return series *after* stripping out effect of time-varying volatility
  - Optimise as you think fit (standard, "robust", Bayesian, Black-Litterman, ...) using adjusted covariance matrix and adjusting the risk aversion/risk budget appropriately
  - Or derive implied alphas using adjusted covariance matrix
  - Implicitly assumes that all (adjusted) return series exhibit same degree of fat-tailed behaviour
- Or continue to treat output from optimisers with some caution!



## Other approaches suggested by some commentators

- Mixtures – e.g. mixtures of multivariate normal distributions
  - This is how we alighted on time varying volatility as a possible source of fat tails
  - But even a mixture of just two multivariate Normal distributions has twice as many covariance terms to estimate, making parameter estimation correspondingly less reliable
  - And results of optimisation exercises were already notoriously sensitive to input assumptions!
- Lower partial moments
  - Any return = threshold + upside + downside
  - Non-quadratic utility will give greater weight to downside rather than upside and will in general also depend on higher moments
  - For single series defined as: $lpm(K,m)=E[\max((r-K)^m,0)]$
  - For multiple return series defined as: $lpm_{i,j}(K,m,n)= E[\max((r_i-K)^m(r_j-K)^n,0)]$
  - I.e. co-skewness, co-kurtosis (or symmetric alternatives)
  - Lots more parameters to estimate
  - Use of skew and kurtosis proved not to be ideal even for single return series



## Summary

- Fat tails are pretty common
  - Both for individual indices and for multiple (joint) distributions
- When analysing fat-tails, treat skew, kurtosis, co-skew, Cornish Fisher etc. with some care
  - As they don't necessarily give appropriate weight to the right observations
  - Better may be to curve-fit the distributional form directly
- Time-varying volatility seems to explain some but not all fat-tailed behaviour
  - Both of individual return series and of joint distributions
  - But still some "unknown unknowns", i.e. "Black swans"
  - Hence merits of stress-testing, implied volatility/correlation analysis, crowded trade analysis
- Portfolio construction
  - Output from optimisers are typically very sensitive to input assumptions
  - Treat pure quant models with scepticism (the fundamental manager's perspective)?
  - Or otherwise use robust optimisation, Bayesian, Black-Litterman, reverse optimisation, …?
  - Try to avoid introducing even more parameters to estimate
  - I.e. keep adjustments for fat-tailed behaviour as simple and as intuitive as possible

# Important Information

Past performance is not a guide to the future. The value of investments and the income from them is not guaranteed and may fall as well as rise and the investor may not get back the original investment. Changes in exchange rates may also cause the value of investments to fall as well as rise.

The mention of any fund does not constitute an offer or invitation to subscribe to shares in such a fund. The research and analysis included in this document has been produced by Threadneedle Asset Management Limited for its own investment management activities and may have been acted upon prior to publication.

In some instances the information contained in this presentation, other than statements of fact, may have been obtained from external sources believed to be reliable but its accuracy or completeness cannot be guaranteed.

Any opinions expressed are as at the date of issue but may be subject to change.

References to target or expected returns are indicative only and are not guaranteed in any way and may be affected by client constraints as well as external factors and management performance.

Performance figures that relate to a fund or a representative account may differ from that of other separately managed accounts due to such differences as cash flows, charges, applicable taxes, and differences in investment strategy and restrictions.

The views represented here are of the individuals and do not necessarily represent those of Threadneedle Asset Management Limited.

For Investment Professional Use Only (not for onward distribution to, or to be relied upon by, private investors)