Extreme Events – Specimen Questions and Answers

[Nematrian website page: ExtremeEventsQuestionsAndAnswers, © Nematrian 2015]

An Appendix to the book Extreme Events: Robust Portfolio Construction in the Presence of Fat Tails contains some specimen questions for students and lecturers on material covered in the book, reproduced courtesy of Nematrian. In the following pages we provide some solution hints/model answers to these questions. Analytical tools available on the Nematrian website that can help with these solutions are referred to in the relevant model solution.

We suggest that you work through these questions and answers approximately sequentially. Techniques and hints given for earlier questions may also be relevant to later ones.

Chapter 2: Fat Tails – In Single (i.e. Univariate) Return Series

[ExtremeEventsQuestionsAndAnswers2]

The exercises in this Section relate to the following hypothetical (percentage) returns on two indices, A and B, over 20 periods:

Period	A (%)	B (%)
1	-25.1	5.0
2	0.1	-16.0
3	24.2	-11.8
4	7.9	-3.0
5	-1.4	-9.3
6	0.5	-0.1
7	-0.6	-7.7
8	1.0	2.8
9	7.3	13.1
10	-12.0	1.5
11	0.7	-10.5
12	0.8	11.7
13	-3.8	17.9
14	2.8	9.7
15	-4.3	8.3
16	-1.2	-6.3
17	-7.0	-7.8
18	6.5	-2.3
19	3.6	5.4
20	1.8	2.8

Questions:

Specimen Question A.2.1

[ExtremeEventsQuestionsAndAnswers2_1q]

You are an investor seeking to understand the behaviour of Index A:

(a) Calculate the mean, (sample) standard deviation, skew and (excess) kurtosis of its log returns over the period covered by the above table.

Answer/Hints

(b) Do the statistics calculated in (a) appear to characterise a fat-tailed distribution if we adopt the null hypothesis that the log returns would otherwise be coming from a normal distribution and we use the limiting form of the distributions for these test statistics (i.e. the form ruling when $n \to \infty$, where n is the number of observations)?

Answer/Hints

(c) Prepare a standardised quantile-quantile plot for Index A. Does it appear to be fat-tailed?

[Answer/Hints]

(d) Does the Cornish-Fisher 4th moment approximation appear to under or overstate the fat-tailed behaviour of this series?

[Answer/Hints]

(e) What other methodologies could you use to formulate a view about how fat-tailed this return series might be if your focus was principally on fat-tailed behaviour around or below the lower 10th percentile quantile level?

[Answer/Hints]

Answers/Hints

A.2.1(a)

[ExtremeEventsQuestionsAndAnswers2_1a]

Q. Calculate the mean, (sample) standard deviation, skew and (excess) kurtosis of [Index A's] log returns over the period covered by the above table.

The logged returns of A, i.e. $x_i = \log(1 + r_i)$ are:

Period	A (logged return)
1	-0.2890163
2	0.0009995
3	0.2167230
4	0.0760347
5	-0.0140989
6	0.0049875
7	-0.0060181
8	0.0099503
9	0.0704585
10	-0.1278334
11	0.0069756
12	0.0079682
13	-0.0387408
14	0.0276152
15	-0.0439519
16	-0.0120726
17	-0.0725707
18	0.0629748
19	0.0353671
20	0.0178399

The mean, (sample) standard deviation, skew and (excess) kurtosis of A's log returns can be calculated in a variety of ways, including:

i. In a Microsoft Excel spreadsheet using the in-built functions AVERAGE, STDEV, SKEW and KURT respectively;

ii. In a Microsoft Excel spreadsheet using the Nematrian web functions MnMean, MnStdev, MnSkew and MnKurt respectively;

iii. In VBA, the built-in programming language packaged with Microsoft Excel, using:

- The VBA equivalents to (i), i.e. Excel.WorksheetFunction.Average, Excel.WorksheetFunction.StDev, Excel.WorksheetFunction.Skew and Excel.WorksheetFunction.Kurt respectively; or

- The VBA equivalents to (ii), which if you have e.g. used the Nematrian website's, automatic code generator or have loaded up a prepopulated spreadsheet from the Nematrian spreadsheet library will also be called MnMean, MnStdev, MnSkew and MnKurt respectively;

iv. Interactively, via the MnMean, MnStdev, MnSkew and MnKurt webpages.

In each case the answers supplied (possibly with some rounding error) are:

Statistic	**Value**
mean	-0.0033204
stdev	0.0954161
skew	-0.9253403
kurt	4.6238918

A.2.1(b)

[ExtremeEventsQuestionsAndAnswers2_1b]

Q. Do the statistics calculated in (a) appear to characterise a fat-tailed distribution if we adopt the null hypothesis that the log returns would otherwise be coming from a normal distribution and we use the limiting form of the distributions for these test statistics (i.e. the form ruling when $n \to \infty$, where n is the number of observations)?

A (univariate) normal distribution is characterised by its mean and standard deviation, so the values of these two statistics cannot be used to differentiate between the normal distribution family and other distributional forms.

However, we can test for normality by reference to the observed skew and kurtosis of the sample. If n is large and if the sample is drawn from a Normal distribution then the skew and (excess) kurtosis are approximately normally distributed with the following distributions, see also: MnConfidenceLevelSkewApproxIfNormal and MnConfidenceLevelKurtApproxIfNormal

$$skew = \gamma_1 \sim \sqrt{6/n}$$
$$kurtosis = \gamma_2 \sim \sqrt{24/n}$$

Suppose we wish to reject the null hypothesis that the sample is coming from a normal distribution with a symmetric two-sided significance level of α (and we assume that $n = 20$ is sufficiently large for the above approximations to apply), then we would reject the null hypothesis if the observed γ_1 and γ_2 are above $N^{-1}(1-\alpha/2)$ or below $N^{-1}(\alpha/2)$ where $N^{-1}(x)$ is the (standardised) inverse normal distribution. $N^{-1}(x)$ can be obtained via the built-in Microsoft Excel worksheet function NORMSINV or via the equivalent Nematrian web function MnInverseNormal.

For example, if $\alpha = 5\%$ then $-N^{-1}(\alpha/2) = N^{-1}(1-\alpha/2) = 1.96$. Thus at this level of significance, index A **does not** appear to be skewed, but **does** appear to be fat-tailed, since the observed value of γ_2 is 4.62 is significantly larger than 1.96.

Other tests for normality that might be used (including ones that can handle small samples and/or focus on just some parts of the overall distributional form) are described in TestsForNormality.

A.2.1(c)

[ExtremeEventsQuestionsAndAnswers2_1c]

Q. Prepare a standardised quantile-quantile plot for Index A. Does it appear to be fat-tailed?

A standardised quantile-quantile plot for Index A can be derived using the following steps:

i. Standardise the data, i.e. if the original logged returns are x_i then calculate:

$$y_i = \frac{x_i - \hat{\mu}}{\hat{\sigma}}$$

where $\hat{\mu}$ is the sample mean and $\hat{\sigma}$ is the sample standard deviation

ii. Order this data, i.e. calculate $y_{(i)}$ where $y_{(1)}$ is the smallest value of the set $\{y_1, \dots, y_n\}$, $y_{(2)}$ is the next smallest etc.

iii. Calculate the 'expected' value that we would expect each of the $y_{(i)}$ to take were they to be coming from a normal distribution, i.e. $q_{(i)} = N^{-1}\left(\left(i - \frac{1}{2}\right)/n\right)$

iv. Plot the $y_{(i)}$ against the $q_{(i)}$, usually with the $q_{(i)}$ along the horizontal axis and the $y_{(i)}$ along the vertical axis.

Steps (i) to (iii) can be all be carried out without too much effort on a case-by-case basis in a spreadsheet system such as Microsoft Excel or using the Nematrian web functions MnNormaliseArray (to carry out steps (i) and (ii) and MnStandardisedNormalQuantilesArray (to carry out step (iii)).

However, it is simpler to use the Nematrian website charting capability. This includes a web function, i.e. MnPlotStandardisedQQ, which both creates (internally) the required data to be viewed in a standardised QQ-plot (i.e. carries out steps (i) to (iii) above) and also creates a temporary Nematrian 'SmartChart', i.e. the plot itself as in step (iv), accessible via a suitable 'SmartCode'. As with other Nematrian web functions it is possible to access the function programmatically (in which case it merely returns the temporary 'SmartCode' assigned to the relevant 'SmartChart'). It can also be accessed interactively, in which case the output on the webpage includes both the 'SmartCode' and the SmartChart.

For users with appropriate permissions, temporary SmartCharts can be converted into permanent SmartCharts, which are stored permanently on the Nematrian website and retain an association between the creator of the chart, its underlying data and its visual form. These can be copied and pasted from the Nematrian website into other venues, e.g. Word documents or Powerpoint presentations. For users without these permissions the image forming the SmartChart can be copied in a manner similar to any other image you might find on a webpage, but what is copied no longer retains an association with its creator and its underlying data.

A permanent SmartChart answering this question is:

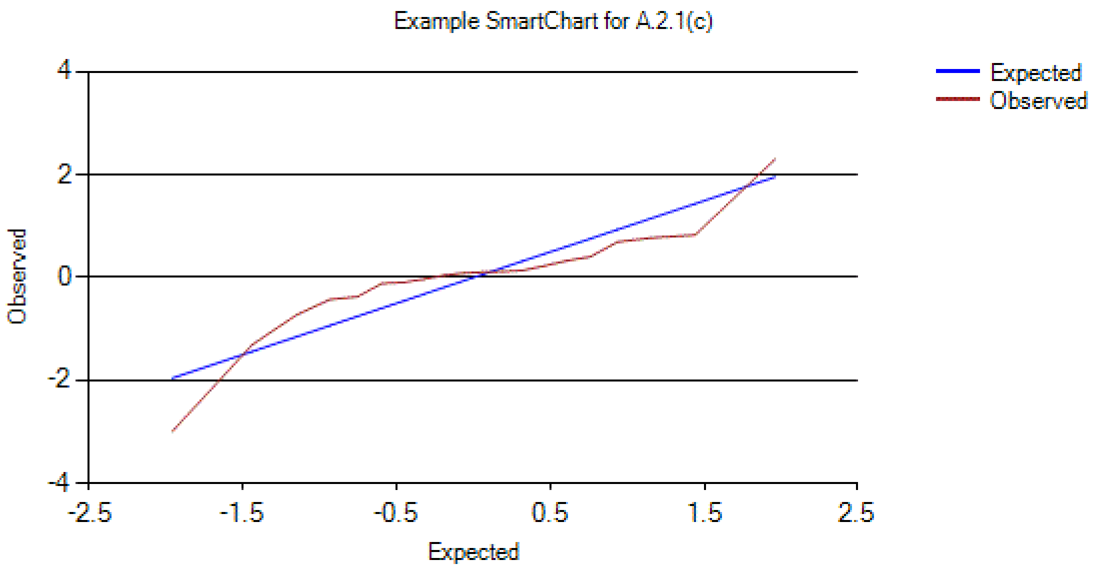

A.2.1(d)

[ExtremeEventsQuestionsAndAnswers2_1d]

Q. Does the Cornish-Fisher 4th moment approximation appear to under or overstate the fat-tailed behaviour of this series?

The Cornish-Fisher 4th moment approximation, see MnCornishFisher4, in effect characterises the form of the QQ-plot by a cubic, which if the sample data has been standardised (using e.g. MnNormaliseArray) involves a plot of the following form:

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

Whilst the series underlying this plot can be created relatively simply in, say, Microsoft Excel (or using MnCornishFisher4, it is again probably simpler to use the Nematrian website charting capability that makes the creation of such a chart very simple, using MnPlotStandardisedQQWithComparisons.

The Cornish-Fisher 4th moment approximation appears in this instance to be providing a reasonable representation of the distributional form towards the tails of the distribution, if anything perhaps understating the extent of fat-tailed behaviour.

This is in contrast to its apparent tendency to overstate fat-tailed behaviour in some of the index series analysed in the book Extreme Events. The last observation in either tail in the small sample analysed here is actually only at the 2.5%/97.5% quantile level, so is not particularly far into the tail relative to some of the larger samples analysed in that book.

A.2.1(e)

[ExtremeEventsQuestionsAndAnswers2_1e]

Q. What other methodologies could you use to formulate a view about how fat-tailed this return series might be if your focus was principally on fat-tailed behaviour around or below the lower 10th percentile quantile level?

With a sample size of 20 there are only two data points in the sample around or below the lower 10th percentile quantile level. This is too few to permit any meaningful statistical analysis to be carried out in this instance.

However, if the sample size were much larger, then we could carry out the following:

(a) We could define a weighting 'schema', i.e. w_i, to apply to the ordered data series, $y_{(i)}$. In this case the form of the question might lead us to use:

$$w_i = \begin{cases} 1 & \text{if } i \leq n/10 \\ 0 & \text{if } i > n/10 \end{cases}$$

(b) We could calculate the normal distribution that best fitted the observed data sample, but giving weight w_i to the observations, i.e. here only taking into account data where $i \leq n/10$. To do so we would use functions as set out in the Nematrian web page on Weighted Moments and Cumulants.

(c) We could now compare the ordered observed data versus that 'expected' were the normal distribution in (ii) to have applied, perhaps visually and/or perhaps fitting suitable curves through this comparison, e.g. a variant of the approach underling the Nematrian's standard weighted cubic curve fit.

Formal tests for non-normality could then be carried out using suitable refinements to standard test methodologies, see TestsForNormality.

Specimen Question A.2.2

[ExtremeEventsQuestionsAndAnswers2_2q]

You are an investor seeking to understand the behaviour of Index B:

(a) You think that the returns shown for Index B may exhibit a material element of smoothing. What sorts of assets might lead to this type of behaviour?

Answer/Hints

(b) Using a first order autoregressive model, de-smooth the observed returns for Index B to derive a return series that you think may provide a better measure of the underlying behaviour of the relevant asset category.

Answer/Hints

(c) Prepare a standardised quantile-quantile plot for this underlying return series. Does it appear to exhibit fat-tailed behaviour?

Answer/Hints

Answers/Hints

A.2.2(a)

[ExtremeEventsQuestionsAndAnswers2_2a]

Q. You think that the returns shown for Index B may exhibit a material element of smoothing. What sorts of assets might lead to this type of behaviour?

Assets that seem to exhibit smoothing in practice are typically ones that take a relatively long time to buy or sell or are otherwise less liquid.

A classic example is direct property, i.e. real estate, given the very wide variety of forms that it can take and the often lengthy legal processes and negotiations that are necessary to buy or sell it.

However, many other types of asset can also be relatively illiquid, or proved to be less liquid than investors had hoped during the 2007-2009 credit crisis. For example, some times of complicated structured investments proved very illiquid during the credit crisis and could only be traded, if at all, with very wide bid-offer spreads.

The equities of smaller companies are often less liquid than their larger brethren.

A.2.2(b)

[ExtremeEventsQuestionsAndAnswers2_2b]

Q. Using a first order autoregressive model, de-smooth the observed returns for Index B to derive a return series that you think may provide a better measure of the underlying behaviour of the relevant asset category.

The approach for de-smoothing (or 'de-correlating') such returns suggested in the book Extreme Events involves assuming that there is some underlying 'true' return series, $\tilde{r}_t$, and that the observed series, r_t, derives from it via a first order autoregressive model, $r_t = (1-\rho)\tilde{r}_t + \rho\tilde{r}_{t-1}$. We will assume that the autoregressive model actually applies to the logged returns, which are given in ExtremeEventsQuestionsAndAnswers2_1a.

If we assume that the corresponding $\tilde{r}_t$ are independent, identically distributed (normal) random variables with standard deviation s (and mean 0 given that we have standardised the data) then the (expected) variance of the series $\{r_1, \dots, r_n\}$, will be $V = ((1-\rho)^2 + \rho^2)s^2$ and the (expected) covariance of the series $\{r_2, \dots, r_n\}$ with the series $\{r_1, \dots, r_{n-1}\}$ will be $C = \rho(1-\rho)s^2$.

One way of proceeding would be to estimate ρ as the solution to the equation:

$$\frac{\rho(1-\rho)}{((1-\rho)^2 + \rho^2)} = \frac{C}{V}$$

More precisely, since the above does not differentiate between ρ and $1-\rho$ we would probably choose the ρ closest to zero (and, ideally, we would expect it to be positive and smaller than 0.5, which also requires C to be positive, as a value of ρ outside this range would be implausible).

In this instance:

Statistic (using logged returns)	**Value**
(Sample*) Variance of $\{r_1, \dots, r_n\}$ (=V)	0.008575828
(Sample*) Covariance of $\{r_2, \dots, r_n\}$ with $\{r_1, \dots, r_{n-1}\}$ (=C)	0.002896449
Ratio (=C/V)	0.337746
Estimated ρ	0.279955

* Ideally, the Variance and Covariance should both include the same small sample size adjustment, i.e. both be “sample” or both be “population” estimates. The Microsoft Excel functions, VAR, VARP and COVAR, are somewhat confusing in this respect, since COVAR is calculated using a multiplier $1/n$, and is therefore properly a “population” statistic and consistent with VARP, whilst VAR is calculated using a multiplier of $1/n-1$, and so is a “sample” statistic. The Nematrian website’s web functions are clearer as it provides both a MnCovariance function and a MnPopulationCovariance function.

Ignoring small sample size adjustments, estimates of $\tilde{r}_t$ can be then be found using:

$$\tilde{r}_t = \begin{cases} r_t \quad \text{if } t = 1 \\ \dfrac{r_t - \rho\tilde{r}_{t-1}}{1-\rho} \quad \text{if } t > 1 \end{cases}$$

However, if we do this in practice we find that there are second order effects that mean that estimating $\tilde{r}_t$ as above still leaves some residual autocorrelation:

Period	r_t (logged returns of B)	$\tilde{r}_t$ (first pass)
1	0.0487902	0.0487902
2	-0.1743534	-0.2611119
3	-0.1255632	-0.0728617
4	-0.0304592	-0.0139731
5	-0.0976128	-0.1301320
6	-0.0010005	0.0492060
7	-0.0801260	-0.1304104
8	0.0276152	0.0890558

9	0.1231022	0.1363395
10	0.0148886	-0.0323317
11	-0.1109316	-0.1414913
12	0.1106465	0.2086780
13	0.1646666	0.1475549
14	0.0925792	0.0712046
15	0.0797350	0.0830516
16	-0.0650720	-0.1226626
17	-0.0812101	-0.0650933
18	-0.0232686	-0.0070071
19	0.0525925	0.0757649
20	0.0276152	0.0088945

Variance (=V)	0.0085758	0.0138041
Covariance (=C)	0.0028964	0.0007054
Ratio (=C/V)	0.3377457	0.0511003
Estimated ρ	0.2799546	

Better, therefore, is to use a root search algorithm in which we explicitly search for the ρ (ideally between 0 and 0.5) for which $\tilde{r}_t$ has zero autocorrelation. This can be done using the Nematrian MnDesmooth_AR1 or MnDesmooth_AR1_rho functions (the former returns the desmoothed series, the latter returns the value of ρ for which $\tilde{r}_t$ has zero autocorrelation). Given the form of the problem given here, these provide the following de-smoothed series (with rho equal to 0.31094682):

Period	r_t (logged returns of B)	$\tilde{r}_t$ (de-smoothed)
1	0.0487902	0.0487902
2	-0.1743534	-0.2750507
3	-0.1255632	-0.05810445
4	-0.0304592	-0.01798381
5	-0.0976128	-0.13354672
6	-0.0010005	0.05881321
7	-0.0801260	-0.14282465
8	0.0276152	0.10452905
9	0.1231022	0.13148365
10	0.0148886	-0.03772687
11	-0.1109316	-0.14396646
12	0.1106465	0.22554488
13	0.1646666	0.13719425
14	0.0925792	0.07244591
15	0.0797350	0.08302433
16	-0.0650720	-0.13190296
17	-0.0812101	-0.05833410
18	-0.0232686	-0.00744471
19	0.0525925	0.07968530
20	0.0276152	0.00411769

Variance (=V)	0.0085758	0.01487681
Covariance (=C)	0.0028964	0
Ratio (=C/V)	0.3377457	0

Note:

(a) In general de-smoothing *increases* the variance of the return series being analysed. Here it has gone from 0.0926 to 0.1219.

(b) Whilst the problem would usually be stated as shown, it perhaps makes more sense not to make the arbitrary implicit assumption that the smoothing is around a mean of zero, but around some mean (that is, for example, estimated from the data), i.e. as if the model was $r_t - \mu = (1-\rho)(\tilde{r}_t - \mu) + \rho(\tilde{r}_{t-1} - \mu)$.

(c) We might not necessarily want to give equal weight to each observation. This is possible using the Nematrian MnWeightedDesmooth_AR1 and MnWeightedDesmooth_AR1_rho functions.

A.2.2(c)

[ExtremeEventsQuestionsAndAnswers2_2c]

Q. Prepare a standardised quantile-quantile plot for this underlying return series. Does it appear to exhibit fat-tailed behaviour?

Again, the simplest approach is to use the Nematrian charting facility e.g. MnPlotStandardisedQQ, applied to, say, the series in the final column in the last table of A.2.2(b). The resulting chart does not appear to exhibit material fat-tailed behaviour.

The skew and (excess) kurtosis of the eventual data series are -0.322 and -0.089 which are both close to zero (given the small sample size in question), so they too are not indicative of non-normality.

Specimen Question A.2.3

[Nematrian website page: ExtremeEventsQuestionsAndAnswers2_3q, © Nematrian 2015]

You discover that the periods being used for Index A are quite long (e.g. yearly), sufficiently long for secular change to make the relevance of data from some of the earlier periods suspect. You decide to exponentially weight the data using a half-life of 10 periods, i.e. the weight given to period t (for $t = 1, \dots, 20$) is $w(t) = e^{-(20-t)\log(2)/10}$.

(a) Recalculate the mean, standard deviation, skew and kurtosis weighting the data as above. Are they still suggestive of fat-tailed behaviour?

Answer/Hints

(b) Prepare a quantile-quantile plot of the weighted data. Is it also suggestive of fat-tailed behaviour? Hint: the 'expected' values for such plots need to bear in mind the weight given to the observation in question.

Answer/Hints

Answers/Hints

A.2.3(a)

[ExtremeEventsQuestionsAndAnswers2_3a]

Q. Recalculate the mean, standard deviation, skew and kurtosis weighting the data as above. Are they still suggestive of fat-tailed behaviour?

The proposed exponential weighting approach gives the following weights to the observations:

Period	Weight	A (logged return)
1	0.2679434	-0.2890163
2	0.2871746	0.0009995
3	0.3077861	0.2167230
4	0.3298770	0.0760347
5	0.3535534	-0.0140989
6	0.3789291	0.0049875
7	0.4061262	-0.0060181
8	0.4352753	0.0099503
9	0.4665165	0.0704585
10	0.5	-0.1278334
11	0.5358867	0.0069756
12	0.5743492	0.0079682
13	0.6155722	-0.0387408
14	0.6597540	0.0276152
15	0.7071068	-0.0439519
16	0.7578583	-0.0120726
17	0.8122524	-0.0725707
18	0.8705506	0.0629748
19	0.9330330	0.0353671
20	1	0.0178399

The ordinary mean of a set of numbers is $\bar{x} = \sum x_i / n$, where the x_i are the observations and there are n of them. The weighted mean, if each observation is given a weight w_i is $\tilde{x}$, say, where:

$$\tilde{x} = \frac{\sum w_i x_i}{\sum w_i}$$

We see immediately that $\tilde{x} = \bar{x}$ if the observations are given equal weight, i.e. if all the w_i are equal.

It is relatively simple to calculate this statistic from the above data using Microsoft Excel (or using the Nematrian web function MnWeightedMean), giving an answer of -0.0007543

However, identifying exactly how to calculate the corresponding weighted standard deviation, skew and kurtosis is less simple, particularly if our focus is on small samples.

For 'population' statistics, i.e. those applicable if we have large observation sets, the calculations involved appear to be relatively unambiguous, see WeightedMomentsAndCumulants

However, for 'sample' statistics, there appear to be differences in opinion in how, precisely, to adjust 'population' statistics to allow for the different degrees of freedom that are present. Given this apparent disagreement between commentators, the Nematrian website provides web functions for calculating the weighted population mean, standard deviation, skew and excess (excess) kurtosis, but only corresponding weighted sample mean, standard deviation and skew (i.e. not yet weighted sample excess kurtosis). Rimoldini (2013) proposes some formulae for these statistics, which appear to match the Nematrian formulae.

Formulae used by the Nematrian website can be accessed using the following web functions: MnWeightedPopulationStdev, MnWeightedPopulationSkew, MnWeightedPopulationKurt, MnWeightedStdev and MnWeightedSkew and give the following results for the above data:

Statistic	Unweighted (population)	Unweighted (sample)	Weighted (population)	Weighted (sample)
Mean	-0.0033204	-0.0033204	-0.0007543	-0.0007543
Standard deviation	0.0930001	0.0954161	0.0741301	0.0763675
Skew	-0.791166	-0.9253403	-0.8903571	-0.9761716
(Excess) kurtosis	2.6500366	4.6238918	5.0848899	N/A

These statistics are still strongly suggestive of fat-tailed behaviour.

A.2.3(b)

[ExtremeEventsQuestionsAndAnswers2_3b]

Q. Prepare a (standardised) quantile-quantile plot of the weighted data. Is it also suggestive of fat-tailed behaviour? Hint: the 'expected' values for such plots need to bear in mind the weight given to the observation in question.

We need to:

(a) Scale the weights so that they add to unity

(b) Standardise the observations, so that they have weighted mean equal to zero and weighted (sample) standard deviation equal to unity

(c) Sort the data

(d) Work out the quantile plots corresponding to each data point (e.g. assume the data point corresponds to the cumulative (scaled) weight of smaller points plus one-half of the (scaled) weight of the observation in question

(e) Identify standardised inverse normal values corresponding to the quantile points in (d)

(f) Plot 'observed' vs 'expected', the latter being the values from (e)

Steps (a) and (b) give:

Period	Scaled Weight	Standardised Observation
1	0.0239245	-3.7746674
2	0.0256416	0.0229646
3	0.0274820	2.8477710
4	0.0294545	1.0055183
5	0.0315686	-0.1747424
6	0.0338343	0.0751858
7	0.0362627	-0.0689278
8	0.0388654	0.1401715
9	0.0416550	0.9325003
10	0.0446447	-1.6640468
11	0.0478490	0.1012191
12	0.0512833	0.1142167
13	0.0549640	-0.4974175
14	0.0589090	0.3714857
15	0.0631371	-0.5656546
16	0.0676687	-0.1482089
17	0.0725255	-0.9404055
18	0.0777309	0.8345045
19	0.0833099	0.4729935
20	0.0892893	0.2434824

Steps (c), (d) and (e) give:

Scaled Weight	Observation (i.e. 'Observed')	Quantile Point	'Expected'
0.0239245	-3.7746674	0.0119622	-2.2583397
0.0446447	-1.6640468	0.0462468	-1.6823879
0.0725255	-0.9404055	0.1048319	-1.2544904
0.0631371	-0.5656546	0.1726632	-0.9436933
0.054964	-0.4974175	0.2317138	-0.7332146
0.0315686	-0.1747424	0.2749801	-0.5978199
0.0676687	-0.1482089	0.3245987	-0.4548775
0.0362627	-0.0689278	0.3765644	-0.3145165
0.0256416	0.0229646	0.4075166	-0.233938
0.0338343	0.0751858	0.4372545	-0.1579336
0.047849	0.1012191	0.4780962	-0.0549323
0.0512833	0.1142167	0.5276623	0.0693948
0.0388654	0.1401715	0.5727367	0.1833459
0.0892893	0.2434824	0.6368141	0.3499558
0.058909	0.3714857	0.7109132	0.5560546
0.0833099	0.4729935	0.7820227	0.7790426
0.0777309	0.8345045	0.8625431	1.0918164
0.041655	0.9325003	0.922236	1.4202738
0.0294545	1.0055183	0.9577907	1.7256048
0.027482	2.847771	0.986259	2.2046022

The following chart plots the observed vs expected values, using Microsoft Excel. It is also suggestive of fat-tailed behaviour:

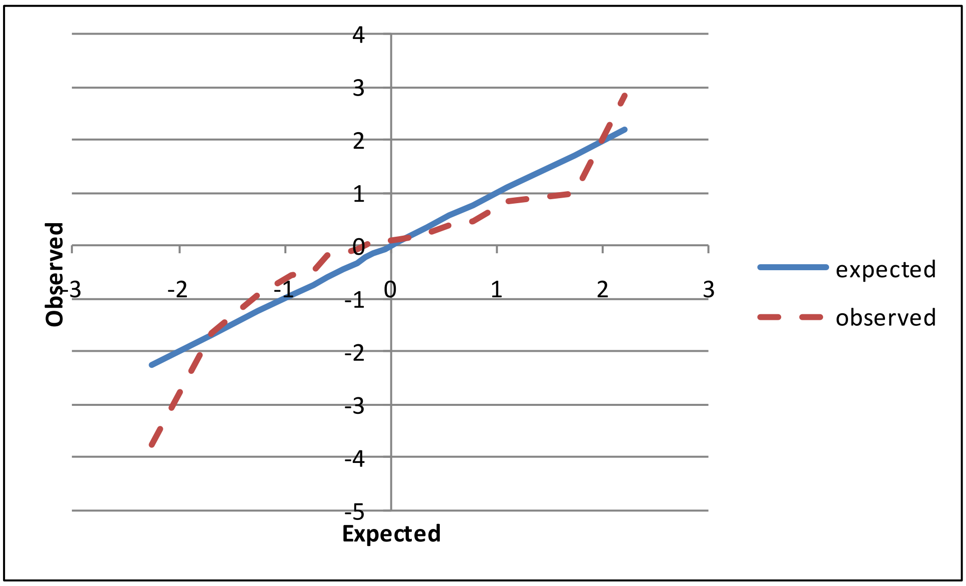

However, in practice it is easier to use the Nematrian Charting facility, i.e. MnPlotWeightedStandardisedQQ which can do all of these steps simultaneously.

Chapter 3: Fat Tails – In Joint (i.e. Multivariate) Return Series

[Nematrian website page: ExtremeEventsQuestionsAndAnswers3, © Nematrian 2015]

You have the following (percentage) return information on three different indices, A, B and C as follows (indices A and B are as per A.2). Index C has not been around as long as Index A and B.

Period	A (%)	B (%)	C (%)
1	-25.1	5.0	N/A
2	0.1	-16.0	N/A
3	24.2	-11.8	N/A
4	7.9	-3.0	N/A
5	-1.4	-9.3	N/A
6	0.5	-0.1	N/A
7	-0.6	-7.7	N/A
8	1.0	2.8	3.1
9	7.3	13.1	4.6
10	-12.0	1.5	-7.5
11	0.7	-10.5	-10.5
12	0.8	11.7	3.9
13	-3.8	17.9	5.2
14	2.8	9.7	0.1
15	-4.3	8.3	3.5
16	-1.2	-6.3	-2.9
17	-7.0	-7.8	-10.8
18	6.5	-2.3	2.5
19	3.6	5.4	6.3
20	1.8	2.8	2.4

Questions:

- Question A.3.1
- Question A.3.2

Specimen Question A.3.1

[ExtremeEventsQuestionsAndAnswers3_1q]

(a) Explain the advantages and disadvantages of attempting to back-fill data for Index C for the first seven periods from data applicable to either Index A or Index B when creating a model that jointly describes the behaviour of all three indices.

Answer/Hints

(b) If you had to select between either Index A or Index B to backfill data for Index C as per (a), which would you use? Why?

Answer/Hints

(c) What are the advantages and disadvantages of using a linear combination of Index A and Index B to backfill data for Index C, rather than using an either/or approach as per (b)?

Answer/Hints

Answers/Hints

A.3.1(a)

[ExtremeEventsQuestionsAndAnswers3_1a]

Q. Explain the advantages and disadvantages of attempting to back-fill data for Index C for the first seven periods from data applicable to either Index A or Index B when creating a model that jointly describes the behaviour of all three indices.

In general, the advantage of back-filling data in the proposed manner is so that we do not throw away information available on the (joint) behaviour of A and B merely because we do not have information available on C for the relevant period. If we have many assets with varying history lengths then not doing so can result in only a small number of overlapping periods, resulting in us throwing away most of the available data.

The disadvantages of back-filling data in this manner are:

i. The linkage between C, had it existed, and A and B during the time that data was not available for C may not bear much relationship to its linkage when there was overlapping data, i.e. we are implicitly making an assumption about time stationarity that may not be reasonable.

ii. Back-filling data merely using either A or B (or a combination) implicitly eliminates the idiosyncratic risk that C might have had during the period being back-filled. If the period being backfilled is large compared to the overall data set then this may result in the overall idiosyncratic characteristics of C being understated. Thus there is a risk that we fool ourselves into believing that C is more like A or B (or a combination) than it really is, merely because for convenience we have 'assumed' that it is for some of the data set.

iii. Point (b) could have important implications for the fine structure of asset allocations deemed optimal, since selection of C versus A or B (or a combination) depends heavily on the assumed characteristics of idiosyncratic risk expressed by each asset class.

A.3.1(b)

[ExtremeEventsQuestionsAndAnswers3_1b]

Q. If you had to select between either Index A or Index B to backfill data for Index C as per (a), which would you use? Why?

The natural one to use is Index B, since it has a much higher correlation with Index C than Index A over the period when both are present. The logged returns over these periods and correlation coefficients are:

Period	A (logged return)	B (logged return)	C (logged return)

8	0.0099503	0.0276152	0.0305292
9	0.0704585	0.1231022	0.0449734
10	-0.1278334	0.0148886	-0.0779615
11	0.0069756	-0.1109316	-0.1109316
12	0.0079682	0.1106465	0.0382587
13	-0.0387408	0.1646666	0.0506931
14	0.0276152	0.0925792	0.0009995
15	-0.0439519	0.0797350	0.0344014
16	-0.0120726	-0.0650720	-0.0294288
17	-0.0725707	-0.0812101	-0.1142891
18	0.0629748	-0.0232686	0.0246926
19	0.0353671	0.0525925	0.0610951
20	0.0178399	0.0276152	0.0237165
Correlation with C	0.16	0.78	

A.3.1(c)

[ExtremeEventsQuestionsAndAnswers3_1c]

Q. What are the advantages and disadvantages of using a linear combination of Index A and Index B to backfill Index C, rather than using an either/or approach as per (b)?

The main advantage is that Index C may have some similarities with both A and B, and so some blend of the two may better characterise C than either in isolation.

The main disadvantages are:

(a) We are making our model for C more complicated. This naturally means that we should expect it to provide a better fit (as there are more degrees of freedom now available to us). Use of techniques such as the Akaike Information Criterion, that trade off model complexity against goodness of fit, might be used to mitigate this issue.

(b) Adjusting the model in this manner implicitly involves assumptions about idiosyncratic risks applicable to C, see A.3.1(a).

(c) As with Question A.3.1(a), point (b) could have important implications for the fine structure of asset allocations deemed optimal, since selection of C versus A or B (or a combination) depends heavily on the assumed characteristics of idiosyncratic risk expressed by each asset class.

Specimen Question A.3.2

[ExtremeEventsQuestionsAndAnswers3_2q]

Plot an empirical two-dimensional quantile-quantile plot as per Section 3.5 characterising the joint distribution of Index A and Index B.

Answer/Hints

Answers/Hints

A.3.2

[ExtremeEventsQuestionsAndAnswers3_2a]

Q. Plot an empirical two-dimensional quantile-quantile plot as per Section 3.5 characterising the joint distribution of Index A and Index B.

See Figure 3.21 in Section 3.5 of Extreme Events for more details of what a 2-d quantile-quantile plot might look like. It is not currently possible to use Nematrian charting facility to create such a chart, although the Nematrian website does have a facility to plot an 'upward' one-dimensional quantile-quantile plot i.e. a cross-section through such a 2-d plot.

Chapter 4: Identifying Factors That Significantly Influence Markets

[Nematrian website page: ExtremeEventsQuestionsAndAnswers4, © Nematrian 2015]

Questions:

- Question A.4.1
- Question A.4.2

Specimen Question A.4.1

[ExtremeEventsQuestionsAndAnswers4_1q]

You are an investor trying to understand better the behaviour of Index B in A.2.1. You think that it is likely to be best modelled by an AR(1) autoregressive model along the lines of $y_t - \mu = c(y_{t-1} - \mu) + w_t$ with random independent identically distributed normal error terms w_t.

(a) Estimate the value of c 16 times, the first time assuming that you only have access to the first 5 observations, the next time you only have access to the first 6 observations, etc.

Answer/Hints

(b) Do these evolving estimates of appear to be stable? How would you test such an assertion statistically?

Answer/Hints

Answers/Hints

A.4.1(a)

[ExtremeEventsQuestionsAndAnswers4_1a]

Q. Estimate the value of c 16 times, the first time assuming that you only have access to the first 5 observations, the next time you only have access to the first 6 observations, etc.

In each estimation, the value of c can be estimated as the slope of the best fit line in a standard linear regression (e.g. using the function SLOPE in Microsoft Excel or the MnSlope Nematrian web function). [N.B. Some refinement of this would be needed if the model were the simpler one $y_t = cy_{t-1} + w_t$ since this is equivalent to the constrained regression where the line of best fit is forced to pass through the origin.]

Using logged returns, these values of c are:

Estimation up to and including period	Estimated value of c
5	-0.3354190
6	-0.4051569
7	-0.3408570

8	-0.3844843
9	0.0505541
10	0.1635878
11	0.1143870
12	-0.0408473
13	0.2354016
14	0.3187719
15	0.3548701
16	0.3061011
17	0.3293394
18	0.3233795
19	0.3163910
20	0.3218388

A.4.1(b)

[ExtremeEventsQuestionsAndAnswers4_1b]

Q. Do these evolving estimates of c appear to be stable? How would you test such an assertion statistically?

The estimates derived in A.4.1(a) do not appear to be stable – early ones are generally negative, whilst later ones are generally positive.

Whilst it is possible to create analytical statistical tests for many problems, it is often easier to carry out a Monte Carlo simulation, in which we simulate the outcomes assuming that some prior model is correct and we work out the proportion of times that outcomes as extreme as observed outcome arise in the simulation. This, of course, still requires us to identify significance levels etc. as would be the case with any other type of statistical technique

Leaving aside generic issues to do with simulation techniques (such as numbers of simulations to carry out, see e.g. Section 6.11 of the book Extreme Events), the main challenges with applying such a methodology to this type of problem are:

(a) Defining the right prior distribution and adjusting the problem to take account of degrees of freedom introduced by parameter estimation. In this particular case the form of the prior is well defined, but there is flexibility over the selected value of c. We cannot assume, say, that the 'true' model involves $c = 0.3218388$. This value was itself estimated. So instead, we might carry out simulations as if $c = 0.3218388$ but then include an adjustment to the elements of each separate simulation forcing the results always to correspond to this value (in effect a 'constrained' simulation). Imposing a constraint in this manner can be done in several different ways, each of which is implicitly adjusting somewhat the prior distribution we are implicitly using in our testing, so we need to take this into account in our end conclusions

(b) Defining how to measure how far away from the 'expected' are the actual observations. This problem is a generic one whenever we have several different observations within the overall observation set. We need to take a view on whether we are most interested in the spread of differences, the most extreme difference etc. Some of the issues are explored further in pages on the website relating to tests for normality.

Specimen Question A.4.2
[ExtremeEventsQuestionsAndAnswers4_2q]

You are an investor trying to understand better the joint behaviour of Indices A and B in Exercise A.2.1:

(a) Identify the series corresponding to the principal components of Indices A and B.

Answer/Hints

(b) Given a linear combination of Index A and B what is the maximum possible kurtosis of a linear combination of A and B?

Answer/Hints

(c) More generally, can you identify two series where the linear combination of the series with the least variance is also the one with the maximum kurtosis? Hint: try identifying two series with very few terms in them as it simplifies the relevant mathematics.

Answer/Hints

(d) What lessons might you draw from (c) in terms of use of variance or variance-related risk statistics when used to estimate the likelihood of extreme events?

Answer/Hints

Answers/Hints

A.4.2(a)
[ExtremeEventsQuestionsAndAnswers4_2a]

Q. Identify the series corresponding to the principal components of Indices A and B

There are several different ways of identifying the series corresponding to the principal components of A and B.

The first point to note is that these series are only computable up to a mean drift term and up to a scalar multiple, although it is conventional to arrange for the series to have zero mean and for them to be scaled in magnitude according to some suitable convention.

Thus, it is easiest to work with the following adjusted data, which involve application of constant adjustments to the original data so that they now have zero means:

Period	Adjusted logged return A (a_i, say)	Adjusted logged return B (b_i, say)
1	-0.2856959	0.0511585
2	0.0043199	-0.1719851

3	0.2200434	-0.1231949
4	0.0793551	-0.0280909
5	-0.0107785	-0.0952445
6	0.0083080	0.0013678
7	-0.0026977	-0.0777577
8	0.0132707	0.0299835
9	0.0737789	0.1254705
10	-0.1245130	0.0172569
11	0.0102960	-0.1085632
12	0.0112886	0.1130148
13	-0.0354204	0.1670349
14	0.0309356	0.0949475
15	-0.0406315	0.0821033
16	-0.0087522	-0.0627037
17	-0.0692503	-0.0788417
18	0.0662952	-0.0209003
19	0.0386876	0.0549608
20	0.0211603	0.0299835

The first principal component can be found by finding the θ $(0 \leq \theta < 2\pi)$ which maximises the standard deviation of $a_i \cos\theta + b_i \sin\theta$. The weights to give to *A* and *B* in constructing the principal component series are then $\cos\theta$ and $\sin\theta$ respectively. Subsequent principal components can be found by Gram-Schmidt orthogonalisation. This is possible to do in Microsoft Excel, e.g. using the Solver Add-in, but is rather convoluted. With this data the θ (in radians) corresponding to the first principal component is -0.7136 and that for the second (i.e. in this case, other) principal component is 0.8572

Simpler is to use a standard statistics package that identifies the principal component weights directly from the underlying data or to use the corresponding web service function available via the Nematrian website, i.e. MnPrincipalComponentsWeights which using this data returns an array:

Principal Component	Multiplier to apply to a_i	Multiplier to apply to b_i
1	0.756	-0.655
2	0.655	0.756

The Nematrian website also provides a web service function which calculates the principal component series directly, i.e. applies these weights to the underlying (adjusted) data series. This is MnPrincipalComponents.

A.4.2(b)

[ExtremeEventsQuestionsAndAnswers4_2b]

Q. Given a linear combination of Index A and B what is the maximum possible kurtosis of a linear combination of A and B?

The linear combination of the form $a_i \cos\theta + b_i \sin\theta$ with the highest kurtosis can be found by finding the θ $(0 \leq \theta < 2\pi)$ which maximises the kurtosis of this combination. As kurtosis is scale invariant, this will also provide the maximum possible kurtosis. As *A* has a much higher kurtosis than *B* (4.62 vs -0.84) we expect that the resulting linear combination will be more biased towards A than was the one in (a) that maximised standard deviation and hence characterised the first principal component.

As in (a) we can find the solution using Microsoft Excel, e.g. using the Solver Add-in. It has θ = -0.088, i.e., as expected, close to A. The kurtosis of the resulting combination is 4.71.

A.4.2(c)

[ExtremeEventsQuestionsAndAnswers4_2c]

Q. More generally, can you identify two series where the linear combination of the series with the least variance is also the one with the maximum kurtosis? Hint: try identifying two series with very few terms in them as it simplifies the relevant mathematics.

Two simple series that have this property are:

Period	A	B
1	0.25	1
2	0.25	1
3	0.25	-1
4	0.25	-1
5	-1	0
Mean	0	0
Standard deviation	0.56	1
(Excess) kurtosis	5	-3

These series both have zero mean, and are chosen so that the main contributor to kurtosis in series A is an observation which is zero in Series B. We find that the linear combination with the maximum kurtosis is the same as the linear combination with the minimum standard deviation, i.e. 100% series A. This would remain true even if we multiplied B by anything greater than approximately 0.6.

A.4.2(d)

[ExtremeEventsQuestionsAndAnswers4_2d]

Q. What lessons might you draw from (c) in terms of use of variance or variance-related risk statistics when used to estimate the likelihood of extreme events?

The main lesson that can be drawn from (c) is that if observations are likely to be drawn from quite fat-tailed distributions then the dominant drivers of likelihood of extreme events occurring may relate to the extent to which a possible contribution is fat-tailed, rather than the extent to which the contribution has high variability. This is essentially the same point as was noted in Chapter 4 of the book Extreme Events.

Answers/Hints

Chapter 5: Traditional Portfolio Construction Techniques

[Nematrian website page: ExtremeEventsQuestionsAndAnswers5, © Nematrian 2015]

Questions:

- Question A.5.1
- Question A.5.2

Specimen Question A.5.1

[ExtremeEventsQuestionsAndAnswers5_1q]

You are an asset allocator selecting between five different asset categories A1 to A5 using mean-variance optimisation. Your expected future returns, standard deviations of returns and correlations for the asset categories are as follows:

			Expected correlation coefficients				
	Expected return (%pa)	Expected standard deviation (%pa)	A1	A2	A3	A4	A5
A1	3.0	2	1				
A2	5.0	4	0.4	1			
A3	6.0	8	-0.6	-0.5	1		
A4	7.0	14	0.0	-0.4	0.2	1	
A5	7.5	15	-0.4	-0.4	0.6	0.3	1

(a) Plot the efficient frontier and the asset mixes making up the points along the efficient frontier, assuming that risk-free is to be equated with zero volatility of return and that no non-negative holdings are allowed for any asset category.

Answer/Hints

(b) Show how the efficient frontier and the asset mixes making up the points along the efficient frontier would alter if risk-free is equated with 50% in Asset A1 and 50% in Asset A2.

Answer/Hints

(c) In what circumstances might a mixed minimum risk portfolio as per (b) apply? Give examples of the types of asset that might then be A1 and A2.

Answer/Hints

Specimen Question A.5.2

[ExtremeEventsQuestionsAndAnswers5_2q]

A colleague has a client for which the mixed minimum risk portfolio as per A.5.1(b) applies. She has invested the client's portfolio as follows:

	Portfolio mix (%)
A1	20
A2	20
A3	20
A4	20
A5	20

(a) You know that she has also used mean-variance optimisation techniques and adopted the same expected covariances as you would have done in A.5.1. What return assumptions might she have adopted when choosing her portfolio mix? Specify mathematically all possible sets of return assumptions she could have adopted and still reached this answer.

Answer/Hints

(b) Suppose that she is just about to adjust her portfolio mix so that it includes no holding in A3 and with the amounts that were invested in A3 in the above table redistributed equally between the remaining four asset categories. What return assumptions might she be adopting when choosing her new portfolio mix? Can you specify mathematically all possible sets of return assumptions she could have adopted and still reached this answer?

Answer/Hints

Answers/Hints

A.5.1(a)

[ExtremeEventsQuestionsAndAnswers5_1a]

Q. Plot the efficient frontier and the asset mixes making up the points along the efficient frontier, assuming that risk-free is to be equated with zero volatility of return and that no non-negative holdings are allowed for any asset category.

The efficient frontier can be found using standard constrained quadratic optimisation techniques, constrained quadratic optimisation techniques and then plotting the results.

This can be done using the Solver add-in that comes as standard with Microsoft Excel, but rather simpler (in our opinion) is to make use of the Nematrian online toolkit or an equivalent. The Nematrian toolkit provides three different tools, each one of which can be used to solve this particular problem:

(a) The asset mix (and corresponding risk and expected return) corresponding to a *single* point along the efficient frontier (i.e. for a specific λ, i.e. risk-reward trade-off parameter) can be found interactively using Nematrian's Example Quadratic Portfolio Optimiser page.

(b) In this case we want to plot the efficient frontier, i.e. we want the risk and expected returns for a range of points along the efficient frontier (and we also want to plot the corresponding asset mixes). We could run (a) for several different efficient frontiers, but this would be quite laborious. Simpler than (a) (particularly if you are likely to carry out several similar exercises) may be to call the corresponding Nematrian MnConstrainedQuadraticPortfolioOptimiser web service function many times using VBA.

(c) Perhaps best in this instance is to piggy-back off of the Nematrian website's SmartChart facility and to use the (standard) Nematrian web function plotting equivalents to (b), i.e. MnPlotQuadraticEfficientFrontier and MnPlotQuadraticEfficientPortfolios respectively. The web function itself returns a string corresponding to the SmartCode of a suitable Smart Chart. The interactive variant returns an entire (temporary) Nematrian SmartChart (which incorporates this SmartCode).

Even simpler, if it exists, is to use a previously established spreadsheet that embeds the relevant Nematrian web functions. For this particular purpose one does exist and can be found here. Other spreadsheets that simultaneously illustrate a selection of related Nematrian web functions can be found here.

Permanent SmartCharts corresponding to (c) are shown below. To equate risk free with zero volatility the minimum risk portfolio should be input as {0,0,0,0,0}. To bar non-negative holdings, place a lower limit of zero on each holding, i.e. the Lower Bounds should be input as {0,0,0,0,0}. Asset weights must add to unity, which can be achieved by suitable choice of Constraint Matrix, Constraint Limits and Constraint Types. A full symmetric correlation matrix needs to be input into Forecast Correlations, i.e. here:

	A1	A2	A3	A4	A5
A1	1	0.4	-0.6	0.0	-0.4
A2	0.4	1	-0.5	-0.4	-0.4

A3	-0.6	-0.5	1	0.2	0.6
A4	0.0	-0.4	0.2	1	0.3
A5	-0.4	-0.4	0.6	0.3	1

Some playing around with the input lambda values may be needed to get a set that span the entire range of efficient portfolios. For the charts below we have used the following input values of lambda:

0.1
1
2
3
4 5
7
10
12
14
16
18
20
25
30
35
40
45
50
60
70
80
90
100
125
150
175
200
300
400
500
600
700
900

Efficient frontier:

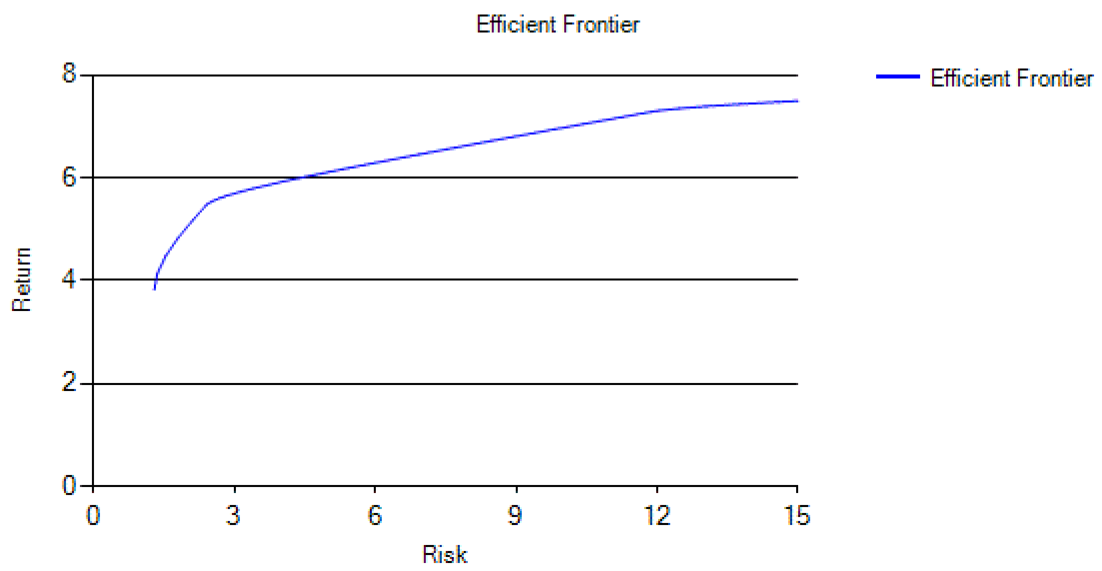

Portfolios forming the efficient frontier

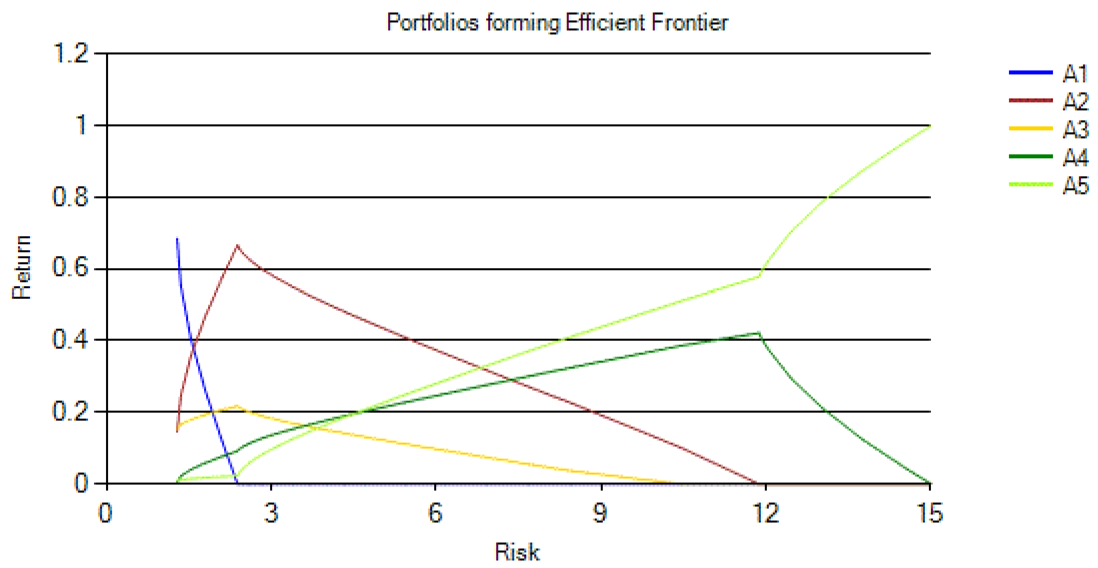

A.5.1(b)

[ExtremeEventsQuestionsAndAnswers5_1b]

Q. Show how the efficient frontier and the asset mixes making up the points along the efficient frontier would alter if risk-free is equated with 50% in Asset A1 and 50% in Asset A2.

Any of the approaches used to answer Question A.5.1(a) can be used to answer this question. If you are using the Nematrian online toolkit (perhaps in conjunction with its SmartChart facility) then the only thing you need to change is to alter the Minimum Risk Portfolio so that it is now {0.5, 0.5, 0, 0, 0}, in which case the answer is:

Efficient frontier:

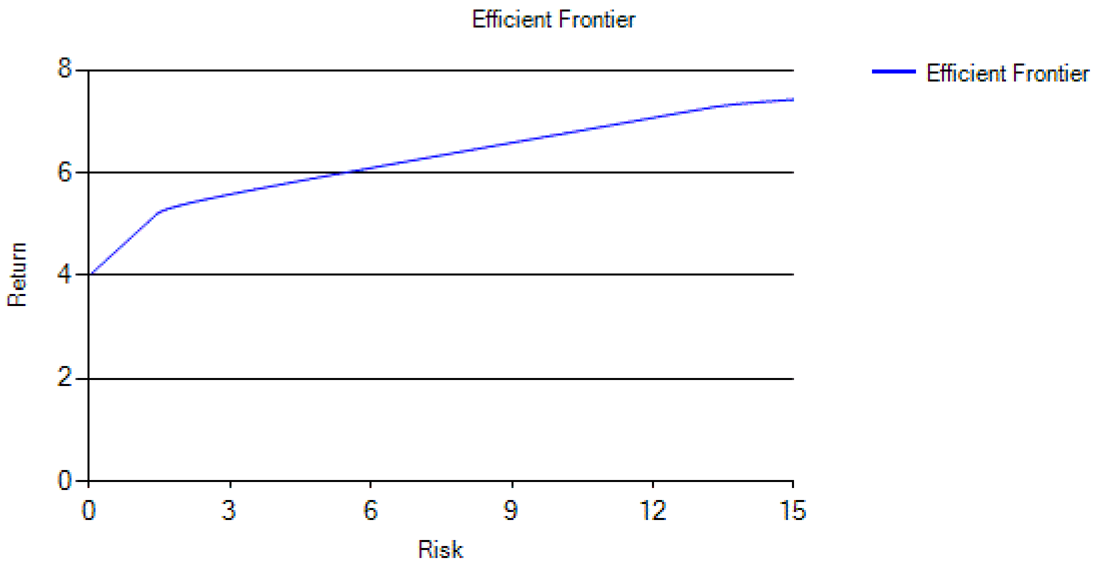

Portfolios forming the efficient frontier

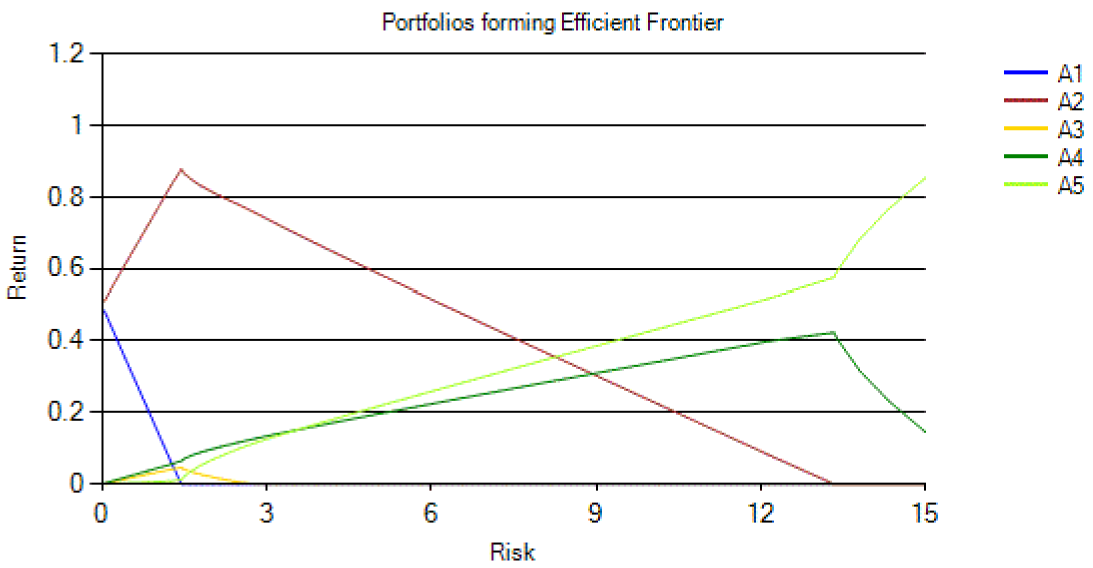

A.5.1(c)

[ExtremeEventsQuestionsAndAnswers5_1c]

Q. In what circumstances might a mixed minimum risk portfolio as per (b) apply? Give examples of the types of asset that might then be A1 and A2.

The most obvious circumstance is where we are investing a portfolio with liabilities that do not look much like cash. For example, a pension fund might have liabilities that are partly fixed in nature (but payable many years into the future) and partly inflation-linked. A suitable minimum risk portfolio might then be a suitable mixture of fixed interest securities and index-linked securities.

A.5.2(a)

[ExtremeEventsQuestionsAndAnswers5_2a]

Q. You know that she has also used mean-variance optimisation techniques and adopted the same expected covariances as you would have done in A.5.1. What return assumptions might she have adopted when choosing her portfolio mix? Specify mathematically all possible sets of return assumptions she could have adopted and still reached this answer.

The return assumptions that your colleague might have used can be derived using implied alphas. These are the (mean) returns that your colleague needs to have assumed will apply for the portfolio in question to be deemed efficient.

No constraints are biting in relation to the specific portfolio mix in question (apart from the constraint that all weights add to unity). If returns assumed for each asset category are characterised by the vector $\mathbf{r}$ and if the portfolio weights are characterised by the vector $\mathbf{x} = (0.2,0.2,0.2,0.2,0.2)^T$ then we need $r - \lambda x^T x$ to be at a maximum, subject to the constraint that $\sum x_i = 1$. Using Lagrange multipliers as per Section 5.10 of Extreme Events, this implies that:

$$\mathbf{r} = 2\lambda\mathbf{V}\mathbf{x} - L.\mathbf{1}$$

It is relatively straightforward to calculate an example $\mathbf{r}$ that satisfies this equation using Microsoft Excel. Alternatively, you could use the Nematrian web function, MnReverseQuadraticPortfolioOptimiser, which assumes $L = 1$ but allows an arbitrary choice of λ (the input parameter is there called the 'TradeOffFactor'). Choosing $\lambda = 0.1$ gives the following possible values for r_i:

i	r_i
1 (i.e. A1)	-0.288
2 (i.e. A2)	-0.864
3 (i.e. A3)	2.656
4 (i.e. A4)	5.18
5 (i.e. A5)	6.48

A.5.2(b)

[ExtremeEventsQuestionsAndAnswers5_2b]

Q. Suppose that she is just about to adjust her portfolio mix so that it includes no holding in A3 and with the amounts that were invested in A3 in the above table redistributed equally between the remaining four asset categories. What return assumptions might she be adopting when choosing her new portfolio mix? Can you specify mathematically all possible sets of return assumptions she could have adopted and still reached this answer?

A similar approach as used in Question A.5.2(a) can be used here, but with $\mathbf{x} = (0,0.25,0.25,0.25,0.25)^T$ rather than $\mathbf{x} = (0.2,0.2,0.2,0.2,0.2)^T$ and with one further important difference. This is that it is impossible to identify precisely the magnitude of the implied alpha for the first asset class relative to the others – we merely know that it is worse than a certain value, since the no short-selling constraint ($x_1 \geq 0$ applies to it.

To solve this question we therefore need to select a very negative return for r_1 and find implied alphas for $r_2, \dots, r_5$ as if we only had a 4 asset problem. We would then carry out optimisation exercises using these values for r_i but selectively increasing r_1 until x_1 just starts to become positive for the chosen value of λ. This has introduced a further parameter in the specification of possible asset mixes, but one that is constrained to fall below a given value.

N.B. Conceptually the same sort of approach can be used if we have more than one inequality biting although the trial and error approach used above would soon become relatively impractical.

Chapter 6: Robust Mean-Variance Portfolio Construction

[Nematrian website page: ExtremeEventsQuestionsAndAnswers6, © Nematrian 2015]

Questions:

- Question A.6.1
- Question A.6.2

Specimen Question A.6.1

[ExtremeEventsQuestionsAndAnswers6_1q]

A daily series that you are analysing seems to have a small number of extreme movements that look suspiciously like errors to you.

(a) To what extent should you exclude such observations when developing a robust portfolio construction algorithm?

Answer/Hints

(b) What sorts of circumstances (applying to what sorts of financial series) might lead to extreme movements that are not actually errors?

Answer/Hints

(c) What other sorts of observations arising in financial series might be ones that you would question?

Answer/Hints

Specimen Question A.6.2

[ExtremeEventsQuestionsAndAnswers6_2q]

You are an asset allocator selecting between five different asset categories as per A.5.1 and you believe that the covariances between the asset categories are as set out in A.5.1. The 'market' involves the following asset mix and is viewed as implicitly involving a minimum risk portfolio which is 100% invested in asset class A1. The mandate does not allow short sales.

	Market mix (%)	Your views relative to those implicit in the market mix
A1	10	-0.5
A2	15	0.0
A3	20	+0.5
A4	30	-0.5
A5	15	+0.5

(a) Set out how you would use the Black-Litterman approach to identify a robust optimal asset mix for the portfolio.

Answer/Hints

(b) What additional information would you need before you could decide what is the most suitable asset allocation for this portfolio?

Answer/Hints

(c) Making plausible assumptions about this additional information, suggest a suitable asset allocation for this portfolio. Ideally create a spreadsheet that takes this information as an input and selects the most suitable asset allocation given this information.

Answer/Hints

(d) Set out a series of return and covariance assumptions which results in the same mean-variance asset allocation as in (c) but without using the Black-Litterman methodology.

Answer/Hints

Answers/Hints

A.6.1(a)

[ExtremeEventsQuestionsAndAnswers6_1a]

Q. To what extent should you exclude such observations when developing a robust portfolio construction algorithm?

If the observations really are errors then they should be excluded from the underlying return series. In this respect 'robust' portfolio construction is not really like 'robust' regression in which we (usually) give reduced (but not nil) weight to outliers. Instead, with robust portfolio construction we should clean the dataset as far as possible, and include outliers if they really are valid data points but exclude them if they are not.

A.6.1(b)

[ExtremeEventsQuestionsAndAnswers6_1b]

Q. What sorts of circumstances (applying to what sorts of financial series) might lead to extreme movements that are not actually errors?

Examples of extreme movements that are not actually errors include:

- Large price movements arising from one off market events such as take-over offers

- Large price movements that reflect drops in value due to payment of dividends or coupons to investors

- One-off movements that correspond to 'new' news being revealed about the investment (e.g. a fraud, an announcement of signing of a major contract or discovery of a mineral find, or some change in the tax position)

It is often difficult for anyone other than a market professional specialising in the relevant market area to spot whether such movements are 'genuine'. Data providers often respond to this challenge by seeking data from a variety of sources and comparing them against each other. However, even this may not be fully robust. All the available data feeds may merely repeat the same 'error' coming from further upstream. If the apparent discrepancy is large enough then this can lead to review by the original data providers themselves and/or by the regulator (e.g. to see if there is some evidence of market abuse or other type of failure in market price formation).

A.6.1(c)

[ExtremeEventsQuestionsAndAnswers6_1c]

Q. What other sorts of observations arising in financial series might be ones that you would question?

These would include ones that appear intrinsically implausible in the context of how effective price formation 'ought' to work. Perhaps the most important of these is if the prices show an unreasonably smooth progression. This might indicate that the prices are 'stale', i.e. not current, particularly if the price series is constant. Or, it might indicate that the prices were fictitious and perhaps fraudulent, if the price and/or return series exhibits an implausibly high risk/return trade-off.

A.6.2(a)

[ExtremeEventsQuestionsAndAnswers6_2a]

Q. Set out how you would use the Black-Litterman approach to identify a robust optimal asset mix for the portfolio.

An example of how to apply the Black-Litterman approach is included in an example spreadsheet that illustrates a wide range of Nematrian web functions linked to mean-variance portfolio optimisation, see here.

A.6.2(b)

[ExtremeEventsQuestionsAndAnswers6_2b]

Q. What additional information would you need before you could decide what is the most suitable asset allocation for this portfolio?

The most important additional pieces of information that you would need are:

- The degree of credibility to give to the 'market' in such a computation; and

- The risk aversion of the investor, in effect the value of λ (the risk/return trade-off factor) that should be used in any subsequent optimisation exercise for the client in question.

A.6.2(c)

[ExtremeEventsQuestionsAndAnswers6_2c]

Q. Making plausible assumptions about this additional information, suggest a suitable asset allocation for this portfolio. Ideally create a spreadsheet that takes this information as an input and selects the most suitable asset allocation given this information.

There is no real way to make plausible assumptions about the additional information set out in A.6.2(b), except to note that, for example:

i. The degree of credibility to give to the market in effect corresponds to how large should be the adjustment to the implied alphas for a unit sized investment view. The conventional way of applying Black-Litterman is to identify suitable return assumptions that result in the 'market portfolio' being efficient. These assumptions are usually framed so that the spread of mean returns is broadly in line with the spread of long-term average returns seen in practice, e.g. perhaps a c. 3-5% difference between the expected return for the lower returning asset class and that for the highest returning asset class. Although no scale is given for the views in the question, we might implicitly assume that +1 or -1 corresponded to a strong view, e.g. perhaps a view similar in magnitude to this range; and

ii. The Black-Littermann approach is most usually applied to equity portfolios. A suitable default value for λ (the risk/reward trade-off) might be one that corresponds to a portfolio that expresses approximately the same overall risk (versus the minimum risk position) as the 'market' portfolio.

A.6.2(d)

[ExtremeEventsQuestionsAndAnswers6_2d]

Q. Set out a series of return and covariance assumptions which results in the same mean-variance asset allocation as in (c) but without using the Black-Litterman methodology.

This is merely another example of calculation of implied alpha. It can therefore be achieved by running the optimisation outputs provided by the answer to A.6.2(c) through a reverse optimiser.

Chapter 7: Regime Switching and Time-Varying Risk and Return Parameters

[Nematrian website page: ExtremeEventsQuestionsAndAnswers7, © Nematrian 2015]

In Section 7.2 we set out formulae for the means and covariance matrices for the conditional probability distributions involved in a RS model that involved just two multivariate normal regimes.

Questions:

- Question A.7.1
- Question A.7.2
- Question A.7.3

Specimen Question A.7.1

[ExtremeEventsQuestionsAndAnswers7_1q]

Suppose we had K regimes rather than just 2. How would the formulae in Equations 7.1 to 7.6 generalise in such circumstances?

Answer/Hints

Specimen Question A.7.2

[ExtremeEventsQuestionsAndAnswers7_2q]

Suppose we revert to the 2 regime case and we also have just two assets. Derive formulae for the skew and kurtosis of the conditional probability distributions.

Answer/Hints

Specimen Question A.7.3

[ExtremeEventsQuestionsAndAnswers7_3q]

Suppose that the regimes in A.7.2 have the following distributional characteristics and transition probabilities:

Distributional characteristics

	Regime 1			Regime 2		
	Means	Covariances		Means	Covariances	
Asset		A	B		A	B
A	0.03	0.01	0.005	0.01	0.02	0.02
B	0.06	0.005	0.02	-0.01	0.02	0.03

Transition probabilities

	State at start of next period	
State at start of this period	Regime 1	Regime 2

Regime 1	p	$1-p$
Regime 2	$1-q$	q

where $p = 0.2$ and $q = 0.3$.

(a) What are the conditional means and covariance matrices of the distributions for the next period if the world is in (i) Regime 1, (ii) Regime 2?

Answer/Hints

(b) What in broad terms is the impact on optimal portfolios of increasing p and q by equal amounts, i.e. the likelihood that we switch states over the coming period whatever the state of the world we are currently in?

Answer/Hints

Answers/Hints

A.7.1

[ExtremeEventsQuestionsAndAnswers7_1a]

Q. Suppose we had K regimes rather than just 2. How would the formulae in Equations 7.1 to 7.6 generalise in such circumstances?

[Page under development]

A.7.2

[ExtremeEventsQuestionsAndAnswers7_2a]

Q. Suppose we revert to the 2 regime case and we also have just two assets. Derive formulae for the skew and kurtosis of the conditional probability distributions.

[Page under development]

A.7.3(a)

[ExtremeEventsQuestionsAndAnswers7_3a]

Q. What are the conditional means and covariance matrices of the distributions for the next period if the world is in (i) Regime 1, (ii) Regime 2?

[Page under development]

A.7.3(b)

[ExtremeEventsQuestionsAndAnswers7_3b]

Q. What in broad terms is the impact on optimal portfolios of increasing p and q by equal amounts, i.e. the likelihood that we switch states over the coming period whatever the state of the world we are currently in?

[Page under development]

Chapter 8: Stress Testing

[Nematrian website page: ExtremeEventsQuestionsAndAnswers8, © Nematrian 2015]

Questions:

- Question A.8.1
- Question A.8.2
- Question A.8.3

Specimen Question A.8.1

[ExtremeEventsQuestionsAndAnswers8_1q]

The wording of the Solvency II Directive indicates that in the Solvency II Standard Formula SCR the operational risk charge should be added to the charge for all other risk types without any diversification offset. This means that there is little scope in practice for a less onerous approach to be adopted by any particular national insurance regulator.

(a) You are an EU insurance regulator who does not wish to rely merely on the wording of the Directive to justify no diversification offset between operational risk and other types of risk. What other arguments might you propose for such an approach?

Answer/Hints

(b) Conversely, you are an insurance firm which is attempting to persuade the regulator to allow it to use an internal model and you would like to be able to incorporate a significant diversification offset between operational risk and other elements of the SCR. What arguments might you propose to justify your position?

Answer/Hints

(c) You are another EU insurance firm. Why might you have an interest in how successful or unsuccessful the firm in (b) is at putting forward its case for greater allowance for diversification between operational and other risks?

Answer/Hints

Specimen Question A.8.2

[ExtremeEventsQuestionsAndAnswers8_2q]

Summarise the main risks to which the following types of entity might be most exposed (and which it would be prudent to provide stress tests for if you were a risk manager for such an entity):

(a) A commercial bank

Answer/Hints

(b) A commercial bank

Answer/Hints

(c) A life insurance company

Answer/Hints

(d) A non-life insurance company

Answer/Hints

(e) A pension fund

Answer/Hints

Specimen Question A.8.3
[ExtremeEventsQuestionsAndAnswers8_3q]

You are a financial services entity operating in a regulated environment in which your capital requirements are defined by a nested stress test approach as described in Section 8.3.3. Your investment manager has approached you with a new service which will involve the manager explicitly optimising your investment strategy to minimise your regulatory capital requirements and is proposing that you pay them a performance related fee if they can reduce your capital requirements. Set out the main advantages and disadvantages (to you) of such a strategy.

Answer/Hints

Answers/Hints

A.8.1(a)

[ExtremeEventsQuestionsAndAnswers8_1a]

Q. You are an EU insurance regulator who does not wish to rely merely on the wording of the Directive to justify no diversification offset between operational risk and other types of risk. What other arguments might you propose for such an approach?

Three possible lines of argument are:

i. The correct level at which to set regulatory capital depends heavily on the magnitude and drivers of relatively extreme events. Different risk factors tend to be more correlated in such circumstances (relative to their behaviour under more usual outcomes). Adopting a high correlation in such a computation implicitly recognises the lack of risk diversification that typically applies in such circumstances.

ii. A more sophisticated variant of (i) might observe that the key statistic that ought ideally to be most focused on in regulatory solvency computations (as far as the regulator is concerned) is Expected Shortfall. This metric is typically even more tail dependent than Value-at-Risk, the metric most commonly used in regulatory computations at present. Moreover, if essentially all cases where Expected Shortfall is large include an operational risk failure then the relevant Expected Shortfalls may be approximately additive, even if the actual tail dependency takes a more complex form.

iii. A very prudent approach to incorporating operational risk may increase the incentives on firms to minimise this type of risk. This may be considered desirable by regulators either because they in general think that such risks are underemphasised by firms or because they think that it is particularly desirable to incentivise firms to tackle these risks. Some types of risk can be expected to be compensated for by additional reward, but this is arguably less likely to be the case with operational risks, which are often asymmetric and very largely only downside orientated. Mitigating operational risk, if it is not too expensive, may therefore be viewed as close to providing firms with the opportunity to benefit from a 'free lunch'.

A.8.1(b)

[ExtremeEventsQuestionsAndAnswers8_1b]

Q. Conversely, you are an insurance firm which is attempting to persuade the regulator to allow it to use an internal model and you would like to be able to incorporate a significant diversification offset between operational risk and other elements of the SCR. What arguments might you propose to justify your position?

Essentially, your line of reasoning would need to be one that came up with the opposite conclusion to that set out in the answer to A.8.1(a). This might involve attempting to reason as follows:

i. Different risks do not all in practice hit at the same time. Most business failures involve a complex interplay of factors, and although operational risk failings are a common contributory factor they are often poorly correlated with other types of risk present in the SCR computation.

ii. A regulatory framework that makes overly cautious assumptions about correlations between different risk factors is likely to be overly prudent. This would discourage use of efficient business structures and could lead to excessive capital charges, both of which may result in businesses exiting the market and ultimately probable delivery of poorer value-for-money to end customers.

A.8.1(c)

[ExtremeEventsQuestionsAndAnswers8_1c]

Q. You are another EU insurance firm. Why might you have an interest in how successful or unsuccessful the firm in (b) is at putting forward its case for greater allowance for diversification between operational and other risks?

i. You may also want to argue for reduced regulatory capital requirements for your own business

ii. Conversely, if industry-wide support mechanisms are ultimately paid for by the rest of the industry, then you may not want other firms to be unduly weakly capitalised, because you might then end up contributing to any bail-outs that they might benefit from or you might be worried about systemic reputational issues that might hit the industry as a whole.

   N.B. This line of argument seems less prevalent than (i), perhaps because there is less perceived incentive on firms to pursue it. This may be because stronger firms take the view that failure of weak firms might reduce capacity and hence increase the profitability of survivors. Alternatively, they may believe that if the stress is severe enough then losses will eventually be picked up primarily by the public purse rather than by other industry participants.

Answer A.8.2(a)

[ExtremeEventsQuestionsAndAnswers8_2a]

Q. Summarise the main risks to which the following types of entity might be most exposed (and which it would be prudent to provide stress tests for if you were a risk manager for such an entity): (a) A commercial bank

Commercial banks are exposed to a variety of risks, but the most important are usually:

i. Exposure to interest rate movements. Commercial banks carry out maturity transformation as they typically borrow short and lend long. Substantial changes to yield levels (and to the shape of the yield curve) can give them major headaches depending on their aggregate cash flow profile.

ii. Exposure to liquidity squeezes. This was a particular issue for commercial banks during the 2007-2009 credit crisis. Banks reliant on particular markets for sources of funding can run into trouble if these markets dry up. When designing suitable stresses it is also worth considering the precise nature of the funding and the extent to which off-balance sheet arrangements might cease to be off-balance sheet in a stressed scenario. For example, several banks that ran into trouble during the 2007-09 credit crisis were heavy users of 'shadow banking' structures, which ostensibly moved their funding exposures to third party entities. However the liquidity squeeze during the 2007-2009 credit crisis was so severe that these banks often found that they had to support their own special purpose vehicles despite the aim being that these vehicles would be ring fenced away from the bank's own balance sheet in adverse circumstances.

iii. Credit risk and bad debts. A particular issue here is that the credit exposures may turn out not be as diversified as the bank might have hoped. For example, residential and commercial mortgage business may ostensibly appear to involve a very diversified client base. However, it may actually turn out to be less diversified than expected if there is a major economic downturn that coincides with a major decline in general property values.

iv. Operational risk. Like other financial services, commercial banks are exposed to operational risks. These might include fraud (by the company's own employees or directors), or inappropriate incentive elements for employees or within the products it is selling that lead to unhelpful aggregate customer behaviours.

P.S. Similar types of risk (but in other guises) usually arise with other types of financial services entities, which may be one contributory factor in the increasing popularity of the discipline of Enterprise Risk Management.

A.8.2(b)

[ExtremeEventsQuestionsAndAnswers8_2]

Q. Summarise the main risks to which the following types of entity might be most exposed (and which it would be prudent to provide stress tests for if you were a risk manager for such an entity): (b) An investment bank

- Market risk. Investment banks commonly have significant market exposures, either in the form of inventory that they are using to provide market making services, or in the form of proprietary positions. The types of market risk (e.g. equity, credit, commodity, interest rate) that any particular bank is exposed to can vary significantly both by bank and through time. Investment banks may also be particularly exposed to warehousing risk during M&A transactions.

- Credit risk. Investment banks may nowadays trade credit exposures using CDS etc. These sorts of credit risk exposures would usually be thought of as a form of 'market risk', being just another type of market making activity. Investment banks may also have more traditional forms of counterparty credit exposure via their trading activities (which they may aim to mitigate using collateralisation techniques).

- Liquidity risk. Investment banks fund their businesses in a variety of ways, many of which explicitly or implicitly require the instruments they have on their balance sheet to be acceptable collateral to others. If liquidity dries up then (and particularly if the instruments they might otherwise have then been relying on to provide them with access to alternative funding sources at the same time prove difficult to value) then investment banks can become very exposed to liquidity squeezes.

- Credit rating downgrade. Some of the funding sources that an investment bank may be relying on can be sensitive to the credit rating assigned to the bank (e.g. because this interacts with the collateralisation processes applicable on derivative contracts it may have entered into). Thus a significant credit rating downgrade can in effect freeze availability of fund from one type of source, which can lead to a loss of confidence on the part of others and difficulties accessing other sources as well.

- Operational risk. The large sizes and speed of transactions that an investment bank might enter into seem to make investment banks particularly prone to very large operational risk losses even if they previously thought that they had well managed staff and business processes.

P.S. Similar types of risk (but in other guises) usually arise with other types of financial services entities, which may be one contributory factor in the increasing popularity of the discipline of Enterprise Risk Management.

A.8.2(c)

[ExtremeEventsQuestionsAndAnswers8_2c]

Q. Summarise the main risks to which the following types of entity might be most exposed (and which it would be prudent to provide stress tests for if you were a risk manager for such an entity): (c) A life insurance company

i. Market risk. Many lines of business in life insurance can be heavily exposed to market risk. However, in many cases much of this risk may be borne by policyholders rather than shareholders (if any). For example, with unit-linked business the benefits payable to policyholders may rise and fall as the underlying asset values rise and fall, although there may also be guarantees that may not be fully hedged and the future revenue streams the firm earns on in-force policies may be reduced if asset values decline.

ii. Credit risk. Some elements of life insurance can create material credit exposures, e.g. credit risk in reinsurance contracts that are not collateralised or otherwise passed through to policyholders.

iii. Liquidity risk. Life insurers have historically been viewed as less exposed to liquidity risk than most other financial services entities, if anything being seen as likely providers of liquidity rather than likely consumers. However this is not necessarily the complete picture, see e.g. Liquidity Risk – Its Relevance To Actuaries).

iv. Insurance risk. The nature of insurance involves the assumption and pooling of risk. Many life insurance contract types, e.g. unit-linked savings vehicles, in effect repackage most of the risks and pass them back to policyholders. But some life insurance contract types, e.g. mortality or disability protection business and arguably annuity business may involve a greater proportion of the risk implicit in the contracts remaining with the shareholders (if any). Insurance risk can come in many different forms, many of which may be largely unhedgeable except via (re)insurance contracts.

v. Operational risk. Like other financial services organisations, life insurers are also exposed to operational risks. Given their often substantial interaction with members of the public, and given how life insurance business is often sold, they have perhaps been more exposed than most other financial services entities to mis-selling risk. Some of this risk may reflect changing expectations regarding standards of care that an insurer might owe to its customers.

P.S. Similar types of risk (but in other guises) usually arise with other types of financial services entities, which may be one contributory factor in the increasing popularity of the discipline of Enterprise Risk Management.

A.8.2(d)

[ExtremeEventsQuestionsAndAnswers8_2d]

Q. Summarise the main risks to which the following types of entity might be most exposed (and which it would be prudent to provide stress tests for if you were a risk manager for such an entity): (d) A non-life insurance company

i. Market risk. Non-life insurers tend to invest less in equities than life insurers and tend to have shorter duration fixed income portfolios, given the typically shorter-term nature of their liabilities. They are therefore usually less exposed to market risk than is usually the case with life insurance. However, there are exceptions, e.g. some long-tail business lines or ones in which the precise timing of payments is particularly important.

ii. Credit risk. Most life insurers have a relatively diversified policyholder base. This may be less true for non-life insurers. They can often also be quite dependent on the continued creditworthiness of reinsurers.

iii. Liquidity risk. As with life insurance, liquidity risk is usually thought of as less relevant to general insurers than to most other types of financial services entities. However, some business activities that non-life insurers can be involved with can be more sensitive to liquidity risk, particularly if other market participants are requiring that the insurer posts collateral (plus haircuts) to protect them against possible default by the insurer. This can potentially require the insurer to have funding lines in place in order to be able to fund delivery of collateral if needed, and if these lines dry up then the insurer can be left in an exposed position.

iv. Insurance risk. The nature of insurance involves the assumption and pooling of risk. Unlike life insurance, most such risks in non-life insurance end up falling on the shareholder (only typically obliquely falling on policyholders and then only in a generalised kind of fashion via the workings of the insurance cycle). If the firm prices these risks wrongly, or if strong selection effects mean that the firm ends up insuring the least profitable market segments, then this can rapidly undermine the firm's business model. Many lines of business also have a 'catastrophe' component (i.e. some likelihood of very large adverse claims, even if likelihood is small), so variability of outcomes can also be a problem for inadequately capitalised companies that do not effectively hedge such risks (e.g. by using appropriate reinsurance programmes).

v. Operational risk. Like other financial services organisations, non-life insurers are also exposed to operational risks. It is perhaps more common for these to involve systematic failures in business model design and execution although outright fraud has felled some such insurers.

P.S. Similar types of risk (but in other guises) usually arise with other types of financial services entities, which may be one contributory factor in the increasing popularity of the discipline of Enterprise Risk Management.

A.8.2(e)

[Nematrian website page: ExtremeEventsQuestionsAndAnswers8_2e, © Nematrian 2015]

Q. Summarise the main risks to which the following types of entity might be most exposed (and which it would be prudent to provide stress tests for if you were a risk manager for such an entity): (e) A pension fund

[Page under development]

A.8.3

[ExtremeEventsQuestionsAndAnswers8_3a]

Q. You are a financial services entity operating in a regulated environment in which your capital requirements are defined by a nested stress test approach as described in Section 8.3.3. Your investment manager has approached you with a new service which will involve the manager explicitly optimising your investment strategy to minimise your regulatory capital requirements and is proposing that you pay them a performance related fee if they can reduce your capital requirements. Set out the main advantages and disadvantages (to you) of such a strategy.

Advantages include:

- The strategy should result in reduced regulatory capital and hence higher risk-adjusted return on capital, RAROC (if the firm's overall capital base is adjusted accordingly). If the firm is particularly capital constrained then any mitigation of capital requirements may be attractive to it.

Disadvantages include:

- The answers are likely to be very sensitive to the precise structure of the capital requirements, and hence may change significantly if these requirements change (as generally seems to happen through time). More specifically, it is quite likely that the optimisation process will disproportionately favour elements of the existing capital framework that are most out-of-line with what might turn out to be applicable over the longer term. Thus the adage that unless carefully done optimisation can merely involve error maximisation rather than return maximisation is potentially particularly appropriate here.

- An investment strategy that minimises regulatory capital requirements at the expense of everything else is likely to give insufficient weight to return in any risk/return trade-off.

- Arguably, the problem can be converted into a mathematical exercise the optimal answer for which is known in advance, but with you having insufficient analytical tools to uncover it yourself. Arguably, this is not the sort of exercise for which performance related fees are ideal, since the main justification for such fee arrangements are that they provide better alignment of incentives between the manager and the client.

- Financial services entities are usually nowadays (at least in the UK) expected to determine what they think is an intrinsically appropriate amount of capital to hold irrespective of any particular regulatory capital computations specified by the regulator. The proposed service may result in the company unduly focusing on the regulatory capital requirement (given the proposed fee structure) which means that it runs the risk of giving insufficient emphasis to risks to which it might be exposed but which do not figure prominently in its regulatory capital computation.

Chapter 9: Really Extreme Events

[Nematrian website page: ExtremeEventsQuestionsAndAnswers9, © Nematrian 2015]

Questions:

- Question A.9.1
- Question A.9.2
- Question A.9.3

Specimen Question A.9.1

[ExtremeEventsQuestionsAndAnswers9_1q]

Set out the main types of risk to which a conventional asset manager managing funds on behalf of others might be exposed. Which of these risks is likely to be perceived to be most worth rewarding by its clients?

Answer/Hints

Specimen Question A.9.2

[ExtremeEventsQuestionsAndAnswers9_2q]

You work for an insurance company and you are attempting to promote the use of enterprise risk management within it. Describe some ways in which risks might be better handled when viewed holistically across the company as a whole.

Answer/Hints

Specimen Question A.9.3

[ExtremeEventsQuestionsAndAnswers9_3q]

You work for a bank and have become worried that different business units might be focusing too little on the liquidity needs that their business activities might be incurring. How might you set an appropriate 'price' for the liquidity that the bank as a whole is implicitly providing to its different business units?

Answer/Hints

Answers/Hints

A.9.1

[ExtremeEventsQuestionsAndAnswers9_1a]

Q. Set out the main types of risk to which a conventional asset manager managing funds on behalf of others might be exposed. Which of these risks is likely to be perceived to be most worth rewarding by its clients?

The types of risk that a conventional asset manager might face in this context are many and varied, and readers are advised to consult an expert for more details or to refer to appropriate books/material covering this topic (e.g. by searching third party content referred to in Nematrian's Reference Library).

They include:

(a) *It could suffer from adverse market movements*. Most asset managers are remunerated on an 'ad valorem' basis, i.e. as a percentage of funds under management. Its future revenue stream is therefore dependent on their (market) value.

(b) *It could be carrying on business unprofitably*. Asset managers compete with each other. This competition includes competition on price. Most (good) asset managers are profitable, although severe declines in asset values can drag them into loss, and if it is a new business there may also be start-up costs.

(c) *It could suffer operational failures that result in it needing to compensate clients or otherwise incurring reputational damage*. This is perhaps the most obvious risk that an asset manager faces on its own account (and the one that 'risk managers' in such a firm might often concentrate on), but is not necessarily as large a risk in financial terms as (a) (or (b)).

Operational failures could include pricing errors, incorrect booking of trades, misunderstanding of tax position of clients, staff fraud, etc.

Risks can also interact. For example, the investment manager might fail to invest assets in line with client restrictions or other relevant documentation (a type of operational risk). However, risk of breach of client mandates might increase in volatile market conditions. Moreover, in times of market distress clients and their lawyers may be more creative and more focused on linking losses they have suffered to supposed failures in business processes, moving the loss from being an investment one borne by the client to an operational one borne by the fund manager.

(d) *Its outsourcing arrangements could prove flawed*. For example, the entities to which it has outsourced could fail to provide it with an adequate service leaving the asset manager itself to compensate its clients or to incur the expense necessary to rectify the issue.

(e) *Its investment performance could deteriorate*, leading to client defections and/or failure to win new business.

(f) *There may be weaknesses in non-investment related aspects of its business activities, e.g. client servicing, or deficiencies in contractual arrangements*. Asset managers may offer ancillary services to clients, e.g. member record keeping, which create risks not directly related to investment management. There are also servicing requirements for their core business activities, and these could be weak and in extremis result in loss of clients and/or compensation payments.

(g) *It could grow too rapidly, particularly if the lines of business are capital intensive*. Rapid growth may place strains on investment and business processes. Some types of activities that conventional asset managers carry out may also be more capital intensive (in terms of systems and staffing requirements and/or in terms of regulatory capital).

(h) *It might become overly complicated*. More complicated business models may be more difficult to manage and more prone to operational risk than simpler business models. Of course, they may also be more remunerative.

(i) *It might have inappropriate sales or staff incentives* (including remuneration structures) leading to mis-selling or inappropriate communication of product characteristics, or other behaviours that benefit the staff involved but create risk for the firm.

(j) *It might have inappropriate credit and/or liquidity exposures*, either for itself (e.g. unpaid fees payable by clients) or for its clients for which it is then found liable to compensate them for. Certain types of fund structures can have specific credit risk exposures which may need managing appropriately.

(k) *Its product structures may be rendered redundant due to changes in external factors such as tax and regulatory frameworks.*

(l) *Its funds could suffer a liquidity squeeze, and it might need to buy in the assets in question onto its own balance sheet*.

(m) *It might find itself contributing to industry wide compensation schemes even though it has not itself run into difficulties.* More generally, its business will be influenced by sentiment elsewhere in the industry – if clients are generally disinvesting from an asset class then asset managers in general will struggle to buck this trend.

(n) *It will be subject to a large number of other types of risk more generally applicable to businesses operating in the financial sector, including:*
- *It may have inadequate management information systems and the like to manage its business effectively*.
- *It may be exposed to terrorist risk, IT systems failure, breach of competition rules, health and safety rules etc.*

A.9.2

[ExtremeEventsQuestionsAndAnswers9_2a]

Q. You work for an insurance company and you are attempting to promote the use of enterprise risk management within it. Describe some ways in which risks might be better handled when viewed holistically across the company as a whole.

Possible ways in which risks might be better handled when viewed holistically across the company as a whole include:

(a) A holistic approach should provide a better appreciation of the ways in which different risks might interact. Extreme events often involve more than one risk appearing simultaneously.

(b) A firm should define its overall risk appetite based on its available capital resources (and its access to further capital should it need some). The precise cause of a loss may be less relevant to this than the magnitude of the loss and its financial consequences. Managing risks holistically should allow the firm to adopt a more standardised, less silo orientated approach to risk management and mitigation, making it more likely that proportionate effort will be applied to each possible source of risk to which the firm might be exposed.

(c) Effective holistic risk management should allow a firm to choose more effectively between the different possible strategies it might adopt for risk management, e.g. retention, hedging mitigation or transfer, enhancing shareholder value. It should also provide a better platform for taking full account of different risks in product pricing and business strategy.

(d) A silo approach might result in lessons from one business area not being applied in another, leading to sub-optimal management.

(e) Some types of risk, e.g. business continuity, may naturally span multiple business lines.

(f) Adoption of a holistic approach may confer advantage in relation to external parties. For example, rating agencies may focus on the extent to which the firm appears to be adopting an 'Enterprise Risk Management' approach to risk management. Regulators may specifically require the firm to adopt such approaches and may penalise the company in terms of extra regulatory capital requirements if it is not seen to be viewing and managing risks in a holistic fashion.

(g) The discipline of viewing and considering different risks in coherent holistic manner should result in improved governance and business management.

A.9.3

[ExtremeEventsQuestionsAndAnswers9_3a]

Q. You work for a bank and have become worried that different business units might be focusing too little on the liquidity needs that their business activities might be incurring. How might you set an appropriate 'price' for the liquidity that the bank as a whole is implicitly providing to its different business units?

Liquidity risk is a risk that was arguably underappreciated and under catered for in regulatory frameworks in the run-up to the 2007-09 credit crisis, see e.g. Market Consistency. Some changes have been mandated or proposed since then by regulators seeking to address some of the issues that were uncovered during this crisis.

However, relying exclusively on regulator mandated criteria potentially risks failing to heed one of the lessons of this crisis. This is that personnel (even up to Board level) may often be overly willing to design business strategies around what might look good from a regulatory perspective rather than what a more robust analysis of business needs might dictate.

The most effective way of setting a price for liquidity is likely to be to identify what the 'market' would charge if the business line in question had to source the liquidity externally. Any other value potentially runs the risk of undercharging for the implicit support that the business unit might be receiving from the corporate centre (if too low), or stifling activity and potential business opportunity (if too high). Superimposed on this are likely to be refinements set at a corporate level explicitly designed to reduce or increase the tendency of each business unit to market products that 'consume' liquidity, to control the overall level of liquidity risk that the firm is likely to be running.

In practice, this requires proxies capable of providing a guide as to the amount of 'liquidity' that any particular product consumes. Liquidity risk has the characteristic that it typically manifests itself with low probability but high severity, i.e. painful outcomes disproportionately correspond to outcomes in the tail of the distribution of potential future outcomes. It may therefore be very hard to estimate precisely. However, this is no reason not to try to do so; indeed if anything it makes the need to do so even greater, even if the resulting proxies may need to rely more on stress test outcomes and other sorts of scenario analyses than many other types of market risk.

One outcome of the 2007-09 credit crisis was the development of a 'term structure' to liquidity risk, with interest rate swap rates differentiating according to the 'refresh' frequency implicit in the swap in question. Thus swap rates involving exchange of fixed for floating payments became differentiated according to whether the reference rate underlying the floating leg was overnight, 1 month, 3 month, 6 month Libor etc. (for the same overall swap maturity date). The shorter the period between successive resets, the easier it is for an investor to move his money away from a bank that appears to be heading for default before it gets there. This could be used as a way of identifying the market-implied price of the liquidity benefit that a bank might gain from locking up (floating rate) funding sources for longer time periods or might expect to be compensated for if it does the same in reverse.