The Actuarial Profession
making financial sense of the future

Seth Eshun, Will Machin, James Sharpe
Members of Extreme Events Working Party

Extreme Value Theory
For a 1-in-200 event

26th January 2012

Contents

- Introduction to EVT
- How to fit EVT to a given set of data
- Case study – EVT fitted to equity returns
- Comparison of EVT to other methods of determining 1 in 200
- Simulation based exercise to validate EVT

Introduction to EVT

- Extreme Value Theory(EVT) is a statistical approach that allows a practitioner to model the occurrence of extreme events with relatively small amounts of extreme data.
- A key difference between EVT and other statistical approaches is that, in EVT we fit a distribution to a subset of the available data, whiles in other statistical approaches, we fit a chosen distribution to the entire set of data.
- This difference can be viewed positively or negatively. Negatively as judgement is required in selecting the subset of data to use. Positively, as EVT allows us to concentrate on the part of the distribution we are most interested.

2

Introduction to EVT

- There are two EVT approaches:
  - Block Maxima Models (BMM): this is the traditional approach and it involves the following:
    - Grouping the data into samples/blocks;
    - Calculating the maximum observation in each block;
    - Fitting the Generalised Extreme Value (GEV) distribution to the maxima of the blocks; and
    - Estimating the risk measure we are interested in from the fitted Generalised Extreme Value (GEV) distribution.
  - The Peak over Threshold: this is a more recent technique which involves the following:
    - Selecting a threshold that defines which observation are included in modelling;
    - Calculate the exceedances (this is the excess of the observations over the threshold);
    - Fit the Generalised Pareto Distribution (GPD) to the exceedances or use the Hill estimators to fit the excedances ; and
    - Compute the measure of risk that is desired.
- The POT is preferred over the BMM because data is used more efficiently.

3

Introduction to GPD

- The GPD distribution is based on a concept similar to that of the central limit theorem.
- It assumes that the maximum values of a sample of independent and identically distributed random variable would approximately follow a GPD distribution.
- Thus GPD is commonly used to model the tail of other distributions. It is defined by three parameters and its cumulative distribution function is:

$$F_{(\xi,\mu,\sigma)}(x) = \begin{cases} 1 - \left(1 + \frac{\xi(x-\mu)}{\sigma}\right)^{-1/\xi} & \text{for } \xi \neq 0, \\ 1 - \exp\left(-\frac{x-\mu}{\sigma}\right) & \text{for } \xi = 0. \end{cases}$$

  - μ is the location parameter;
  - σ > 0 the scale parameter; and
  - ξ is the shape parameter.

4

Introduction to GPD

- The following are examples of tail behaviours where the GPD can be used as an approximation:
  - Decreasing Exponential: such as the gamma, normal and log-normal. Achieved by setting the shape parameter of the GPD to zero;
  - Decreasing Polynomial: such as Student's t and Cauchy. Achieved by setting the shape parameter of the GPD to a positive number;
  - Finite: such as the beta distribution. Achieved by setting the shape parameter of the GPD to a negative number.
- As the name suggests, the GPD simplifies to other distributions under specific conditions. For example:
  - When ξ =0, it simplifies to an exponential distribution;
  - When ξ >0, it simplifies to an ordinary Pareto distribution; and
  - When ξ <0, it simplifies to a Pareto type II distribution.

5

How to apply EVT to a set of data

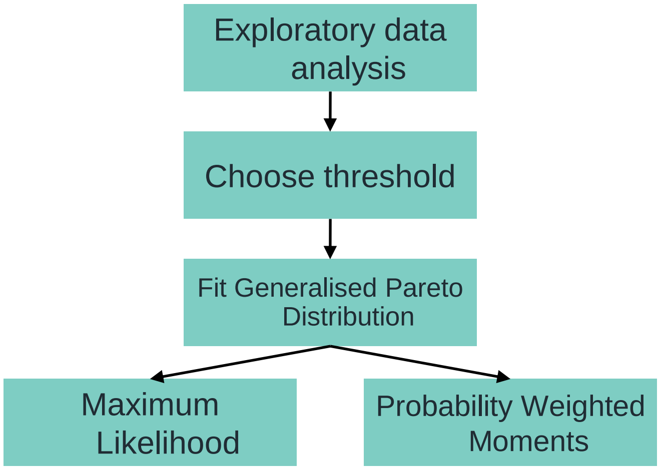

6

How to apply EVT to a set of data

- The first step is to analyse the available data to assess whether the data, is fat tailed enough for us to apply EVT to it. The key tools used in this analysis are the Q-Q plots and the mean excess plots.
  - The Q-Q plot is a graphical means of assessing whether a sample data follows a given probability distribution.
  - The mean excess function express the mean of the excess over a threshold as a function of the threshold. The mean excess function for the GPD is a linear function of the threshold u.
- The next step in applying EVT is selecting an appropriate threshold.
  - A lower threshold reduces the variance of the estimates of the GPD model. However, a lower threshold can introduce bias in the data.
  - A higher threshold reduces bias but increase the volatility of the estimate of the GPD distribution.
  - The mean excess analysis may be used to select an optimum threshold.
  - An alternative approach is to fit the GPD to the data using different thresholds.
- Finally the parameters of the GPD can be estimated by any of the the following approaches:
  - Maximum likelihood;
  - Probability-weighted moments.

7

Case Study - Fitting GPD to Equity returns

- This section of the presentation looks at applying the Generalised Pareto distribution (GPD) to equity returns
- Two data sets were used:

Data	Range
DMS	Annual equity returns 1900-2008
MSCI	Monthly year-on-year returns 1969-2011

Results – GPD distribution

- Estimates of the 99.5% quantile, using a threshold of 9%

Data	DMS estimate	MSCI estimate
UK	-42%	-47%
USA	-43%	-40%
Belgium	-52%	-66%
Denmark	-38%	-42%

Exploratory data analysis – QQ plots of negative annual returns vs. Exponential (GPD ξ=0) UK

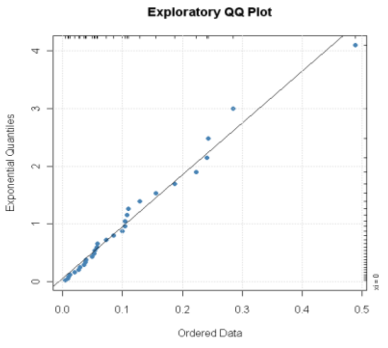

Figure 1: QQ plot of UK simple returns

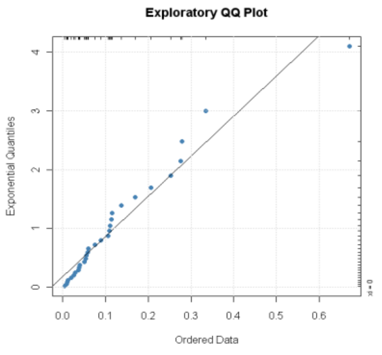

Figure 2: QQ plot of UK log returns

QQ plots - Interpretation

- QQ plot with a straight diagonal line of data points from the bottom left of the chart to the top right indicates an exponential distribution is a relatively good fit to the tail of this data
- A concave (ie starting bottom left and curving round in the top half of C shape to a horizontal line) shape to the QQ plot indicates fatter tail than the exponential distribution; so would suggest fitting a GPD with a $\xi > 0$
- A convex shape (ie starting bottom left and curving round in the right half of a U shape to a vertical line) to the QQ plot indicates a thinner tail than the exponential distribution with $\xi < 0$

QQ plots - Interpretation – UK data

- From the QQ plots, the exponential distributions appears a relatively good fit to the negative annual returns
- This indicates that annual returns in the UK from 1900 - 2008 have not been particularly fat tailed
- Similar results seen for other major EU equity returns, with only German log returns indicating a tail potentially fatter than the exponential

Exploratory data analysis – Mean Excess plots of negative annual returns

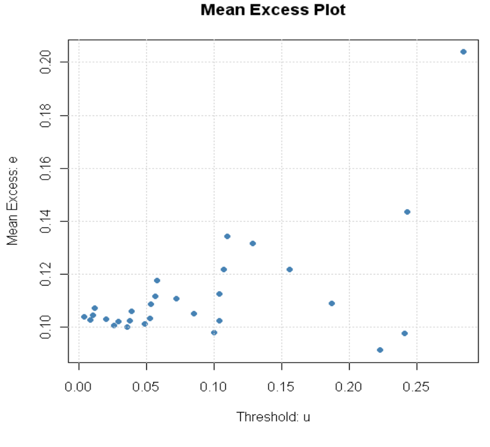

Figure 3: QQ plot of UK simple returns

Mean Excess Plots - Interpretation

- The mean excess plot is a tool used to aid the choice of threshold and also to determine the adequacy of the GPD model in practice.
- A characteristic of a fat tailed GPD type distribution with positive shape parameter is a straight line from bottom left to top right of the mean excess plot
- A mean excess plot with a downwards sloping line from top left to bottom right indicates thin tailed behaviour. A straight horizontal line indicates exponential type behaviour
- There is some significant up and down behaviour in the UK mean excess plots, with a broad upwards trend

Generalised Pareto Distribution – Calibrated to UK annual returns

- The GPD is calibrated to the annual returns based on the points over a certain threshold
- Calibrations are made at each threshold
- Five plots are presented:
  1. The number of points above each threshold
  2. The 99.5th percentile for each calibration at each threshold
  3. The shape parameter ξ for each calibration at each threshold
  4. The standard error for the shape parameter for each calibration at each threshold
  5. The maximum log likelihood for each calibration at each threshold

Generalised Pareto Distribution – Calibrated to UK annual returns – Maximum Likelihood Estimate

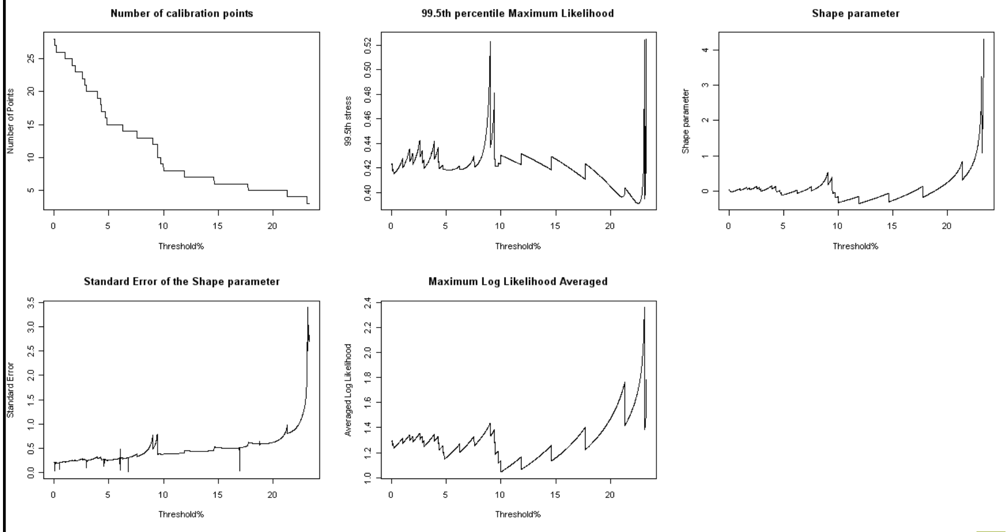

16

Generalised Pareto Distribution – Calibrated to UK annual returns – Maximum Likelihood Estimate

- Starting in the top left, the first plot shows how the number of points in the tail decreases as the calibration moves further into the tail
- The second plot shows the calibration of the 99.5th percentile at each of the thresholds. We can see how this varies as the calibration moves into the tail. There is a significant spike at the -10% threshold where a number of points are clustered. Otherwise the calibration of the 99.5th percentile is fairly stable with changes in threshold
- The top right plot shows the shape parameter for the calibration at each threshold. This stays around the 0 level for most of the thresholds, spiking upwards as the number of points in the calibration falls
- The bottom left plot shows the standard error of the shape parameter estimate. The standard error rises as the number of points in the calibration falls
- The final plot shows the maximum log likelihood, which may indicate better calibration for higher values

17

Generalised Pareto Distribution – Calibrated to UK annual returns – Probability Weighted Moments

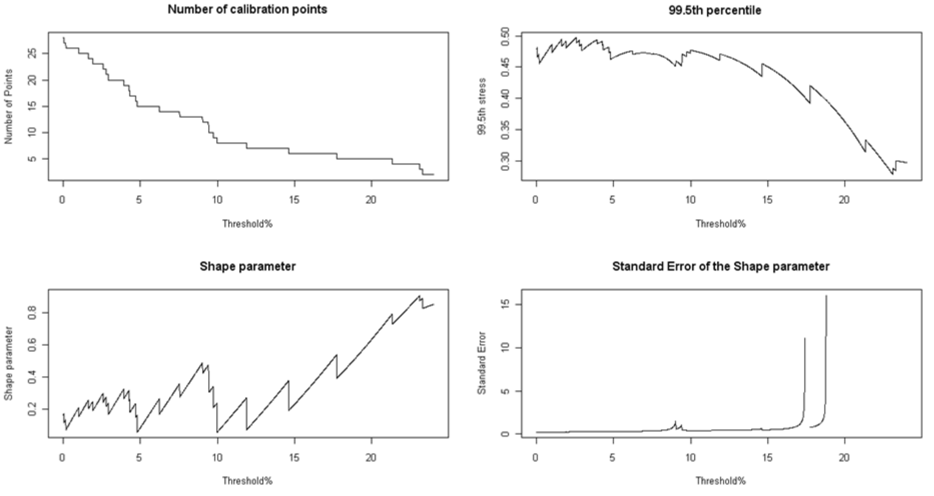

Generalised Pareto Distribution – Calibrated to UK annual returns – Probability Weighted Moments

- Starting in the top left, the first plot shows how the number of points in the tail decreases as the calibration moves further into the tail
- The second plot shows the calibration of the 99.5th percentile at each of the thresholds. We can see how this falls from around 45% to around 30% as the threshold increases
- The bottom left plot shows the shape parameter for the calibration at each threshold. This is generally over 0 for most of the thresholds.
- The bottom right plot shows the standard error of the shape parameter estimate. The standard error rises as the number of points in the calibration falls

Confidence intervals around estimated 99.5th percentiles – UK annual simple returns

- Using three different thresholds with the maximum likelihood calibration method, confidence intervals have been estimated for the 99.5th percentile stress level

	Thresholds		
UK 99.5th percentiles	**-1%**	**-5%**	**-10%**
Lower CI	-29.2%	-28.6%	-27.2%
Best Estimate	-42.2%	-42.9%	-51.3%
Upper CI	-97.7%	-97.7%	-97.7%

- Confidence intervals are very wide reflecting significant uncertainties in calibration of the 99.5th percentile stress
- Confidence intervals created using profile likelihood method [McNeil et al ., 2005]

Software used for calibration

- Statistics package R has been used using the QRM library
- Other R libraries such as the fExtremes library were also investigated
- The work was partly re-created in excel as a check on the results
- The R libraries are freely available online and the code used to produce the results in this presentation and excel checking tool will be made available

Questions or comments?

Expressions of individual views by members of The Actuarial Profession and its staff are encouraged.

The views expressed in this presentation are those of the presenter.

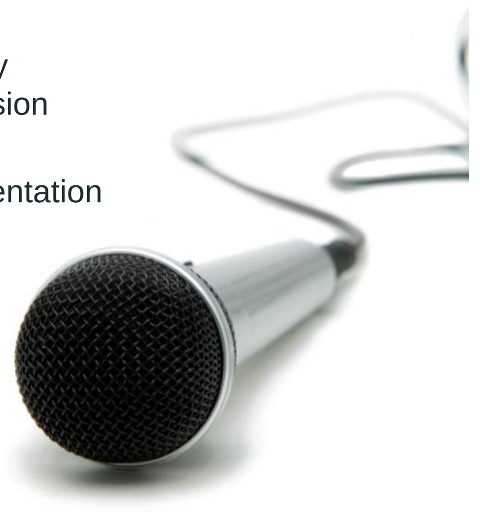

References

McNeil, A., Frey, R., and Embrechts, P. (2005).
Quantitative Risk Management: Concepts, Techniques and Tools.
Princeton University Press, Princeton.