Exponential Spectral Risk Measures

Kevin Dowd and John Cotter***

Spectral risk measures are attractive risk measures as they allow the user to obtain risk measures that reflect their subjective risk aversion. This paper examines spectral risk measures based on an exponential utility function, and finds that these risk measures have nice intuitive properties. It also discusses how they can be estimated using numerical quadrature methods, and how confidence intervals can be estimated for them using a parametric bootstrap. Illustrative results suggest that estimated exponential spectral risk measures obtained using such methods are quite precise in the presence of normally distributed losses.

Introduction

One of the most interesting and potentially most promising recent developments in the financial risk area has been the theory of Spectral Risk Measures (SRMs), recently proposed by Acerbi (2002 and 2004). SRMs belong to the family of coherent risk measures proposed by Artzner *et al.* (1997 and 1999), and therefore possess the highly desirable property of subadditivity.[1] It is also well-known by now that the most widely used risk measure, the Value-at-Risk (VaR), is not subadditive, and the work by Artzner *et al.,* and Acerbi has shown that many (if not most) of the inadequacies of VaR as a risk measure can be traced to its non-subadditivity.

One of the nice features of SRMs is that they relate the risk measure itself to the user's subjective risk aversion—in effect, the spectral risk measure is a weighted average of the quantiles of a loss distribution, the weights of which depend on the user's risk-aversion function. SRMs enable us to link the risk measure to the user's attitude towards risk, the underlying objective being to ensure that if a user is more risk averse, other things being equal, then that user should face a higher subjective risk, which is what the SRM measures. This implies that two different investors might have the same portfolios and share the same set of forecasts, but their subjective risk measures will still be different if one of them is more risk averse than the other.

* Professor of Financial Risk Management, Centre for Risk and Insurance Studies, Nottingham University Business School, Jubilee Campus, Nottingham, UK. E-mail: Kevin.Dowd@nottingham.ac.uk

** Associate Professor in Finance and Director, Centre for Financial Markets, School of Business, University College Dublin, Carysfort Avenue, Dublin, Ireland. E-mail: john.cotter@ucd.ie

[1] More formally, if $\rho(\cdot)$ is a measure of risk, and A and B are any two positions, subadditivity means that $\rho(A+B)\leq \rho(A)+\rho(B)$. Subadditivity is a crucial condition because it ensures that our overall risks do not increase when we put them together. As Acerbi and others have pointed out, any risk measure that does not satisfy subadditivity has no real claim to be regarded as a 'respectable' risk measure at all (*see*, *e.g.*, Acerbi, 2004, p. 150).

In principle, SRMs can be applied to any problems involving risky decision-making. Acerbi (2004) suggests that, amongst many other possible applications, they can be used to set capital requirements or to obtain optimal risk-expected return tradeoffs in portfolio analysis, and Cotter and Dowd (2006) suggest that SRMs could be used by futures clearinghouses to set margin requirements that reflect their corporate risk appetites.

This paper investigates the SRMs further. In particular, it focuses on SRMs based on an underlying exponential utility function. The objective here is twofold. First, the paper seeks to establish some of the properties of SRMs to see how intuitive and 'well-behaved' they might be. The second objective is computational: the paper discusses how these SRMs might be estimated and how the confidence intervals can be estimated for them. Next, the paper sets out the essence of Acerbi's theory of spectral risk measures. Then, it discusses the SRMs based on exponential utility functions and their estimation. Finally, it discusses the estimation of confidence intervals for them and concludes.

Spectral Risk Measures

Following Acerbi (2004), consider a risk measure, M_ϕ, defined as:

$$M_\phi = \int_0^1 \phi(p) q_p dp \qquad \text{...(1)}$$

where q_p is the p^{th} loss quantile and $\phi(p)$ is a user-defined weighting aversion function with weights defined over p, where p is a continuous range of cumulative probabilities $p \in [0,1]$. We can think of M_ϕ as the class of quantile-based risk measures, where each individual risk measure is defined by its own particular weighting function. Two well-known members of this class are the VaR and the Expected Shortfall (ES). The VaR at α confidence level is:

$$VaR_\alpha = q_\alpha \qquad \text{...(2)}$$

The VaR places all its weight on a single quantile that corresponds to a chosen confidence level, and places no weight on any others, i.e., with the VaR risk measure, $\phi(p)$ takes the degenerate form of a Dirac delta function that gives the outcome $p = \alpha$, an infinite weight, and gives every other outcome a zero weight. For its part, the ES at the confidence level α is the average of the worst $1-\alpha$ of losses and (in the case of a continuous loss distribution) is:

$$ES_\alpha = \frac{1}{1-\alpha} \int_\alpha^1 q_p dp \qquad \text{...(3)}$$

With ES, $\phi(p)$ gives tail quantiles a fixed weight of $\frac{1}{1-\alpha}$ and gives non-tail quantiles a weight of zero.

A drawback with both of these risk measures is that they are inconsistent with risk aversion in the traditional sense. This can be illustrated in the context of the theory of lower partial moments (Bawa, 1975; Fishburn, 1977; and Grootveld and Hallerbach 2004). Given a set of returns r and a target return r^*, the lower partial moment of order $k \geq 0$ around r^* is equal to $E\{[\max(0, r^* - r)]^k\}$. The parameter k reflects the user's degree of risk aversion, and the user is risk averse if $k > 1$, risk neutral if $k = 1$ and risk loving if $0 < k < 1$. It can then be shown that the VaR is a preferred risk measure only if $k = 0$, i.e., the VaR is the preferred risk measure only if the investor is very risk loving. The ES would be the preferred risk measure if $k = 1$, this implies that ES is the preferred risk measure only if the user is risk neutral between better and worse tail outcomes.

A user who is risk averse might prefer to work with a risk measure that takes account of his/her risk aversion, and this takes us to the class of Spectral Risk Measures (SRMs). In loose terms, an SRM is a quantile-based risk measure that takes the form of (1) where $\phi(p)$ reflects the user's degree of risk aversion. More precisely, we can consider SRMs as the subset of M_ϕ that satisfy the following properties of positivity, normalization and increasingness, due originally to Acerbi[2]:

- Positivity: $\phi(p) \geq 0$
- Normalization: $\int_0^1 \phi(p)dp = 1$
- Increasingness: $\phi'(p) \geq 0$

The first coherent condition requires that the weights are weakly positive and the second requires that the probability-weighted weights should sum to 1, but the key condition is the third one. This condition is a direct reflection of risk aversion, and requires that the weights attached to higher losses should be no less than the weights attached to lower losses. Typically, we would also expect the weight $\phi(p)$ to rise with p.[3] In a 'well-defined' case, we would expect the weights to rise smoothly, and the more risk averse the user, the more rapidly we would expect the weights to rise.

A risk measure that satisfies these properties is attractive not only because it takes account of the risk-aversion of the user, but also because such a risk measure is known to be coherent.[4] However, there still remains the question of how to specify $\phi(p)$, and perhaps the most natural way to obtain $\phi(p)$ is from the user's utility function.[5]

---

2 *See* Acerbi (2002 and 2004). However, it is worth pointing out that he deals with a distribution in which profit outcomes have a positive sign, whereas we deal with a distribution in which loss outcomes have a positive sign. His first condition is therefore a negativity condition, whereas ours is a positivity condition, but this difference is only superficial and there is no substantial difference between his conditions and ours.

3 The conditions set out allow for the degenerate limiting case where the weights are flat for all p values, and such a situation implies risk-neutrality and is therefore inconsistent with risk aversion. However, we shall rule out this limiting case by imposing the additional (and in the circumstances very reasonable) condition that $\phi(p)$ must rise over point at least as p increases from 0 to 1.

4 The coherence of SRMs follows from Acerbi (2004, *Proposition* 3.4).

5 *See also* Bertsimas *et al.* (2004).

Exponential Spectral Risk Measures

As an utility function is required to be specified, the common choice is the exponential utility function. This utility function is defined conditional on a single parameter, the coefficient of absolute risk aversion. The exponential utility function defined as follows, over outcomes x:

$$U(x) = -e^{-ax} \qquad \text{...(4)}$$

where $a > 0$ is the Arrow-Pratt coefficient of Absolute Risk Aversion (ARA). The coefficients of absolute and relative risk aversion are:

$$R_A(x) = -\frac{U''(x)}{U'(x)} = a \qquad \text{...(5a)}$$

$$R_R(x) = -\frac{xU''(x)}{U'(x)} = xa \qquad \text{...(5b)}$$

We now set:

$$\phi(p) = \lambda e^{-a(1-p)} \qquad \text{...(6)}$$

where λ is an unknown positive constant. This clearly satisfies the first and the third properties, and we can easily show (by integrating $\phi(p)$ from 0 to 1, setting the integral to 1 and solving for λ) that it also satisfies the second property if we set:

$$\lambda = \frac{a}{1-e^{-a}} \qquad \text{...(7)}$$

Hence, substituting Equation 7 into Equation 6 gives us the exponential weighting function (or risk aversion function) corresponding to Equation 4:[6]

$$\phi(p) = \frac{ae^{-a(1-p)}}{1-e^{-a}} \qquad \text{...(8)}$$

This risk aversion function is illustrated in Figure 1 for two alternative values of a, the ARA coefficient. Observe that this weighting function has a nice shape: for the higher p values associated with higher losses, we get bigger weights for greater degrees of risk aversion. In addition, as p rises, the rate of increase of $\phi(p)$ rises with the degree of risk aversion.

The SRM based on this risk aversion function—the exponential SRM—is then found by substituting Equation 8 into Equation 1, viz.:[7]

[6] Strictly speaking, Acerbi's *Proposition* 3.19 in Acerbi (2004, p. 182) defines his weighting function in terms of a parameter $\gamma > 0$, but his weighting function and Equation 7 are equivalent subject to the proviso that $\gamma = 1/a$.

[7] Estimates of Equation 9 were obtained using Simpson's rule numerical quadrature with p divided into N = 10,000,001 'slices'. The calculations were carried out using the CompEcon functions in MATLAB given in Miranda and Fackler (2002). As the results suggest, this value of N suffices effectively the exact estimates of SRMs. We use N = 10,000,001 rather than the more obvious N = 10,000,000 because the Simpson's rule algorithm requires N to be odd.

Figure 1: Exponential Risk Aversion Functions

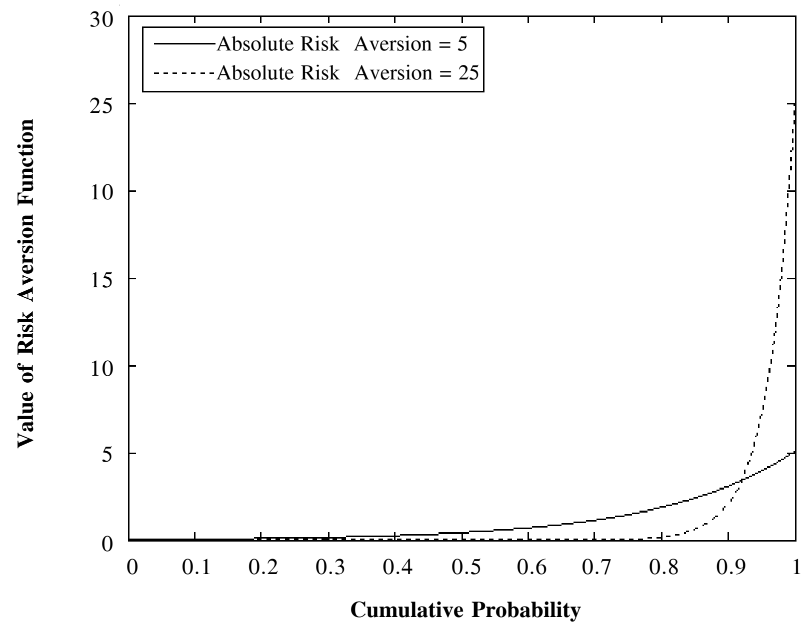

Note: Weights are based on the exponential risk aversion function, Equation 8.

$$M_{\phi} = \int_0^1 \phi(p)\, q_p\, dp = M_{\phi} = \frac{a}{1-e^{-a}} \int_0^1 e^{-a(1-p)} q_p\, dp \qquad ...(9)$$

We also find that the risk measure itself rises with the degree of risk aversion, and some illustrative results are given in Table 1. For example, if losses are distributed as standard normal and we set $a = 5$, then the spectral risk measure is 1.0816. But if we increase a to 25, the measure rises to 1.9549, that is, the greater the risk aversion, the higher is the exponential spectral risk measure.

Table 1: Values of Exponential Spectral Risk Measure with Standard Normal Losses

Coefficient of Absolute Risk Aversion	Exponential Spectral Risk Measure
1	0.2781
5	1.0816
25	1.9549
100	2.5055

Note: Estimates of Equation 9 are obtained using Simpson's rule numerical quadrature with p divided into N = 10,000,001 'slices'. The calculations were carried out using the Miranda and Fackler (2002), CompEcon functions in the 2007a version of MATLAB on a Pentium4 desktop computer.

The relationship of the exponential SRM and the coefficient of absolute risk aversion is illustrated further in Figure 2. We see that as expected, the risk measure rises smoothly as the coefficient of risk aversion increases.

Figure 2: Plot of Exponential Spectral Risk Measure Against the Coefficient of Absolute Risk Aversion: Standard Normal Loss Distribution

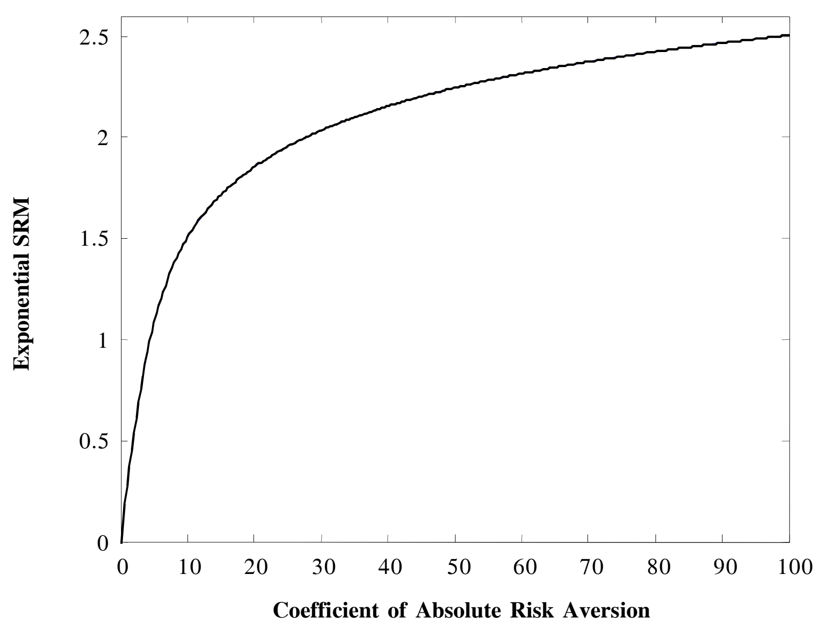

Note: Estimates of Equation 9 are obtained using Simpson's rule numerical quadrature with p divided into N = 10,000,001 'slices'. The calculations were carried out using the Miranda and Fackler (2002), CompEcon functions in the 2007a version of MATLAB on a Pentium4 desktop computer.

These results suggest that exponential SRMs are 'well-behaved' and have nice intuitive properties.

Estimating Spectral Risk Measures

The equation for the SRM, Equation 9, indicates that solving for the SRM involves integration. Special cases aside, this integration would need to be carried out numerically rather than analytically. This means that the estimation of SRMs requires a suitable numerical integration or numerical quadrature method. Such methods estimate the integral from a numerical estimate of its discreditized equivalent in which p is broken into a large number of N 'slices'.[8] This, in turn, raises the question of how different quadrature methods can be compared. Furthermore, since quadrature methods depend on a specified value of N, it also raises the issue of how estimates might depend on the value of N.

[8] For more on these methods, *see*, *e.g.*, Borse (1997, Chapter 7), Kreyszig (1999, Chapter 17), or Miranda and Fackler (2002, Chapter 5).

To investigate these issues further, Figure 3 provides some plots showing how estimates of standard normal SRMs vary with different quadrature methods and different values of *N*. These plots are based on an illustrative ARA coefficient equal to 5, but we get similar plots for other values of this coefficient. The methods examined are those based on the trapezoidal rule, Simpson's rule, and Niederreiter and Weyl quasi-Monte Carlo procedure. As we might expect, all four quadrature methods give estimates that converge to their true values as *N* gets larger. However, the trapezoidal and Simpson's rule methods produce estimates that converge smoothly as *N* gets larger, whereas the two quasi-Monte Carlo methods produce estimates that converge more erratically as *N* gets larger. The plots also suggest that the first two methods are usually more accurate for any given value of *N*, and that the method based on Simpon's rule is marginally better than that based on the trapezoidal rule.

Figure 3: Estimates of Standard Normal Spectral Risk Measure for Various Numerical Quadrature Methods and Values of *N* (Coefficient of Absolute Risk Aversion = 5)

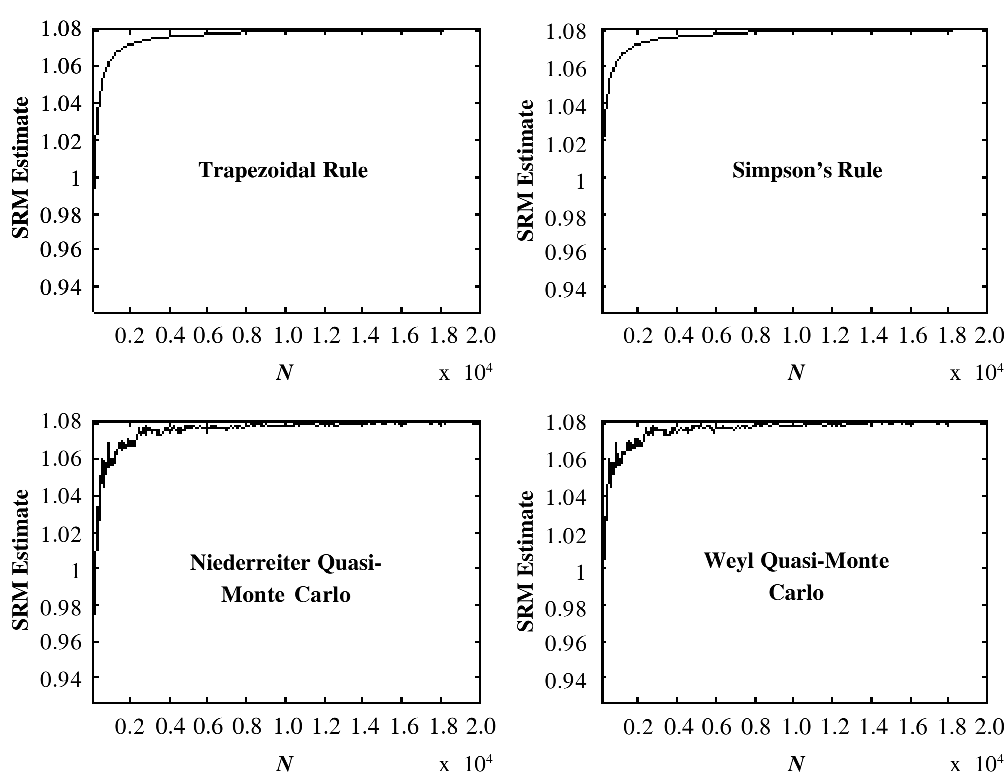

Note: As per Table 1. Each chart gives plots of the estimated SRM predicated on ARA = 5 for each of the four numerical quadrature methods: the trapezoidal rule, Simpson's rule, the Niederreiter quasi-Monte Carlo and Weyl quasi-Monte Carlo. Estimates of Equation 9 are obtained for values of *N* ranging from 100 to 20,000.

In addition, these plots show that all methods produce estimates of SRMs that have a small downward bias. If we wish to get accurate estimates of SRMs, it is therefore important to choose a value of N, large enough to make this bias negligible. To investigate further, Table 2 reports results for the accuracy and calculation times of the same SRM estimated using the Simpson's rule but different values of N. This shows that, for N = 1,001, we get an estimate with an error of –1.55% and this takes 0.0027 seconds to calculate using the latest version (Version 2007a) of MATLAB using a Pentium 4 desktop computer. For N = 10,001, the error is –0.18% and this estimate takes 0.0096 seconds to calculate. The accuracy and calculation times increase as N gets larger, and when N = 10,000,001, we get an error of –0.00% and a calculation time of a little over 8 seconds. Bearing in mind that real-world risk measurement is subject to many different sources of error that are beyond the control of the risk manager, it is pointless to go for spuriously accurate estimates that ignore these other sources of error. We would therefore suggest that a value of N = 10,001 is accurate enough for the needs of risk managers in the real world for practical purposes.[9]

Table 2: Percentage Errors and Calculation Times for Estimates of Standard Normal Exponential Spectral Risk Measure Obtained Using the Simpson's Rule (Coefficient of Absolute Risk Aversion = 5)

N	Percentage Error	Calculation Time (in seconds)
1,001	–1.55%	0.0027
10,001	–0.18%	0.0096
100,001	–0.03%	0.0818
1,000,001	–0.01%	0.8689
10,000,001	0.00%	8.3775

Note: As per the Note to Table 1, but with the specified value of N.

Estimating Confidence Intervals for Spectral Risk Measures

Given that our estimates are prone to many sources of error, it is a good practice to estimate some precision metrics to go with our estimated risk measures, and perhaps the best such metrics are estimates of their confidence intervals.

We can easily obtain such estimates using a parametric bootstrap, and this can be implemented using the following procedure:

- Given N, for each of b bootstrap trials, we simulate a set of N loss values from our assumed distribution, in this case a standard normal. We order these simulated losses from lowest to highest, to obtain a set of simulated quantiles, $\tilde{q}_p$. We then apply our chosen quadrature method to Equation 9 with q_p replaced by $\tilde{q}_p$ to obtain a bootstrap estimate of the SRM.
- We repeat this step 'b' times and obtain an estimate of the confidence interval from the distribution of the simulated SRM estimates.

[9] These results are obtained assuming that losses are normal and were also obtained using a good numerical integration routine. In practice, risk managers would be well advised to check the accuracy of their numerical integration routines before using them and also check the extent to which similar findings hold or do not hold for the distributions they are dealing with.

Some illustrative results are shown in Table 3. This shows estimates of the 90% confidence intervals for values of ARA coefficient equal to 5 and 100. The Table also shows estimates of the standardized confidence intervals, which are equal to the 90% confidence intervals with the bounds divided by the mean of the bootstrap SRM estimates. All results are based on N = 10,001 and b = 1,000. We can see that for ARA = 5, the 90% confidence is [1.0591, 1.1012] and its equivalent for ARA = 100 is [2.4075, 2.5380]. For ARA = 5, the standardized interval is [0.9805, 1.0195] and for ARA = 100 the standardized interval is [0.9739, 1.0267]. It is striking to note that how narrow these intervals are: for ARA = 5, the width of the interval is only 3% of the estimated mean SRM, and for ARA = 100, the width is only a little over 5% of the estimated mean SRM. These narrow confidence intervals indicate that the SRM estimates are very precise.

Table 3: Illustrative Estimates of 90% Confidence Intervals for Spectral Risk Measures

ARA = 5	ARA = 100
90% Confidence Interval	
[1.0591, 1.1012]	[2.4075, 2.5380]
Standardized 90% Confidence Interval	
[0.9805, 1.0195]	[0.9739, 1.0267]

Note: Estimates of Equation 9 are obtained using Simpson's rule numerical quadrature with N = 10,001 and 1,000 resamples from a parametric bootstrap. The SRM estimates were obtained using the Miranda and Fackler (2002), CompEcon functions in MATLAB on a Pentium4 desktop computer. The underlying loss distribution is assumed to be standard normal. Each confidence interval took just over 14 seconds to calculate. The standardized 90% confidence intervals are equal to the associated 90% confidence interval with each bound divided by the mean of the bootstrap SRM estimates.

Conclusion

This paper has examined Spectral Risk Measures (SRMs) based on an exponential utility function. The results reveal that the exponential utility function leads to risk aversion functions and SRMs with intuitive and nicely behaved properties. These exponential SRMs are easy to estimate using numerical quadrature methods, and accurate estimates can be obtained very quickly in real time. It is also easy to estimate confidence intervals for these SRMs using a parametric bootstrap. Illustrative results suggest that these confidence intervals are surprisingly narrow, and this indicates that SRM estimates are quite precise. Of course, the results presented here are based on an assumed normal distribution, and further work is needed to establish results for other distributions.[10] ❖

Acknowledgment: *The authors would like to thank Carlo Acerbi and Dirk Tasche for fruitful conversations on the subject. Cotter's contribution to the study has been supported by a University College Dublin School of Business research grant.*

[10] However, a start has already been made, for example, Acerbi (2004, pp. 202-205) provides some results for lognormal and power-law distributions, and Cotter and Dowd (2006) provide results for generalized Pareto distributions.

References

1. Acerbi C (2002), "Spectral Measures of Risk: A Coherent Representation of Subjective Risk Aversion", *Journal of Banking and Finance*, Vol. 26, pp. 1505-1518.

2. Acerbi C (2004), "Coherent Representations of Subjective Risk Aversion", in Szegö G (Ed.) *Risk Measures for the 21st Century*, pp. 147-207, John Wiley & Sons, New York.

3. Artzner P, Delbaen F, Eber J-M and Heath D (1997), "Thinking Coherently", *Risk*, Vol. 10, pp. 68-71, November.

4. Artzner P, Delbaen F, Eber J-M and Heath D (1999), "Coherent Measures of Risk", *Mathematical Finance*, Vol. 9, pp. 203-228.

5. Bawa V S (1975), "Optimal Rules for Ordering Uncertain Prospects", *Journal of Financial Economics*, Vol. 2, pp. 95-121.

6. Cotter J and Dowd K (2006), "Extreme Spectral Risk Measures: An Application to Futures Clearinghouse Margin Requirements", *Journal of Banking and Finance*, Vol. 30, pp. 3469-3485.

7. Fishburn P C (1977), "Mean-Risk Analysis with Risk Associated with Below-Target Returns", *American Economic Review*, Vol. 67, pp. 116-126.

8. Grootveld H and Hallerbach W G (2004), "Upgrading Value-at-Risk from Diagnostic Metric to Decision Variable: A Wise Thing to Do?", in Szegö G (Ed.) *Risk Measures for the 21st Century*, pp. 33-50, John Wiley & Sons, New York.

Bibliography

1. Bertsimas D, Lauprete G J and Samarov A (2004), "Shortfall as a Risk Measure: Properties, Optimization and Applications", *Journal of Economic Dynamics and Control*, Vol. 28, pp. 1353-1381.

2. Borse G J (1997), *Numerical Methods with MATLAB: A Resource for Scientists and Engineers*, PWS Publishing Company, Boston.

3. Dowd K (2005), *Measuring Market Risk*, 2nd Edition, John Wiley & Sons, Chichester.

4. Kreyszig E (1999), *Advanced Engineering Mathematics*, 8th Edition, John Wiley & Sons, New York.

5. Miranda M J and Fackler P L (2002), *Applied Computational Economics and Finance*, MIT Press, Cambridge, MA and London.

Reference # 42J-2007-12-03-01