



Munich Personal RePEc Archive

Extreme Spectral Risk Measures: An Application to Futures Clearinghouse Margin Requirements

Cotter, JOhn and Dowd, Kevin

2006

Online at <https://mpra.ub.uni-muenchen.de/3505/>
MPRA Paper No. 3505, posted 13 Jun 2007 UTC

Extreme Spectral Risk Measures: An Application to Futures Clearinghouse Margin
Requirements

By

John Cotter and Kevin Dowd*

Abstract

This paper applies the Extreme-Value (EV) Generalised Pareto distribution to the extreme tails of the return distributions for the S&P500, FT100, DAX, Hang Seng, and Nikkei225 futures contracts. It then uses tail estimators from these contracts to estimate spectral risk measures, which are coherent risk measures that reflect a user's risk-aversion function. It compares these to more familiar VaR and Expected Shortfall (ES) measures of risk, and also compares the precision and discusses the relative usefulness of each of these risk measures.

Keywords: Spectral risk measures, Expected Shortfall, Value at Risk, Extreme Value, clearinghouse.

JEL Classification: G15

February 21 2005

* John Cotter is at the Centre for Financial Markets, Smurfit School of Business, University College Dublin, Carysfort Avenue, Blackrock, Co. Dublin, Ireland; email: john.cotter@ucd.ie. Kevin Dowd is at the Centre for Risk and Insurance Studies, Nottingham University Business School, Jubilee Campus, Nottingham NG8 1BB, UK; email: Kevin.Dowd@nottingham.ac.uk. Cotter's contribution to the study has been supported by a University College Dublin Faculty of Commerce research grant. Dowd's contribution was supported by an Economic and Social Research Council research fellowship on 'Risk measurement in financial institutions', and he thanks the ESRC for their financial support.

1 INTRODUCTION

One of most important functions of a futures clearinghouse is to act as counterparty to all trades that take place within its exchanges. This ensures that individual traders do not have to concern themselves with credit risk exposures to other traders, because the clearinghouse assumes all such risk itself. However, it also means that the clearinghouse has to manage this risk, and perhaps the most important way it can do so is by setting margin requirements to protect itself against possible default by investors who suffer heavy losses. But how should clearinghouses set their margin requirements?

A good starting point is that investor defaults are due to large – that is to say, extreme – price movements that are best analysed using some form of Extreme-Value (EV) theory. A number of papers have followed this line of inquiry (e.g., Longin (1999, 2000) and Booth *et al* (1997)). Typically, extremes are modelled by applying a Generalized Pareto Distribution (GPD) to exceedences X over a high threshold u . The application of the GPD is justified by theory that tells us that exceedences should follow a GPD in the asymptotic limit as the threshold gets bigger. Once the GDP curve is fitted to the data, it can then be extrapolated to give us estimates of any quantiles or tail probabilities we choose. Where we are primarily interested in the tails of the distribution in modeling margin requirements, the GPD is far superior to alternatives such as a normal (Gaussian) distribution (as in Figlewski (1984), which tends to under-estimate the heaviness of futures tail returns and is, in any case, inconsistent with any of the distributions that EV theory tells us to expect. Margin setting has also relied on historical distributions (see Edwards and Neftci (1988) and Warshawsky (1989)), but historical approaches are unable to provide very low probability estimates due to insufficient data. Alternatively, time varying measures could be developed using, for example, GARCH type process (see Giannopoulos and Tunaru (2004 and Cotter (2001)), but would involve continuous updating of futures margins on a daily basis.

These previous studies applying EV theory or other statistical models, have only applied margin requirements as a quantile or VaR. However, the VaR has been heavily criticised as a risk measure on the grounds that it does not satisfy the properties of coherence and, most particularly, because the VaR is not subadditive

(Artzner *et al.* 1999; Acerbi, 2004). The failure of VaR to be subadditive can then lead to strange and undesirable outcomes: in the present case, the use of the VaR to set margin requirements takes no account of the magnitude of possible losses exceeding VaR, and can therefore leave the clearinghouse heavily exposed to very high losses exceeding the VaR.¹

This paper first contributes to the literature by using an alternative risk measure that could but has yet to be applied to margin setting, namely the Expected Shortfall (ES). The ES is the average of the worst α of losses, where α is the confidence level. This measure is closely related, but not identical to, the Tail Conditional Expectation, which is the probability-weighted average of losses exceeding VaR.² Unlike the VaR, the ES is coherent (and hence subadditive as well) and so satisfies many of the properties we would desire a priori from a ‘respectable’ risk measure.³ The ES is (generally) bigger than the VaR and, more importantly, takes account of the magnitude of losses exceeding the VaR. These attractions suggest that the ES would provide a better basis for margin requirements than the VaR. However, both VaR and ES margin requirements would depend on the choice of a confidence level, and there is no a priori ‘obvious’ confidence level to choose. In addition, the use of VaR or ES to set clearinghouse margin requirements has the undesirable implication that the clearinghouse is either risk-loving or risk-neutral, and this sits uncomfortably with the common-sense notion that any clearinghouse should be risk-averse.

¹ One important consequence of using a non-subadditive risk measure like the VaR to set margin requirements is that investors might break up their accounts to reduce overall margin requirements, and in so doing leave the clearinghouse exposed to a hidden residual risk against which the clearinghouse has no effective collateral from its investors. This type of problem does not arise with subadditive risk measures such as the other ones considered in this paper.

² For more on these risk measures and their distinguishing features, see Acerbi and Tasche (2001) or Acerbi (2004). We don’t consider the TCE further in this paper because it is equivalent to the ES where the density function is continuous, and where it differs from the ES, it is not coherent.

³ Loosely speaking, let X and Y represent any two portfolios’ P/Ls (or future values, or the portfolios themselves) over a given forecast horizon, and let $\rho(\cdot)$ be a measure of risk. The risk measure $\rho(\cdot)$ is subadditive if it satisfies $\rho(X + Y) \leq \rho(X) + \rho(Y)$. Subadditivity is the most important criterion we would expect a ‘respectable’ risk measure to satisfy. It can be demonstrated that VaR is not subadditive unless we impose the empirically implausible requirement that returns are elliptically distributed. Given the importance of subadditivity, the VaR’s non-subadditivity makes it very difficult to regard the VaR as a ‘respectable’ measure of risk.

This paper further adds to the literature on margin setting by providing risk measures that specifically incorporate an agent's degree of risk aversion. In theory, the clearinghouse should use a risk measure that takes account of the nature and extent of its risk aversion. Thus a clearinghouse that is more risk averse would have a higher estimated risk measure and impose a higher margin requirement, other things being equal. Such risk measures have recently been proposed by Acerbi (2002, 2004). These measures are known as spectral risk measures because they relate the risk measure directly to the user's risk spectrum or risk-aversion function. 'Well-behaved' spectral risk measures are a subset of the family of coherent risk measures, and therefore have the attractions of coherent risk measures as well. One attractive type of spectral risk measure is based on an exponential risk aversion function. A nice feature of this type of spectral risk measure is that the extent of risk aversion depends on a single parameter γ : the lower is γ , the more risk-averse the user. In principle, once a clearinghouse chooses the value of γ that reflects its own attitude to risk, it can then obtain an 'optimal' risk measure that directly reflects its risk aversion. So, whereas the VaR, previously applied in the literature, or even ES, are contingent on the choice of an arbitrary parameter, the confidence level, whose 'best' value cannot easily be determined, a spectral-coherent risk measure is contingent on a parameter whose 'best' value can in principle be ascertained by the clearinghouse that uses it.

This paper provides a variety of alternative estimates of VaR, ES and spectral risk measures for each of 12 different types of contract, these being long and short positions in each of the S&P500, the FTSE100, the DAX, the Hang Seng, and the Nikkei225 indexes. It also compares these different estimates to each other. Bearing in mind that the value of any estimated risk measures depends crucially on their precision, the paper also examines alternative methods of estimating their precision. Our findings on that front suggest the (controversial) conclusion that (at least in this context) estimates of ES are generally a little more precise than estimates of VaR, whereas estimates of the spectral risk measures are somewhat less precise than estimates of either of the other two risk measures. The former finding is encouraging in that it helps strengthen the argument that risk practitioners should 'upgrade' their risk measures from the VaR to the ES; however, the latter finding is a little

disappointing in that it weakens a little the otherwise (to us compelling) argument that clearinghouses would do better still to upgrade to spectral risk measures.

This paper is organised as follows. Section 2 reviews the risk measures to be examined. Section 3 then reviews the extreme-value (EV) theory to be applied: the Peaks-Over-Threshold (POT) theory based on the Generalised Pareto distribution (GPD) applied to exceedances over a high threshold. Section 4 introduces the data and provides some preliminary data analysis on both long and short positions in five representative futures contracts. Section 5 then estimates VaR and ES, and section 6 estimates the spectral risk measures. Each of these sections also examines the precision of these estimated risk measures. Section 7 discusses these results and compares the suitability of each type of risk measure for futures clearinghouse margin requirements. Section 8 concludes.

2. MEASURES OF RISK

We are interested in three main risk measures, the first two (VaR and ES) mainly for comparison, and the third (spectral or, more properly, extreme-value coherent spectral-exponential risk measures) because we believe these are superior in principle and potentially more useful for determining how clearinghouse margin requirements should be set.

Let X be a random loss variable (which gives losses a positive sign and profits a negative one) over a daily horizon period on a futures contract position (which might itself be long or short in the underlying index). If the confidence level is α , the VaR at this confidence level is:

$$VaR_{\alpha} = q_{\alpha} \tag{1}$$

where q_{α} is the relevant quantile of the loss distribution. We have already noted some of the problems with the VaR as a risk measure.⁴ Viewed as a function of the

⁴ There are also other problems. The VaR is not consistent with expected utility maximisation, except in the very unusual case where risk preferences are lexicographic (Grootveld and Hallerbach, 2004, p. 33). More insight into the limitations of VaR comes from the perspective of the downside risk literature (see, e.g., Bawa (1975) and Fishburn (1977)). These papers suggest that we can think of downside risk

quantiles of the loss distribution, it is also useful to note here that the VaR places all its weight on a single quantile that corresponds to a chosen confidence level, and it places no weight on any others.

Our second risk measure is the ES, which is the average of the worst $1 - \alpha$ of losses. In the case of a continuous loss distribution, the ES is given by:

$$ES_\alpha = \frac{1}{1 - \alpha} \int_\alpha^1 q_p dp \quad (2)$$

Using an ES measure implies taking an average of quantiles in which tail quantiles have an equal weight and non-tail quantiles have a zero weight. However, the fact that the ES gives all tail losses equal weights suggests that a user who uses this measure is risk-neutral at the margin between better and worse tail outcomes, and this is inconsistent with risk-aversion.⁵

It is also possible to relate coherent risk measures to a user's risk aversion, and, indeed, to tailor coherent measures to fit a user's risk aversion. Let us define more general risk measures M_ϕ that are weighted averages of the quantiles q_p :

$$M_\phi = \int_0^1 \phi(p) q_p dp \quad (3)$$

where the weighting function, $\phi(p)$, known as the risk spectrum or risk-aversion function, still remains to be determined.⁶

in terms of lower-partial moments (LPMs), which are probability-weighted deviations of returns r from some below-target return r^* : more specifically, the LPM of order $k \geq 0$ around r^* is equal to $E[\max(0, r^* - r)^k]$. The parameter k reflects the degree of risk aversion, and the user is risk-averse if $k > 1$, risk-neutral if $k = 1$, and risk-loving if $0 < k < 1$. However, we would choose the VaR as our preferred risk measure only if $k = 0$ (Grootveld and Hallerbach, 2004, p. 35), and this suggests that the use of the VaR as a preferred risk measure indicates negative risk aversion of a fairly extreme sort.

⁵ This interpretation is also confirmed from the downside risk literature. From that perspective, the ES is the ideal risk measure if $k=1$, and this implies that the user is risk-neutral (Grootveld and Hallerbach, 2004, p. 36).

⁶ The spectral risk measure (3) also includes both the VaR and the ES as special cases. The VaR implies a $\phi(p)$ function that takes the degenerate form of a Dirac delta function that gives the outcome

We are interested in the broader class of coherent risk measures, and want to know the conditions that $\phi(p)$ must satisfy to make M_ϕ coherent. The answer is the class of (non-singular) spectral risk measures, in which $\phi(p)$ takes the following properties:⁷

- *Non-negativity*: $\phi(p) \geq 0$ for all p belong in the range $[0,1]$.
- *Normalization*: $\int_0^1 \phi(p) dp = 1$.
- *Weakly increasing*: $\phi(p_1) \leq \phi(p_2)$ for all $0 \leq p_1 \leq p_2 \leq 1$.

The first condition requires that the weights are non-negative, and the second requires that the probability-weighted weights should sum to 1. Both of these are obvious. However, the third condition is a direct reflection of risk-aversion, and requires that the weights attached to higher losses should be bigger than, or certainly no less than, the weights attached to lower losses. This implies that the key to coherence is that a risk measure must give higher losses at least the same weight as lower losses. This explains why the VaR is not coherent and the ES is, and it also tells us that the VaR's most prominent inadequacies are closely related to its failure to satisfy the weakly increasing property.

If a user has a 'well-behaved' risk-aversion function, then the weights will rise smoothly, and the rate at which they do so is related to the degree of risk aversion: the more risk-averse the user, the more rapidly the weights will rise. This implies that there is an optimal risk measure for each user, which depends on the user's risk aversion function. Thus, two users might have identical portfolios, but they

$p=\alpha$ an infinite weight, and every other outcome a zero weight, and the ES implies a discontinuous $\phi(p)$ that takes the value 0 for profits or small losses and takes a constant value for high losses. However, these are not 'well-behaved' spectral risk measures, because they are inconsistent with risk aversion..

⁷ For more on these, see Acerbi (2004, proposition 3.4). There is also a good argument that spectral measures are the only really interesting coherent risk measures. Kusuoka (2001) and Acerbi (2004, pp. 180-182) show that all coherent risk measures that satisfy the two additional properties of comonotonic additivity and law invariance are also spectral measures. The former condition is that if two random variables X and Y are comonotonic (i.e., always move in the same direction), then $\rho(X+Y) = \rho(X) + \rho(Y)$; comonotonic additivity is an important aspect of subadditivity, and represents the limiting case where diversification has no effect. Law-invariance is equivalent to the (in practice essential) requirement that a measure be estimable from empirical data. Both conditions are important, and coherent risk measures that do not satisfy them are probably not worth considering further.

will only have the same (coherent) risks if they also have exactly the same risk-aversion.

To obtain a spectral risk measure, the user must specify a particular form for their risk-aversion function. This decision is subjective, but can be guided by the economic literature on utility-function theory. A user would pick some ‘reasonable’ utility function whose risk-aversion properties seem to reflect a suitable attitude toward risk. The chosen utility function could then be transformed into a risk-aversion function suitable for risk measurement purposes. A plausible candidate is a standard exponential utility function, which can be transformed into an exponential risk-aversion function defined by

$$\phi_{\gamma}(p) = \frac{e^{-p/\gamma}}{\gamma(1 - e^{-1/\gamma})} \quad (4)$$

where $\gamma \in (0, \infty)$ (see Acerbi, 2004, p. 178). This function satisfies the conditions required of a spectral risk measure, and is also attractive because it is a simple function that depends on a single parameter, the value of which reflects the risk aversion of the user. A spectral risk-aversion function is illustrated in Figure 1. This shows how the weights rise with the cumulative probability p , and the rate of increase depends on γ : the more risk-averse the user, the more rapidly the weights rise.

Insert Figure 1 here

It is also curious to note that γ plays a role in spectral measures similar to the role that the confidence level α plays in the VaR and ES. More specifically, if we think in loose terms of a higher confidence level reflecting a greater concern with higher losses – which might reflect increasing risk-aversion in a crude sense – then this is comparable to a falling γ in a spectral risk measure. However, whereas a falling γ reflects a well-defined sense of increasing risk aversion, the choice of confidence level is arbitrary.

To obtain our spectral measure M_ϕ , we choose a value of γ and substitute $\phi(p)$ and $q_p(X)$ into M_ϕ to get:

$$M_\phi = \int_0^1 \phi(p)q_p dp = \int_0^1 \frac{e^{-(1-p)/\gamma}}{\gamma(1-e^{-1/\gamma})} q_p dp \quad (5)$$

(5) gives us our third risk measure, the spectral risk measure, contingent on a chosen value of γ .

Of course, with all these risk measures we still face the problem of how to obtain the quantile q_p , and it is to this problem that we now turn.

3. THE PEAKS OVER THRESHOLD (GENERALISED PARETO) APPROACH

As we are particularly interested in the extreme risks faced by the clearinghouse, we model extreme returns using an Extreme Value (EV) approach. Perhaps the most suitable of these for our purposes is the Peaks over Threshold (POT) approach based on the Generalized Pareto distribution (GPD).⁸ This approach focuses on the realisations of a random variable X over a high threshold u . More particularly, if X has the distribution function $F(x)$, we are interested in the distribution function $F_u(x)$ of exceedances of X over a high threshold u :

$$F_u(x) = P\{X - u \leq x | X > u\} = \frac{F(x+u) - F(u)}{1 - F(u)} \quad (6)$$

As u gets large (as would be the case for the thresholds relevant to clearinghouses), then the distribution of exceedances tends to a GPD:

⁸ Alternatively, extreme tail returns could be modelled by Generalised Extreme Value (GEV) theory, which deals with the distribution of the sample maxima. The GEV and POT approaches are analogous in the limit, but we prefer to use the POT approach because it (generally) uses one less parameter, and because the GEV approach does not utilise all extreme returns if extremes occur in clusters.

$$G_{\xi, \beta}(x) = \begin{cases} 1 - (1 + \xi x / \beta)^{-1/\xi} & \text{if } \xi \geq 0 \\ 1 - \exp(-x / \beta) & \text{if } \xi < 0 \end{cases} \quad (7)$$

where

$$x \in \begin{cases} [0, \infty) & \text{if } \xi \geq 0 \\ [0, -\beta / \xi] & \text{if } \xi < 0 \end{cases}$$

and the shape ξ and scale $\beta > 0$ parameters are estimated conditional on the threshold u (Balkema and de Haan (1974); Embrechts *et al.*, 1997, pp. 162-164). Note that the shape parameter ξ sometimes appears in GPD discussions couched in terms of its inverse, a tail index parameter α given by $\alpha = 1/\xi$.

The GPD parameters can be estimated by maximum likelihood methods. The log likelihood function of the GPD for $\xi \neq 0$ is:

$$l(\xi, \beta) = -n(\ln(\beta) - (1 + 1/\xi) \sum_{i=1}^n \ln(1 + \xi x_i / \beta)) \quad (8)$$

where x_i satisfies the constraints specified for x . If $\xi = 0$, the log likelihood function is:

$$l(\beta) = -n(\ln(\beta) - \beta^{-1} \sum_{i=1}^n x_i) \quad (9)$$

ML estimates are then found by maximising the log-likelihood function using suitable (e.g., numerical optimisation) methods.

The behaviour of the GPD tail depends on the parameter values, and the shape parameter is especially important. A negative ξ is associated with very thin-tailed distributions that are rarely of relevance to financial returns, and a zero ξ is associated with other thin tailed distributions such as the normal. However, the most relevant for our purposes are heavy-tailed distributions associated with $\xi > 0$. The tails of such distributions decay slowly and follow a ‘power law’ function. Moreover the number of finite moments is ascertained by the value of ξ : if $\xi \leq 0.5$ (or, equivalently,

$\alpha \geq 2$) we have infinite second and higher moments; if $\xi \leq 0.25$ (or $\alpha \geq 4$), we have infinite fourth and higher moments, and so forth. α thus indicates the number of finite moments. Evidence generally suggests that the second moment is probably finite, but the fourth moment is more problematic (see, e.g., Loretan and Phillips, 1994).

Assuming that u is sufficiently high, the distribution function for exceedances is given by:

$$F_u(x) = 1 - \frac{N_u}{n} \left(1 + \xi \frac{x-u}{\beta} \right)^{-\frac{1}{\xi}} \quad (10)$$

where n the sample size and N_u is the number of observations in excess of the threshold (Embrechts *et al.*, 1997, p. 354). The p^{th} quantile of the return distribution - which is also the VaR at the (high) confidence level p - can then be obtained by inverting the distribution function:

$$q_p = \text{VaR}_p = u + \frac{\beta}{\xi} \left\{ \left(\frac{n}{N_u} p \right)^{-\xi} - 1 \right\} \quad (11)$$

The ES is then given by:

$$ES_p = \frac{q_p}{1-\xi} + \frac{\beta - \xi u}{1-\xi} \quad (12)$$

To obtain more general spectral risk measures, we substitute (11) into (5) to obtain:

$$M_\phi = \int_0^1 \phi(p) q_p(X) dp = \int_0^1 \frac{e^{-(1-p)/\gamma}}{\gamma(1-e^{-1/\gamma})} \left[u + \frac{\beta}{\xi} \left\{ \left(\frac{n}{N_u} p \right)^{-\xi} - 1 \right\} \right] dp \quad (13)$$

Estimates of our risk measures are then obtained by estimating/choosing the relevant parameters and plugging these into the appropriate risk measure equation

(i.e., (11), (12) or (13)). This is straightforward where our risk measures are the VaR and the ES; where our risk measures are spectral, we can solve (13) using a suitable numerical integration method (e.g., a trapezoidal rule, Simpson's rule, Monte Carlo, etc.).⁹

4. PRELIMINARY DATA ANALYSIS

Our data set consists of daily geometric returns (taken as the difference between the logs of respective end-of-day daily prices) for the most heavily traded index futures – that is, the S&P500, FTSE100, DAX, Hang Seng and Nikkei 225 futures – between January 1 1991 and December 31 2003.

As a preliminary, Figure 2 shows QQ plots for these contracts' empirical return distributions relative to a normal (or Gaussian) distribution. If the normal distribution is an adequate fit, then the QQ plot should be approximately a straight line. However, in each case, we find that the QQ plot is approximately straight only in the central region, but not for the tails. For their part, the tails – on both left-hand and right-hand sides – of the QQ plot show steeper slopes than the central observations, indicating that the tails exhibit excess kurtosis (or tail heaviness) relative to the normal distribution. These findings are consistent with the widely held perception that both long and short futures positions are heavy-tailed.

Insert Figure 2 here

In addition, the points where the QQ plots change shape provide us with natural estimates of tail thresholds. These lead us to select thresholds of 2 for the S&P, DAX, Hang Seng and Nikkei indices, and to select a threshold of 1.5 for the FTSE.

This choice of threshold values is also consistent with the tail index plots – plots of the estimated tail index α and its 95% confidence interval against the number of exceedances – shown in Figure 3. The number of exceedances reflects the choice of threshold, a smaller number reflecting a higher threshold. In each case the

⁹ More details on such methods can be found in standard references (e.g., Miranda and Fackler (2002, chapter 5).

estimated tail index is stable over a wide range of exceedance numbers (or threshold size, if you prefer). This tells us that the estimated indices are stable over the thresholds selected.

Insert Figure 3 here

We now assume that the distributions of exceedances take the form of GPDs, and ML estimates of the GPD parameters are given in Table 1 for both long and short trading positions. The tail indices are positive except for the Nikkei and the estimated scale parameters fluctuate around 1. All of these estimates are plausible and in line with those reported from other studies. The Table also gives the assumed thresholds u , the associated numbers of exceedances (N_u) and the observed exceedance probabilities (*prob*). The numbers and probabilities of exceedances vary somewhat, but all confirm that the chosen thresholds are in the stable tail-index regions identified in Figure 2.

Insert Table 1 here

To check that the GPD provides an adequate fit, Figure 4 shows empirical exceedances fitted to the GPDs based on the parameter estimates given in Table 1. The GPD provides a good fit in all cases, and this confirms both the selection of the GPD as an appropriate distribution and the parameter estimates on which the fitted distribution is predicated.

Insert Figure 4 here

5. ESTIMATION OF VAR AND EXPECTED SHORTFALL

Estimates of VaRs and ESs based on these parameters are shown in Table 2 and illustrated in Figures 5 and 6. These risk measures are based on extremely high confidence levels and reflect the clearinghouses' concerns with very high trading losses and the associated possibilities of investor default. The ESs are notably larger than the VaRs: they are typically 20-30% larger, but in some cases considerably more. However, the two risk measures otherwise behave in similar ways. Both risk measures

increase as the confidence levels get bigger, and the rates of increase tend to be exponential. Estimated risks are lowest for the S&P and FTSE contracts and highest for the Hang Seng; and there is relatively little difference between the estimated risk measures of short and long positions.

Insert Table 2 here

Insert Figure 5 here

Insert Figure 6 here

These risk estimates are only useful if we have some idea of their precision, and one common way to assess precision is through estimates of standard errors. Table 3 presents some estimates of their standard errors based on a parametric bootstrap. The results presented in this Table show that the ES standard errors are higher than the VaR standard errors for all contracts except the Nikkei 225.

Insert Table 3

However, it can be argued that a straight comparison of standard errors is likely to be unfair, because the ratio of the ES to its standard error (coefficients of variation) is generally larger than the ratio of the VaR to its standard error. A comparison of these coefficients of variations would then suggest that where the ES is in relative terms more precisely estimated than the VaR. To investigate this issue further, Table 4 presents estimates of the 90% confidence intervals for the two risk measures. To facilitate their comparison, these confidence intervals are standardised (i.e., divided) by the corresponding mean risk measures. So, for example, for a confidence level of 98%, the first result on the top row of Table 4 tells us that the 90% confidence interval spans the range 0.9476 to 1.0560 of the mean (or expected) VaR. The results presented here show that the ES confidence intervals are narrower than the VaR ones, and confirm that the ES is estimated relatively more precisely than the VaR.

Insert Table 4 here

These results also indicate a second interesting finding: for relatively low confidence levels they are more or less symmetric, but for higher confidence levels (and especially a confidence level of 0.999), the confidence intervals are notably asymmetric: in particular, the right bound of the confidence level is further away from the mean than the left bound. This tells us that we cannot use symmetric textbook confidence intervals based on the underlying central limit theorem assumptions when dealing with risk measures at very extreme confidence levels. The explanation for this asymmetry is that at these very high confidence levels the data are notably more sparse on the right-hand side than on the left-hand side of the confidence interval, and this greater sparsity pulls the right-hand bound of the interval further away from the mean.

6. ESTIMATION OF SPECTRAL RISK MEASURES

As noted already, to apply a spectral risk measure requires that we first choose a suitable value for the risk-aversion parameter γ . In principle, the value of this parameter is a matter of choice, subject only to the constraint that γ be positive. However, in the present context where we are primarily concerned with extremes, we need a relatively low γ value (i.e., a high-degree of risk-aversion) for the application of an EV approach to make sense. The reason why can be seen in Figure 1. If we have a relatively low degree of risk aversion (e.g., such as that implied by $\gamma = 0.05$), then the risk-aversion weights applied to extreme losses are not much higher than those for somewhat lower losses; in such circumstances, we would be concerned with losses over a relatively wide range of confidence levels and we would have no reason to apply an EV approach in the first place. On the other hand, with a very high degree of risk-aversion (e.g., $\gamma = 0.005$), then the weights attached to extreme losses are much higher than those attached to smaller losses – indeed, the weights attached to smaller losses are negligible – and in these circumstances we have little real choice but to apply an EV approach. A casual inspection of Figure 1 thus suggests that in the present context we would want to work with a γ value well below 0.05.

Having chosen a γ value, we calculate the integral (13), which we can do using numerical integration. This requires that we approximate the continuous integral by a discrete equivalent: we discretise the continuous variable p into a number N of discrete slices, where the approximation gets better as N gets larger. We also have to choose a suitable numerical integration method, and some obvious ones are trapezoidal and Simpson's rules, and quasi-Monte Carlo methods (e.g., using Niederreiter and Weyl algorithms).¹⁰

To evaluate the accuracy of these methods, Table 5 gives estimates of the spectral-exponential risk-measure approximation errors generated by these alternative numerical integration methods based on various values of N and a plausible set of benchmark parameters (i.e., more specifically, the mean long-position parameters in Table 1 with $\gamma = 0.01$). These results indicate that all methods have a negative bias for relatively small values of N , but the bias disappears as N gets large. For example, on average, approximation errors are of the order of -2% to -3% for $N = 10000$ and of the order of -0.05% for $N = 1$ million. They also show that the Simpson's and trapezoidal methods are a little more accurate than the quasi-methods. This perception is confirmed by the plots in Figure 7, which show how rapidly the different integration methods converge as N gets large: the Simpson and trapezoidal methods converge most rapidly, and the quasi methods a little less so. Thus, any of these methods ought to give reliable results provided that one is able to choose a sufficiently high value of N .

Insert Table 5 here

Insert Figure 7 here

For the purposes of the remaining estimations, we selected a benchmark method consisting of the trapezoidal rule calibrated with $N=1$ million, and the results just examined suggest that this benchmark should deliver highly accurate estimates.

Estimates of the spectral-exponential risk measures themselves are given in Figure 8 and Table 6. The Figure shows plots of estimated spectral-exponential

¹⁰ One might also use pseudo-MC methods too, but results for pseudo methods are not reported because they are considerably less accurate than the methods whose results are reported here.

measures against $1 - \gamma$, which we can take as a proxy for risk-aversion: the risk measures increase as $1 - \gamma$ (or risk aversion) increases. They also increase in a similar way to the exponential increases exhibited by the estimated VaR and ES measures shown in Figures 5 and 6. The estimated spectral measures are also similar to the earlier ones in that they are lowest for the S&P and FTSE and highest for the Hang Seng, and show relatively little asymmetry across short and long positions.

Insert Figure 8 here

Insert Table 6 here

It is also interesting to compare the rough magnitudes of the spectral and earlier risk measures. For example, the VaR at the 0.995 confidence level, the ES at the 0.99 confidence level, and the spectral risk measure with $\gamma = 0.01$ are all of much the same size. This type of comparison gives us a sense of the relative sizes of the different risk measures: it tells us that a spectral risk measure with $\gamma = 0.01$ gives us comparable margin requires as the VaR at the 0.995 confidence level; this implies that a spectral risk measure with $\gamma < 0.01$ gives us margin requirements greater than the VaR at the 0.995 confidence level, and so on.

As with the earlier risk measures, it also important to gauge the precision of these estimates. Accordingly, Table 7 presents estimates of the standard errors of the spectral risk measures based on a parametric bootstrap similar to our earlier ones.¹¹ These results show us that the estimated standard errors are broadly similar across futures contracts, but increase as γ falls: this makes sense because a falling γ implies that we are placing ever more weight on fewer observations; this suggests that the effective sample size is falling and a smaller effective sample size implies a higher standard error. However, this increase in standard errors is quite substantial: for example, for the long S&P500 contract, $\gamma = 0.05$ produces a standard error of 0.1575, whereas $\gamma = 0.005$ produces a standard error of 0.8862. In terms of a rough order of

¹¹ As with the earlier bootstrap, this involves resampling from the estimated distribution function. However, in doing so we also have to restrain the number of slices N to the sample size: the parametric bootstrap therefore involves $N=3392$.

magnitude, these results also suggest that a spectral risk measure with $\gamma = 0.01$ usually has a standard error larger than those of the VaR or ES at the 0.995 confidence level, but less than those the VaR or ES at the 0.999 confidence level.

Table 8 presents the standardised confidence intervals for the spectral risk measures (i.e., confidence intervals divided by bootstrapped means). Again, these are similar across different futures positions, but the confidence intervals also expand markedly as γ falls, again reflecting the fact that the effective sample size falls as γ gets smaller. For example, for the long S&P500 contract considered earlier, $\gamma = 0.05$ produces a confidence interval of [0.8895 1.1143], whereas $\gamma = 0.005$ produces a standard error of [0.6769 1.3806]. The confidence intervals for very low γ values also show a small asymmetry similar to those we got earlier with VaR and ES confidence intervals predicated on extremely high confidence levels.

However, when we compare the magnitudes of the spectral confidence intervals with the earlier ones, we see that the confidence intervals of the spectral risk measures are considerably wider. For example, the standardised confidence intervals for the spectral risk measures for $\gamma = 0.01$ are notably wider than those for the VaR or ES at the 0.999 confidence interval: as a rough average, the former are about [0.77 1.25], whereas the latter are about [0.84 1.19]. Naturally, the confidence intervals for spectral risk measures with $\gamma = 0.005$ are even wider. Thus, our results present clear evidence that estimates of spectral-exponential risk measures are less precise than estimates of VaR or ES. The explanation for the wider confidence intervals of the spectral measures would again appear to be related to effective sample size: to see this, note that the ES has a fixed sample size given by the number of tail observations, whereas a spectral risk measure predicated on a high degrees of risk aversion effectively puts a lot of weight on a small number of tail observations, and therefore operates with fewer effective observations.

7. DISCUSSION

All our estimated risk measures show considerable similarity. They all agree that the S&P and FTSE contracts are the least risky indices, and that the Hang Seng is the most risky. The use of any of these measures for setting margin requirements would

therefore lead to the former ones having the lowest margin requirements and to the Hang Seng having the highest. All the estimated risk measures also agree that there is only mild asymmetry across long and short positions, and this suggests that there should be only small differences across the margin requirements of long and short positions.

It is also interesting that all three types of risk measure react in similar ways as the appropriate key parameter (i.e., the confidence level in the case of the VaR and ETL, and γ in the case of the spectral measures) changes. In all cases, the estimated risk measure depends critically on the value of the key parameter, and the relevant plots (of risk measure against key parameter) are all close to exponential. However, whilst these plots look broadly similar, the different risk measures vary considerably in terms of their interpretation and usefulness:

- The VaR is of limited use for setting margin requirements because a VaR as such gives the clearinghouse no indication of how big its losses might be in the event that an investor suffers losses that exhaust its margin. As far as the clearinghouse is concerned, the only practically useful VaR information relates to the probability of default associated with any given VaR: if clearinghouse risk managers have a VaR calculation engine, then they can use it to work out the default probabilities associated with particular margin requirements interpreted as VaRs. VaR information is also limited in so far as it is contingent on the confidence level, but there is little a priori guide to tell the clearinghouse what particular confidence level it should work with.
- The ES is more useful to the clearinghouse than the VaR because it does, and the VaR does not, take account of the sizes of losses higher than the VaR itself. It also has the helpful interpretation that it tells the clearinghouse the loss an investor can expect to make conditional on it experiencing a loss that exceeds a chosen VaR threshold. So if the clearinghouse sets a VaR-based margin requirement, then the ES tells the clearinghouse the expected default loss conditional on the investor experiencing a loss that exceeds its margin. This said, the ES measure is also contingent on a confidence level, and there is little a priori guide to tell the clearinghouse what that should be.

- The spectral risk measures are in principle the most useful, because they alone take account of the user's (i.e., clearinghouse's) degree of risk aversion: the more risk averse the user, the greater the risk measure. This gives clearinghouse risk managers an opportunity to select a γ value that reflects the clearinghouse's corporate risk aversion. So, whereas the VaR or the ES are contingent on a key parameter whose value is to a greater or lesser extent arbitrary, spectral measures are contingent on a key parameter whose value can be determined from the clearinghouse's risk appetite. Thus, spectral measures are generally better *both* because they take account of the clearinghouse's risk aversion, and because their key parameter can be chosen from it. Spectral risk measures can be higher or lower than the earlier risk measures, depending on the sizes of the two key parameters, but our confidence-interval results also suggest that spectral risk measures have a tendency to be less precisely estimated than the VaR or the ES – which is unfortunate given the advantages of spectral risk measures over earlier ones.

8. SUMMARY AND CONCLUSIONS

By acting as a counterparty in all trades, a clearinghouse relieves individual traders of credit risk concerns but acquires credit-risk exposures of its own. It then seeks to manage these exposures by imposing margin requirements. However, even with margin requirements the clearinghouse is still exposed to the risk of loss arising from investor defaults triggered by extreme price movements located in the tail of a distribution. This paper has sought to model these risks using an EV approach, where margin requirements might be set using one of three different financial risk measures: the VaR, the Expected Shortfall (ES), and spectral-coherent risk measures.

This paper compares the three approaches for margin setting as to whether risk aversion is explicitly incorporated into the risk measures. Both VaR and ES do not incorporate risk aversion, and as a consequence result in inappropriate attitudes to risk, being either risk lovers (VaR) or risk neutral (ES). In contrast, spectral risk measures allow clearinghouses specifically exhibit their degree of risk aversion, and importantly attach greater weights to the more extreme price movements, those associated with investor default. The paper illustrates the relative merits of the

approaches by noting, for example, that spectral measures are not only coherent like ES measures, but also allow for unique margins with varying risk aversion across contract type and trading position.

Furthermore, the paper examines the empirical features of the alternative approaches to margin setting. Comparing the margin requirements based on VaR and ES measures, we see two desirable results for the latter, namely, they take some account of losses beyond the various quantiles, and are also more precise than VaR measures. Whilst the spectral measures incorporate varying degrees of risk aversion by giving higher weights to the more extreme returns unlike the equal weighting scheme of ES measures, the paper finds that the spectral measures are generally less precise. Given the obvious benefits that spectral measures offer economic agents, it would be interesting to determine if the choice of utility function is driving the lack of precision and this will be investigated in future work.

REFERENCES

- Acerbi, C., Tasche, D., 2001. Expected shortfall: a natural alternative to value at risk. *Economic Notes* 31: 379-388.
- Acerbi, C., 2002. Spectral measures of risk: a coherent representation of subjective risk aversion. *Journal of Banking and Finance* 26: 1505-1518.
- Acerbi, C., 2004. Coherent representations of subjective risk-aversion, in G. Szegö (Ed), *Risk Measures for the 21st Century*, Wiley, New York, pp. 147-207.
- Artzner, P., F. Delbaen, J.-M. Eber, D. Heath. 1999. Coherent measures of risk. *Mathematical Finance* 9, 203-228.
- Balkema, A. A., de Haan, L., 1974. Residual lifetime at great age. *Annals of Probability* 2, 792-804.
- Bawa, V. S., 1975. Optimal rules for ordering uncertain prospects. *Journal of Financial Economics* 2: 95-121.
- Booth, G. G., Broussard, J.P., Martikainen, T., Puttonen, V., 1997. Prudent margin levels in the Finnish stock index futures market. *Management Science* 43, 1177-1188.
- Broussard, J. P. 2001. Extreme-value and margin setting with and without price limits, *Quarterly Review of Economics and Finance*, 41, 365-385.

- Cotter, J. 2001. Margin exceedences for European stock index futures using extreme value theory. *Journal of Banking and Finance* 25, 1475-1502.
- Edwards, F. R., Neftci, S. N., 1988. Extreme price movements and margin levels in futures markets. *Journal of Futures Markets* 8, 639-655.
- Embrechts, P., Kluppelberg C., Mikosch, T., 1997. *Modelling Extremal Events for Insurance and Finance*. Springer Verlag, Berlin.
- Figlewski, S. 1984. Margins and market integrity: margin setting for stock index futures and options. *The Journal of Futures Markets*, 4, 385–416.
- Fishburn, P. C. 1977 Mean-risk analysis with risk associated with below-target returns. *American Economic Review* 67: 116-126.
- Giannopoulos, K. Tunaru, R., 2004. Coherent risk measures under filtered historical simulation. *Journal of Banking and Finance*, Forthcoming.
- Grootveld, H., Hallerbach, W. G., 2004. Upgrading value-at-risk from diagnostic metric to decision variable: a wise thing to do?, in G. Szegö (Ed.) *Risk Measures for the 21st Century*. Wiley, New York, pp. 33-50.
- Kusuoka, S., 2001. On law invariant coherent risk measures. *Advances in Mathematical Economics* 3: 83-95.
- Longin, F., 1999. Optimal margin levels in futures markets: extreme price movements. *Journal of Futures Markets*, 19, 127-152.
- Longin, F., 2000. From value at risk to stress testing: The extreme value approach. *Journal of Banking and Finance*, 24, 1097-1130.
- Loretan, M., Phillips, P. C. B., 1994. Testing the covariance stationarity of heavy-tailed time series. *Journal of Empirical Finance*, 1, 211-248.
- Miranda, M. J., Fackler, P. L., 2002. *Applied Computational Economics and Finance*. MIT Press, Cambridge MA and London.
- Warshawsky, M. J. 1989. The adequacy and consistency of margin requirements: the cash, futures and options segments of the equity markets. *The Review of Futures Markets* 8, 420–437.

TABLES

Table 1: GPD Parameters for Futures Indexes

Futures index	Long position					Short position				
	u	$prob$	N_u	Tail $\hat{\xi}$	Scale $\hat{\beta}$	u	$prob$	N_u	Tail $\hat{\xi}$	Scale $\hat{\beta}$
S&P500	2.00	0.04	130	0.18 (0.10)	0.60 (0.08)	2.00	0.03	118	0.13 (0.15)	0.76 (0.13)
FTSE100	1.50	0.07	250	0.10 (0.08)	0.71 (0.07)	1.50	0.08	276	0.02 (0.07)	0.73 (0.07)
DAX	2.00	0.07	235	0.01 (0.05)	1.19 (0.10)	2.00	0.07	237	0.05 (0.07)	1.00 (0.10)
Hang Seng	2.00	0.10	353	0.13 (0.06)	1.18 (0.10)	2.00	0.11	367	0.14 (0.05)	1.15 (0.09)
Nikkei 225	2.00	0.08	277	-0.01 (0.06)	0.89 (0.07)	2.00	0.08	255	-0.07 (0.05)	1.04 (0.08)

Notes: The Table presents estimates of the GPD parameters for long and short futures positions in the five contracts shown. The sample size n is 3392, the threshold is u , the probability of an observation in excess of u is $prob$, the number of exceedences in excess of u is N_u , the estimated tail parameter is $\hat{\xi}$ and the estimated scale parameter is $\hat{\beta}$. The numbers in brackets are the estimated standard errors of the parameters concerned. The thresholds u are chosen as the approximate points where the QQ plots in Figure 2 change slope.

Table 2: Estimates of GPD VaRs and Expected Shortfalls for Futures Positions

Futures index	Long position				Short position			
	$\alpha=0.98$	$\alpha=0.99$	$\alpha=0.995$	$\alpha=0.999$	$\alpha=0.98$	$\alpha=0.99$	$\alpha=0.995$	$\alpha=0.999$
	VaR							
S&P500	2.414	2.912	3.476	5.092	2.436	3.029	3.677	5.428
FTSE100	2.489	3.070	3.692	5.315	2.539	3.063	3.594	4.857
DAX	3.488	4.326	5.170	7.152	3.291	4.042	4.819	6.731
Hang Seng	4.171	5.231	6.392	9.526	4.190	5.250	6.419	9.611
Nikkei 225	3.243	3.850	4.452	5.833	3.315	3.957	4.568	5.877
	Expected Shortfall (ES)							
S&P500	3.237	3.844	4.532	6.503	3.375	4.056	4.801	6.813
FTSE100	3.388	4.033	4.725	6.527	3.305	3.840	4.382	5.670
DAX	4.705	5.551	6.404	8.406	4.411	5.202	6.020	8.033
Hang Seng	5.851	7.070	8.404	12.007	5.884	7.117	8.475	12.188
Nikkei 225	4.112	4.712	5.308	6.677	4.201	4.801	5.372	6.595

Notes: Estimates in daily % return terms based on the parameter values shown in Table 1, where α is the confidence level and the holding period is 1 day.

Table 3: Standard Errors for VaRs and Expected Shortfalls

Futures index	Long position				Short position			
	$\alpha=0.98$	$\alpha=0.99$	$\alpha=0.995$	$\alpha=0.999$	$\alpha=0.98$	$\alpha=0.99$	$\alpha=0.995$	$\alpha=0.999$
	VaR							
S&P500	0.0811	0.1311	0.2028	0.6386	0.0977	0.1500	0.2331	0.6555
FTSE100	0.0954	0.1448	0.2195	0.5909	0.0882	0.1309	0.1830	0.4210
DAX	0.1438	0.2030	0.2916	0.6629	0.1278	0.1842	0.2724	0.6552
Hang Seng	0.1738	0.2667	0.4147	1.1749	0.1735	0.2700	0.4201	1.2130
Nikkei 225	0.1037	0.1490	0.2095	0.4546	0.1551	0.1522	0.2049	0.4079
	Expected Shortfall							
S&P500	0.0976	0.1598	0.2498	0.7789	0.1110	0.1742	0.2633	0.7440
FTSE100	0.1089	0.1609	0.2406	0.6581	0.0906	0.1312	0.1914	0.4321
DAX	0.1462	0.2112	0.2921	0.6795	0.1335	0.1999	0.2811	0.6875
Hang Seng	0.2025	0.3069	0.4775	1.3617	0.2017	0.3197	0.4932	1.4061
Nikkei 225	0.1036	0.1465	0.2047	0.4478	0.1071	0.1446	0.1934	0.3851

Notes: Estimates in daily % return terms based on a parametric bootstrap with 5000 resamples using the parameter values shown in Table 1. α is the confidence level and the holding period is 1 day.

Table 4: 90% Confidence Intervals for VaRs and Expected Shortfalls

	$\alpha=0.98$	$\alpha=0.99$	$\alpha=0.995$	$\alpha=0.999$
	VaR of long position			
S&P500	[0.9476 1.0560]	[0.9294 1.0769]	[0.9072 1.1025]	[0.8243 1.2253]
FTSE100	[0.9384 1.0651]	[0.9252 1.0805]	[0.9082 1.1015]	[0.8413 1.2015]
DAX	[0.9327 1.0691]	[0.9245 1.0797]	[0.9107 1.0965]	[0.8602 1.1638]
Hang Seng	[0.9346 1.0731]	[0.9179 1.0859]	[0.9001 1.1140]	[0.8250 1.2214]
Nikkei 225	[0.9479 1.0533]	[0.9376 1.0659]	[0.9258 1.0810]	[0.8825 1.1386]
	VaR of short position			
S&P500	[0.9348 1.0693]	[0.9224 1.0870]	[0.9002 1.1120]	[0.8318 1.2177]
FTSE100	[0.9434 1.0580]	[0.9310 1.0724]	[0.9202 1.0880]	[0.8729 1.1574]
DAX	[0.9351 1.0646]	[0.9271 1.0785]	[0.9106 1.0959]	[0.8557 1.1731]
Hang Seng	[0.9325 1.0713]	[0.9192 1.0882]	[0.8959 1.1137]	[0.8222 1.2269]
Nikkei 225	[0.9440 1.0587]	[0.9385 1.0647]	[0.9263 1.0762]	[0.8935 1.1186]
	Expected Shortfall of long position			
S&P500	[0.9519 1.0515]	[0.9338 1.0711]	[0.9141 1.0985]	[0.8334 1.2221]
FTSE100	[0.9476 1.0536]	[0.9377 1.0671]	[0.9198 1.0903]	[0.8542 1.1776]
DAX	[0.9499 1.0517]	[0.9399 1.0654]	[0.9295 1.0805]	[0.8821 1.1413]
Hang Seng	[0.9450 1.0592]	[0.9315 1.0746]	[0.9111 1.0978]	[0.8366 1.2073]
Nikkei 225	[0.9583 1.0420]	[0.9509 1.0527]	[0.9382 1.0653]	[0.8994 1.1156]
	Expected Shortfall of short position			
S&P500	[0.9482 1.0543]	[0.9312 1.0731]	[0.9139 1.0944]	[0.8450 1.1996]
FTSE100	[0.9559 1.0469]	[0.9456 1.0581]	[0.9313 1.0752]	[0.8878 1.1366]
DAX	[0.9516 1.0501]	[0.9398 1.0662]	[0.9278 1.0831]	[0.8737 1.1536]
Hang Seng	[0.9469 1.0580]	[0.9297 1.0777]	[0.9072 1.1011]	[0.8363 1.2112]
Nikkei 225	[0.9587 1.0426]	[0.9521 1.0508]	[0.9426 1.0623]	[0.9107 1.1018]

Notes: Estimates in daily % return terms based on a parametric bootstrap with 5000 resamples using the parameter values shown in Table 1. α is the confidence level and the holding period is 1 day. Bounds of confidence intervals are standardised (i.e., divided) by the means of the bootstrapped estimates.

Table 5: Approximation Errors (%) of Numerical Integration Estimates of Spectral-Exponential Risk Measure

Numerical integration method	Number of slices (N)				
	1000	10000	100000	1000000	10000000
Trapezoidal rule	-16.38	-2.48	-0.34	-0.04	0.00
Simpson's rule	-16.67	-2.51	-0.34	-0.04	0.00
Niederreiter quasi MC	-14.27	-3.49	-0.56	-0.07	0.00
Weyl quasi MC	-14.27	-3.49	-0.56	-0.07	0.00

Notes: Estimates are based on the mean long-position parameters in Table 1 (i.e., $\beta = 0.914$, $\xi = 0.082$, threshold = 1.9, and $N_u = 249$), and $\gamma = 0.01$. Errors are assessed against a 'true' value of 4.595 obtained using the trapezoidal rule with $N = 20$ million. Estimates of pseudo MC errors are standard derivations of sample pseudo risk estimates based on samples of size 100.

Table 6: Estimates of Spectral-Exponential Risk Measures for Futures Positions

Futures index	$\gamma = 0.05$	$\gamma = 0.01$	$\gamma = 0.005$
Spectral-exponential risk of long position			
S&P500	2.2965	3.5143	4.156
FTSE100	2.2871	3.6629	4.326
DAX	3.0894	5.0365	5.884
Hang Seng	3.8460	6.3850	7.651
Nikkei 225	2.9378	4.3428	4.940
Spectral-exponential risk of short position			
S&P500	2.2549	3.6731	4.380
FTSE100	2.2973	3.5165	4.053
DAX	2.9767	4.7331	5.533
Hang Seng	3.8804	6.4284	7.713
Nikkei 225	2.9355	4.4180	5.006

Notes: Estimates based on the parameter values shown in Table 1, using the trapezoidal integration method with $N=1$ million.

Table 7: Standard Errors for Spectral-Exponential Risk Measures

Futures index	Long position			Short position		
	$\gamma = 0.05$	$\gamma = 0.01$	$\gamma = 0.005$	$\gamma = 0.05$	$\gamma = 0.01$	$\gamma = 0.005$
S&P500	0.1575	0.5273	0.8862	0.1662	0.5636	0.9247
FTSE100	0.1626	0.5405	0.8960	0.1538	0.5009	0.7988
DAX	0.2226	0.7363	1.1901	0.2117	0.7103	1.1483
Hang Seng	0.2809	0.9724	1.6352	0.2866	0.9934	1.6845
Nikkei 225	0.1950	0.6018	0.9576	0.1969	0.6173	0.9702

Notes: Estimates in daily % return terms based on a parametric bootstrap with 5000 resamples using the parameter values shown in Table 1. The holding period is 1 day.

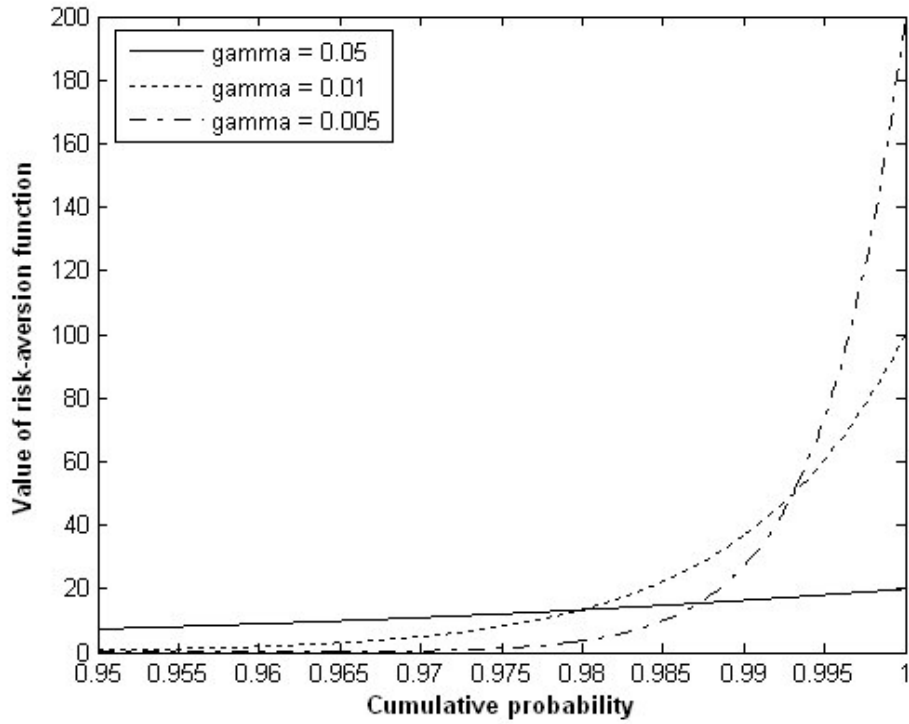
Table 8: 90% Confidence Intervals for Spectral-Exponential Risk Measures

Futures index	$\gamma = 0.05$	$\gamma = 0.01$	$\gamma = 0.005$
	Long position		
S&P500	[0.8895 1.1143]	[0.7682 1.2576]	[0.6769 1.3806]
FTSE100	[0.8860 1.1221]	[0.7704 1.2505]	[0.6822 1.3594]
DAX	[0.8844 1.1231]	[0.7697 1.2512]	[0.6834 1.3597]
Hang Seng	[0.8824 1.1235]	[0.7661 1.2610]	[0.6742 1.3758]
Nikkei 225	[0.8938 1.1123]	[0.7783 1.2314]	[0.7004 1.3378]
	Short position		
S&P500	[0.8825 1.1257]	[0.7622 1.2654]	[0.6768 1.3737]
FTSE100	[0.8931 1.1126]	[0.7711 1.2454]	[0.6963 1.3394]
DAX	[0.8854 1.1167]	[0.7648 1.2554]	[0.6816 1.3556]
Hang Seng	[0.8818 1.1258]	[0.7623 1.2660]	[0.6602 1.3816]
Nikkei 225	[0.8924 1.1117]	[0.7790 1.2424]	[0.6958 1.3277]

Notes: Estimates in daily % return terms based on a parametric bootstrap with 5000 resamples using the parameter values shown in Table 1. The holding period is 1 day. Bounds of confidence intervals are standardised (i.e., divided) by the means of the bootstrapped estimates.

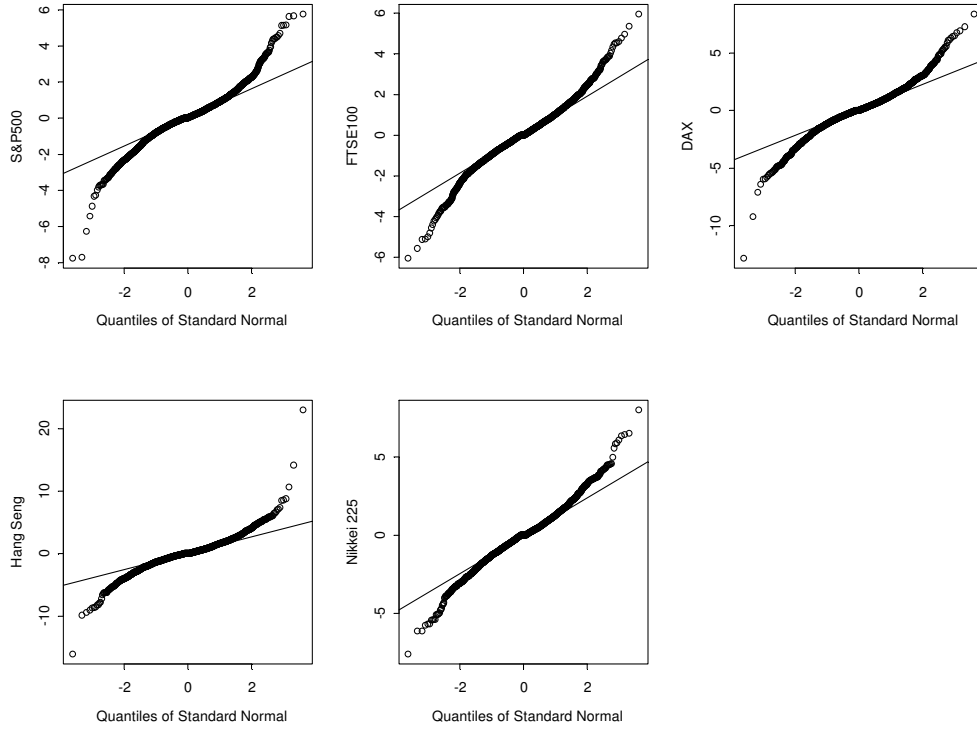
FIGURES

Figure 1: Exponential Risk-Aversion Function for Various Values of γ



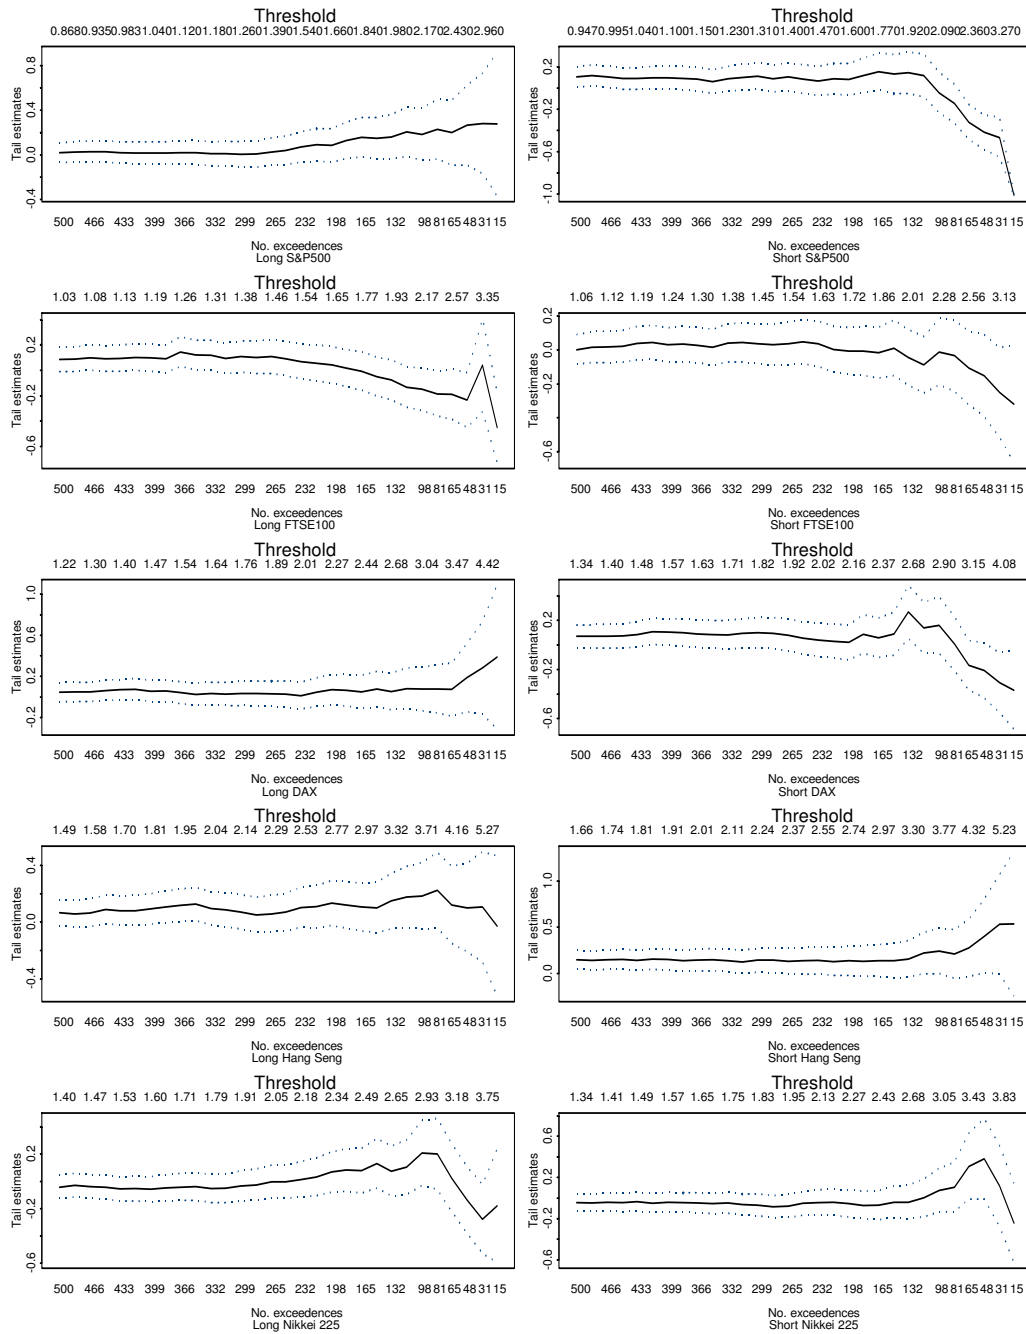
Notes: Based on equation (4) in the text, for stated γ values.

Figure 2: QQ Plots for Futures Return Indexes



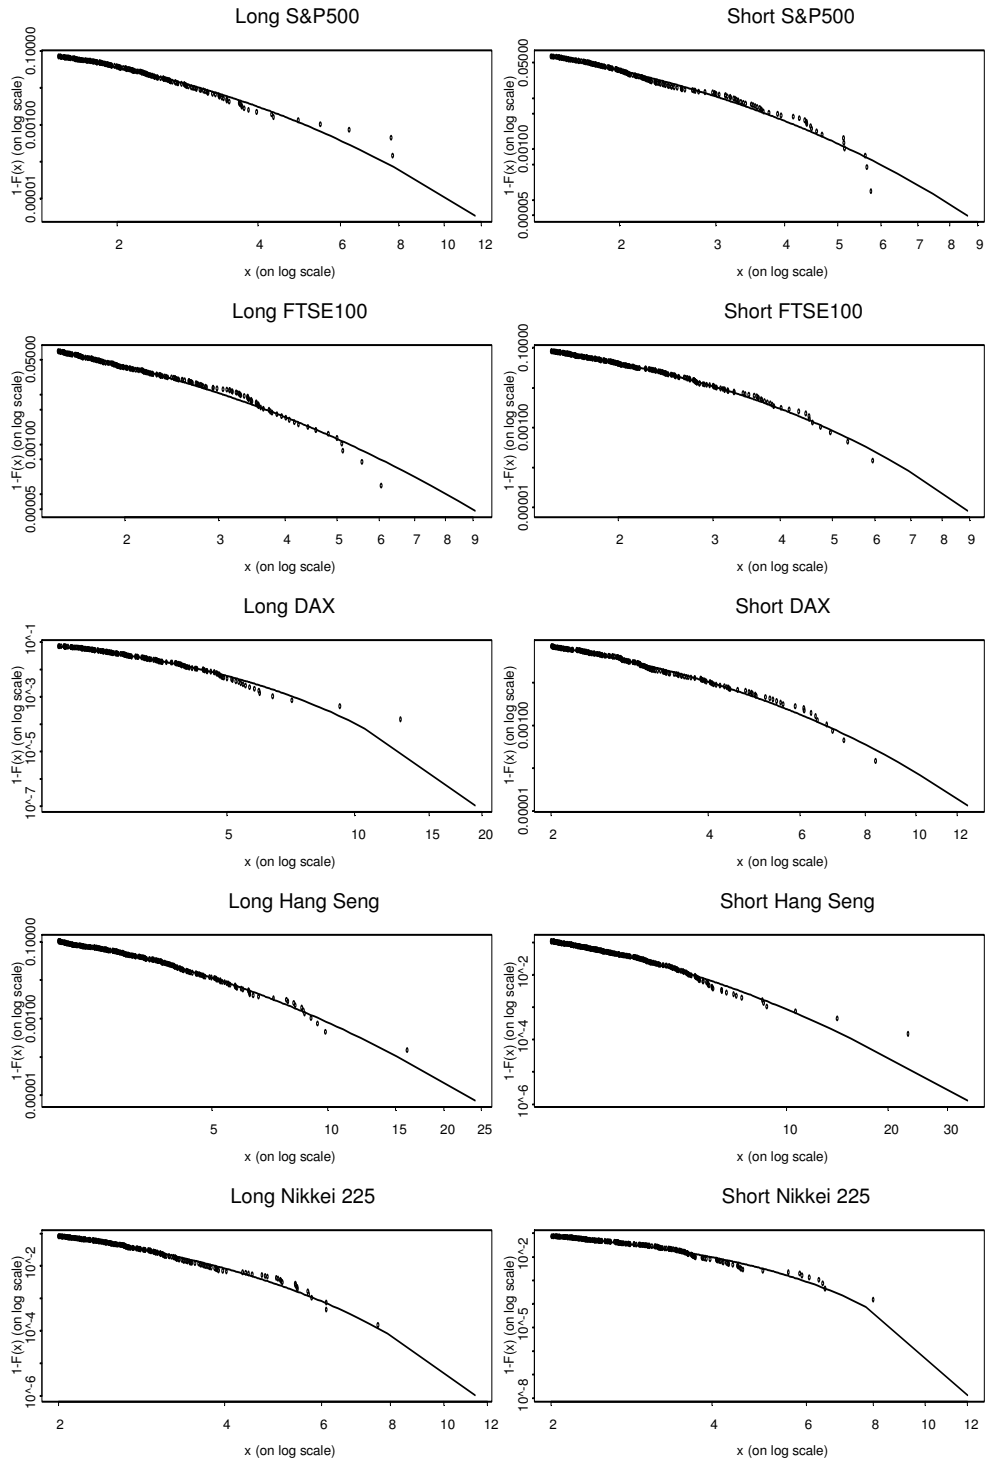
Notes: Quantiles of the respective empirical return distribution against those of normal distributions.

Figure 3: Tail Index Plots as Functions of Numbers of Exceedances



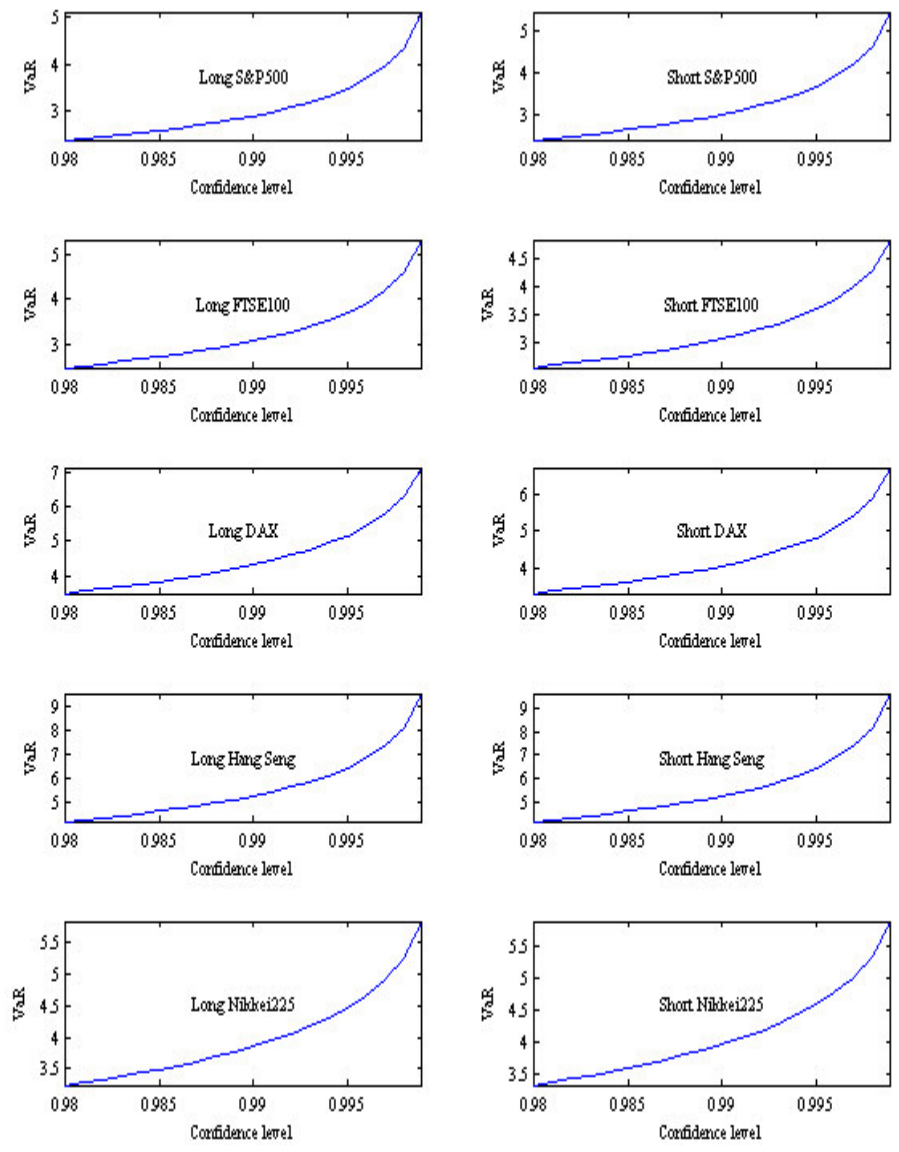
Notes: Tail estimates with 95% confidence bands are presented as a function of threshold size and number of exceedances.

Figure 4: Exceedances Fitted to GPD




Notes: Tail distribution based on the parameter values given in Table 1.

Figure 5: Generalised Pareto VaRs of Futures Positions at Extreme Confidence Levels



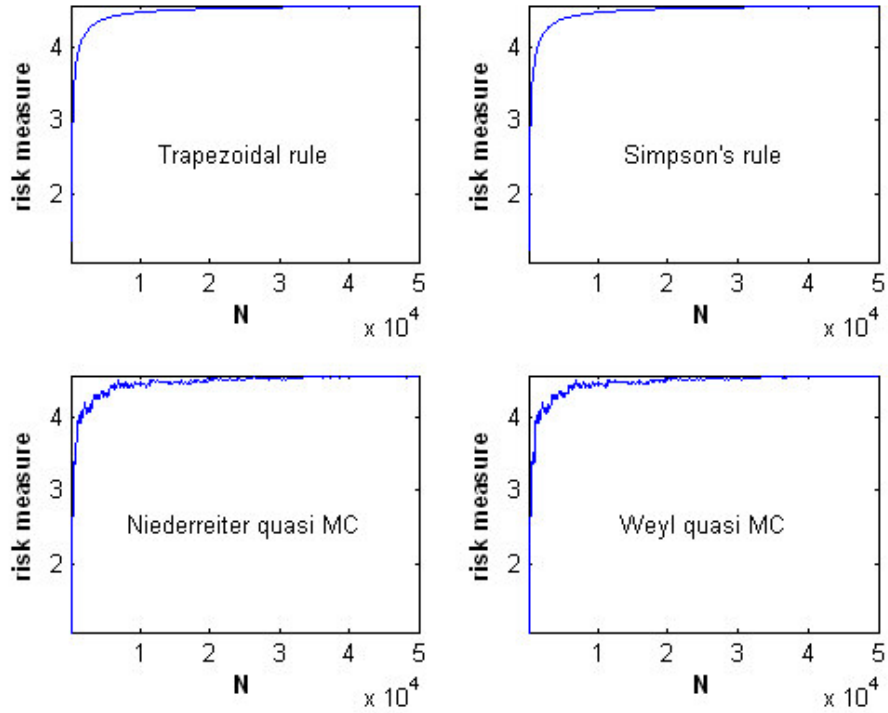
Notes: Based on the parameter values given in Table 1.

Figure 6: Generalised Pareto Expected Shortfalls of Futures Positions at Extreme Confidence Levels




Notes: Based on the parameter values given in Table 1.

Figure 7: Plots of Estimated Spectral-Exponential Risk Measures Against the Number of Slices, N



Notes: Each plot shows the estimated spectral-exponential risk measure against N , where N covers the range 100 to 50000 in steps of 100, obtained using the numerical integration routines shown on each plot. Estimates are based on the mean long-position parameters in Table 1 (i.e., $\beta = 0.914$, $\xi = 0.082$, threshold=1.9, and $N_u = 249$) along with $\gamma = 0.01$.

Figure 8: Spectral-Exponential Risk Measures of Futures Positions



Notes: Based on the parameter values given in Table 1.