

A comparison of extreme value theory approaches for determining value at risk

Article

Accepted Version

Brooks, C. ORCID: <https://orcid.org/0000-0002-2668-1153>,
Clare, A. D., Dalle Molle, J. W. and Persaud, G. (2005) A
comparison of extreme value theory approaches for
determining value at risk. *Journal of Empirical Finance*, 12 (2).
pp. 339-352. ISSN 0927-5398 doi:
10.1016/j.jempfin.2004.01.004 Available at
<https://centaur.reading.ac.uk/20555/>

It is advisable to refer to the publisher's version if you intend to cite from the
work. See [Guidance on citing](#).

Published version at: <http://dx.doi.org/10.1016/j.jempfin.2004.01.004>

To link to this article DOI: <http://dx.doi.org/10.1016/j.jempfin.2004.01.004>

Publisher: Elsevier

All outputs in CentAUR are protected by Intellectual Property Rights law,
including copyright law. Copyright and IPR is retained by the creators or other
copyright holders. Terms and conditions for use of this material are defined in
the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

NOTICE: this is the author's version of a work that was accepted for publication in the *Journal of Empirical Finance*. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in the *Journal of Empirical Finance*, 12.2 (2005), DOI: 10.1016/j.jempfin.2004.01.004

A Comparison of Extreme Value Theory Approaches for Determining Value at Risk

by

C. Brooks, A. D. Clare, J.W. Dalle Molle and G. Persand*

December 2003

Abstract

This paper compares a number of different extreme value models for determining the value at risk of three LIFFE futures contracts. A semi-nonparametric approach is also proposed where the tail events are modeled using the Generalised Pareto Distribution and normal market conditions are captured by the empirical distribution function. The value at risk estimates from this approach are compared with those of standard nonparametric extreme value tail estimation approaches, with a small sample bias-corrected extreme value approach, and with those calculated from bootstrapping the unconditional density and bootstrapping from a GARCH(1,1) model. The results indicate that for a hold-out sample, the proposed semi-nonparametric extreme value approach yields superior results to other methods, but the small sample tail index technique is also accurate.

Keywords: Bootstrap, Value at Risk (VaR), Generalised Pareto Distribution, Parametric, Semi-nonparametric and Small Sample Bias Corrected Tail Index Estimators, GARCH models.

JEL Classifications: C14, C15, G13

* We would like to thank the Editor, Franz Palm, an Associate Editor of this journal, and an anonymous referee for useful comments on a previous version of this paper, and Salih Neftci for useful conversations. The usual disclaimer applies. Brooks (corresponding author) is with the ISMA Centre, University of Reading, Whiteknights Park, PO Box 242, Reading, RG6 6BA, United Kingdom, tel: (+44) 118 931 67 68; fax: (+44) 118 931 47 41; e-mail: C.Brooks@rdg.ac.uk ; Clare is with Legal and General Asset Management, London; Dalle Molle is from Massey University, New Zealand; and Persand is with the University of Southampton, UK.

1. Introduction

The calculation of a financial institution's Value at Risk (VaR) has now become the standard approach to the determination of appropriate levels of bank capital. For example, under the European Union's second Capital Adequacy Directive (CAD II), the use of internal risk management models (IRMM) is now permitted as long as the institutions can demonstrate that the model and the operational procedures relating to the model are "sound".

The standard ("delta-normal") Value at Risk methodology requires that the underlying returns generating distribution for the security in question is normally distributed, with moments which can be estimated using historical data and are time-invariant (see Danielsson and DeVries (2000) or Neftci (2000)). However, the stylised fact that returns are fat-tailed is likely to lead to under-prediction of both the size of extreme market movements and the frequency with which they occur. These under-predictions have potentially serious solvency implications in the context of future margin systems. Margin setting in futures markets is known to be sensitive to the occurrence of large price changes. Margin committees and brokers in future markets must deal with the following trade-off in the process of setting the margin level: A high margin level protects brokers against insolvent customers, which in turn reinforces market integrity, but it also increases the costs that must be supported by investors, which in the end makes the market less attractive. Various approaches have been used to derive the margin level for a given probability of margin violation desired by margin committees or brokers. Longin (1994) proposed a new method to set margins that takes into account the appropriate proportion of extremes in the distribution of price changes and provides a simple analytical formula to compute the "*optimal*" margin level. Broussard (2001) and Broussard and Booth (1998) also analyze future margins series using such methods.

Neftci (2000) argues that it is likely that extreme events are "structurally" different from the return generating process under normal market conditions. An obvious response to this problem is to employ a methodology that explicitly allows for the fat-tailed nature of return distributions, such as those based on extreme value theory (EVT). Although there are a number of extreme value approaches available, little empirical work to date has conducted a comparative analysis of the various methods. Moreover, extant approaches using EVT focus on the tails only and have nothing to say concerning how observations in the centre of the distribution should be used.

By contrast, the approach advocated in this paper makes use of information from both the tails and the centre of the distribution, but treats each separately. In this paper we calculate a number of VaR estimates for three of London International Financial Futures Exchange's (LIFFE's) most popular derivatives contracts.

We use an unconditional model, a GARCH(1,1) model and a bootstrapping approach based on a combination of a Generalised Pareto Distribution (GPD) and the empirical distribution of the returns. These models are compared with standard nonparametric tail index estimation methods and with an approach recently proposed in Huisman, Koedijk, Kool, and Palm (2001). Our main finding is that out-of-sample tests of the calculated VaRs show that the proportion of exceedences produced by the extreme value semi-nonparametric approach, which separately models the tail and central regions, are considerably closer to the nominal probability of violations than competing approaches which fit a single model for the whole distribution.

The remainder of this paper is organised as follows: in Section 2 we present the data sets; in Section 3 we present extreme value theory; in Section 4 we present the proposed semi-nonparametric method for estimating VaR, while nonparametric tail index estimators for comparison are described in Section 5. Section 6 displays and analyses the results and finally, Section 7 concludes.

2. Data

In this study we calculate VaRs for three futures contracts traded on LIFFE - the FTSE-100 Index Futures Contract, the Long Gilt Futures Contract and the Short Sterling Interest Rate Futures Contract - based upon their daily settlement prices¹. Although it has lost some prominence in recent years, the London futures market is still among the largest in the world, and we examine data on three of its most heavily traded contracts. This should provide an interesting parallel with the majority of existing studies that focus exclusively on US markets. Our data were collected from Primark Datastream, and span the period 24 May 1991 to 3 September 1997. Sample observations from when LIFFE was closed were deleted from the data set to avoid the incorporation of spurious zero returns, leaving 1344 observations in the in-sample estimation period (to 16 September 1996. Observations from 17 September 1996 to 3 September 1997 are reserved for out-of-sample testing). In the empirical work below, we use the daily log return of the original price series.

Table 1 exhibits that all three returns series show strong evidence of skewness – the FTSE-100 and Short Sterling contract returns are positively skewed while the returns on the Long Gilt contract are negatively skewed. They are also highly leptokurtic. In particular, the Short Sterling series has a coefficient of excess kurtosis of nearly 200. The Jarque-Bera test statistic consequently rejects normality for all three log-return series. The extreme fat-tailed nature of the three series provides a strong motivation for the estimation methodologies employed in this paper that specifically focus on the tails.

3. The Excess Distribution Function and the Generalized Pareto Distribution

The tail region of a distribution can be modeled using the conditional excess distribution function, which describes the conditional distribution of the exceedences (or excesses) over a given threshold level. The GPD is a flexible family of distributions applicable to the approximation of the conditional distribution of scaled exceedences – see Pickands (1975)². The distribution for the excesses over a high threshold converges to one of three different extreme value distributions: the exponential, Pareto, and Beta distributions, which correspond to the Gumbel, Fréchet and Weibull distributions, respectively. Introducing some notation, let $\{x_1, x_2, \dots, x_n\}$ denote the log-returns of the collection of prices $\{P_0, P_1, P_2, \dots, P_n\}$ for $t = 1, 2, \dots, n$. Let $F_X(x_t)$ denote the “unknown” distribution function of the returns. Next, let x_U and x_L represent the upper and lower thresholds of the tails respectively, such that $x_t > x_U > 0$ and $x_t < x_L < 0$ lie in the two tails of the distribution $F_X(x_t)$.

The log likelihood function for estimating the generalised Pareto distribution parameters σ_u and γ at an upper tail threshold, U , is given by

$$\ell(\gamma_U^{ML}, \sigma_U, k) = -k \ln(\sigma_U) + \left(\frac{1}{\gamma_U^{ML}} - 1 \right) \sum_{i=1}^k \ln \left(1 - \frac{\gamma_U^{ML} v_i}{\sigma_U} \right) \quad (1)$$

¹ Since these contracts expire 4 times per year - March, June, September and December - to obtain a continuous time series we use the closest to maturity contract unless the next closest has greater volume, in which case we switch to this contract.

² Also see Smith (1987), Embrechts, Kluppelberg, and Mikosch (1997) and Neftci (2000) for further discussions on the conditions required for the use of the GPD approximation.

where k is the number of exceedences in a sample of n observations and γ_U^{ML} is the maximum likelihood tail index estimator, v_i is the exceedence of the threshold for observation i and $\sigma_U > 0$ is the scale parameter. The log-likelihood function can similarly be derived for the lower threshold, L . In this case, the number of sample extremes (k) is obtained by first estimating the standard deviation of the entire sample of returns. Second, using a normality approximation for the centre of the distribution that incorporates 90% of the data (corresponding to two 5 % tail regions), selecting all returns in the sample that are greater than (less than minus) 1.96 times the standard deviation to represent the upper (lower) extremes.

The results for the sample number included in the tail (k), the sample normalising coefficient (σ) and the coefficient determining the sample fatness of the tail region (γ) are given in Table 2. The number of extremes (k) for the upper tail is higher than those of the lower tail, except for the Long Gilt contract where the number of extremes is 44 in the lower tail compared to 29 in the upper tail. As expected, σ_U is positive for all three contracts, and is highest for the FTSE-100 index contract, followed by the Long Gilt and then the Short Sterling contracts. The result is quite similar for the lower tail: σ_L is positive for all the contracts, and is highest for the FTSE-100 index contract, followed by the Short Sterling and then the Long Gilt contracts. Whereas the parameter γ is positive in the lower tail for all three contracts (the highest being for the Long Gilt contract, followed by the Short Sterling and FTSE-100 Index contracts), it is negative for the FTSE-100 Index and Long Gilt contracts in the upper tail.

The next step is to estimate the upper and lower VaR threshold levels, $T_U = T_U(1-q)$ and $T_L = T_L(q)$, which are the q^{th} and $(1-q)^{\text{th}}$ percentiles, respectively. Following the definition of x_U and x_L

$$T_U \gg \max[|x_L|, |x_U|] \text{ and } T_L \ll \max[|x_L|, |x_U|] \quad (2)$$

since for the calculation of the VaRs, it is important that the VaR threshold levels are points far from where the tail starts. Following Nefci (2000), the estimator of the tail probability is

$$1 - F_x^{ML}(x_U + v_t) \cong \frac{k}{n} \left(1 - \frac{\gamma^{ML} v_t}{\sigma_U^{ML}} \right)^{1/\gamma^{ML}} \quad (3)$$

where, σ_U^{ML} and γ^{ML} are the maximum likelihood estimates of σ_U and γ respectively. Denoting the VaR threshold probability by the percentile $q = F_x(x_U + \nu_t)$, we can obtain the VaR for the upper tail:

$$T_U = \frac{\sigma_U^{ML}}{\gamma^{ML}} \left[1 - \left(\frac{qn}{k} \right)^{\gamma^{ML}} \right] \quad (4)$$

Recall the results for both the upper and lower VaR threshold values are presented in Table 2. For the upper tail, the threshold (i.e. the start of the tail) is set at 0.01664 for the FTSE-100 Index contract, at 0.01003 for the Long Gilt contract, and at 0.00325 for the Short Sterling contract. Thus, the threshold is further in the tail for the FTSE-100 Index, followed by the Long Gilt and the Short Sterling contracts, respectively. The same result is obtained for the lower tail, with the threshold being 0.018 for the FTSE-100 Index contract, 0.00983 for the Long Gilt contract and 0.000189 for the Short Sterling contract. The threshold is higher in the lower tail for the FTSE-100 Index contract compared to the upper tail. On the other hand, the threshold is higher in the upper tail for the Short Sterling contract compared with its lower tail.

4. A Semi-nonparametric Methodology for Estimating VaR

VaR is estimated for 1 day, 1 week, 1 month and 3 month investment horizons by simulating the conditional densities of price changes, using the Efron (1982) bootstrapping methodology. The simulation study is conducted for the generalised Pareto model by bootstrapping from both the two fitted tails and from the empirical distribution function derived from the log returns.

For the GARCH model, since the standardised residuals from these models are *iid* (according to the BDS test - see Brooks, Clare and Persaud (2000)), samples are drawn randomly, with replacement, from the sample standardised residuals and a path of future x_t 's can be generated, using the estimates of the GARCH(1,1) model parameters from the sample and multi-step ahead forecasts of the conditional variance.

In the case of the generalised Pareto model, the path for future prices is simulated as follows: (1) draw samples, with replacement, of the x_t 's from the empirical distribution $F^{ML}(x_t)$, (2) if $x_t < T_L$, then draw from the generalised Pareto distribution fitted to the lower tail, (3) however, if $x_t > T_U$, then draw from the generalised Pareto distribution fitted to the upper tail, and (4) on the other hand, if x_t falls in the middle of the empirical

distribution, i.e. $T_L < x_t < T_U$, then x_t is retained. The number of draws of x^t is equal to the length of the investment horizon of 1 day, 1 week, 1 month and 3 months. This procedure can be considered as a type of structured Monte Carlo study, where particular attention is paid to the extreme returns in the tails of the distribution. It is these extreme returns that strongly influence the value of the VaR, and hence most influence the likelihood of financial distress. The reason that the generalised Pareto distribution is used for the tails rather than simply using the empirical distribution throughout is that the number of observations in the tails may be insufficient to obtain accurate results without using an appropriate fitted distribution.

To calculate the appropriate VaR, a securities firm would have to estimate the maximum loss that its trading positions might experience over the proposed holding period and at a specified coverage level. For example, by tracking the daily value of a long futures position and recording its lowest value over the sample period, the firm can report its maximum loss per contract for this particular simulated path of futures prices. Repeating this procedure for 20,000 simulated paths generates an empirical sampling distribution for this maximum loss. The expression for the maximum loss of a short position is analogously calculated.

Hsieh (1993) assumed that prices are lognormally distributed, i.e. that the lowest (highest) of the log ratios of the simulated prices over the holding period, i.e. $x_l = \ln(P_l / P_0)$ ($x_h = \ln(P_h / P_0)$), is normally distributed. However, in this paper (and focusing on long positions in the following exposition), we do not impose this restriction, but instead the distribution of x_l is transformed into a standard normal distribution by matching the moments of distribution of simulated values of x_l to one of a set of possible distributions collectively known as the Johnson system of distributions (see Johnson (1949) or Kendall, Stuart, and Ord (1987)). Matching moments to the family of Johnson distributions requires a specification of the transformation that maps the distribution of x_l to a distribution that has a standard normal distribution. In this case, matching moments implies finding a distribution whose first four moments are known, i.e. one that has the same mean, standard deviation, skewness and kurtosis as the distribution of the samples of x_l . For all three contracts, the distributions of the x_l 's were found to match the unbounded Johnson distribution. Therefore, the estimated 5th percentile of the simulated distribution of x_l fitted to the Johnson system is based on the following transformation:

$$X_{l,5}^* = \sinh\left(\frac{-1.645 - a}{b}\right) \cdot d + c \quad (5)$$

where $X_{l,5}^*$ is 5th percentile of the Johnson distribution derived from the samples of x_l , and a , b , c and d are estimated parameters whose values are determined by the first 4 sample moments of the simulated distribution of x_l . (For the short position, we use the 95th percentile of the distribution of simulated values of x_l and analogously use the 95th percentile $X_{h,95}^*$ derived from the Johnson distribution fitted to the samples of x_h).

It can be shown that the distribution of Q/P_0 where Q is the maximum loss, will depend on the distribution of P_l/P_0 . Hence, the first step is to find the 5th percentile of the distribution of x_l

$$\frac{x_l - \mu_l}{\sigma_l} \cong X_{l,5}^* \quad (6)$$

where $X_{l,5}^*$ is the 5th percentile of the resulting Johnson distribution for x_l , μ_l is the expected value of the simulated distribution of x_l and σ_l is the standard deviation of the simulated distribution of x_l . An expression for Q/P_0 and P_l/P_0 can be found by rearranging equation (6) and exponentiating both sides of the resulting expression

$$\frac{P_l}{P_0} = \exp\left((X_{l,5}^* \times \sigma_l) + \mu_l\right) \quad (7)$$

and therefore

$$\frac{Q}{P_0} = 1 - \exp\left((\pm X_{l,5}^* \times \sigma_l) + \mu_l\right) \quad (8)$$

VaR is also calculated using a semi-nonparametric procedure based on a GARCH(1,1) model. Here, the model is estimated for the in-sample returns and the bootstrapping is conducted on the standardised residuals from the estimated model. The GARCH equations are then used to construct a simulated path of returns of the required length. We also use the unconditional density to calculate VaRs in order to make a direct comparison between this and the two other approaches since this simpler approach ignores the non-linear dependence in the

conditional volatility (which would be captured by the GARCH formulation). To use the unconditional density, the x_t 's are drawn randomly, with replacement, from the in-sample returns.

5. Comparison of the Semi-nonparametric Approach with Nonparametric Estimators

The tail index can be estimated parametrically or nonparametrically. The tail index can be estimated parametrically using the maximum likelihood method, as was the case for the tail indices used in the semi-nonparametric bootstrapping procedure in this investigation. As a comparison, the results of three common nonparametric tail index estimators and a recently proposed small-sample approach will be examined. In the nonparametric estimation of the tail index of a GPD, one chooses a high threshold and then fits the tail index to the realizations in the corresponding tail region.

There are a large number of nonparametric tail index estimators that have been developed. Chapters 2 and 3 of Beirlant, Teugels, and Vyncker (1996), section 6.4 of Embrechts, Kluppelberg, and Mikosch (1997) and Pictet, Dacorogna and Müller (1998) contain discussions of a large number of these nonparametric estimators. The nonparametric estimators that are used in the comparison are those of De Haan and Resnick (1980), Hill (1975) and Pickands (1975). The Hill estimator is an estimator for the Pareto index, γ of the heavy-tailed distributions that are in the maximum domain of attraction of the heavy tailed Fréchet distribution, i.e. for positive values of the Pareto index. The Pickands and the De Haan and Resnick tail index estimators can be used over the entire range of the tail index for all three limiting extreme value distributions.

Let k be the number of upper order statistics to include in the estimator out of a sample of size n where $n \gg k$. The De Haan and Resnick (1988), Hill (1975), and Pickands (1975) estimators are based on the ordered sample, $x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(k)} \leq x_{(k+1)} \dots \leq x_{(n-1)} \leq x_{(n)}$, of the observed sample of log returns $\{x_1, x_2, \dots, x_n\}$, where $\{x_{(j)} : j = 1, 2, \dots, n\}$ denote the order values of $\{x_1, x_2, \dots, x_n\}$. The Hill (1975) estimator is given by

$$\gamma^H(k, n) = -\frac{1}{k-1} \sum_{i=1}^{k-1} \ln(x_{(i)}) - \ln(x_{(k)}) \quad (9)$$

The Hill estimator can be interpreted as the average vertical excess of the log-transformed data above a given threshold. The Pickands (1975) estimator is given by

$$\gamma^P(k, n) = \frac{1}{\ln(2)} \ln \left(\frac{X_k - X_{(2k)}}{X_{(2k)} - X_{(4k)}} \right), \quad (10)$$

and the DeHaan and Resnick (1988) estimator is given by

$$\gamma^{DR}(k, n) = \frac{\ln(X_{(1)}) - \ln(X_{(k)})}{\ln(k)} \quad (11)$$

Typically, extreme value theory for tail estimators is based on large samples but in practice, the limited number of data points in the tails leads to small sample biases. In order to address this problem, Huisman, Koedijk, Kool and Palm (2001) proposed a robust small sample bias-corrected estimator that is based on the linear regression of set of Hill tail estimates where each estimate is conditioned on a different number of observations included in the tail region. The Huisman, Koedijk, Kool and Palm estimator can be viewed as a modified Hill estimator that is based on the following expression:

$$\gamma^{HKKP}(k, n) = \beta_0 + \beta_1 k + \varepsilon(k) \quad (12)$$

for $k = k(n) = 1, 2, \dots, \text{int}[(n+0.5)/2]$. Huisman, Koedijk, Kool and Palm (2001) provide a generalised least squares estimator that corrects for the correlation and heteroscedasticity present in the residual series $\{\varepsilon(i) : i = 1, 2, \dots, k\}$. The optimal estimate for the tail index using this method is the intercept β_0 as $k \rightarrow 0$.

The following estimator of the scale parameter σ^F of a sample of k independent observations $\{X_i, i = 1, 2, \dots, n\}$ from a Fréchet distribution will be used to estimate σ_U :

$$\sigma^F = \left(\frac{1}{k} \sum_{i=1}^k X_i^{1/\gamma^F} \right)^{\gamma^F} \quad (13)$$

where γ^F is a tail index estimate corresponding to a Fréchet distribution including the γ^{DR} , γ^H , γ^{HKKP} , and γ^P estimators, and k is the number of exceedences used in the estimation of the tail index.

The results for each of the three nonparametric tail index estimation procedures, the corresponding estimated threshold and scaling factor in equation (13) are summarized in Table 3. For reasons of comparison, the three

nonparametric tail index estimation procedures were applied to the same section of the tail as was used above, i.e. the same values of k for the upper and lower tails and for $q = 0.01$.

To determine the accuracy of this methodology, we compared the actual daily profits and losses of the three futures contracts with their daily value at risk estimates. Our measure of model performance is a count of the number of times the VaR “underpredicts” realised losses over the out-of-sample period of 250 days from 17th September 1996 to 3rd September 1997. The Basel Committee requires the use of a 1-trading year “back-test” sample of returns in order to evaluate the suitability of the model.

6. Results

The VaRs for the three contracts based semi-nonparametric bootstrapping the unconditional density, the GARCH(1,1) and EVT models are presented together with the results of direct calculations from the nonparametric tail index estimators in Table 4. An examination of the results reveals that the VaR estimates are always higher for short compared with long futures positions, particularly as the investment horizon is increased. This is because the distribution of log-price changes is not symmetric: there is a larger probability of a price rise in all three futures contracts than a price fall over the sample period, indicating that there is a greater probability that a loss will be sustained on a short relative to a long position. For example, the VaR for a long Short Sterling position, calculated using the GARCH(1,1) model and held for three months is 3.627%, but is 5.798% for a short position.

The VaRs based upon the GARCH(1,1) model are always higher than for the unconditional density and extreme value bootstrap methods of calculation. This result highlights the excess volatility persistence implied in the GARCH(1,1) model (see Brooks, Clare and Persaud, 2000 for a discussion of this issue). A higher degree of persistence implies that a large innovation in contract returns (of either sign) causes volatility to remain high for a relatively long period, and therefore the capital level required is also higher.

Comparing among the VaRs calculated directly from the three traditional nonparametric tail index estimators, there is little to choose between them. There appears to be a tendency for the Pickands estimator to generate slightly smaller VaRs, except for the upper tail of the Short Sterling contract. The three nonparametric tail estimators typically lead to VaRs that are smaller than both the GARCH and EVT semi-nonparametric

bootstrapping approaches, but larger than those resulting from bootstrapping from the returns themselves. The VaRs calculated using the modified Hill estimator of Huisman, Koedijk, Kool, and Palm (2001) are close to those of the GARCH semi-nonparametric bootstrap approach, and are also higher than those of the other nonparametric tail estimators, particularly in the case of the Short Sterling series.

Considering how the VaRs increase with horizon, it should be noted that the 1-week, 1-month and 3-month VaRs are scaled using the square root of time rule for the nonparametric tail index estimators. This leads them to increase more slowly with horizon than the bootstrap-based approaches, where the longer period VaRs are calculated directly by drawing more bootstrapped observations and simulating over a longer time interval.

The percentages of days that the VaRs were exceeded by actual trading losses are given in Table 5 for a 250-day out of sample period for each model. The nominal probability of a violation is 5%, but some of the models show considerable deviation from this level. Whilst a “good” model would be one that generated a percentage of exceedences close to the nominal 5% value, an inadequate coverage is likely to be far more serious for the firm than having too much capital. Considering first the FTSE contract, the traditional nonparametric tail estimators lead to too low a VaR for long positions and slightly over-estimated VaRs for short positions. The procedures based on the bootstrap and the modified Hill estimator all perform better, with very few or no exceedences for the GARCH and EVT bootstrap approaches and for the modified Hill approach. The percentage of exceedences is close to 5% for bootstrapping from the unconditional distribution. The VaRs from all 6 models appear to yield adequate coverage for the long gilt since none generate any exceedences at all in the out of sample period. However, the conventional nonparametric tail estimators and the unconditional bootstrap generate insufficient VaRs on close to 40% of days for both the long and short positions in the Short Sterling futures contract. It is only the bootstrap with GARCH and EVT, and the small-sample approach to measuring the tail index that give reasonable coverage in this case. In fact, the proportion of observations in the tails fell for the Long Gilt series, while the distribution of returns on the Short Sterling changed dramatically between the in-sample and out-of-sample periods. Short-term interest rates were more volatile during the out-of-sample period and sterling was at a higher level than at any time during the in-sample period. The average percentage error between the in-sample and out-of-sample percentiles was about 2.5 times greater for the Short Sterling than for the other two series. Whilst the regulatory focus on the first percentile of returns together with

the scaling factor would considerably reduce the number of Short Sterling exceedences, this does suggest that the Basel Committee recommendation of 1-year back-tests could be insufficient.

6. Conclusions

Under the EU's CAD II, investment firms and banks in Europe are now permitted to use their own internal risk management models to calculate the required capital to cover losses in their trading positions. Given that such models are now in widespread usage, it is crucial that a body of research is generated that compares between different approaches to computing value at risk. To this end, our paper has sought to propose a new semi-nonparametric method for calculating value at risk, based on bootstrapping and extreme value theory. This approach has been applied to three of the most heavily traded LIFFE futures contracts, and was compared to the VaRs obtained from bootstrapping from the actual returns and from a GARCH model. We also examined the performance of VaRs calculated directly using three nonparametric tail index estimators. It was observed that the semi-nonparametric procedure generated more accurate VaRs than any other method based on a holdout sample of returns, although a modified Hill estimator that is explicitly designed for use with small samples also performed well.

While all three series employed in our analysis were asymmetric, the approaches proposed will also work for symmetric distributions since the left and right tails are modelled separately. Also, leaving the Short Sterling series aside, there appears to be nothing in the properties of our series that would suggest that the semi-nonparametric approach based on structured Monte Carlo should not be just as applicable to other financial time series. Within the structured Monte Carlo approach, we used both a conditional (GARCH) and an unconditional (EVT) model, and there was little to choose between them. The important issue, therefore, seems to be to treat the tails as being distinct from the rest of the distribution, and to model them separately but to incorporate information from both. The choice between conditional and unconditional models appears to be of secondary concern. Clearly, setting appropriate capital requirements represents a delicate trade off between ensuring the safety of the banking system on one hand and not tying up firms' resources unnecessarily on the other and further research in this area is warranted. The approaches proposed in this paper may also be used in risk-based margin setting systems for time when markets are stressed.

References

- Beirlant, J. J.L. Teugels, and P. Vyncker, P., 1996. Practical Analysis of Extreme Values. Leuven University Press, Leuven, Belgium.
- Brooks, C., Clare, A.D. and Persaud, G., 2000. A Word of Caution on Calculating Market-based Minimum Capital Risk Requirements, *Journal of Banking and Finance* 14, 1557-1574.
- Broussard, J.P., 2001. Extreme-value and Margin Setting With and Without Price Limits *Quarterly Review of Economics and Finance* 41, 365-385.
- Broussard, J-P. and Booth, G.G., 1998. The Behaviour of Extreme Values in Germany's Stock Index Futures: An Application to Intra-daily Margin Setting *European Journal of Operational Research* 104, 393-402.
- Daniélsson, J. and de Vries, C.G., 2000. Value-at-Risk and Extreme Returns, Embrechts, P. *Extremes and Integrated Risk Management*, London, England, RISK Books 85-106.
- De Haan, L. and Resnick, S.I., 1980. A Simple Estimate for the Index of a Stable Distribution, *Journal of the Royal Statistical Society B* 83-87.
- Efron, B., 1982, *The Jackknife, the Bootstrap, and Other Resampling Plans*, Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Embrechts, P., Kluegel, C. and Mikosch, T., 1997. *Modelling Extreme Events for Insurance and Finance*, Berlin: Germany, Springer-Verlag.
- Hill, B.M., 1975. A Simple General Approach to Inference About the Tail of a Distribution, *The Annals of Statistics* 3, 1163-1174.
- Huisman, R., Koedijk, K.G., Kool, C.J.M. and Palm, F., 2001. Tail Index Estimates in Small Samples, *Journal of Business and Economic Statistics* 19, 208-216.
- Hsieh, D.A., 1993. Implications of Nonlinear Dynamics for Financial Risk Management, *Journal of Financial and Quantitative Analysis* 28, 41-64.
- Johnson, N.L., 1949. Systems of Frequency Curves Generated by Methods of Translations, *Biometrika* 36 149-176.
- Kendall, M.G., Stuart, A, and Ord, J.K., 1987. *Kendall's Advanced Theory of Statistics: Volume 1*. Charles Griffin and Company Limited, London.
- Longin, F., 1994. Optimal Margin Level in Futures Markets: A Parametric Extreme-Based Method IFA Working Paper 192-94, London Business School, UK.

- Neftci, S.N., 2000. Value at Risk Calculations, Extreme Events, and Tail Estimation, *The Journal of Derivatives* 7, 23-37.
- Pickands III, J., 1975. Statistical Inference Using Extreme Order Statistics, *The Annals of Statistics* 3, 119-131.
- Pictet, O.V., Dacorogna, M.M. and Müller, U.A., 1998. Hill, Bootstrap and Jackknife Estimators for Heavy Tails, Adler, R.J., Feldman, R.E. and Taqqu, M.S. *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, Boston, Massachusetts: Birkhäuser Boston Inc, 283-310.
- Smith, R.L., 1987. Estimating Tails of Probability Distributions, *The Annals of Statistics* 15, 117.

Table 1: Summary Statistics of Derivative Returns

Futures Contracts	FTSE-100	Long Gilt	Short Sterling
Mean	0.00034	0.00013	0.00004
Variance	8.283E-5	2.654E-5	1.680E-6
Skewness	0.29556*	-0.09153*	8.55407*
Kurtosis	2.73215*	3.43428*	199.165*
Normality Test Statistic†	484.2252*	639.9767*	2223267*

Notes: * represents significance at the 5% level (2 tailed-test); † Jarque-Bera test

Table 2: Number of Extremes, Maximum Likelihood Parameters of the Generalised Pareto Distribution and the Threshold Level

	FTSE-100		Long Gilt		Short Sterling	
	<i>Upper Tail</i>	<i>Lower Tail</i>	<i>Upper Tail</i>	<i>Lower Tail</i>	<i>Upper Tail</i>	<i>Lower Tail</i>
k	28	19	29	44	19	15
σ^{ML}	0.02246	0.05232	0.01243	0.01324	0.00667	0.01773
γ^{ML}	-0.02521	0.03680	-0.12329	0.86250	0.15124	0.54101
$T(\gamma^{ML})$	0.01664	0.01800	0.01003	0.00983	0.00325	0.00189

Table 3: Nonparametric Tail Index Estimates of the Generalised Pareto Distribution and the Corresponding Threshold Levels

	FTSE-100		Long Gilt		Short Sterling	
	<i>Upper Tail</i>	<i>Lower Tail</i>	<i>Upper Tail</i>	<i>Lower Tail</i>	<i>Upper Tail</i>	<i>Lower Tail</i>
k	28	19	29	44	19	15
σ^{DR}	0.02463	0.02425	0.01345	0.01337	0.004171	0.003851
σ^H	0.02465	0.02366	0.01349	0.01335	0.004284	0.003849
σ^P	0.02488	0.02363	0.01358	0.01340	0.004408	0.004072
γ^{DR}	-0.3164	-0.30742	-0.3769	-0.20986	-0.8292	-0.4749
γ^H	-0.2997	0.27259	-0.2860	-0.26638	-0.59753	-0.4799
γ^P	0.08703	0.587591	0.08193	-0.12873	-0.37696	0.1305
γ^{HKKP}	0.4139	0.4738	0.4024	0.6184	0.5302	0.6224
$T(\gamma^{DR})$	0.020351	-0.008633	0.01200	-0.01800	0.001673	-0.000434
$T(\gamma^H)$	0.02024	-0.008588	0.01160	-0.01862	0.001648	-0.000434
$T(\gamma^P)$	0.01886	-0.007597	0.01078	-0.01717	0.00163	-0.000444
$T(\gamma^{HKKP})$	0.04776	-0.04211	0.02658	-0.01967	0.00707	-0.00564

Note: DR, H, P and HKKP denote the DeHaan-Resnick, Hill, Pickands, and Huisman, Koedijk, Kool, and Palm estimators respectively.

Table 4: Value at Risk Calculated by a Semi-nonparametric Bootstrap using Unconditional Density, a GARCH(1,1) model, and Extreme Value Theory, and also Calculated Directly from Nonparametric Tail Index Estimators

Panel A: Long Positions

	Nonparametric Tail Estimators				Semi-nonparametric Bootstrapping		
Horizon	DR	H	P	HKKP	Unconditional	GARCH	EVT
FTSE 100							
3 months	6.685	6.654	5.886	32.618	12.775	25.498	20.391
1 month	3.859	3.842	3.399	18.832	7.954	10.417	13.369
1 week	1.930	1.921	1.699	9.416	3.272	6.031	5.600
1 day	0.863	0.859	0.760	4.211	1.392	4.275	2.340
Long Gilt							
3 months	13.943	14.423	13.300	15.236	7.906	12.028	4.954
1 month	8.050	8.327	7.679	8.796	4.855	7.305	3.672
1 week	4.025	4.164	3.839	4.398	2.007	4.653	2.506
1 day	1.800	1.862	1.717	1.967	0.849	2.932	1.152
Short Sterling							
3 months	0.333	0.333	0.341	4.369	1.643	3.627	2.810
1 month	0.192	0.192	0.196	2.522	0.986	2.377	2.001
1 week	0.096	0.096	0.098	1.261	0.348	1.423	1.555
1 day	0.043	0.043	0.044	0.564	0.127	0.903	0.753

Panel B: Short Positions

		Nonparametric Tail Estimators			Semi-nonparametric Bootstrapping		
Horizon	DR	H	P	HKKP	Unconditional	GARCH	EVT
FTSE 100							
3 months	15.793	15.678	14.609	36.995	21.102	32.540	30.820
1 month	9.101	9.052	8.434	21.359	10.782	14.567	19.763
1 week	4.550	4.526	4.217	10.679	3.845	7.905	5.998
1 day	2.035	2.024	1.886	4.776	1.419	5.570	3.161
Long Gilt							
3 months	9.295	8.985	8.350	20.589	10.906	14.070	5.489
1 month	5.367	5.188	4.821	11.887	5.623	9.833	4.010
1 week	2.683	2.594	2.410	5.943	2.090	5.378	3.005
1 day	1.200	1.160	1.078	2.658	0.898	3.276	1.413
Short Sterling							
3 months	1.294	1.278	1.263	5.476	3.061	5.798	4.320
1 month	0.747	0.738	0.729	3.162	1.237	4.008	3.010
1 week	0.373	0.369	0.364	1.581	0.382	2.799	2.004
1 day	0.167	0.165	0.163	0.707	0.130	1.437	0.975

Note: DR, H, P and HKKP denote the DeHaan-Resnick, Hill, Pickands, and Huisman, Koedijk, Kool, and Palm estimators respectively. VaR is expressed as a percentage of the initial value of the position.

Table 5: Out-of-Sample Tests: Realised Percentages of Daily VaR Violations

	Nonparametric Tail Estimators				Semi-nonparametric Bootstrapping		
	DR	H	P	HKKP	Unconditional	GARCH	EVT
<i>Panel A: Long Position</i>							
FTSE Index	11.6	11.6	15.2	0.0	4.0	0.0	0.8
Long Gilt	0.0	0.0	0.0	1.6	0.0	0.0	0.0
Short Sterling	40.8	40.8	40.8	18.4	35.6	2.4	6.4
	Nonparametric Tail Estimators				Semi-nonparametric Bootstrapping		
	DR	H	P	HKKP	Unconditional	GARCH	EVT
<i>Panel B: Short Position</i>							
FTSE Index	0.8	1.2	8.0	0.0	4.8	0.0	0.0
Long Gilt	0.0	0.0	0.0	0.8	0.0	0.0	0.0
Short Sterling	38.4	38.4	38.4	21.2	39.6	0.4	2.4

Note: the nominal probability of VaR violations was set at 5% (see text for more details). Note: DR, H, P and HKKP denote the DeHaan-Resnick, Hill, Pickands, and Huisman, Koedijk, Kool, and Palm estimators respectively.