

# Assessing central counterparty margin coverage on futures contracts using GARCH models

*Raymond Knott* \*  
*and*  
*Marco Polenghi* \*\*

\* Financial Market Operations, Bank of England, Threadneedle Street, London, EC2R 8AH.

E-mail: [raymond.knott@bankofengland.co.uk](mailto:raymond.knott@bankofengland.co.uk)

\*\* Financial Stability, Bank of England.

E-mail: [marco.polenghi@bankofengland.co.uk](mailto:marco.polenghi@bankofengland.co.uk)

The results and the analysis contained in this paper should only be cited with the express written consent of the authors.

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We would like to thank Stephen Millard for helpful comments on an earlier draft. Any remaining errors or omissions are, of course, our own.

Word count: 9,708



## Contents

Abstract	4
Summary	5
1 Introduction	9
1.2 Previous studies of CCP margin coverage	10
1.3 Conditional vs. unconditional estimates of margin coverage	11
1.4 Conditional expected losses as a complementary measure of CCP exposure	13
2 Method	14
2.1 Modelling the returns process	14
2.2 Computing conditional expected losses	16
2.3 Estimation of the tails under EVT	17
2.4 Backtesting to select the best-performing innovation distribution	18
3 Data and preliminary analysis	19
4 Results	20
4.1 Unconditional analysis	20
4.2 Conditional analysis	21
4.3 Conditional expected losses	22
5 Discussion and conclusions	23
Tables	26
Charts	29
Appendix A: Relevant results from EVT	34
Appendix B: Hill's estimator small sample properties	35
References	37



## **Abstract**

This study considers how the probability of exceeding central counterparty (CCP) initial margin levels can be estimated, in order to provide a timely and informative measure of risk coverage. Previous studies of CCP margining have largely focussed on the unconditional distribution of returns, estimating margin coverage on a long-term average basis. The present study extends previous work by estimating conditional margin coverage using a GARCH (1,1) model, so that variations in coverage can be tracked over a much shorter time-frame. The model is applied to estimating non-coverage probabilities for two heavily-traded derivatives contracts, the Brent and FTSE 100 futures. To account for the well-documented fat-tailed characteristics of futures returns distributions, several variants of the GARCH model are estimated, assuming innovations are distributed according to Student t, Extreme Value Theory or historical distributions, and these are compared against the more usual assumption of normally distributed innovations. A backtesting procedure described by McNeil and Frey (2000) is used to select the best performing variant. Finally, the model is applied to estimating a measure of conditional expected losses, defined in this context as the expected liquidity demands on the market of a price move that exceeds initial margin levels.

Key words: EVT, GARCH, margin, futures, central counterparty clearinghouse

*JEL* classification: G28, C53



## Summary

In order to reduce the risk of non-performance for market participants, futures and other exchange-traded derivatives contracts are typically guaranteed against counterparty failure by a central counterparty clearing house (CCP). In providing this service, however, CCPs themselves become exposed to the risk of counterparty default. To protect themselves, they have developed a system of collateralisation, or margining, by which members of the CCP are required to place a sum of initial margin in a CCP account when they register positions. The initial margin is designed to provide protection against potential changes in the market value of a member's positions over a time-horizon of one or more days.

CCPs typically select the appropriate level of initial margin by inspecting the historical distribution or price movements, focussing particularly on recent price changes. After being set, however, initial margin levels will often remain static for some period of time. During this time, the coverage they provide changes dynamically, varying according to market conditions.

This paper describes a model that can be used to assess the statistical coverage provided by initial margins. Previous studies have largely concentrated on assessing long-run average coverage levels. The present paper shows how margin can be assessed on a day by day basis. In order to measure variations in margin coverage, we apply a GARCH framework to modelling the returns of two heavily traded derivatives contracts, the Brent and the FTSE 100 futures. To account for the well-documented fat-tailed characteristics of futures return distributions, several variants of the GARCH model are estimated, assuming innovations which are distributed according to Student t, Extreme Value or historical distributions, and these are compared against the more usual assumption of normally distributed innovations. A backtesting procedure described by McNeil and Frey (2000) is used to select the best performing variant. As expected, the normal distribution typically underestimates appropriate margin levels, because it fails to account for the fat-tailed properties of the return distribution. But an Extreme Value Theory variant, which is explicitly designed to model fat-tails, does not produce reliably better results than a Student t or a historical distribution, even at high coverage levels (99.85%). Overall, across all coverage levels, the Student t and historical distribution variants offer the best fit.

The model is applied to estimating a measure of conditional expected losses, defined in this context as the expected liquidity demands on the market of a price move that exceeds initial margin levels. The results quantify how the tail-shape of the return distribution can play an important role in determining the required market liquidity.





## 1 Introduction

### *1.1 Margining and CCP risk management*

Central counterparty clearing houses (CCPs) play an important role in mitigating credit risk for market participants. For exchange-traded derivatives transactions, they typically form a core part of the market infrastructure, and in OTC markets, they now play an increasingly significant role.

Through a process known as ‘novation’, CCPs interpose themselves in transactions, such that the original bilateral contracts between market participants are extinguished and replaced by new contracts with the CCP. As a result, bilateral counterparty risks of variable quality are replaced with a high-quality counterparty risk against the CCP. This has benefits both for individual traders and for markets as a whole. Since CCP members need no longer be concerned about the creditworthiness of their counterparties, they need to hold less capital against the risk of default. As a wider spectrum of counterparties now becomes available for each market participant, markets become more liquid.

In assuming responsibility for contract performance, however, CCPs themselves become exposed to counterparty risk. Moreover the funnelling of market activity through one institution concentrates this risk and the responsibility for its management in the CCP. Many CCPs clear cash and derivatives markets, and if such a CCP were to fail, activities in a wide range of markets might face disruption, with the CCP acting as a possible channel of contagion. For these reasons, central banks and regulators who are concerned with systemic stability require CCPs to maintain high standards of risk management.

To reduce the risk of counterparty default, and ensure that if a default does occur it can be absorbed with the minimum loss, CCPs have evolved a variety of risk management safeguards. Primary protection against counterparty default is provided by initial margin. When traders register positions with the CCP, they are required to place a sum of initial margin, intended to cover the market risk of the positions, in a CCP account. To prevent margin from becoming gradually eroded by market movements, positions are marked to market on a daily basis, and members are periodically ‘called’ to replenish their accounts.

Should a trader become financially distressed, and be unable to meet margin calls, the CCP has the legal right to use the margin in the CCP account to settle the trader’s outstanding obligations. Margin therefore acts as a performance bond making good the traders promise to meet any losses, and provided the initial margin level is set at a sufficiently high level, a default which results in losses for the CCP becomes effectively impossible.<sup>(1)</sup>

---

<sup>1</sup> As a further layer of protection, CCPs usually require clearing members to hold a minimum level of capital and satisfy specific credit-quality requirements. The CCP may also have access to additional default resources, which may be used if margin proves insufficient to meet losses, and may maintain a mutual guarantee, or default, fund to which members make a contribution when joining the CCP. Insurance policies may provide further cover, and some CCPs have the power to assess members for funds if other default resources prove insufficient. Margin however, provides the primary layer of protection for most CCPs. Other default resources would only be used as a last resort, when margin is exhausted.

CCPs usually set initial margins at a level that will provide them with coverage against market movements over a time horizon of one or two days. Under normal market conditions, these margin levels are quite adequate.<sup>(2)</sup> In stressed markets, however, margins can become rapidly eroded, potentially leading to uncovered losses for the CCP. Post-mortem analyses of the 1987 Crash have suggested that initial margins immediately prior to the Crash period may have been too low to provide some US CCPs with adequate protection (Warshawsky and Dietrich (1989); Bates and Craine (1999); Dewachter and Gielens (1999)). Certainly the immediate post-Crash response of some US CCPs was to raise initial margin levels, indicating that increased protection was perceived to be desirable. Several CCPs had also experienced acute liquidity problems, which may have been exacerbated by margin levels which, in hindsight, provided insufficient protection. Bernanke (1990) suggests that these liquidity problems were only resolved by the timely intervention of the Federal Reserve Banks.

## 1.2 Previous studies of CCP margin coverage

Studies which have attempted to estimate CCP margin coverage have broadly adopted three types of approach, (a) statistical, (b) optimisation, and (c) option-based.

- Statistical approaches usually assume that initial margin should be set at a level that produces a pre-specified and acceptably small probability of exhaustion, over a time horizon that reflects the period of potential exposure for the CCP. The coverage level is set exogenously in these approaches.
- Optimisation models, by contrast, attempt to minimise the total sum of margin, settlement, default and other costs to which the clearing house and its members are potentially exposed. Unlike the statistical models which pre-specify acceptable coverage levels, an appropriate coverage level emerges endogenously in optimisation approaches (see for example, Fenn and Kupiec (1993)).
- Option pricing models take a strategic approach, assuming that, at least in theory, clearing members may choose to default if the reduction in value of their positions is greater than their posted margin. Since the CCP does not charge counterparties for this default option, initial margins should be set high enough to ensure that it is valueless (Kupiec (1997)).

The approach taken in the present paper is most closely related to the statistical models. See Knott and Mills (2002) for a more general review of other modelling approaches.

---

<sup>2</sup> The one or two-day time horizon that CCPs employ when setting initial margin levels implicitly assumes that the CCP can close out or transfer a defaulting member's positions within this time-frame. If not, the CCP becomes exposed to potential loss if the member's margin becomes exhausted. In the present study, we avoid making any assumptions about the appropriate time-horizon for setting margin. A time-horizon of one day is simply used as a convenient time scale over which to *measure* risk coverage provided by initial margin levels, and to compare coverage against what is typically claimed by CCPs. We make no assumptions about how rapidly a CCP could close-out or transfer positions in practice.

A variety of statistical studies have tried to quantify the market risks faced by CCPs. Figlewski (1984) estimated the degree of coverage provided by typical initial margins on stock and stock index futures, assuming that returns could be modelled as a geometric Brownian motion process. Gay, Hunter and Kolb (1986) constructed a similar model. These studies made the assumption that either returns or price changes were normally distributed; there is compelling evidence, however, that return distributions, particularly in futures markets, are leptokurtic, exhibiting fatter tails than the normal distribution, and indicating a greater probability of extreme price moves (Cornew, Town and Crowson (1984); Cotter and McKillop (2000); Venkateswaran, Brorsen and Hall (1993)). The potential inaccuracy of assuming a normal distribution, when evaluating margins, is illustrated by a study of Warshawsky and Dietrich (1989) which showed that for margin coverage levels of 99%, Normality assumptions consistently underestimated the required margin level.<sup>(3)</sup>

Recent studies of CCP margining have given greater attention to the fat-tailed characteristics of financial return distributions and have applied Extreme Value Theory (EVT) to estimating margin levels for futures contracts (e.g., Booth, Broussard, Martikainen and Puttonen (1997); Broussard (2001); Cotter (2001); Dewachter and Gielens (1999); and Longin, (1999)) concluding that EVT provides a useful tool for informing margining decisions. Generally, however these studies have concentrated on modelling the unconditional distribution of returns.<sup>(4)</sup>

A main objective of the present study is to show how initial margin levels can be modelled with the *conditional* distribution of returns, to provide an accurate assessment of margin coverage. We also illustrate how EVT or other distributional assumptions can be incorporated into the conditional analysis. The results suggest that an AR-GARCH (1,1) model with appropriate distributional assumptions can provide a simple and parsimonious way of assessing the coverage level provide by CCP initial margins on a day by day basis.

### *1.3 Conditional vs. unconditional estimates of margin coverage*

In calculating margins, two different return distributions can be considered - the unconditional or the distribution conditioned on current market information. Although the majority of research studies that have applied EVT to estimating margin levels have based their analysis on the

---

<sup>3</sup> In setting margins for individual futures positions, CCPs typically aim for levels which provide coverage against 95-99% of one-day price movements. Most make an estimate of the appropriate margin level by inspecting the historical distribution of price movements over recent months, and select the quantile that provides the desired level of coverage. But this method does not necessarily provide an adequate estimate of the margin required to protect against unusually large price movements, particularly if the sample window is small. CCPs will therefore adjust their empirical quantile estimates with subjective judgements on whether previously observed extreme moves (from outside the sample window) are likely to be repeated, and on whether anticipated future events (e.g. interest rate decisions) are likely to have an unusually large impact. Cotter (2001) notes the potential for inconsistency in this approach, advocating instead a fully statistical approach based on Extreme Value Theory (EVT).

<sup>4</sup> Fenn and Kupiec (1993) is a notable exception; it employs a conditional approach and derives implied non-coverage probabilities for historical margin levels.

unconditional distribution of returns, in practice, CCPs recognise the importance of giving additional weight to recent price history, especially current market volatility.<sup>(5)</sup>

McNeil and Frey (2000) note that conditional and unconditional analyses provide different but complementary information. In the unconditional analysis we ‘attempt to assign a magnitude to a specified rare adverse event, such as a five or ten-year loss’. Modelling the tail of the unconditional distribution can therefore be likened to stress-testing. By contrast, for the conditional analysis, we ask a different question, namely, what is the possible loss over the next day given current market conditions.<sup>(6)</sup>

One way to incorporate conditionality into margining is to apply a stochastic volatility model. As Danielsson and Morimoto (2000) and McNeil and Frey (2000) note, however, there are difficulties in applying standard stochastic volatility models to modelling the tails of the return distribution. Such models typically assume conditional normality of returns, which is not usually well-supported by the data. Whereas the central portion of the return distribution can be satisfactorily modelled assuming conditional normality, the tails generally cannot. Danielsson and Morimoto (2000) explain this by observing that a standard AR-GARCH process with normally distributed innovations gives the same weight to observations in the centre and in the tails of the distribution, so that the common events (which represent the middle 95% of the distribution) drive the estimation. To counter this problem, innovations can be modelled using a fat-tailed distribution such as a Student t, or one of the family of Extreme-Value Theory distributions.<sup>(7)</sup>

In the present study, four different alternatives for the distribution of innovations are considered: (1) historical, (2) normal, (3) Student t, and (4) extreme-value. To establish the best fit, a backtesting approach is applied, similar to that described in McNeil and Frey (2000).

---

<sup>5</sup> To take current market conditions into account, many CCPs focus their analysis on price data over the past two to three months. At London Clearing House, for example, for a naked futures position, an initial margin would usually be set to cover 99% (or approximately three standard deviations) of market price movements observed during the most recent three months. Although this falls short of a full conditional analysis, restricting the length of the sampling window to the most recent three months, demonstrates a desire to make margining decisions sensitive to current market conditions.

<sup>6</sup> Although margins set to ‘cover once every ten-year’ events would probably provide a comfortable level of coverage for CCPs, they would almost certainly be rejected by members who have to bear the opportunity cost of posting margin.

<sup>7</sup> A conditional approach to incorporating EVT also provides an additional benefit. The main asymptotic results of EVT require that sample observations be independent and identically distributed (iid). Typically, however, futures returns show volatility clustering. Whilst this is often ignored in EVT studies of margin, it can potentially lead to inaccurate tail estimates. In order to remove volatility clustering and produce an iid series suitable for EVT analysis, Kiesel, Perraudin and Taylor (2000) recommend filtering data using an AR-GARCH process and then fitting the residuals with the EVT distribution. Fitting a stochastic volatility model to the data prior to applying EVT therefore provides the additional benefit of filtering or pre-whitening the data.

#### 1.4 Conditional expected losses as a complementary measure of CCP exposure

In addition to setting margins at an appropriate level, CCPs also need to make an estimate of the resources they need to hold in reserve to meet a potential default. Typically this is done by stress-testing. Most stress-tests, however, suffer from a problem of subjectivity, as it is usually hard to associate a given scenario with a probability of occurrence. An alternative way to estimate the resources potentially required is to calculate the conditional expected losses on a contract – i.e. the expected loss to the CCP conditional on the initial margin level being exceeded – a risk-measure also known as the tail loss.

In the same way that conditional margin levels change on a day-by-day basis, conditional expected losses will also vary. Bates and Craine (1999) offered a graphic illustration of this by examining conditional expected losses on Chicago Mercantile Exchange (CME) contracts before and during the Crash of 1987. They found that, post-Crash increases in margin reduced the probability of a margin exceedance back to pre-Crash levels, but the expected losses faced if a margin exhausting price move were to have occurred were still an order of magnitude higher. This result highlights the importance of calculating other measures of clearing house exposure, and suggests that, when considered alone, margin coverage probabilities may in fact provide a misleadingly comforting picture of CCP exposure.

To summarise, estimates of conditional expected losses can provide an informative supplement to margin exceedance probabilities by providing a forward-looking measure of the economic impact of a margin-exhausting price move<sup>(8)</sup>. As such, they can be used to estimate the liquid capital resources that CCP members' might need to meet losses on their positions, or what the CCP itself might require as back-up resources to meet a potential default

The main goal of this study is to measure the risk coverage provided by margins set on futures contracts. To provide a way of assessing day-by-day changes in risk coverage, we employ, in Section 2, a GARCH model. To incorporate the well-documented fat-tailed characteristics of futures return distributions, conditional Student t and Extreme-Value Theory distributions are integrated into the modelling framework. Variants of the basic model are assessed which assume that residuals from a GARCH (1,1) model are distributed according to conditional normal, historical, Student t or Extreme Value Distributions, respectively, and backtesting is used to select the best-performing model, following a method described by McNeil and Frey (2000). After describing the data in Section 3, we apply our model to calculating conditional margins for the Brent and FTSE 100 futures that provide a pre-specified level of statistical coverage similar to that used in practice by CCPs, in Section 4.

The application of the model is illustrated by calculating the conditional exceedance probabilities for initial margins set historically by the London Clearing House, i.e., the implied probability that, on any given day, the margin level set by LCH would have been exceeded by market

---

<sup>8</sup> Compared to simple quantile measures, conditional expected losses (CEL) also show a number of desirable theoretical properties, in particular sub-additivity i.e.  $CEL(A+B) \leq CEL(A) + CEL(B)$  (see Artzner, Delbaen and Heath (1997)).

movements. As an additional measure of CCP exposure, market liquidity requirements are also estimated with the model, assuming a scenario where these historic margin levels were, in fact, exhausted. Section 5 concludes.

## 2 Method

This section summarises the econometric approach. A model which is designed to capture the volatility structure of futures returns is first described. Quantiles of the return distribution are estimated from this model, and hence margin levels that provide a pre-specified level of coverage. One of a number of contrasting distributional assumptions can be selected when formulating the model. The implications of these different distributional assumptions are compared, and the model is selected which provides the best fit to the empirical return distribution, using the backtesting method outlined by McNeil and Frey (2000). Finally, the model is applied to estimating conditional expected losses for individual futures contracts, and hence to calculating the aggregate market liquidity required to meet a margin exhausting price move.

### 2.1 Modelling the returns process

When modelling futures price returns, two main empirical regularities need to be accounted for: the presence of fat-tails, and volatility clustering. To incorporate these features, we adopt the following model in which futures returns can be described by a strictly stationary process  $\{X_t\}_{t \geq 0}$  with the following dynamic structure:

$$\begin{aligned} X_t &= a + bX_{t-1} + \sigma_t Z_t, \\ \sigma_t^2 &= \alpha + \beta\sigma_{t-1}^2 + \gamma\varepsilon_{t-1}^2, \\ \varepsilon_t &= \sigma_t Z_t, \quad Z_t \sim \text{IID}(0,1), \quad P(Z \leq z) = F_Z(z). \end{aligned} \tag{1}$$

The innovations  $\{Z_t\}_{t \geq 0}$  are assumed to have zero mean, unit variance and to be independently and identically distributed (i.e., IID(0,1)) according to a distribution function  $F_Z(z)$ .

With this structure, returns have a simple AR(1) recursive dependency, and volatility follows a GARCH (1,1) process such that the volatility at time  $t$  is dependent on shocks and volatility at time  $t-1$ . This model provides a simple and parsimonious way to take into account volatility clustering.<sup>(9)</sup>

From (1), the time  $t$  quantile,  $x_q$ , can be derived for a given probability level  $q \in (0,1)$  and conditional on the information set available at time  $t-1$ . First, the quantile,  $x_q$ , is defined such that

---

<sup>9</sup> A more complicated volatility structure could be assumed, but it is not clear this would significantly improve the model's forecasting ability. Hansen and Lunde (2001), for example, estimate more than 300 volatility models concluding that there is no model that provides a clearly superior forecast to the GARCH(1,1) framework. Recent studies by Barone-Adesi, Giannopoulos and Vosper (2000), Kiesel, Perraudin and Taylor (2000) and McNeil and Frey (2000) adopt a similar model.

$$q = P_{t-1}(X_t \leq x_q).$$

Using **(1)** this can be rewritten as,

$$\begin{aligned} q &= P_{t-1}(a + bX_{t-1} + \sigma_t Z_t \leq x_q) \\ &= P_{t-1}\left(Z_t \leq \frac{x_q - a - bX_{t-1}}{\sigma_t}\right) = F_Z\left(\frac{x_q - a - bX_{t-1}}{\sigma_t}\right). \end{aligned} \quad (2)$$

If the quantile associated with the distribution function  $F_Z(z)$  is denoted as  $z_q$ , the quantile  $x_q$  can be rewritten as,

$$x_q = a + bX_{t-1} + \sigma_t z_q. \quad (3)$$

Equation **(3)** says that given the information set available at time  $t-1$ , the quantile  $x_q$  for the returns is a linear function of the innovation quantile and is proportional to the volatility level, implying that the quantile  $x_q$  and the margin estimates derived from it, will inherit fluctuations in the estimate of  $\sigma_t$ .<sup>(10)</sup>

This model can be used to derive the next day conditional distribution and the non-coverage (or tail) probability implied by a pre-specified margin level. The probability that  $X_t$  will be less than a pre-specified level  $x^{mar}$  is equal to

$$\begin{aligned} P_{t-1}(X_t \leq x^{mar}) &= P_{t-1}(a + bX_{t-1} + \sigma_t Z_t \leq x^{mar}) \\ &= P\left(Z_t \leq \frac{x^{mar} - a - bX_{t-1}}{\sigma_t}\right) = F_Z\left(\frac{x^{mar} - a - bX_{t-1}}{\sigma_t}\right). \end{aligned} \quad (4)$$

$x^{mar}$  can be taken as the maximum return permitted by initial margin. With this interpretation, **(4)** gives the coverage level, on a specific date, provided by margin,  $x^{mar}$ .<sup>(11)</sup>

To compute actual exceedence probabilities, a functional form must be provided for  $F_Z(z)$ . GARCH models typically assume the distribution of innovations to be normal. McNeil and Frey (2000), however, note that when modelling a fat-tailed return distribution it may be appropriate to choose a different innovation distribution. Here, four different alternatives are considered: (i) historical, (ii) normal, (iii) Student t and (iv) extreme value distributions. The historical distribution is derived from the standardised residuals of an AR(1)-GARCH(1,1) process with normal distribution.

Where innovations are assumed to be distributed according to the normal or historical distributions, the derivation of quantiles is straightforward. Suppose  $T$  is a Student-t distributed random variable with  $v$  degrees of freedom, then  $E[T]=0$  and  $V[T] = v/(v-2)$ ,  $v>2$ . Standard

---

<sup>10</sup> See Danielsson and Morimoto (2000).

<sup>11</sup> Note that these expressions are valid only one-step ahead.

assumptions require that innovations have unit variance, so the process for  $Z$  in **(1)** is scaled as  $Z = \sqrt{(v-2)/v} * T$ . The quantile,  $z_q$ , can then be expressed in terms of the quantile of the Student  $t$  random variable as  $z_q = \sqrt{(v-2)/v} * F_T^{-1}(q)$  where  $F_T^{-1}(q)$  is the  $q$ -quantile of  $T$ .

Finally, innovations can be assumed to follow an extreme value theory (EVT) distribution. The reader is referred to Appendix A for technical details on EVT. Here we simply note that if the distribution function can be assumed to be in the domain of attraction of a Frechet distribution,  $F \in \text{MDA}(\Phi_\alpha)$ , then the desired quantile can be estimated by  $x_q = (n/k(1-q))^{-1/\alpha} X_{n-k:n}$ , where  $\alpha$  is the tail parameter and  $X_{n-k:n}$  is the  $k^{\text{th}}$  order statistic.

## 2.2 Computing conditional expected losses

In addition to non-coverage probabilities, measures of conditional expected loss, i.e., the expected losses conditional on margin being exceeded, can provide important information on CCP exposure. Using the model described in the previous section, it is possible to derive the conditional expected return,  $\text{CER}_t$ , defined as the expected value of the return,  $X_t$ , given that it has exceeded a certain level  $u$ .

$$\text{CER}_t \equiv E_{t-1}[X_t | X_t > u]$$

Substituting for  $X_t$  from (1) gives,

$$\text{CER}_t \equiv a + bX_{t-1} + \sigma_t E[Z | Z > z] \quad (5)$$

where  $z = (u - a - bX_{t-1})/\sigma_t$  and  $E[Z | Z > z]$  depend on the chosen distribution for the innovations and can be summarised in the following way,

$$E[Z | Z > z] = \begin{cases} \frac{\phi(z_q)}{1 - \Phi(z_q)} \approx z_q (1 + z_q^{-2}), & Z \sim N(0,1), \\ z_q \left[ \frac{v}{v-1} (1 + o(v^{-1})) \right], & z_q = \sqrt{\frac{v-2}{v}} t_q, \quad t_q = F_T^{-1}(q), \\ \frac{1}{1 - \xi} z_q, & \text{EVT} \end{cases}$$

$\phi$  and  $\Phi$  are, respectively, the density and the distribution for a standardised normal random variable.

Equation **(5)** gives the *expected return* given that a certain threshold level return has been exceeded. A more conventional measure of conditional *expected loss* can be derived from this. Suppose a conditional distribution is obtained for date  $t+1$ . Knowing the maximum return permitted by initial margin at date  $t+1$ ,  $r_{t+1}^{\text{mar}}$ , the expected price variation, given that this margin level has been exceeded, can be computed. Given the futures closing price  $P_t$  at time  $t$ , the additional expected financial resources required per unit, conditional on margin being exceeded is  $|\exp(\text{CER}_{t+1}) - \exp(r_{t+1}^{\text{mar}})|P_t$ . Multiplying this expression by the contract size and the open



interest figure then provides a rough estimate of the total liquidity that would be required in the market, in the event that initial margin is exceeded.

### 2.3 Estimation of the tails under EVT

Estimation of the AR(1)-GARCH(1,1) model shown in (1) is relatively straightforward for both normal, Student t and historical innovations. The estimation for EVT distributed innovations is more complex, but essentially involves two steps: (i) estimate the model with normal innovations, and (ii) apply EVT to the standardised residuals. It can be shown (see Gouriéroux (1997)) that step (i) yields asymptotically consistent estimates of the model parameters. The intuition is that the AR and the GARCH parts of the model fit the central part of the distribution (i.e., the conditional mean and variance) leaving the tails unaffected. The tails can then be estimated in step (ii) using an appropriate estimation technique. In the present paper, a version of the Hill estimator suitable for small samples is applied. We describe this estimation in what follows.

If the innovation distribution can be assumed to be fat-tailed, EVT tells us that the distribution of the sample maxima can be modelled as a Fréchet distribution. More correctly, the (unknown) distribution function,  $F$ , belongs to the domain of attraction of the Fréchet distribution,  $F \in \text{MDA}(\Phi_\alpha)$ . If this is the case, the complementary distribution function  $\bar{F}(x) \equiv 1 - F(x)$  can be approximated as  $\bar{F}(x) = Cx^{-\alpha}$ , where  $C = u^\alpha$ ,  $x \geq u > 0$ . The maximum likelihood estimator for  $\alpha$  can be shown to be,

$$\hat{\alpha} = \left( \frac{1}{n} \sum_{j=1}^n \log X_{n-j+1:n} - \log u \right)^{-1},$$

where  $X_{n-j+1:n}$  is the  $(j+1)^{\text{th}}$  upper order statistic<sup>(12)</sup>. In the general case, the threshold  $u$  is unknown, but can be approximated by the  $k^{\text{th}}$  order statistics. Then, conditional on the number  $k$  of exceedances of the threshold, the maximum likelihood estimator (MLE) becomes

$$\hat{\alpha}^{(H)} = \hat{\alpha}_{k,n}^{(H)} = \left( \frac{1}{n} \sum_{j=1}^n \log X_{n-j+1:n} - \log X_{n-k:n} \right)^{-1}, \quad (6)$$

The MLE for the constant,  $C$ , is  $\hat{C}_{k,n} = X_{n-k:n}^{\hat{\alpha}_{k,n}} \frac{k}{n}$ .

Equation (6) is called the Hill Estimator. The quality of the estimate depends upon the choice of  $k^{\text{th}}$  order statistics.

Several recent papers have discussed how to select an optimal value for  $k$ . Drees and Kaufmann (1998) find  $k(n)$  using an iterative procedure. Beirlant, Vynckier and Teugels (1996) use a regression approach based on a Pareto quantile plot. Danielsson, de Vries, de Haan and Peng

---

<sup>12</sup> See Appendix B for a definition of upper order statistics.

(1999) use a two step subsample bootstrap method. A drawback of these methods is the very large sample size required which makes them unsuitable for many applications.

A somewhat simpler alternative, appropriate for small samples, has been recently suggested by Huisman, Koedijk, Kool and Palm (1997). A weighted average of a set of Hill estimators is calculated, with weights obtained by using simple least square methods. Huisman *et al.* (1997) show that this approach produces unbiased estimates for the tail-index in small samples<sup>(13)</sup>. We adopt this approach for three reasons, (i) it removes the need to choose the optimal threshold, (ii) it yields acceptable estimates even for small samples (ca. 250), and (iii) it provides a clear criterion for selecting the best Hill estimate<sup>(14)</sup>.

#### 2.4 Backtesting to select the best-performing innovation distribution

The previous section, described how to estimate the AR(1)-GARCH(1,1) model under different distributional assumptions, but left unresolved the question of which distribution should be selected as the best model. To choose amongst the four possible alternatives, a backtesting procedure was applied, in which the model was estimated for a large number of subsamples, so the predictive ability of each alternative distribution could be assessed, and the best performing distribution selected. This procedure is described in more detail below.

Suppose the sample size is equal to  $N$  and a subsample size is fixed equal to  $n < N$ . The model is estimated using the first  $n$  observations. Parameter estimates are then used to compute the next day quantile  $x_{n+1|n}^q$  at a confidence level  $1 - q$  which can be compared with the actual observed return for day  $n+1$ ,  $x_{n+1}$ . If  $|x_{n+1}| > |x_{n+1|n}^q|$ , an exceedance is said to have occurred. The sample window is then shifted forward by one observation, and the model is re-estimated using the subsample from  $t=2$  to  $t=n+1$ . In parallel fashion, the predicted quantile  $x_{n+2|n+1}^q$  for date  $n+2$  is compared with the actual realisation  $x_{n+2}$  and it is noted whether or not an exceedance has occurred. The process continues in this way until all the subsamples are exhausted.

If the non-coverage probability has been fixed at  $q$ ,  $qN$  exceedances would be expected on average. For a specific distribution, if there are fewer exceedances than the expected number, the predicted quantiles  $x_{n+1|n}^q$  of the distribution may be too large. If there are more exceedances than expected, the quantiles may be too small. The distribution is selected that gives a number of exceedances closest to the expected number.

To implement the backtesting procedure, a subsample size of 500 observations was employed (i.e.  $n = 500$ ). This is sufficiently small to be comparable to the data windows often used by

---

<sup>13</sup> See Appendix B for details.

<sup>14</sup> An alternative approach for estimating the tails of an EVT distribution is the so-called peaks over threshold method (POT), which assumes that the tail of the distribution can be approximated by a Generalized Pareto distribution (GPD). To apply the model it is necessary to fix a threshold above which it is assumed that the GPD provides a good approximation to the data generation process. McNeil and Frey (2000) and Kiesel *et al.* (2000) both regard the GPD approach as a superior method for estimating tails. In the backtesting exercise, however, this would require us to fix the number of order statistics to be the same in each subsample. This would not always produce an optimal fit.

CCPs for margining, yet sufficiently large to permit empirical quantiles to be derived for very small probabilities (up to 0.2%) and the small-sample EVT estimation techniques (described in section 5.3) to be applied.

The model can also be applied to estimating the coverage provided by historical initial margin levels. If the initial margin for date  $n+1$  is known, equation (4) can be used to estimate coverage probabilities. Implied coverage probabilities for initial margins set historically by LCH on the Brent and FTSE contracts are reported in section 5.

### **3 Data and preliminary analysis**

The data set consists of daily closing prices for the Brent and FTSE 100 futures contracts. Both are high volume contracts of considerable importance to the exchanges that trade them (the London International Financial Futures and Options Exchange, and the International Petroleum Exchange, respectively). Closing prices were used because the obligations of the CCP are based on close of day positions.

#### *(i) Time series of futures prices*

To create a time-series of futures prices sufficiently long for analysis, individual contracts must be rolled-over as they approach maturity. In order to link individual contract prices, two choices must be made: (1) *when* to switch from one contract to the next (i.e. how to choose the *roll-over date*) and (2) *how* to correct for any discrete change in price (i.e. any *price adjustment* on the roll-over date). A variety of approaches have been taken in the research literature. Dewachter and Gielens (1999) use the spot price as a proxy for the futures price, sidestepping the problem of how to link contract prices. Longin (1999) uses futures prices but does not amend the time series on the roll-over dates creating the possibility of artificial jumps on the roll-over dates. Cornew, Town and Crowson (1984) roll over futures contracts by selecting a date two weeks before the first notice day. Cotter and McKillop (2000) change contracts on the last trading day before the delivery month making no correction, but recognising the existence of artificial jumps.

In order to prevent artificial jumps from arising, the following procedure was adopted here. The front contract was rolled to the next on the first day of the delivery month. Both contracts are liquid at this point and by switching on the first day, a spot-month charge levied by the CCP at the start of the delivery month can be ignored. Nevertheless, switching contracts on the first day of the delivery month can introduce an artificial jump into the time series of futures prices because of differences in the individual contract prices. Since these jumps are artefacts of the method used to construct the time series, which are not economically relevant, the rollover dates, and so any jumps associated with them, have been omitted. Table A presents descriptive statistics for returns on the two contracts, which are discussed further in (iii).

#### *(ii) Times series of initial margins*

Two time series of initial margins for the Brent and FTSE 100 futures contracts were obtained from London Clearing House covering the periods 4/1/99 to 11/2/02 and 6/1/98 to 11/2/02,

respectively.<sup>15</sup> Initial margins for long and short positions are set equally. Charts 1 and 2 show that initial margins remained constant for significant portions of each period. The temporary increases visible for the Brent contract usually coincide with UK public holidays where protection would be required over a longer time-horizon than normal. For example, a temporary increase in initial margin occurs at the end of 1999, to provide additional protection during the Millennium period.

The charts also show close-to-close prices superimposed on initial margin levels. A visual comparison suggests that during the sample period initial margins on the Brent contract were challenged or exceeded much more frequently than the FTSE 100 contract. In fact, there are no initial margin exceedances for the FTSE 100 contract, compared to nine for the Brent contract (see Table B).

### *(iii) Descriptive statistics*

From the time series of closing prices, log-returns for each contract were computed. Descriptive statistics (Table A) show that for both contracts the return distributions were approximately symmetric. Jarque-Bera and the Kolmogorov-Smirnov tests, however, indicate deviations from normality. QQ plots, illustrated in Chart 3B and Chart 4B, confirm fat-tailed behaviour. Return plots (Chart 3A and Chart 4A) suggest the presence of heteroscedasticity (i.e. time varying volatility), which is confirmed by the autocorrelation function for the squared returns (Chart 3D and 4D). Finally, there is little evidence of correlation in the return levels (Charts 3C and 4C), and the standard augmented Dickey-Fuller and the Phillips-Perron tests indicate that both series are stationary.

Overall, the descriptive analysis indicates heteroscedasticity and leptokurtosis (i.e. fat-tails) to be the two main empirical regularities which need to be taken into account.

## **4 Results**

### *4.1 Unconditional analysis*

Charts 5A-D show unconditional margins at a succession of coverage levels, for each of the four distributions. Inspection shows that there is generally a fairly consistent rank ordering of margin levels with  $EVT > \text{Historical} > \text{Student } t > \text{Normal}$ . EVT margins are clearly greater than margins calculated under other distributional assumptions at all coverage levels, with this pattern being most marked for the Brent contract, and particularly for its left tail which represents large price falls. The large EVT margins can be explained by the fact that, unlike the other distributional

---

<sup>15</sup> Like many other CCPs, London Clearing House uses the SPAN system (Systematic Portfolio Analysis of Risk) to calculate initial margins for portfolios of futures and other exchange-traded derivatives contracts. Initial margin for a single unhedged futures contract is represented by a SPAN parameter known as the *scanning range* which is periodically updated by CCPs as market conditions change. Although a single contract usually consists of a fixed number of units of the underlying commodity, for simplicity, the analysis here considers the margin required for each individual unit. For example, each Brent futures contract is for 1,000 barrels. The margin per lot set on 30/9/02 was \$2,000, equivalent to \$2 per barrel.

approaches, the EVT analysis only considers extremes (i.e. observations actually in the tail). These margins therefore tend to be more sensitive to large tail observations. Even at the lowest coverage level of 95%, the EVT analysis still leads to the highest margins, a finding which differs from some previous studies, which have only observed reliable differences between EVT and the historical distribution at coverage levels greater than 99%. In the present dataset, it is differences between margins based on the normal, historical and t distributions, which only begin to emerge at the 99% level.

The most significant feature of Charts 5A-D is the distribution by coverage level interaction. As higher coverage levels are approached, there is a proportionate increase in the EVT margins compared with other distributions. Compared with the normal distribution margins, the historical distribution and Student t margins show a similar but less marked interaction.

To assess the coverage that would have been provided by these theoretically derived margins they can be compared against historical daily price movements. Table C reports the number of margin exceedances observed. For the Brent contract, there are a total of 3034 observations. At a coverage level of 99%, 30 violations would therefore be expected. The historical margins are set in order to produce this expected number. Results from the three other distributions are compared in Table C. The Student-t distribution produces a number of exceedances closest to the expected number. The third and fourth columns of the table show that the dollar margin derived from the Student-t distribution is almost identical to the historical distribution.

#### 4.2 *Conditional analysis*

Charts 6A and 6B depict conditional margins derived from the four distributions as a function of time. Margins are calculated for the right tail of the returns distribution assuming a required coverage level of 99.5% (99%, two-tailed). The main features to note are that both the Brent and the FTSE 100 contracts, the conditional margins show considerable variability over time, largely because the model is driven by market volatility. For the Brent contract, there is nevertheless a fairly consistent rank ordering of the margin levels which mirrors the results of the unconditional analysis i.e.  $EVT > \text{historical} > \text{Student } t > \text{normal}$ . This ordering breaks down, however, after October 2000.

For the FTSE 100 contract, there is no consistent rank ordering of conditional margins, and in general the margins calculated on the basis of the four distributions are similar. The absence of a rank ordering seems most likely attributable to the fact that, compared with the Brent contract, the FTSE 100 contract is less fat-tailed, as indicated by the QQ-plot in Chart 4B. The FTSE 100 distribution shows fewer extreme values, and as a result, the impact of the choice of innovation distribution is less marked.

Chart 7A and B show the Student-t and normal distribution models applied to calculating non-coverage probabilities for the initial margin levels set historically by LCH. The model suggests that the initial margins on the FTSE contract provided an average coverage level of 99.86% during the sample period. The average coverage level for the Brent contract is similar at 98.96%. These levels are consistent with overall number of exceedances of initial margin for each contract

reported in Table B. Despite the high average levels, both charts show that coverage can fall rapidly when more volatile market conditions are encountered.

#### 4.2.1 Backtesting results

Tables D and E display the number of margin exceedances observed in backtesting for the four conditional distributions at four different non-coverage levels. The expected number of exceedances at each non-coverage level is shown in the bottom line of each table.

Consider the Brent contract first (Table D). In seven out of eight cases, margins based on the thin-tailed normal distribution lead to too many margin exceedances. The one exception occurs, as might be expected, at the highest non-coverage level of 5%, where the error introduced by assuming a normal distribution is least pronounced. At the other extreme, EVT margins are observed to be much too conservative at the 5% level, leading to far too few exceptions. At the higher coverage levels, however, performance of the EVT distribution improves significantly. Overall, and taking into account performance at all coverage levels, the historical and t distributions provide the best fit.

FTSE 100 back-test results offer a broadly similar picture. At the 5% non-coverage level, the EVT margins are too stringent, though performance improves substantially at higher coverage levels. For this contract, the normal distribution leads to margins that are too low in five out of eight cases, but the margin levels are generally closer to those predicted by other distributions. This result might be expected given that the FTSE 100 distribution is less fat-tailed than the Brent distribution (compare Charts 3B and 4B).

#### 4.3 Conditional expected losses

This section reports conditional expected losses for the Brent and FTSE 100 contracts, in order to assess the potential liquidity demands on the market of margins on either of the two contracts being exhausted. Section 2.2 introduced the notion of conditional expected return, i.e., the expected return given that the maximum (or minimum) return permitted by margin has been exceeded. A measure of the extra liquidity required to meet outstanding obligations can be derived from this. The procedure is as follows. Using the returns up to date  $t$ , the model parameters are estimated and the conditional distribution of returns for date  $t+1$  is obtained. Given the maximum return permitted by margin for date  $t+1$ ,  $r_{t+1}^{\text{mar}}$ , the conditional expected return,  $\text{CER}_{t+1}$ , is calculated. If  $P_t$  is the closing price at date  $t$ , then  $|\exp(\text{CER}_{t+1}) - \exp(r_{t+1}^{\text{mar}})|P_t$  represents the conditional expected loss, or equivalently, the extra liquidity needed to cover the margin violation.

Tables F and G summarise the expected extra liquidity required to meet a margin exceedance. The Student-t model was used here, and prices were averaged across the whole sample. To calculate a figure for required *market* liquidity (Table F) the contract size was multiplied by the average open interest. For the Brent contract, an extra \$50 million would be required on average, and the maximum expected loss for the market, based on the sample period, was \$90 million. For

the FTSE 100 contract, market liquidity requirements were significantly larger, partly because the contract value was bigger but also because the open interest was double that of the Brent contract. The required market liquidity varied substantially between a minimum of £14 million and a maximum of £250 million.

To put these figures into perspective, they can be converted into a measure of the extra liquidity required as a percentage of the margin already posted (Table G). From Table G, it is apparent that coping with a margin violation on the Brent contract would have placed higher demands on the market than a margin exceedance on the FTSE 100 contract. A 7% minimum additional margin would have been required for the Brent contract, compared with only 1% for the FTSE 100 contract. Looking at the averages, the Brent contract would have required an 8-9% additional margin, double that required for the FTSE 100. Thus, although the FTSE 100 contract would have required more initial margin than the Brent contract, a margin exceedance on the Brent contract would have had a more serious affect on the liquidity position of market participants. This can be largely attributed to the fatter-tails of the Brent return distribution.

## **5 Discussion and conclusions**

This study has described a method for assessing the conditional coverage provided by CCP initial margins, illustrating its application to two heavily traded futures, the Brent and FTSE 100 contracts. Previous work on margining has largely focussed on assessing margin coverage unconditionally. This study suggests that estimating the conditional distribution of returns using a GARCH model can potentially provide more timely, price-sensitive information on the risk coverage offered by CCP margins.

Previous studies have advocated the use of EVT to compute initial margins in order to take into account the well-documented fat-tailed characteristics of futures return distributions. Using the approach developed by McNeil and Frey (2000), we have extended previous work on margining by constructing a conditional model of margin coverage that integrates fat-tailed characteristics and stochastic volatility within the same framework. By exploiting the additional information present in the stochastic volatility structure of returns, this model permits the calculation of estimates of the probability of margin exceedance, on a day-by-day basis.

The model was applied estimating the coverage provided by initial margins set by LCH on the Brent and FTSE 100 contracts. During the sample period, margins were found to provide a coverage level generally in excess of 99%, over a 1-day time horizon. Long-run coverage levels were also independently assessed using the observed frequency of initial margin exceedances, with the two analyses producing similar results. It was noted, however, that the coverage probability implied by the model could fall rapidly in more volatile markets, underlining the need for CCPs to carefully monitor margin levels and adjust them as necessary to take account of changing market conditions.

Backtesting indicated that an AR(1)-GARCH(1,1) process provided a simple but effective model for the conditional mean and volatility of futures returns. Danielsson and Morimoto (2000) have

previously argued that pre-whitening the data using this approach removes only volatility clustering and not extreme dependence, i.e. there is no guarantee that GARCH filtering will remove all the structure in the data. Our backtesting results, however, suggest that, at least in this instance, an AR(1)-GARCH(1,1) model with innovations based on the historical or t distributions produced a reasonably good fit to the tails of the return distribution. This is a potential avenue for future research.

The choice of innovation distribution did make a marked difference to the quality of the fit. Normally distributed innovations generally led to margin levels that were too low to provide the required level of coverage. At high confidence levels (greater than 99%), it might have been expected that the extreme value distribution, which explicitly models only the extremes, would provide the best fit. Even for the fat-tailed Brent return distribution, however, the EVT fit was not significantly better than one based on t distributions, and across all coverage levels (including 95%), the historical and t distributions produced superior performance.<sup>(16)</sup>

One reason why the EVT approach may have failed to outperform other models is that it implicitly assumes the innovation distribution has fat-tails. When estimating the AR(1)-GARCH(1,1) model with t distributed innovations, however, it was observed that the degrees of freedom changed frequently, sometimes reaching values greater than 20, indicating that the tails of the sample distribution were not always fat-tailed. In contrast to the EVT approach, the historical and Student t models have the flexibility to accommodate these variations in tail-thickness.

Overall, it can be concluded that: (1) a thin-tailed distribution, such as the normal, does not provide an adequate fit to the innovation distribution. (2) In a conditional context, a fat-tailed EVT distribution performs better than the normal distribution but does not necessarily outperform other simpler models, (3) for the coverage levels considered here, the historical and Student t distributions generally provide satisfactory estimates of the appropriate margin. If very high coverage levels are sought, however, or estimates of conditional expected losses, it is probably preferable to use an EVT or a Student t distribution. (4) When all non-coverage levels are considered, the Student t distributions probably provide the most flexible all-round model.

In the final part of the paper, a conditional expected loss measure was calculated to provide a further indicator of potential CCP exposure, i.e. an estimate of the additional liquidity required if the initial margin on a member's positions was to be exceeded. As might be expected, the tail-shape of the return distribution was found to be an important determinant of the amount of liquidity required. In terms of the percentage of liquidity already posted, the results indicate that an initial margin exceedance for the Brent contract would put a considerably greater strain on the market than a similar exceedance for the FTSE 100 contract. This can be attributed, in part, to the

---

<sup>16</sup> In some circumstances, the relatively small sample size employed (500 observations) might be considered a problem in applying EVT estimation techniques. In the present study, however, we have addressed this issue by using a variant of the Hill estimator which produces unbiased estimates for the tail index, without loss of efficiency, for samples as small as 250 observations (see Huisman *et al.* (1997) for further details).



lower coverage provided by the Brent margins during the sample period, but more importantly to the marked leptokurtic characteristics of the Brent return distribution.

## Tables

**Table A:** Descriptive statistics for the Brent and the FTSE 100

	FTSE100	Brent
Mean	1.51	6.55
St.Dev	113.06	215.44
Min	-608.23	-1451.62
Max	637.27	1315.06
Skewness	-0.03	-0.08
Excess Kurtosis	2.20	5.52
Kolmogorov-Smirnov	0.04	0.07
Jarque-Bera	632.8	3860.6
Ljung-Box	53.85	62.17
#Observations	3147	3034
Dates	2/1/90-13/08/02	2/1/90-13/08/02

Descriptive statistics for log returns are expressed in basis points. The 5% significance values for the Kolmogorov-Smirnov statistic, the Ljung-Box with 30 lags and the Jarque-Bera test are 0.02, 43.77 and 5.99, respectively.

**Table B:** Initial margin changes and exceedances for the Brent and the FTSE 100 contracts

	Start	End	# obs.	# changes	# exceed.
Brent	4-01-99	11-02-02	749	24	9
FTSE	6-01-98	11-02-02	1026	5	0

An initial margin exceedance is said to occur if the absolute value of the price change between two consecutive days is greater than the initial margin, i.e.,  $|\Delta P_t| > M_t$ , where  $M_t$  is the margin level for day  $t$  and  $\Delta P_t = P_t - P_{t-1}$ .

**Table C: Unconditional margins and exceedances for the Brent contract**

	Right(%)	Left(%)	Margin(right) (USD)	Margin(left) (USD)	# exceed. Right	# exceed. Left
Hist.	6.06	-5.51	1.25	1.07	30	29
Normal	5.08	-4.95	1.04	0.96	50	48
Stud-t	5.56	-5.43	1.14	1.06	39	32
EVT	8.71	-11.49	1.82	2.43	7	4

Theoretical initial margins are reported for both long and short positions. For the EVT case, left and right tails are estimated separately. The tail-shape parameter,  $\xi$ , is equal to 0.29 for the left tail and 0.16 for the right tail. Initial margins are estimated to provide a 2% non-coverage probability (1% probability in each tail). At this coverage level, the expected number of exceedances is 30.

**Table D: Backtesting performance of the four distributions for the Brent contract**

	Left tail				Right tail			
	5%	1%	0.5%	0.135%	5%	1%	0.5%	0.135%
Hist.	137	35	18	9	129	29	15	8
Normal	116	43	31	21	140	50	30	10
Stud-t	125	27	20	2	155	27	10	5
EVT	79	26	16	3	83	30	17	9
<b>Expected</b>	127	25	13	3	127	25	13	3

Margin exceedances are shown for the Brent contract assuming theoretical margins based on each of the four distributions. Four non-coverage probabilities are assumed: 5%, 1%, 0.5% and 0.135%. The sample size is  $n = 2534$ .

**Table E: Backtesting performance of the four distributions for the FTSE 100 contract**

	Left Tail				Right Tail			
	5%	1%	0.5%	0.135%	5%	1%	0.5%	0.135%
Hist.	132	28	15	7	141	26	14	7
Normal	109	24	15	7	141	34	25	13
Stud-t	118	18	10	3	142	29	18	8
EVT	54	21	12	7	81	24	15	8
<b>Expected</b>	132	26	13	4	132	26	13	4

Margin exceedances are shown for the FTSE 100 contract assuming theoretical margins based on each of the four distributions. Four non-coverage probabilities are assumed: 5%, 1%, 0.5% and 0.135%. The sample size is  $n = 2646$ .

**Table F: Additional liquidity required by counterparty conditional on margin being exceeded**


	FTSE100 (£mn)		Brent (\$mn)	
	Left	Right	Left	Right
Min	13.9	13.2	33.5	37.4
Max	221.6	250.2	73.7	90.0
Average	53.9	59.7	41.7	46.4
Open Interest at 02/02	464,600		254,400	
Margin rate	£350		US\$2	
Average price	£41,290		US\$12,01	

**Table G: Additional liquidity required as a proportion of margin already posted**


	FTSE100		Brent	
	Left	Right	Left	Right
Min (%)	1	1	7	7
Max (%)	14	15	14	18
Average (%)	3	4	8	9
Total margin	£1,626mn		US\$509mn	

## Charts

**Chart 1: Brent margin and price variation (4-1-99 to 11-2-02)**




**Chart 2: FTSE 100 margins and price variation (6-1-98 to 11-2-02)**

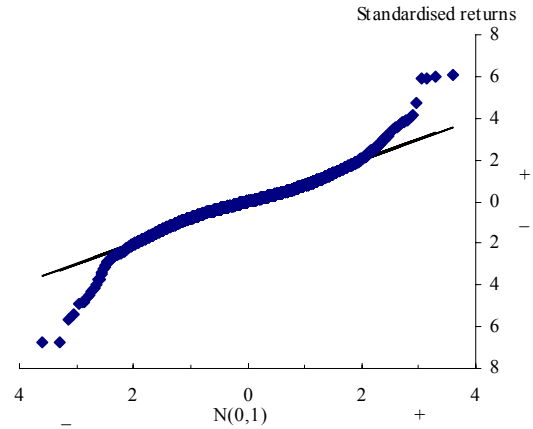


### Chart 3: The Brent contract

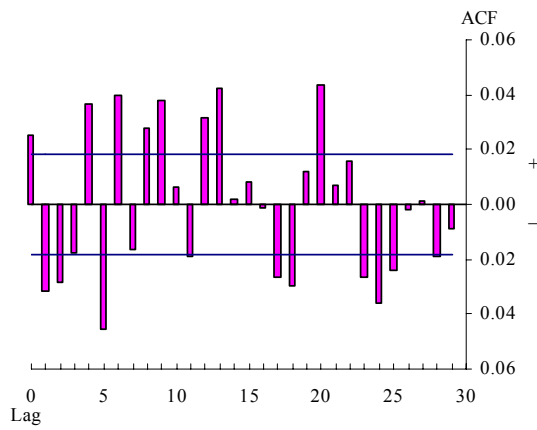
A: Daily ln returns (in bp)




B: QQ plot



C: Autocorrelation function for return levels

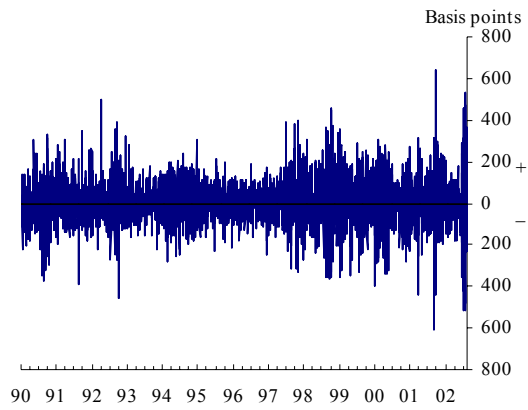


D: Autocorrelation function for squared returns




### Chart 4: The FTSE 100 contract

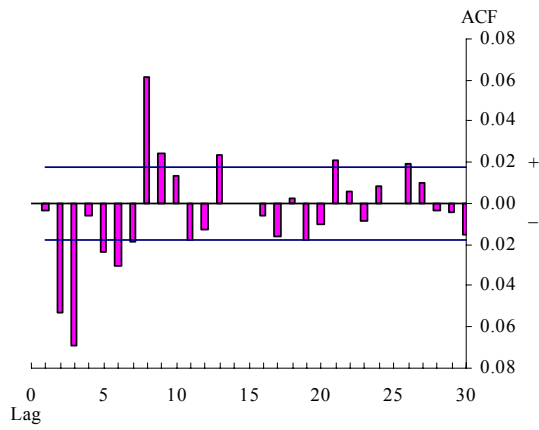
A: Daily ln returns (in bp)




B: QQ plot



C: Autocorrelation function for return levels

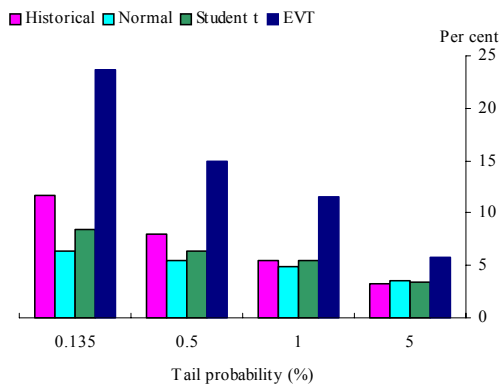


D: Autocorrelation function for squared returns

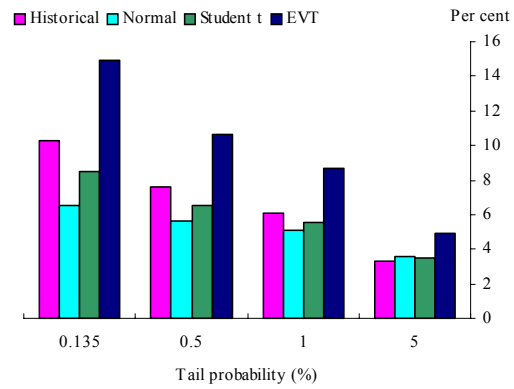


### Chart 5: Unconditional margins

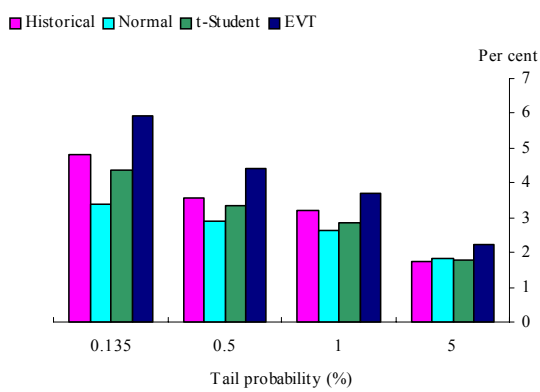
A: Brent - left tail



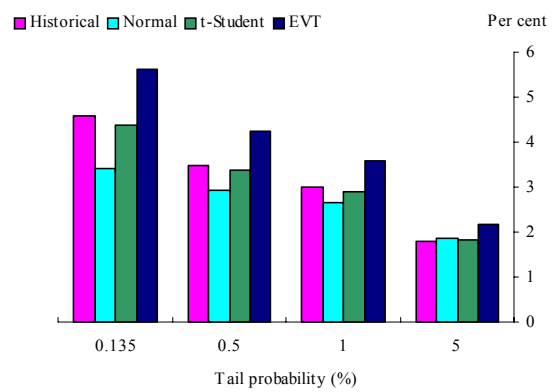
B: Brent - right tail



C: FTSE 100 - left tail




D: FTSE 100 - right tail



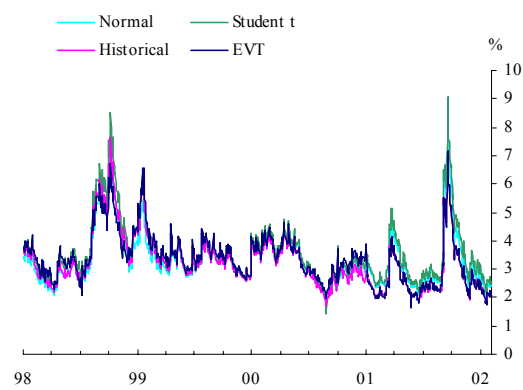


**Chart 6: Theoretical conditional margin at 99.5% for positive changes in return**

Panel A: Brent contract




Panel B: FTSE 100 contract




**Chart 7: Non-coverage probabilities**

Panel A: Brent contract



Panel B: FTSE 100 contract



## Appendix A: Relevant results from EVT

Suppose<sup>(17)</sup> we have a sequence of iid random variables  $\{X_1, X_2, X_3, \dots\}$  with common distribution function  $F$ . Define the *sample maxima* as  $M_1 \equiv X_1$ ,  $M_n \equiv \max\{X_1, X_2, X_3, \dots, X_n\}$  for  $n \geq 2$ .

The distribution function of  $M_n$  is given by,

$$P(M_n \leq x) = P(X_1 \leq x, \dots, X_n \leq x) = F^n(x), \quad x \in \mathfrak{R}, n \in \mathbb{N}.$$

It is usually the case that the distribution function  $F$  is unknown. The Fisher-Tippet theorem, however, shows that approximate families for  $F^n(x)$  can be derived by examining extremes.

### Theorem (Fisher-Tippet):

Let  $\{X_n\}$  be a sequence of iid random variables. If there exist norming constants  $c_n > 0$  and  $d_n \in \mathfrak{R}$  and some non-degenerate distribution function  $H$  such that

$$\frac{M_n - d_n}{c_n} \xrightarrow{d} H,$$

then  $H$  belongs to the type of one of the following three distribution functions:

$$\text{Fréchet:} \quad \Phi_\alpha(x) = \begin{cases} 0 & x \leq 0 \\ \exp\{-x^{-\alpha}\} & x > 0 \end{cases} \quad \alpha > 0,$$

$$\text{Weibull:} \quad \Psi_\alpha(x) = \begin{cases} \exp\{-(-x)^{-\alpha}\} & x \leq 0 \\ 1 & x > 0 \end{cases} \quad \alpha > 0,$$

$$\text{Gumbel:} \quad \Lambda(x) = \exp\{-e^{-x}\} \quad x \in \mathfrak{R}.$$

**Proof:** see Embrechts *et al.* (1997) pg. 122.

This results says that the rescaled sample maxima,  $(M_n - d_n)/c_n$ , converge in distribution to one of three types - Fréchet, Weibull or Gumbel - known as extreme value distributions. Moreover, the result is obtained irrespective of the population distribution function,  $F$ .

Each of the three distributions has a characteristic tail shape. The Weibull distribution has a short, truncated tail. The Gumbel distribution has a thin tail. The Fréchet has fat-tails, and is generally found to be most appropriate for modelling financial distributions.

All the three underlying extreme value distributions can be summarised in a simple functional form  $H_\xi(x)$ , known as the Generalised Extreme Value (GEV) distribution.

$$H_\xi(x) = \begin{cases} \exp\left\{-\left(1 + \xi x\right)^{-\frac{1}{\xi}}\right\} & \text{if } \xi \neq 0, \\ \exp\{-e^{-x}\} & \text{if } \xi = 0, \end{cases}$$

where  $1 + \xi x > 0$  and  $\xi \equiv 1/\alpha$ . The parameter  $\xi$  is called the *tail index*.

---

<sup>17</sup> We follow Embrechts, Klüppelberg and Mikosch (1997).

If the distribution function  $F$  is such that  $(M_n - d_n)/c_n$  converges to  $\Phi_\alpha(x)$ , it is said to be in the domain of attraction of the Fréchet distribution, that is,  $F \in \text{MDA}(\Phi_\alpha)$ . For example, the Pareto, the Cauchy and the Burr distributions all belong to the domain of attraction of the Fréchet distribution. The distributions belonging to  $F \in \text{MDA}(\Phi_\alpha)$  all have fat tails and hence are deemed suitable to model financial time series.

Similar domains of attraction are defined for the short-tailed Weibull distribution and the thin tailed Gumbel distribution<sup>(18)</sup>. In the Fréchet case, the complementary distribution function  $\bar{F}(x) \equiv 1 - F(x)$  can be expressed as a simple asymptotic form,  $\bar{F}(x) \approx L(x)x^{-\alpha}$ , where  $L(x)$  is a slowly varying function<sup>(19)</sup>. Importantly, this result shows that, for  $x$  large enough, the tail of the distribution function can be approximated by a power function. In the case of the Gumbel distribution, the asymptotic form has an exponential type behaviour.

In the case  $F \in \text{MDA}(\Phi_\alpha)$ , we can set  $L(x) = u^\alpha$  and write  $\bar{F}(x) \approx Cx^{-\alpha}$  and  $F'(x) = f(x) = \alpha Cx^{-\alpha-1}$ .

Given the above, for  $F \in \text{MDA}(\Phi_\alpha)$ , the expected loss conditional on  $X$  exceeding  $x_q$  can be calculated as:

$$E[X|X > x_q] = \frac{\int_{x_q}^{\infty} x \alpha C x^{-\alpha-1} dx}{Cx_q^{-\alpha}} = \frac{1}{1-\xi} x_q, \quad \xi = 1/\alpha.$$

$\alpha$  is calculated from maximum likelihood estimation (the Hill's estimate).

### Appendix B: Hill's estimator small sample properties

Suppose we are given  $n$  observations  $\{X_1, \dots, X_n\}$  and we denote order statistics by  $X_{1:n} \leq \dots \leq X_{n:n}$ , that is, the  $n^{\text{th}}$  order statistic is simply the largest element in the sample, the  $n-1^{\text{th}}$ , the second largest element, and so on. The Hill estimator based on  $n$  observations and the first  $k$  order statistics is,

$$\hat{\xi} = \hat{\xi}_{k,n}^{(H)} = \frac{1}{k} \sum_{j=1}^k \log X_{n-j+1:n} - \log X_{n-k:n}.$$

Dacorogna, Muller and Pictet (1996) show that the asymptotic expected value and variance of the Hill estimator are given by,

$$E[\hat{\xi}_{k,n}^{(H)}] \approx \frac{1}{\alpha} - \frac{b\beta}{\alpha(\alpha+\beta)} a^{-\frac{\beta}{\alpha}} \left(\frac{k}{n}\right)^{\frac{\beta}{\alpha}}, \quad (1)$$

$$V[\hat{\xi}_{k,n}^{(H)}] \approx \frac{1}{k\alpha^2}, \quad (2)$$

where  $\alpha, \beta > 0$ ,  $a, b \in \mathfrak{R}$  and are such that distribution function  $F(x)$  can be approximated by

<sup>18</sup> See Embrechts, Klüppelberg and Mikosch (1997) section 3.3 for the MDA conditions.

<sup>19</sup> A positive function  $L$  defined on  $(0, \infty)$  is slowly varying if  $\lim_{x \rightarrow \infty} L(tx)/L(x) = 1$ ,  $t > 0$ .

$$F(x) \approx 1 - ax^\alpha(1 + bx^{-\beta}).$$

Note that according to **(1)**, the bias increases with  $k$ , suggesting that a small value should be chosen for  $k$ . **(2)**, however, shows that the variance of the estimator increases for low values of  $k$ .

Huisman *et al.* (1997) suggest linearizing **(1)** with respect to  $k$ , obtaining,

$$\hat{\xi}_{k,n}^{(H)} = \beta_0 + \beta_1 k + \varepsilon(k), \quad k = 1, \dots, \bar{k} \quad (3)$$

for an appropriate value of  $\bar{k} < n$ . We can rewrite **(3)** in matrix form as  $\xi = \mathbf{Z}\beta + \varepsilon$  where  $\mathbf{Z}$  is a  $(\bar{k} \times 2)$  matrix with ones in the first column and the vector  $[1, 2, \dots, \bar{k}]'$  in the second. Using **(3)**, we can apply a weighted least square method (WLS) with weighting matrix  $\mathbf{W}$  with the vector  $[\sqrt{1}, \dots, \sqrt{\bar{k}}]'$  in the main diagonal and zeros elsewhere. The WLS estimator will be then

$$\hat{\xi}_{\text{WLS}} = (\mathbf{Z}'\mathbf{W}'\mathbf{W}\mathbf{Z})^{-1}\mathbf{Z}'\mathbf{W}'\mathbf{W}\xi,$$

and the estimated tail-index  $\xi$  is equal to the first element of the vector  $\hat{\xi}_{\text{WLS}}$ . Kiesel *et al.* (2000) extend this approach by deriving an approximation of the variance/covariance matrix of the vector of innovations  $\varepsilon$ , so that they can apply a generalised least square estimator for the tail-index  $\xi$ . See Appendix A in Kiesel *et al.* (2000) for details.

## References

- Artzner, P, Delbaen, J E and Heath, D (1997)**, 'Thinking Coherently', *Risk*, Vol. 10.
- Barone-Adesi, G, Giannopoulos, K and Vosper, L (2000)**, 'Backtesting the Filtered Historical Simulation', *Università della Svizzera Italiana Working Paper*.
- Bates, D S and Craine, R (1999)**, 'Valuing the Futures Market Clearinghouse's Default Exposure During the 1987 Crash', *Journal of Money, Credit and Banking*, Vol. 31, pages 248-272.
- Bates, D S (2000)**, 'Post-'87 Crash Fears on the S&P 500 Futures Option Market', *Journal of Econometrics*, Vol. 94, pages 181-238.
- Bernanke, B (1990)**, 'Clearing and Settlement during the Crash', *The Review of Financial Studies*, Vol. 3, pages 133-151.
- Beirlant, J, Vynckier, P and Teugels, J (1996)**, 'Excess functions and estimation of the extreme value index', *Bernoulli*, Vol. 2, pages 293-318.
- Booth, GG, Broussard, J P, Martikainen, T and Puttonen, V (1997)**, 'Prudent Margin Levels in the Finnish Stock Index Futures Market', *Journal of Management Science*, Vol. 43, pages 1177-1188.
- Broussard, J P (2001)**, 'Extreme-Value and Margin Setting with and without Price Limits', *Quarterly Review of Economics and Finance*, Vol. 41, pages 365-385.
- Cornew, R W, Town, D E and Crowson, L D (1984)**, 'Stable distributions, futures prices and the measurement of trading performance', *Journal of Futures Markets*, Vol. 4, pages 531-557.
- Cotter, J (2000)**, 'Conditional and unconditional risk management estimates for European stock index futures', *Department of Banking and Finance Working Paper, Graduate School of Business, University College Dublin*.
- Cotter, J and McKillop, D (2000)**, 'The Distributional Properties of Futures Prices', *Journal of Business Finance and Accounting*, Vol. 27, pages 487-510.
- Cotter, J (2001)**, 'Margin Exceedences for European Stock Index Futures using Extreme Value Theory', *Journal of Banking and Finance*, Vol. 25, pages 1475-1502.
- Dacorogna, M M, Muller, U A and Pictet, O V (1996)**, 'Hill, Bootstrap and Jackknife Estimators for Heavy Tails', *Olsen Associates Working Paper, Zurich*.

**Danielsson, J, de Vries, C, de Haan, L and Peng, L (1999)**, 'Using a Bootstrap Method to choose the Sample Fraction in Tail Index Estimation', *Financial Markets Group Working Paper, London School of Economics*.

**Danielsson, J and Morimoto, Y (2000)**, 'Forecasting Extreme Financial Risk: A Critical Analysis of Practical Methods for the Japanese Market', *Journal of Monetary and Economic Studies*, Vol. 18.

**Danielsson, J (2001)**, 'The Emperor has no Clothes: Limits to Risk Modelling', *Journal of Banking and Finance*, Vol. 26, pages 1273-1296.

**Dewachter, H and Gielens, G (1999)**, 'Setting futures margins: the extremes approach', *Journal of Applied Financial Economics*, Vol. 9, pages 173-181.

**Drees, H and Kaufmann, E (1998)**, 'Selecting the optimal sample fraction in univariate extreme value estimation', *Journal of Stochastic Processes and their Applications*, Vol. 75, pages 149-172.

**Embrechts, P, Klüppelberg, C and Mikosch, T (1997)**, *Modelling extremal events for insurance and finance*, Berlin: Springer.

**Fenn, G W and Kupiec, P (1993)**, 'Prudential Margin Policy in a Futures-Style Settlement System', *Journal of Futures Markets*, Vol. 13, pages 389-408.

**Figlewski, S (1984)**, 'Margins and Market Integrity: Margin Setting for Stock Index Futures and Options', *Journal of Futures Markets*, Vol. 4, pages 385-414.

**Gay, G D, Hunter, WC and Kolb, RW (1986)**, 'A Comparative Analysis of Futures Contract Margins', *Journal of Futures Markets*, Vol. 6, pages 307-324.

**Gouriéroux, C (1997)**, *ARCH models and Financial Applications*, New York: Springer Series in Statistics.

**Hansen, P R and Lunde, A (2001)**, 'A Comparison of Volatility Models: Does Anything beat a GARCH?', *Centre for Analytical Finance Working Paper no .84, Aarhus School of Business, University of Aarhus*.

**Huisman, R, Koedijk, K, Kool, C and Palm, F (1997)**, 'Fat Tails in Small Samples', *Limburg Institute of Financial Economics Working Paper, Maastricht University*.

**Jorion, P (2002)**, 'Fallacies about the effects of market risk management systems', *The Financial Stability Review*, Vol. 13, pages 115-127.

**Keppo, J (1997)**, 'Calling for the True Margin', *Journal of Applied Financial Economics*, Vol. 7, pages 207-212.

**Kiesel, R, Perraudin, W and Taylor, A (2000)**, ‘An Extremes Analysis of VaRs for Emerging Market Benchmark Bonds’, *Birkbeck College, University of London Working Paper*.

**Knott, R and Mills, A (2002)**, ‘Modelling risk in central counterparty clearing houses: a review’, *The Financial Stability Review*, Vol. 13, pages 162-174.

**Kupiec, P (1997)**, ‘Margin Requirements, Volatility, and Market Integrity: What Have We Learned Since the Crash?’, *Finance and Economics Discussion Series FEDS no. 97, Board of Governors of the Federal Reserve System*.

**Longin, F M (1999)**, ‘Optimal margin level in futures markets: extreme price movements’, *Journal of Futures Markets*, Vol. 19, pages 127-152.

**Ma, C K, Mercer, J M and Walker, M A (1992)**, ‘Rolling Over Futures Contracts: A Note’, *Journal of Futures Markets*, Vol. 12.

**McNeil, A J and Frey, R (2000)**, ‘Estimation of Tail Related Risk Measures for Heteroscedastic Financial Time Series: an Extreme Value Approach’, *ETH Zurich*.

**Venkateswaran, M, Brorsen, B W and Hall, J A (1993)**, ‘The Distribution of Standardised Futures Price Changes’, *Journal of Futures Markets*, Vol. 13, pages 279-298.

**Warshawsky, M J and Dietrich, E (1989)**, ‘The Adequacy and Consistency of Margin Requirements in the Markets of Stocks and Derivative Products’, *Federal Reserve Working Paper no. 158*.