

REMOVING MATURITY EFFECTS OF IMPLIED RISK NEUTRAL DENSITIES AND RELATED STATISTICS

Amadeo Alentorn^a and Sheri Markose^{a b *}

^a Centre for Computational Finance and Economic Agents (CCFEA)

^b Department of Economics
University of Essex

22 February 2006

When studying a time series of implied Risk Neutral Densities (RNDs) or other implied statistics, one is faced with the problem of maturity dependence, given that option contracts have a fixed expiry date. Therefore, estimates from consecutive days are not directly comparable. Further, we can only obtain implied RNDs for a limited set of expiration dates. In this paper we introduce two new methods to overcome the time to maturity problem. First, we propose an alternative method for calculating constant time horizon Economic Value at Risk (EVaR), which is much simpler than the method currently being used at the Bank of England. Our method is based on an empirical scaling law for the quantiles in a log-log plot, and thus, we are able to interpolate and extrapolate the EVaR for any time horizon. The second method is based on an RND surface constructed across strikes and maturities, which enables us to obtain RNDs for any time horizon. Removing the maturity dependence of implied RNDs and related statistics is useful in many applications, such as in (i) the construction of implied volatility indices like the VIX, (ii) the assessment of market uncertainty by central banks (iii) time series analysis of EVaR, or (iv) event studies.

JEL classification: G10, G14, C14

Keywords: Economic Value-at-Risk, empirical scaling law, term structure of implied RNDs, maturity effect, RND surface.

* Corresponding author: Dr. Sheri Markose, e-mail address scher@essex.ac.uk

We are grateful for comments from Kyriakos Chourdakis, Olaf Menkens, Christian Schlag, Christoph Schleicher and Radu Tunaru. We would also like to thank participants at the following conferences: FMA 2005 in Siena, CEF 2005 in Washington.

1. Introduction

The majority of studies on extracting information implied by traded option prices have focused on the analysis of the implied distributions at a single point in time. These include event studies, as in Bates (1991) for the study of the 1987 crash, Gemmill and Saflekos (2000) for the study of British elections, and Melick and Thomas (1996) for the analysis of oil prices during the Gulf war crisis. It is only recently that the dynamics of the implied distributions has received some attention with the work of Panigirtzoglou and Skiadopoulos (2004). A problem encountered when looking at the dynamics of RNDs, or RND related statistics, is the maturity dependence of information implied by options, given the fixed maturity of option contracts. There are two separate effects: time-to-maturity effect and contract-switch effect (Melick and Thomas, 1998). Implied RNDs are usually constructed using the options with shortest time to maturity. The time to maturity effect is due to the decrease of uncertainty as the expiry date of the options approaches, reducing the time horizon of the RND. This means that both the time horizon of the RND and the degree of uncertainty decrease as the expiry date approaches. The degree of uncertainty jumps up again in the contract-switch point, when the contract with the shortest time to maturity expires, and we replace it with the next expiration date contract. This maturity dependence needs to be removed in order to make information extracted from option prices useful.

We briefly survey some of the proposed methods in the literature for removing this maturity effects from implied RNDs and related statistics. Gemmill and Saflekos (2000) conducted event studies around crisis periods and British elections. They adjusted the RNDs to give the same maturity before and after the event, since as they point out, without an adjustment, there is a narrowing of the distribution as maturity approaches. However, the adjustment they make is based on an *ad-hoc* scaling of the variances with the square root of time, and linear scaling of the mean. Butler and Davies (1998), who used RNDs to assess market views on monetary policy, suggested an alternative approach to this, by using options of two maturities and then synthesizing an implied distribution for a constant time horizon.

Some other studies that looked at time series of implied statistics worked around the problem of maturity dependence by only using one day a month, to ensure that each point in the time series has the same time to maturity. For example, Gemmill and Saflekos (2000) estimated the skewness time series of implied distributions from 1987 to 1997, but were only able to estimate 120 data points (one per month), in order to ensure each of the points in the time series had the same number of days to maturity (forty-five days). If one were able to estimate a constant time horizon RND every day, this type of time series could be build from daily estimates, instead of only monthly, and therefore, a more accurate picture and richer study of the information embedded in option prices could be obtained.

One of the recent attempts to construct a daily time series of implied statistics was done in Vähämaa (2005), who estimated a daily time series of implied skewness with constant maturity of 30

days by linear interpolation between two skewness estimates from adjacent maturities. Malz (2001) used a similar method, performing linear interpolation of option implied volatilities with adjacent maturities to calculate constant time horizon implied volatilities. Similarly, Härdle and Hlávka (2005) found that the variance of implied RND decreases linearly as the option moves closer to its maturity, and suggested that implied RND estimates calculated for neighbouring maturities can be linearly interpolated in order to obtain an implied RND estimate with arbitrary time to maturity.

The issue of removing the time to maturity effect has also been addressed when constructing implied volatility indices such as the VIX (CBOE, 2003). The VIX represents the expected stock market volatility over the next 30 calendar days, implied by S&P 500 index option prices. Its calculation uses the two nearest-term expiration months in order to bracket a 30-day horizon. It calculates the implied volatility for each of the two maturities, and then it interpolates them to arrive at a single value with a constant maturity of 30 days. A constant maturity needs to be used because implied volatilities change as the time to maturity changes (Fleming *et. al.*, 1995).

All these methodologies simply calculate a constant time horizon implied statistic by linear interpolation between only two maturities that bracket the horizon of interest. But in all major markets, there are usually options trading with more than two maturities. Not much attention has been paid at looking how the moments of implied RNDs scale with time along *all* maturities available. An attempt to capture the time dependency of implied volatilities with time to maturity was done in Dumas *et. al.* (1998) by modelling the implied volatility surface as a function of maturity and time. However, no study has looked at how implied RNDs and implied higher moments vary with time to maturity.

The most sophisticated method to date to remove this maturity effect is the one proposed by Clews *et. al.* (2004), who construct constant horizon RNDs. To achieve that, they convert option prices to implied volatilities (IV) in the delta space, and then they interpolate this IV - delta surface at the required time horizon to obtain a set of interpolated data points. Finally, they convert the interpolated IV-deltas back to option prices, and estimate the implied RND that, by construction, will have the required time horizon. They employ a non-parametric method to estimate the implied RND, which consists of calculating the second derivative of the call pricing function using the Breeden and Litzenberger (1978) result. Having obtained a constant time horizon RND, it is easy to construct a time series of constant time horizon implied statistics. For example, this methodology is used by the Monetary Instruments and Markets Division at the Bank of England to report daily estimates of Economic-VaR (EVaR)¹ with a constant 3 month time horizon². The Bank of England also estimates

¹ The term Economic-VaR was coined by Ait-Sahalia and Lo (2000) to distinguish the market implied VaR (EVaR) from the historic backward looking VaR, which they call Statistical-VaR (SVaR).

² For the EVaR calculations, the Bank of England uses a constant time horizon of 3 months for the FTSE 100 index at different confidence levels ranging from 5% to 95%, in intervals of 5%. These BoE EVaR estimates are available at <http://www.bankofengland.co.uk/statistics/IMPLIEDPDFS/>

RNDs for other assets using this method, such as physical commodities, and uses them in some of the Inflation Reports³.

The BoE methodology for removing the time to maturity effects has the following drawbacks. Conversions to and from delta space make this methodology difficult to both understand and implement. Additionally, when interpolating in the delta space, this methodology uses options only from the two maturities that bracket the required horizon, and does not use the option data available for all other maturities. Because of the need to interpolate between two maturities, this methodology is unable to calculate implied RNDs for shorter horizons than the closest maturity, or longer horizons than the longest maturity, given that the implied volatility surface in delta space is non-linear, and therefore not easily extrapolated. For instance, it is not possible to estimate a daily 10 day constant time horizon implied RND, and therefore, it is not possible to report the daily 10 day EVaR, which is one of the most relevant time horizons used in the industry for VaR.

In this paper we introduce two new methods to exclude this systematic impact of time to maturity on the implied RNDs and related statistics. Our methods use option prices for *all* maturities available in the estimation, instead of using only options for two maturities. Additionally, our methods are computationally less arduous than the BoE one. First, we propose a new method to calculate a fixed horizon EVaR, based on an empirical scaling law in the quantile space for parametric based RND extraction methods⁴. The main steps of our method are as follows: using option prices with different maturities, we construct a discrete term structure of implied RNDs, one RND for each maturity. Then, for each RND in the term structure, we obtain the EVaR estimates at different confidence levels (i.e. 99%, 95%, etc). Finally, we discover a linear behaviour of the EVaR values with time to maturity in the log-log scale⁵, which can be exploited to estimate an empirical scaling law for each confidence level. In this paper we demonstrate this linear behaviour of EVaRs in the log-log scale for the following parametric RND extraction methods: Black-Scholes, mixture of lognormals (Ritchey 1990) and GEV model (Markose and Alentorn, 2005).

One of the advantages of this linear relationship, apart from being considerably less complex than the method proposed by the BoE, is that it allows us to both interpolate and extrapolate any holding period. Therefore, we can obtain EVaR values at any time horizon. For example, we can obtain a daily 10 day EVaR regardless of the time to maturity of the closest maturity contracts. When comparing these results with the BoE non-parametric time series, we find that the GEV EVaR is

³ For example, in the May 2001 Inflation Report, changes in the 6 months constant time horizon implied distributions of oil price, derived from option prices for West Texas Intermediate (WTI) crude oil, were used to argue that the market uncertainty about oil prices increased considerably between February and May of 2001 (BoE, 2001)

⁴ When implied VaR has to be evaluated at high confidence intervals, the non-parametric methods fail to give a reliable estimate, given that those methods are unable to describe the tails of the distribution outside the range of available strikes. Therefore, in this paper we use parametric methods.

⁵ Menkens(2004) and Provizionatou *et. al.* (2005) were the first to point out the linear relationship of VaR estimates using the log-log plot in the context of historical VaR. Here, we find a linear relationship for implied VaR.

remarkably close to the series generated by the BoE non-parametric method, whereas the Black-Scholes method yields lower EVaR estimates, and the mixture of lognormals method yields higher EVaR estimates.

In the second part of this paper we extend the notion of a discrete term structure of RNDs to obtain, for a given day, an implied RND surface. This concept is then applied to the GEV model, by re-parameterizing the GEV option pricing model so that the scaling of the implied volatility is a parameter that will be estimated. By making the RND model an explicit function of time and by utilizing all options with different maturities, we are able to obtain for each day, not only the constant horizon implied VaR, but all implied statistics without maturity effects for any time horizon. This implied RND surface has several useful applications, such as in the (i) construction of implied volatility indices like the VIX, (ii) assessment of market uncertainty by central banks (iii) time series analysis of EVaR, and (iv) event studies. In this paper, we briefly discuss these applications.

The rest of the paper is structured as follows. In the next section we present the method to remove maturity effects from EVaR using an empirical scaling law in the quantile space. In Section 3 we extend the original GEV model to obtain an implied RND surface, and briefly discuss practical applications for this method. We conclude in Section 4.

2. Removing maturity effects for Economic VaR

2.1 The term structure of RNDs

In most derivatives markets there are traded option contracts for several different maturities. For example, in the UK there are options on the FTSE 100 index expiring every month for the next three months, and quarterly (March, June, September and December) for the next eight quarters. That is, at any one time, there are around 10 different maturities available. However, not all of these maturities have traded option prices, because option contracts with very long maturities are not traded very often. In fact, the average number of maturities with traded options for the FTSE 100 index between 1997 and 2003 was found to be around 5 maturities, as shown in Table 1 below.

| Year | Average number of maturities available |
|----------------|---|
| 1997 | 3.96 |
| 1998 | 4.57 |
| 1999 | 5.19 |
| 2000 | 5.49 |
| 2001 | 5.84 |
| 2002 | 6.19 |
| 2003 | 6.09 |
| Average | 5.33 |

Table 1: Average daily number of maturities with traded option prices for the FTSE 100 index.

Despite having a considerable number of different maturities with traded option prices, it is common in the RND extraction literature to only use option prices from a single maturity, usually the maturity with the closest expiration date. Here, we propose, on a daily basis, the extraction of an RND for each of the maturities with a sufficient number of traded option prices⁶. Then, using this discrete set of RNDs, each with a different maturity, we can construct what we call a term structure of implied RNDs. This term structure can be visualized as a 3 dimensional chart that displays, for a given day, how the implied RNDs vary across different maturities. For purposes of illustration, Figure 1 below displays the implied RND term structure for a typical day, 21 August 2001, using the GEV model. A comprehensive comparison between the three parametric methods will be made later on in the paper, Note from Figure 1 that the main feature of the term structure, which we find to be independent of the RND extraction method used, is that the peakedness of the RNDs decreases as the time horizon increases. This illustrates the so called maturity effect. This term structure of implied RNDs will be used in the following section to obtain constant time horizon EVaRs.

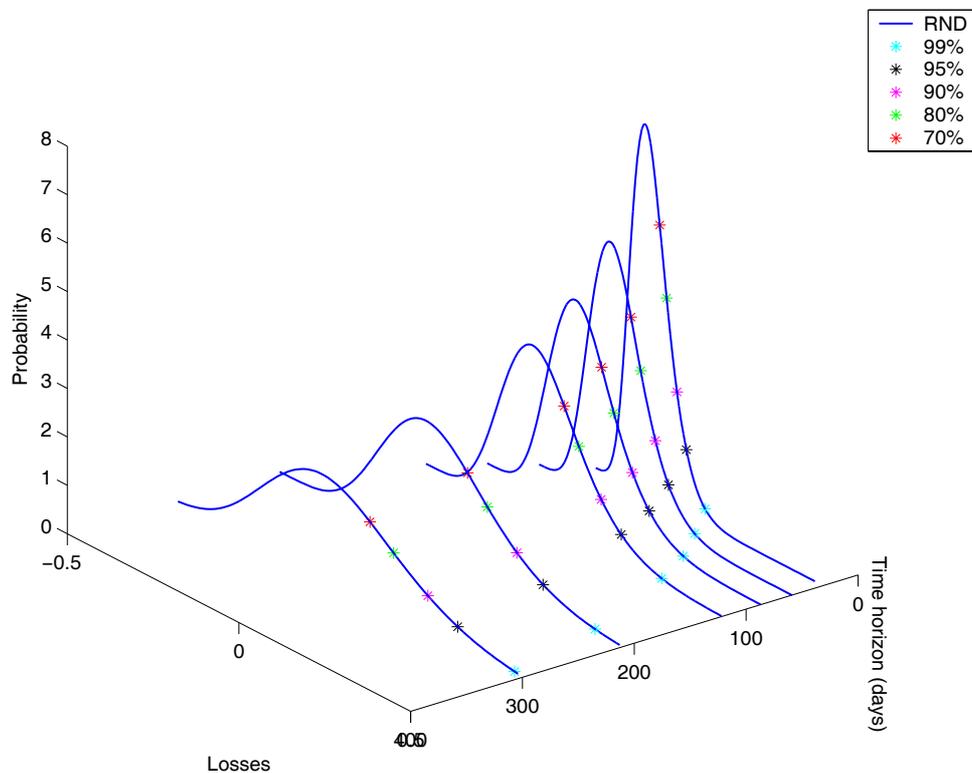


Figure 1: Term structure of GEV based implied RNDs for 21 August 01. The coloured stars on the graph indicate the EVaR values for each RND at different confidence intervals.

⁶ The number of option prices needed to extract the RND must be at least equal to the number of degrees of freedom for the parametric method used. The number of degrees of freedom is equal to the number of parameters that need to be estimated minus the number of constraints. For example, the GEV model has three parameters while the mixture of lognormals have five parameters. We only use one constraint, based on a martingale condition that the mean of the RND has to be equal to the Futures price.

2.2 Economic Value at Risk (EVaR)

The most popular measure for risk management is Value-at-Risk, denoted by $VaR(q, k)$, which is an estimate, with a given degree of confidence q , of how much can be lost from a portfolio over a given time horizon k . An alternative measure of risk is Economic VaR (EVaR), which was proposed by Ait-Sahalia and Lo (2000) and it is calculated under the option-implied risk neutral density. It has been argued that EVaR is a more general measure of risk, since it incorporates the investor's risk preferences, the demand–supply effects, and the probabilities that correspond to extreme losses (Panigirtzoglou and Skiadopoulos 2004). EVaR can be seen as a forward looking measure to quantify market uncertainty about the future course of financial asset prices, whereas Statistical VaR (SVaR) is a backward looking measure, based on the historical distribution of returns.

Ait-Sahalia and Lo (2000) estimated EVaR from RNDs extracted using a non-parametric method. Given that EVaR is calculated as the quantile of the RND at the tails, if one uses a non-parametric distribution, the tails of the RND have to be arbitrarily extrapolated, because there are no options trading at very low or very high strike prices. As pointed out by Neuhaus (2000), in connection with the study of Coopers (2000), there are not always sufficient number of strikes and traded option prices to cover the whole distribution, so if only part of the density function is estimated, it can be difficult to correctly allocate the missing probability mass. Some studies using non-parametric techniques (Anagnou *et. al.* 2002) have explicitly taken this limitation on board by choosing not to model the tails, and instead, estimating only a truncated density for the range of strikes with available traded option prices. Alternatively, when estimating the RND using a parametric approach, the tails of the implied distribution outside the range of available strikes are implicitly given by the distribution function that has been assumed. In particular, the GEV distribution explicitly includes a tail shape parameter. Given the limitations of non-parametric methods, we will be employing parametric methods in what follows.

2.3 Calculating EVaR with parametric RND methods

When extracting RNDs with a parametric method, we can use a non-linear optimization algorithm to estimate the set of method specific parameters θ , which minimizes the sum of squared errors between the traded option prices at different strikes but for the same maturity T , and the prices given by the model:

$$SSE(t) = \min_{\theta} \left\{ \sum_{i=1}^N (C_i(K_i, T) - \tilde{C}_i(K_i, T))^2 \right\} \quad (1)$$

In this paper, we use three different parametric RND extraction methods to study scaling of EVaR. EVaR is obtained by calculating the quantile of the implied RND at the required confidence level q . The first parametric method we use is the Black-Scholes model, which is commonly used as

the benchmark. The Black-Scholes model assumes that log returns are normally distributed, with mean μT and variance $\sigma^2 T$. The EVaR values for this model are obtained using the inverse of the normal cumulative distribution function (cdf) equation. The normal inverse function is defined in terms of the normal cdf H as given in equation (2):

$$EVaR(q, k) = H^{-1}(q | \hat{\mu}T, \hat{\sigma}\sqrt{T}) \quad (2)$$

Here, $EVaR(q, k)$ denotes Economic Value at Risk with a confidence level q , and a time horizon of k days. It is calculated using the implied parameters $\hat{\theta} = \{\hat{\mu}, \hat{\sigma}\}$ obtained from equation (1) above. The second parametric RND extraction method we use is the mixture of lognormals (hereafter MLN). This method was introduced by Ritchey (1990) and has been extensively used in the literature, given that it is very flexible, and allows the modelling of different levels of skewness, as well as bimodal densities. The MLN method models the RND as a weighted sum of two lognormals, and given by:

$$f(S_T) = p h(S_T | \mu_1 T, \sigma_1 \sqrt{T}) + (1-p) h(S_T | \mu_2, \sigma_2 \sqrt{T}) \quad (3)$$

There are five unknown parameters $\theta = \{\mu_1, \mu_2, \sigma_1, \sigma_2, p\}$, the means of each lognormal function μ_1 and μ_2 , the standard deviations σ_1 and σ_2 , and the weighting coefficient p . We obtain the set of implied parameters $\hat{\theta}$ by using equation (1). Then, EVaR is calculated as the quantile of the MLN density, which consists of a weighted sum of the two inverse cdfs, and given by:

$$EVaR(q, k) = \hat{p} H^{-1}(q | \hat{\mu}_1, \hat{\sigma}_1, T) + (1 - \hat{p}) H^{-1}(q | \hat{\mu}_2, \hat{\sigma}_2, T) \quad (4)$$

The last RND extraction method we use is the GEV option pricing model in Markose and Alentorn (2005). This model is based on the Generalized Extreme Value (GEV) distribution, which is a distribution characterized by a set of three parameters $\theta = \{\mu, \sigma, \xi\}$, the location parameter μ , the scale parameter σ , and the tail shape parameter ξ . Again, we estimate the set of implied parameters $\hat{\theta} = \{\hat{\mu}, \hat{\sigma}, \hat{\xi}\}$ using equation (1). The quantile equation of the GEV distribution gives us the EVaR value associated with a given confidence level q , and as shown by Dowd (2004: pp. 274) is a function of the GEV parameters as follow:

$$EVaR(q, k) = \hat{\mu} - \frac{\hat{\sigma}}{\hat{\xi}} \left[1 - (-\log(q))^{-\hat{\xi}} \right] \quad \xi \neq 0 \quad (5)$$

$$EVaR(q, k) = \hat{\mu} - \hat{\sigma} \log[\log(1/q)] \quad \xi = 0 \quad (6)$$

For purposes of illustration, Table 2 below displays the actual EVaR values for 21 August 2001, corresponding to the coloured stars in the implied RND term structure of Figure 1, for the GEV case. Note that a comparison between the three methods will be made later on in the paper. As one would expect, the EVaR values increase both with confidence level and with time horizon. Also, note how the number of options prices available decreases as time to maturity increases, that is, the options with the closest to maturity dates are the ones that have the widest range of traded strikes.

| Expiry month | Days to maturity | Number strikes | EVaR | | | | |
|-----------------|---------------------|-------------------|------|-------|-------|-------|-------|
| | | | 70% | 80% | 90% | 95% | 99% |
| Sep-01 | 31 | 44 | 2.4% | 4.4% | 7.4% | 10.1% | 15.6% |
| Oct-01 | 59 | 31 | 3.1% | 5.9% | 10.2% | 14.0% | 21.7% |
| Nov-01 | 87 | 13 | 3.7% | 7.2% | 12.7% | 17.5% | 27.4% |
| Dec-01 | 122 | 16 | 4.2% | 8.5% | 15.0% | 20.8% | 32.6% |
| Mar-02 | 213 | 13 | 5.7% | 11.4% | 20.1% | 27.7% | 42.8% |
| Jun-02 | 304 | 10 | 6.9% | 13.7% | 23.8% | 32.4% | 49.0% |

Table 2: EVaR values for each available maturity and at different confidence levels

on 21 August 2001, for the GEV model.

2.4 Scaling of EVaR

When calculating EVaR for a given time horizon of k days, ideally, one would use a RND implied by options that mature in k days. For example, to calculate the 10 day EVaR we would use prices of options that mature in 10 days. However, in practice, we only have option prices for a small set of fixed expiration dates. For example, in the above Table 2, for 21st August 2001, a 10 day EVaR can not be readily reported, given that the closest maturity is at 31 days. Hence, to obtain an EVaR for a given time horizon, we need to resort to scaling.

Interest in scaling SVaR has arisen due to the requirements of the Basel accord, which states that banks should report the daily 10 day VaR at the 99% confidence level of their portfolios. However, there are some difficulties on estimating the 10 day VaR, due to the need for a long time series in order to compute the 10 day returns, and then, calculate the quantiles of their distribution. In practice, the square root of time scaling rule is widely used to scale up the 1 day VaR to the 10 day VaR. This scaling rule is only appropriate for time series that have Gaussian properties, but it has been well established in the literature for a long time, such as in Fama (1965) and Mandelbrot (1967), that financial data is non-Gaussian. Following the wide spread use of VaR as a risk measure and reporting requirement, there have been several recent studies that looked at the problem of scaling VaR, such as McNeil and Frey (2000), Kaufmann and Patie (2003), Danielsson and Zigrand (2004), Menkens (2004) and Provizionatou *et. al.* (2005).

When estimating EVaR for a given time horizon k , we are faced with a similar problem, but instead of having to scale up the 1-day VaR to the k -day VaR, we usually need to scale down from the maturities available. Without resorting to a scaling law, we would only be able to calculate the 10 day VaR for only one day each month, the day when there are exactly 10 days to maturity for the closest to

maturity contract. In the case of FTSE 100, a 10 day EVaR can only be obtained around the first Friday of each month, since contracts mature in the third Friday of the month. Following an approach as in Menkens (2004) and Provizionatou *et. al.* (2005), we have identified an empirical scaling law for EVaR with respect to time to maturity which is linear in a log-log plot. In general, we define:

$$EVaR(k,q) = k^{b(q)} EVaR(1,q) \quad (7)$$

where k is the number of days, and $b(q)$ is the scaling parameter for a confidence level q . When applying logarithms to both sides of the above equation, and rearranging, we can write down the linear equation with a slope $b(q)$ and intercept $c(q)$ as shown in equation (8) below. The intercept term $c(q)$ is the logarithm of the 1-day EVaR at a confidence level q .

$$\log(EVaR(k,q)) = b(q) \log(k) + c(q) \quad (8)$$

Note how the two coefficients, the slope $b(q)$ and the intercept $c(q)$, are quantile dependent, and thus written as a function of q . We can estimate these two coefficients for a given day and for a given quantile q , by fitting a linear relationship between the vector of estimated EVaR and the vector of number of days to maturity. These vectors have one element for each of the maturities available in the term structure, as shown in equation (9) below.

$$\log \begin{bmatrix} EVaR(k_1, q) \\ EVaR(k_2, q) \\ \dots \\ EVaR(k_n, q) \end{bmatrix} = b(q) \log \begin{bmatrix} k_1 \\ k_2 \\ \dots \\ k_n \end{bmatrix} + c(q) \quad (9)$$

Here k_i is the number of days left to expiration for maturity $i = 1, \dots, n$. Once the parameters \hat{b} and \hat{c} are estimated for a given day and for a given confidence level q , we can obtain an estimate for the k -day EVaR using equation (10) below:

$$EVaR(k, q) = k^{\hat{b}(q)} \exp(\hat{c}(q)) \quad (10)$$

For illustration purposes, Figure 2 below displays, for the GEV case, the log-log plot of the EVaR values for a typical day (21 August 2001) obtained from the term structure of RNDs in Figure 1. Figure 2 also shows the regression line for each of the quantiles. It is easy to see that the EVaR values increase linearly with time in the log-log plot. As was shown in Table 2, the number of different strikes with options traded decreases as time to maturity increases. Therefore, the implied RNDs for far from maturity dates are usually estimated with a smaller number of option prices, and therefore, EVaR estimates from these RNDs will have a wider confidence interval. Instead of using a simple Ordinary Least Squares (OLS) regression, we found it useful to employ a Weighted Linear Squares (WLS)

regression, where the weights are a function of the number of option prices available in each maturity. More details about this can be found in Appendix A.

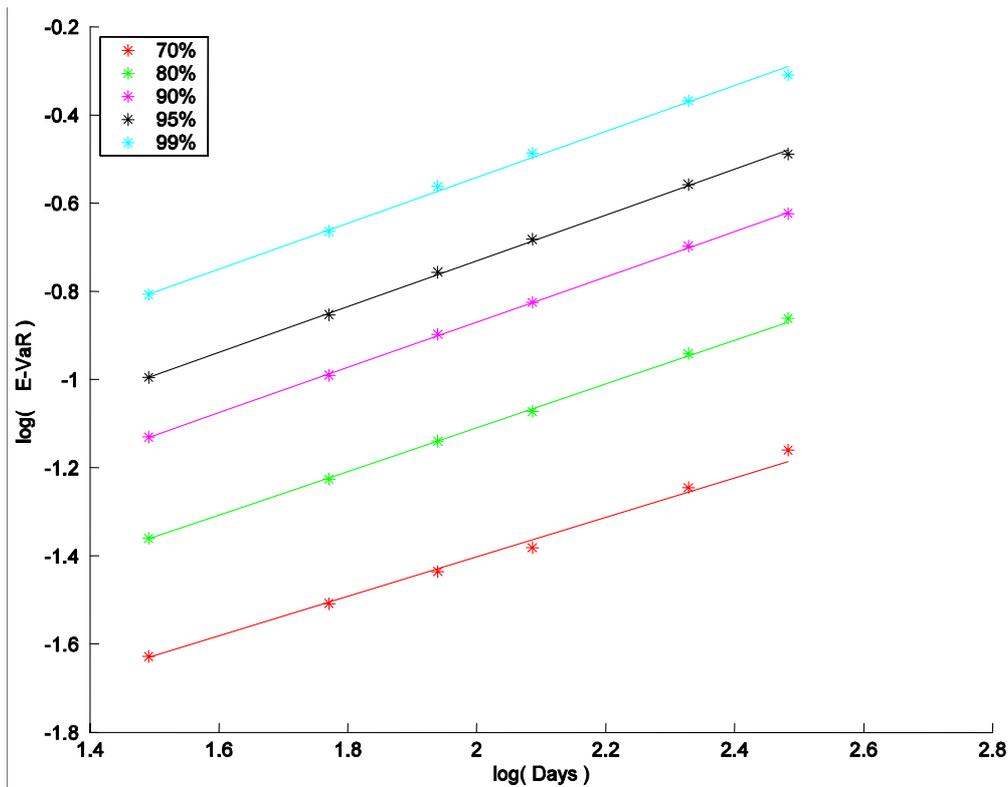


Figure 2: log-log plot of EVaR estimates for 21 August 01, with the estimated linear scaling rule for each confidence level for the GEV model.

2.5 Empirical analysis

The data used in this study are the daily settlement prices of the FTSE 100 index call and put options from 1997 to 2003.⁷ Using option prices for all maturities from the 1733 days in our sample period, we obtain the scaling law coefficient b and the intercept c in equation (8), for different confidence levels. The average values of the regression coefficient b and the exponential of the intercept, $\exp(c)$, are displayed in Table 3 below. We can see that the slope b increases with the confidence level, and that intercept, i.e. the 1 day EVaR also increases with confidence level, as one would expect. We also report the percentage of days where b was found statistically significantly different than one half.⁸ On average, irrespective of the confidence level, we found that in around 50% of the days the scaling was significantly different than 0.5, the scaling implied by the square root of time rule.

⁷ A detailed description of the filtering rules used can be found in Markose and Alentorn (2005).

⁸ Newey-West (1987) heteroskedasticity and autocorrelation consistent standard error was used to test the null hypothesis $H_0 : b=0.5$. We employed this methodology when testing the statistical significance of the estimated slope b , because E-VaR estimates are for overlapping horizons, and therefore are autocorrelated. The Newey-West lag adjustment used was $n - 1$.

| Confidence level q | $b(q)$ | $EVaR(1,q) = \exp(c)$ |
|----------------------|-----------------|-----------------------|
| 70% | 0.41 (51.3%) | 0.6% |
| 80% | 0.48 (51.1%) | 0.9% |
| 90% | 0.51 (52.4%) | 1.4% |
| 95% | 0.53 (52.1%) | 1.9% |
| 99% | 0.56 (52.6%) | 2.7% |

Table 3: Average regression coefficients b and $\exp(c)$ across the 1733 days in the sample, for the GEV case. The percentage number of days where b was found statistically significantly different than 0.5 is indicated in brackets.

Figure 3 below displays the time series of 90 day EVaR estimates at the 95% confidence level for each of the three methods given in equations (2), (4) and (6), together with the estimates from the Bank of England non-parametric method.⁹ The 90 day FTSE returns are displayed in pink. Table 4 below shows the sample mean and standard deviation of each of the four EVaR time series. If we use the BoE values as the benchmark, we can see that on average, the mixture of lognormals method overestimates EVaR, while the Black-Scholes method underestimates it. Among the three parametric methods, the GEV method yields the time series of EVaRs closest to the BoE one. This can be seen both in the Table 4 and also in the Figure 3, where the BoE time series practically overlaps the GEV time series.

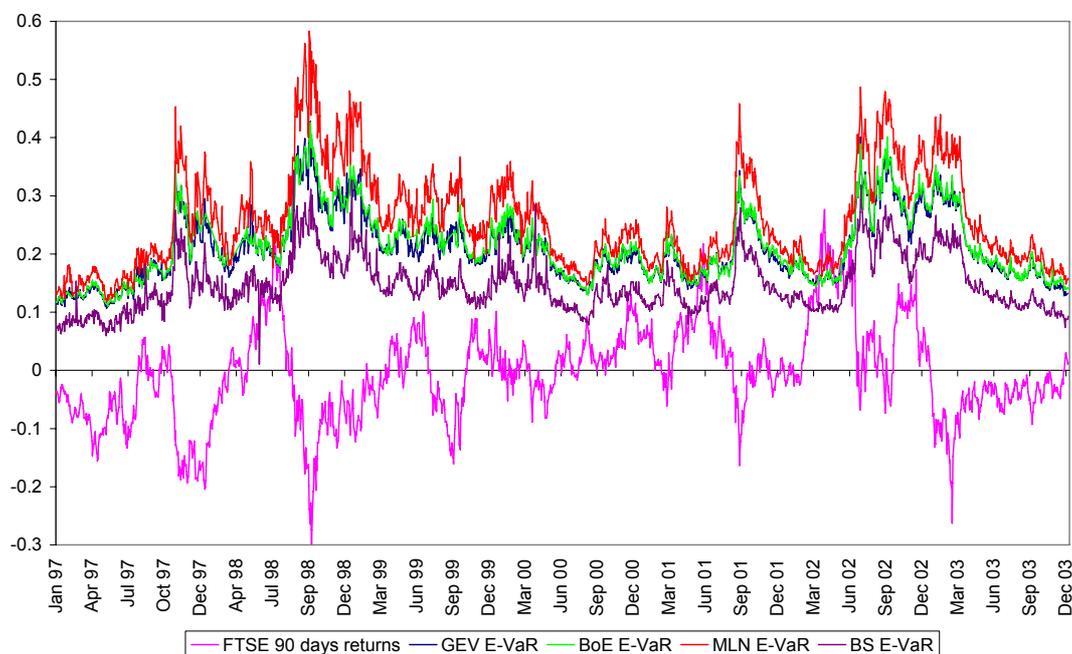


Figure 3: BoE EVaR vs. GEV EVaR vs. Mixture lognormals for 90 days 95%

⁹ A detailed comparison of the backtesting performance of these methods for calculating EVaR can be found in Alentorn and Markose (2005). When comparing the backtesting performance of EVaR with SVaR, it was found that EVaR yields higher and more volatile estimates than SVaR.

| Method | Sample mean | Sample standard deviation |
|--------------------|-------------|---------------------------|
| BoE method | 0.217 | 5.9% |
| GEV model | 0.212 | 6.0% |
| Mixture lognormals | 0.259 | 8.5% |
| Black Scholes | 0.147 | 4.7% |

Table 4: sample mean and standard deviation of the four EVaR time series.

Figure 4 displays the time series of the b estimates (the slope of the scaling law) for the GEV case at the 95% confidence level. Even though the average value of b at this confidence level is 0.53 (see Table 3), it appears to be time varying, taking values that range from 0.2 to 1.

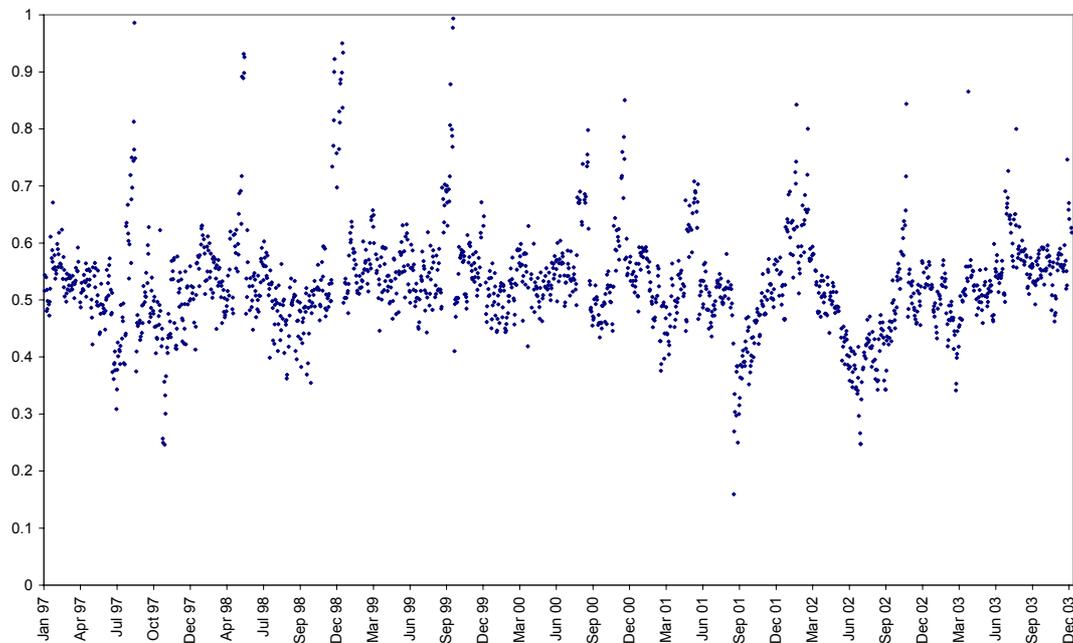


Figure 4: Time series of b estimates for the GEV model at 95% confidence level.

3. Modelling the implied RND surface

In this section, we will construct an implied RND surface to capture the variation of RNDs with time to maturity¹⁰, instead of only considering a discrete term structure of RNDs as was done in the previous section. In the discrete term structure of RNDs, we obtained a different set of implied RND parameters Θ for each of the maturities in a given day. In this section, in order to obtain an implied RND surface, we will constrain Θ to be the same for all maturities in a given day.

For parametric RND models in which volatility is an explicit function of time, such as BS and MLN, this unique set of parameters Θ that fits option prices of all maturities can be readily obtained, without making any changes to the model. However, for the GEV case, as the GEV option pricing equation is not an explicit function of time, we will need to modify the model to make it an explicit function of time, by allowing the mean and volatility to scale with time. For illustration purposes, the implied RND surface obtained from the modified GEV model is plotted below in Figure 5, and it corresponds to the discrete term structure of RND in Figure 1, for 21 August 2001.

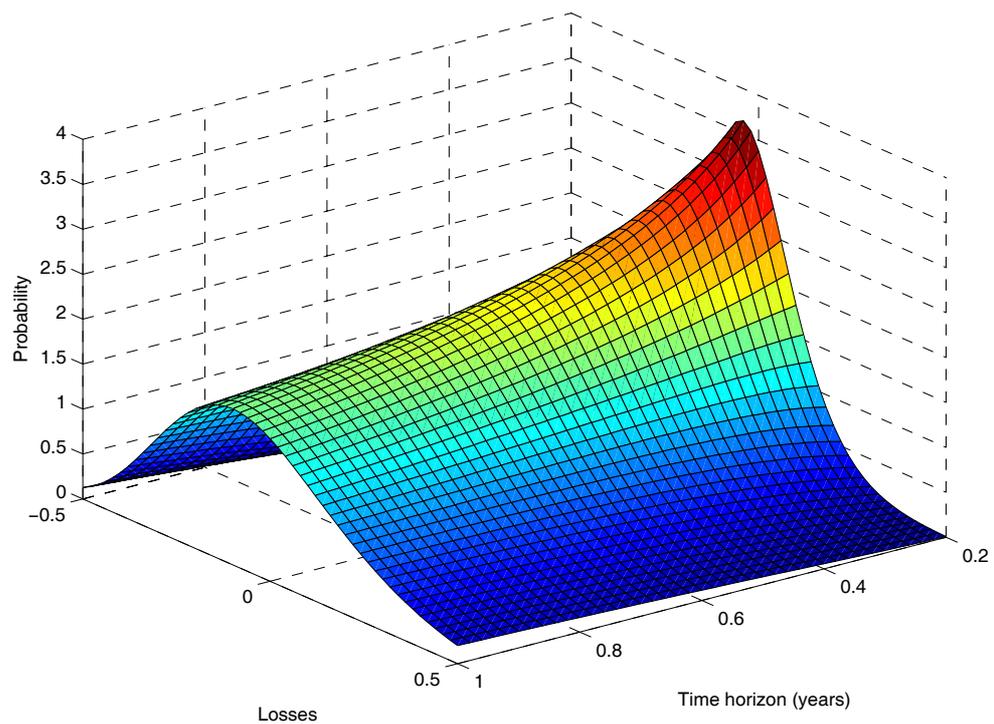


Figure 5: The implied RND surface for the GEV case for 21 August 2001.

¹⁰ Some work on capturing the time to maturity dependence of information embedded in option prices was done in the context of implied volatility surfaces, as in Dumas et al. (1998). They proposed the modelling of the implied volatility surface as a function of both strike and time to maturity, where time to maturity was modelled using a quadratic function. A similar approach was taken by Tompkins (2001), but instead of using a quadratic function of time to maturity, he used a cubic functional form.

3.1 The implied RND surface for the GEV model

The GEV based RND extraction method in Markose and Alentorn (2005) is not an explicit function of time. As a consequence, on any given day the implied GEV parameters are different for each of the available maturities in the term structure of RNDs. In this section we will modify the GEV model and make it a function of time, in such a way that for a given day, it will yield a unique set of parameters consistent with option prices for all maturities.

Note that the GEV option pricing model in Markose and Alentorn (2005) is based on the assumption that negative returns or losses L_T , as defined in equation (11) below, follow a GEV distribution:

$$L_T = -R_T = -\frac{S_T - S_t}{S_t} = 1 - \frac{S_T}{S_t} \quad (11)$$

The standardized GEV distribution, in the form in von Mises (1936) (see, Reiss and Thomas, 2001, p. 16-17), incorporates a location parameter μ , a scale parameter σ , and a tail shape parameter ξ , and it is given by:

$$F_{\xi, \mu, \sigma}(x) = \exp\left(-\left(1 + \frac{\xi}{\sigma}(x - \mu)\right)^{-1/\xi}\right) \quad \text{with} \quad 1 + \frac{\xi}{\sigma}(x - \mu) > 0, \quad \xi \neq 0 \quad (12)$$

and

$$F_{0, \mu, \sigma}(x) = \exp\left(-e^{-\frac{(x - \mu)}{\sigma}}\right) \quad \text{with} \quad \xi = 0. \quad (13)$$

The first step in obtaining an implied RND surface is to allow the mean of the GEV distribution to be a function of time. Reiss and Thomas (1997, pp. 15-18) show that the mean of the GEV distribution is related to the location, scale and tail shape parameters as follows:

$$Mean = \mu + \left[\frac{\Gamma(1 - \xi) - 1}{\xi} \right] \sigma \quad (14)$$

The mean of any RND is constrained by the forward no-arbitrage identity¹¹, which follows from the martingale condition. Using the price at t of a futures contract $F_{t,T}$ that expires at T , and the

¹¹ The forward no-arbitrage identity states that the mean of an RND for a given maturity T has to equal the price at t of a Futures contract with maturity T , and is given by $E_t^Q(S_T) = F_{t,T}$. In the original GEV model this condition was enforced by minimizing the squared of the difference between the mean of the GEV RND and the Futures price as an additional constraint to the non-linear optimization problem. The approach we propose in this paper is more efficient, because by rewriting the GEV distribution as a function of the Futures price, we reduce the number of parameters to be estimated from three to two.

definition of losses in (11) we can impose the martingale condition for the distribution of losses as follows:

$$E_t^Q(L_T) = 1 - \frac{F_{t,T}}{S_t} \quad (15)$$

Equating the right hand sides of equations (14) and (15) and rearranging, we can write the location parameter μ in terms of the futures price $F_{t,T}$ and the other two GEV parameters σ and ξ as follows:

$$\mu = 1 - \frac{F_{t,T}}{S_t} - \left[\frac{\Gamma(1-\xi)-1}{\xi} \right] \sigma \quad (16)$$

The second step in obtaining an implied RND surface for the GEV model consists on making the volatility of the GEV distribution a function of time. Reiss and Thomas (1997, pp. 15-18) show that the volatility of the GEV distribution is related to the scale and tail shape parameters as follows:

$$Volatility = \left[\frac{\sqrt{\Gamma(1-2\xi)-\Gamma^2(1-\xi)}}{\xi} \right] \sigma \quad (17)$$

As can be seen in equation (17), in the GEV model the volatility is not a function of time. But we know that the volatility of an RND function grows with time, because uncertainty increases as the time horizon increases. In order to make our GEV model a function of time, we need to model how the volatility scales with time. Instead of assuming a specific scaling¹², we propose modelling the volatility as a function of T^b as shown below:

$$Volatility(T) = \left[\frac{\sqrt{\Gamma(1-2\xi)-\Gamma^2(1-\xi)}}{\xi} \right] \bar{\sigma} T^b \quad (18)$$

Here, $\bar{\sigma}$ is the annualized GEV scale parameter, and b is a new implied parameter that models how volatility scales with time to maturity. Note that we do not allow ξ to scale with time¹³. Equation (18) above implies the following substitution in the GEV distribution equation:

$$\sigma \equiv \bar{\sigma} T^b \quad (19)$$

¹² The Black-Scholes model assumes that the volatility scales with \sqrt{T} .

¹³ This is based on results not reported in this paper, which showed that ξ does not seem to scale with time, once we correct for the scaling of volatility. When we allowed ξ to scale with T^c , where c was a new implied parameter, it was found that c was insignificantly different than 0 for over 80% of the days in the sample. Additionally, including this scaling parameter for ξ did not improve the explanatory power of the model (in terms of adjusted R^2).

Substituting into the GEV distribution in (12) the equation for μ in (16) and the equation for σ in (19) we obtain a functional form for the GEV distribution, in which by construction, its mean and volatility are a function of time to maturity T :

$$F_{\xi, \mu, \sigma}(x, T) = \exp \left(- \left(1 + \frac{\xi}{\bar{\sigma} T^b} \left(x - \left(1 - \frac{F_{t,T}}{S_t} - \left[\frac{\Gamma(1-\xi)-1}{\xi} \right] \bar{\sigma} T^b \right) \right) \right)^{-1/\xi} \right) \quad (20)$$

Following the same steps as in Markose and Alentorn (2005), we can derive a closed form solution for the European call option price as follows:

$$C_t(K) = e^{-r(T-t)} \left\{ S_t \left(\left(1 - \mu + \frac{\bar{\sigma} T^b}{\xi} \right) e^{-H^{-1/\xi}} - \frac{\bar{\sigma} T^b}{\xi} \Gamma(1-\xi, H^{-1/\xi}) \right) - K e^{-H^{-1/\xi}} \right\} \quad (21)$$

where $H = 1 + \frac{\xi}{\bar{\sigma} T^b} \left(1 - \frac{K}{S_t} - \mu \right)$ and $\mu = 1 - \frac{F_{t,T}}{S_t} - \left[\frac{\Gamma(1-\xi)-1}{\xi} \right] \bar{\sigma} T^b$. A similar expression can be easily obtained for the put option price.

3.2 Empirical analysis

The structural GEV parameters ξ , σ and scaling parameter b can be estimated by minimizing the sum of squared errors (SSE) between the analytical solution of the GEV option pricing equations in (21) and the observed traded option prices for all available strikes K_i , and all available maturities T_j as indicated in (22) below:

$$SSE(t) = \min_{\xi, \sigma, b} \left\{ \sum_{j=1}^M \sum_{i=1}^N \left(C_t(K_i, T_j) - \tilde{C}_t(K_i, T_j) \right)^2 \right\} \quad (22)$$

Note how in the previous section, the SSE was calculated only across different strikes for a given maturity. Here we have a double summation, across both strikes and maturities. In the original GEV study, Markose and Alentorn (2005) tested the empirical performance of the GEV based option pricing model, only using option prices with shortest time to maturity. Here we evaluate the pricing performance of the modified GEV model, which captures option price differences across both strikes and maturities simultaneously, and compare it with the pricing performance of the Black-Scholes and Mixture of Lognormals (MLN) model. The results are displayed in Table 5 below, in terms of RMSE, or average pricing error in pence per option. We can see that, when fitting option prices from several maturities, the GEV RND model, by allowing the scaling of implied volatility to vary, achieves the best pricing performance, even though it has one degree of freedom less than the MLN method.

| Option pricing model | Average pricing error | Number of degrees of freedom |
|-----------------------|-----------------------|------------------------------|
| Black - Scholes | 15.27 | 1 |
| Mixture lognormals | 13.59 | 4 |
| GEV based RND surface | 11.32 | 3 |

Table 5: Average pricing error per option for the GEV RND surface, Black Scholes and Mixture of Lognormals, when pricing both calls and puts, all strikes and maturities simultaneously.

One of the three parameters that we obtain from the estimation of the GEV RND surface is the implied tail shape parameter ξ . The time series of the implied tail shape parameter was first studied in Markose and Alentorn (2005), but the estimator suffered from maturity dependence effects, exhibiting a jump on the contract-switch date, and a downward trend as the time horizon shortened. The implied tail shape parameter estimated using the maturity corrected GEV option pricing model in this paper is more appropriate to analyse its time series behaviour. Recall that when we have a positive ξ the implied RND of the losses is of Fréchet type, and has thicker than normal tails. On the other hand, when we have a negative ξ the implied RND is of Weibull type and has thinner than normal tails. Our period of study, 1997 to 2003, includes some events, such as the Asian crisis, the LTCM crisis and the 9/11 attacks, which resulted in a sudden fall of the underlying FTSE 100 index, and can be useful in analyzing the market reaction. We can see from Figure 6 below that ξ was negative for most part of 1997, and jumped to positive values around the Asian crisis. Similarly, its value considerably increased around the LTCM crisis in September 1998. It stayed negative for most part of 2000 and 2001, at an average of -0.72, except on the days following the 9/11 events, where it increased to around +0.1.

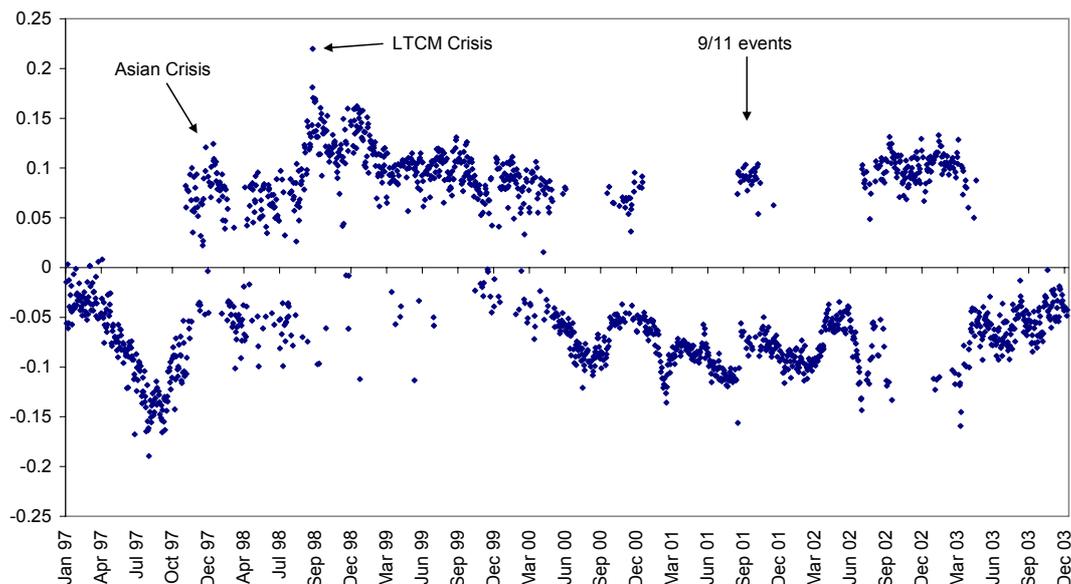


Figure 6: Time series of implied tail shape parameter ξ . Crisis periods where the tail shape parameter substantially increased are indicated.

Another parameter we obtain from the estimation of the GEV RND surface is the implied scaling of volatility b , and Figure 7 below shows its time series. Under the Black-Scholes model, this parameter is assumed to take a value of 0.5 (i.e. scaling volatility with the square root of time). The implied scaling by the GEV RND surface for the period of study is on average 0.52, but its daily time series ranges from 0.24 to 0.74. When testing on a daily basis the null hypothesis:

$$H_0 : b = 0.5$$

to test if b is significantly different than 0.5 we find that we can reject the null hypothesis 86.4 % of the days, at the 95% confidence level.¹⁴

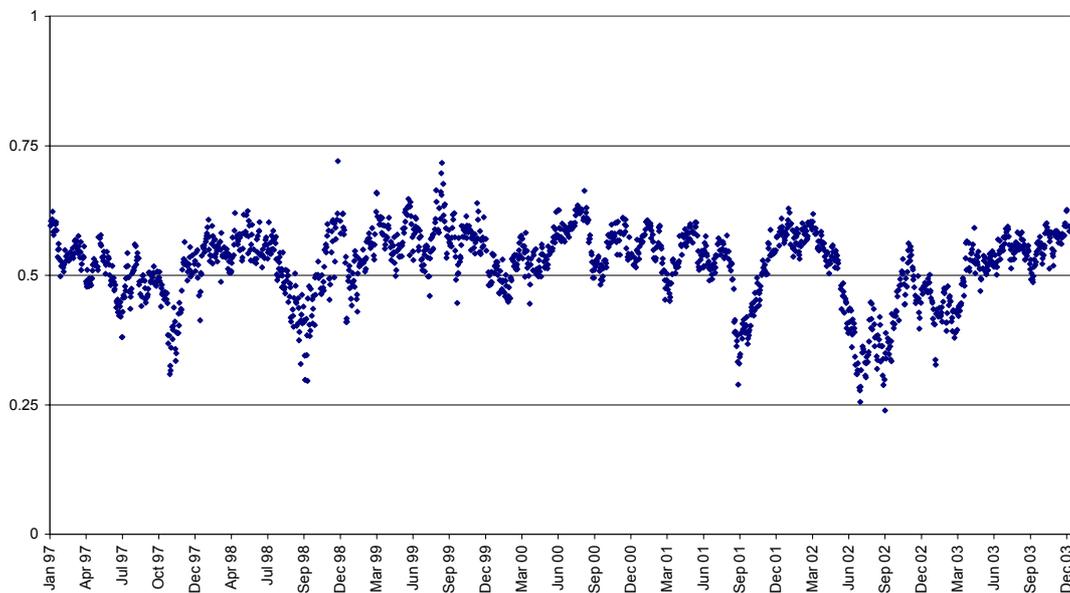


Figure 7: Time series of implied scaling parameter b . This parameter captures how the implied volatility scales with time.

To further test the statistical significance of b , we re-estimated the model by fixing $b = 0.5$ (thus, removing one degree of freedom). The average pricing error was 18.25, a substantial increase from 11.32, the error obtained with the GEV model that allows b to differ from 0.5. The error obtained when fixing $b = 0.5$ is even worse than the Black-Scholes error (see Table 5).

¹⁴ The standard error for the estimator was calculated from the Hessian matrix obtained from the non-linear squares algorithm in Matlab. Following the methodology in Andersson and Lomakka (2005), we analyzed the residuals of our estimation (the pricing errors). When tested by a Bera-Jarque test, we were able to reject normality at the five percent significance level. Additionally, we tested whether the errors were independent of strike price, and correlations within some groups were significant. Therefore, since the errors are neither normally distributed nor independent, the parameter estimates are unbiased, but the estimate of the covariance matrix is inconsistent. Therefore, the results from the hypothesis testing should be interpreted with care.

3.3 Applications of the implied GEV RND surface

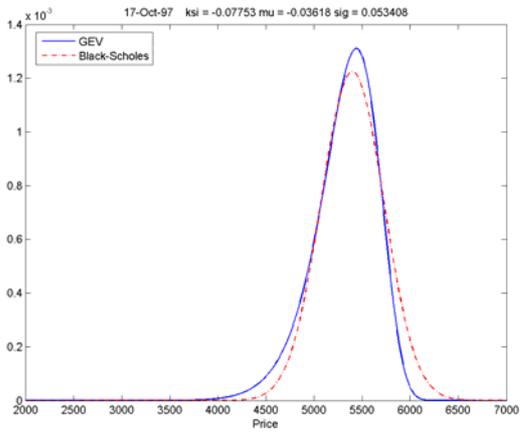
The model for the implied GEV RND surface presented in this paper can be used in applications that need to extract information from option prices. Specifically, applications that need to correct the time to maturity effect, such as event studies, or the construction of a time series of constant time horizon implied statistics.

First, we will look at how the GEV RND surface can be used to remove the maturity effects when conducting before/after type event studies. We can use the GEV RND surface method to obtain two RNDs with the same horizon, one using option prices before the event, and one using option prices after the event. Then, any changes in the RNDs can be attributed to the reaction of the market to the event, rather than having this effect mixed with the change of the RND due to the shortening of time to maturity after the event. Table 6 below displays the GEV implied parameters before and after each of the three crisis events in our period of study, and the implied moments of the returns distribution. As we have previously seen in Figure 6, the implied tail shape parameter ξ increased considerably after the events, in some cases going from negative to positive. The implied volatility, skewness and kurtosis also increased substantially after the events. Note that the implied higher moments of the RND differ substantially from the moments of the normal distribution (skewness of 0 and kurtosis of 3) even before the events.

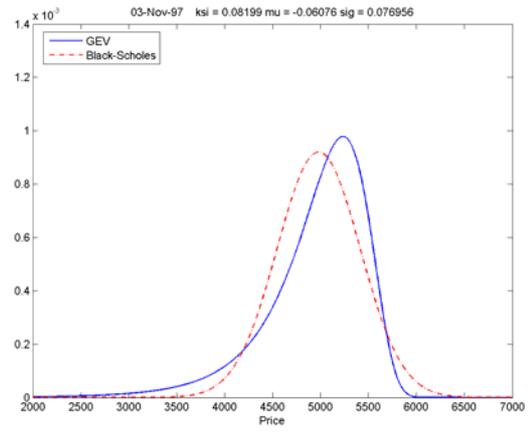
| Event | Date | σ | ξ | b | Implied Volatility | Implied Skewness | Implied Kurtosis |
|--------------|-----------|----------|--------|-------|--------------------|------------------|------------------|
| Asian Crisis | 17-Oct-97 | 0.169 | -0.077 | 0.463 | 19.8% | 0.74 | 3.84 |
| | 3-Nov-97 | 0.208 | 0.082 | 0.400 | 30.1% | 1.74 | 9.36 |
| LTCM | 8-Sep-98 | 0.251 | 0.098 | 0.479 | 37.4% | 1.89 | 10.78 |
| | 23-Sep-98 | 0.232 | 0.219 | 0.424 | 44.5% | 4.12 | 68.53 |
| 9/11 | 7-Sep-01 | 0.214 | -0.113 | 0.492 | 24.2% | 0.58 | 3.43 |
| | 12-Sep-01 | 0.202 | 0.094 | 0.389 | 29.8% | 1.85 | 10.42 |

Table 6: GEV implied parameters and moments of the GEV based RND around crisis events. The implied moments have been calculated using a 1 year horizon, to obtain annualized values.

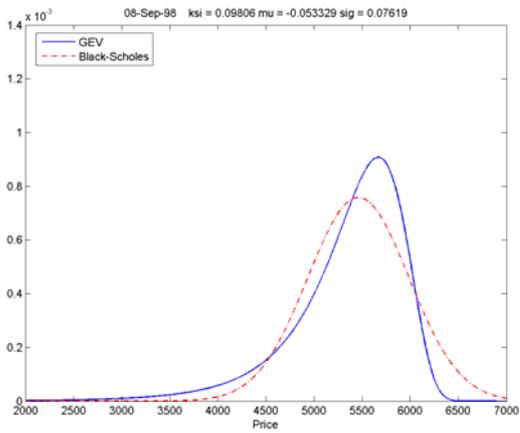
We can see a graphical representation of the changes of the RNDs before and after the events in Figure 8. The blue lines display the GEV based RND for the price distribution, while the red dotted lines display the corresponding Black-Scholes RND for comparison purposes. We can see how the non-gaussian characteristics of the RND become more accentuated after the events, with the left tail becoming thicker, and the distribution more skewed.



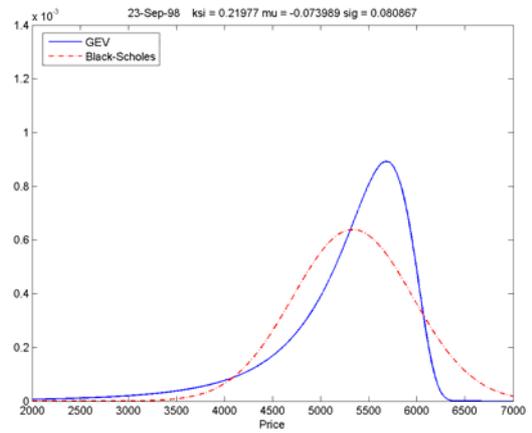
a) Before Asian crisis



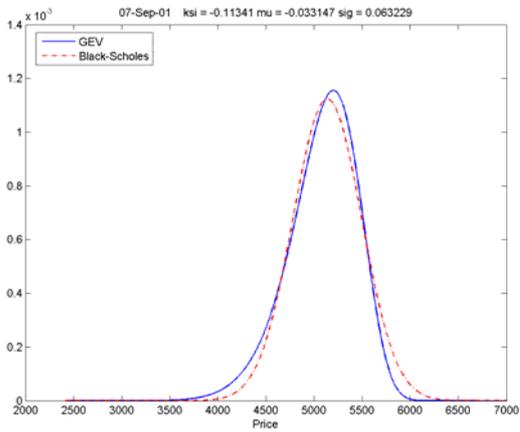
b) After Asian crisis



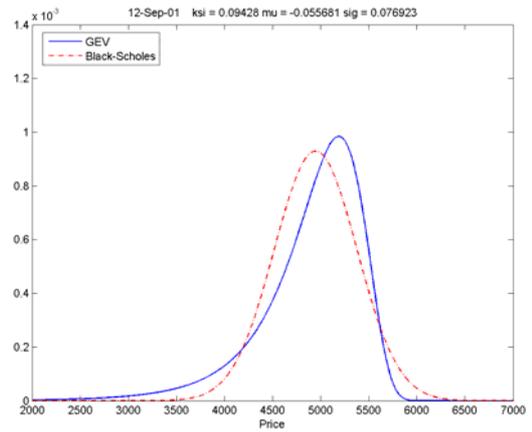
c) Before LTCM crisis



d) After LTCM crisis



e) Before 9/11 events



f) After 9/11 events

Figure 8: Implied risk neutral densities with a 30 day horizon before and after each of the three events.

A second application for the model of the implied GEV RND surface is in calculating constant time horizon EVaR. When comparing the EVaR estimates obtained from the GEV RND surface with the estimates obtained from the scaling law in Section 2, we find that they are very similar to each other, and also very similar to the BoE EVaR time series. The estimates obtained using the GEV RND surface are slightly less volatile than the ones obtained using the empirical scaling law. The reason for this may be that the estimation of the implied parameters for the GEV RND surface is more robust than the estimation of the implied parameters for single maturity RNDs, since there is a larger number of observations to estimate the same number of parameters in the former than in the later.

The last application of the implied GEV RND surface we will consider is in the construction of an implied volatility index, like the VIX. In 1993 the Chicago Board Options Exchange (CBOE) introduced the CBOE Volatility Index (VIX), which represents the expected stock market volatility over the next 30 calendar days, implied by S&P 500 index option prices¹⁵. VIX represents a hypothetical option that is at-the-money and has a constant 22 trading days (30 calendar days) to expiry. The VIX calculation uses the two nearest-term expiration months in order to bracket a 30-day horizon. It calculates the implied volatility for each of the two maturities, and then it interpolates them to arrive at a single value with a constant maturity of 30 days. The VIX has become the benchmark for stock market volatility, and a gauge of market uncertainty, often called “the investor’s fear gauge”. Similar indices also exist for other markets such as the NDX for the NASDAQ, the VDAX for the German Futures and Options Exchange (VDAX) and the VX1 and VX6 for the French Exchange MONEP.

The GEV RND surface can be used as an alternative method of calculating the VIX. Instead of calculating a fixed horizon implied volatility by interpolating between implied volatilities from two maturities, one could obtain the implied volatility for a fixed horizon using equation (18), with the implied parameters of the implied RND surface. Figure 9 below shows the time series of the 30 calendar day implied volatility for the FTSE 100 index obtained using the GEV RND surface. From 1997 to 2003, the implied volatility reached its highest value around the LTCM crisis in September 1998. Other periods of high volatility can be identified around the Asian Crisis (October 1997) and the 9/11 events in September 2001.

¹⁵ Initially the VIX calculation used options on the S&P 100, but in 2003 the calculation algorithm was modified and options on the S&P 500 are used since then.

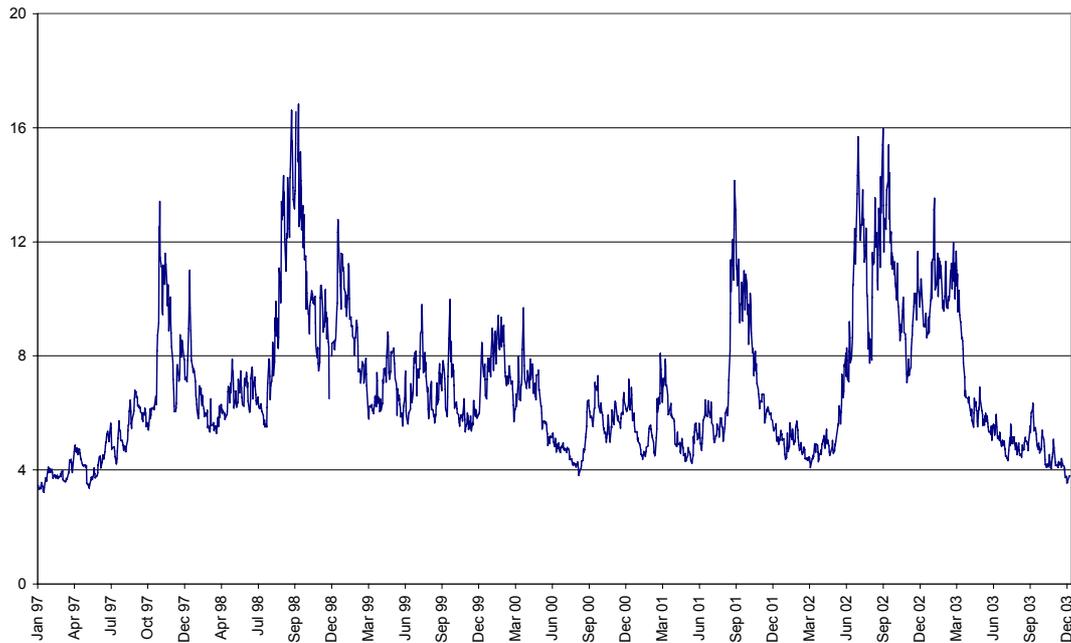


Figure 9: Implied volatility series for a constant time horizon of 30 calendar days implied by FTSE 100 index option prices, using the GEV model

An advantage of using the GEV RND surface to calculate a time series of implied volatilities, over the VIX method is that with the GEV RND surface one can obtain a time series of implied volatilities for any constant time horizon, while the VIX methodology can only calculate it for time horizons equal or greater than the closest maturity date, because it needs options with maturities that bracket the required horizon.

4. Conclusions

In this paper we have proposed two new methods for removing the maturity effects from implied RNDs and related statistics. The first method was based on an empirical scaling law for EVaR in the quantile space, obtained by exploiting the linear behaviour of EVaR with time horizon in a log-log scale. We confirmed this linear behaviour of EVaR for three different parametric RND extraction methods. Comparing the results with the BoE EVaR estimates, the mixture of lognormal method was found to overestimate EVaR, while the Black-Scholes method was found to underestimate it. The EVaR estimates we obtained with the GEV method are very similar to the BoE estimates, even though the estimation techniques are very different. Recall that the BoE is a non-parametric method, and interpolation is done in the delta space, whereas the GEV method is a parametric method, and interpolation is done in the quantile space. These findings are partially in line with the existing literature, such as Melick and Thomas (1997) where it was found that EVaR estimates are very sensitive to the RND extraction technique used. Some advantages of our method with respect to the

BoE one are that our method is simpler; it delivers an empirical scaling law for EVaR, and allows us to extrapolate outside the range of available maturities.

The second method for removing maturity effects was based on modelling an implied RND surface across maturities and strikes. We extended the original GEV model in Markose and Alentorn (2005), to explicitly model the scaling of implied volatility with a new parameter. EVaR estimates from this second method were very close to the ones obtained in the first method. But this second method has many more applications. We showed how the implied tail shape parameter we obtain from this modified model does not suffer from maturity effects, and a time series of this parameter can be used to study how the market views on the probabilities of extreme moves changed around crisis periods or special events. We also showed how the implied GEV RND surface can be used to construct an implied volatility index, similar to the VIX, in event studies and to obtain constant time horizon EVaR estimates.

References

Ait-Sahalia, Y., Lo, A.W., (2000). Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9–51.

Ait-Sahalia, Y., Wang, Y., Yared, F., (2001). Do option markets correctly price the probabilities of movement of the underlying asset? *Journal of Econometrics* 102, 67–110.

Ait-Sahalia and Lo (2000)

Alentorn, A. Markose, S. (2005) “Extreme Economic Value at Risk: EEVaR”, Mimeo, Center for Computational Finance and Economic Agents, University of Essex. Available at: <http://privatewww.essex.ac.uk/~alent/>

Anagnou, I., Bedendo, M., Hodges, S. D. and Tompkins, R. G., "The Relation Between Implied and Realised Probability Density Functions" (February 28, 2002). EFA 2002 Berlin Meetings Presented Paper; University of Warwick Financial Options Research Centre Working Paper.

Bates, D.S., 1991. The Crash of '87: Was it expected? The evidence from options market. *Journal of Finance* 46, 1009–1044.

Bali, T. G. (2003). “The generalized extreme value distribution”. *Economics Letters* 79, 423-427.

Black, F. and M. Scholes (1973). “The Pricing of Options and Corporate Liabilities”. *Journal of Political Economy* 81, 637-659.

Butler, C and H Davies (1998): “Assessing Market Views on Monetary Policy: the Use of Implied Risk Neutral Probability Distributions”. Working Paper, Monetary Instruments and Markets Division, Bank of England.

Chicago Board Options Exchange (CBOE) (2003), “The New CBOE Volatility Index—VIX,” available at www.cboe.com/micro/vix/vixwhite.pdf.

Clews, R., Panigirtzoglou, N., Proudman, J., 2000. Recent developments in extracting information from option markets. *Bank of England Quarterly Bulletin* 40, 50–60.

Cont, R. and da Fonseca, J. (2002) Dynamics of implied volatility surfaces. *Quantitative Finance* 2 45-60

- Daniélsson, J. and C.G. de Vries (1997) "Tail index estimation with very high frequency data", *Journal of Empirical Finance*, 4, 241-257.
- Daniélsson, J. and J. Zigrand (2004). "On time-scaling of risk and the square-root-of-time rule." *FMG Discussion Papers*. dp439. London School of Economics and Political Science
- De Haan, L., D.W. Cansen, K. Koedijk and C.G. de Vries (1994) "Safety first portfolio selection, extreme value theory and long run asset risks. Extreme value theory and applications" (J. Galambos et al., eds.) 471-487. Kluwer, Dordrecht.
- Dowd, K. (2002). *Measuring Market Risk*. Wiley Press.
- Dumas, B., Fleming J. and Whaley R.E. (1998) "Implied Volatility Functions: Empirical Tests", *The Journal of Finance*, Vol LIII, No. 6, 2059-2106
- Embrechts, P., C. Klüppelberg and T. Mikosch (1997) *Modelling Extremal Events for Insurance and Finance*. Berlin: Springer.
- Embrechts P., S. Resnick and G. Samorodnitsky (1999). "Extreme value theory as a risk management tool" *North American Actuarial Journal* 3, 30-41.
- Embrechts, P. (2000) "Extreme Value Theory: Potential and Limitations as an Integrated Risk Management Tool" *Derivatives Use, Trading & Regulation*, 6, 449-456.
- Fama, E. (1965). "The behaviour of stock market prices." *Journal of Business*. **38**. 34-105.
- Fleming, Je, Barbara Ostdiek, and Robert E. Whaley (1995), "Predicting Stock Market Volatility: A New Measure," *Journal of Futures Markets*, vol. 15, 265-302.
- Fleming (1998) "The quality of market volatility forecasts implied by S & P 100 index option prices" *Journal of Empirical Finance* 5 (1998) 317-345
- Gemmil G, and A. Saflekos (2000) "How Useful are Implied Distributions? Evidence from Stock-Index Options" *Journal of Derivatives* 7, 83-98.
- Härdle, W. and Hlávka, Z. (2005) "Dynamics of State Price Densities". *SFB 649 Discussion Paper* 2005-021
- Jackwerth, J. C. (1999). "Option-Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review". *The Journal of Derivatives*, 7, 66-82.
- Kaufmann, R. and P. Patie (2003). "Strategic Long-Term Financial Risks: The One-Dimensional Case." *RiskLab Report*. (ETH Zurich).
- Malz, A. (2001) "Do implied volatilities provide early warning of market stress?" *Journal of Risk* **3**, 2, Winter 2000/2001
- Mandelbrot, B. (1963). "The Stable Paretian Income Distribution when the Apparent Exponent is Near Two." *International Economic Review*. **4**. (1). 111-115.
- Mandelbrot, B. (1967). "The Variation of Some Other Speculative Prices." *Journal of Business*. **40**. 393-413.
- Markose, S. and Alentorn, A. (2005). "The Generalized Extreme Value (GEV) Distribution, Implied Tail Index and Option Pricing". University of Essex, Economics Department Working paper number 594.

- Melick, W. R. (1999). "Results of the estimation of implied PDFs from a common data set". Proceedings of the Bank for International Settlements Workshop on 'Estimating and Interpreting Probability Density Functions', Basel, Switzerland.
- Melick, W. R. and C. P. Thomas (1996). "Using Option prices to infer PDFs for asset Prices: An application to oil prices during the gulf war crisis". International Finance Discussion Paper, No. 541, Board of Governors of the Federal Reserve System.
- Melick, W R and C P Thomas (1998): "Confidence intervals and constant maturity series for probability measures extracted from option prices". Paper presented at the conference 'Information Contained in Prices of Financial Assets', Bank of Canada.
- Menkens, O. (2004). "Value-at-Risk and Self Similarity". Preprint. Centre for Computational Finance and Economic Agents. Department of Economics. University of Essex.
- McNeil, A.J. and R. Frey. (2000). "Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach". *Journal of Empirical Finance*. 7: 271-300.
- Müller, Ulrich A. & Dacorogna, Michel M. & Olsen, Richard B. & Pictet, Olivier V. & Schwarz, Matthia, (1990) "Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis," *Journal of Banking & Finance*, Elsevier, vol. 14(6), pages 1189-1208.
- Newey, W. K. and West, K. D. (1987), "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix", *Econometrica*, 55, 703-708.
- Ncube M. (1996) "Modelling implied volatility with OLS and panel data models", *Journal of Banking and Finance* 20 71-84
- Panigirtzoglou, N. and G. Skiadopoulos (2004) "A New Approach to Modeling the Dynamics of Implied Distributions: Theory and Evidence from the S&P 500 Options" *Journal of Banking and Finance*, 28, 1499-1520.
- Peña, I., Serna, G, and G. Rubio (1999). "Why do we smile? On the determinants of the implied volatility function", *Journal of Banking and Finance* 23 1151-1179
- Ritchey, R. J. (1990) "Call Option Valuation for Discrete Normal Mixtures". *Journal of Financial Research*, 13, 285-296.
- Shiratsuka, S. (2001) "Information Content of Implied Probability Distributions: Empirical Studies on Japanese Stock Price Index Options", IMES Discussion Paper 2001-E-1, Institute for Monetary and Economic Studies, Bank of Japan.
- Tompkins, R., (2001) "Implied volatility Surfaces: Uncovering Regularities for Options on Financial Futures", *European Journal of Finance*, 7, 198-230.
- Vähämaa, S. (2005) "Option-implied asymmetries in bond market expectations around monetary policy actions of the ECB" *Journal of Economics and Business* 57, 23-38.

Appendix A: Using WLS to improve the estimation of the linear scaling law

The linear regression suggested in Section 2 of this paper to scale EVaR can be affected by EVaR values that have been calculated from an RND estimated with very few option prices. The EVaR estimates in such cases will have very wide confidence intervals. As an example, take the 12 Nov 97 (see Figure A.1 below) at the 95% confidence level. The R^2 of the OLS regression was 64.8%, a very poor fit as can be seen from Figure A.1 below. The EVaR value obtained for the furthest away maturity was obtained from an RND estimated using only 4 option prices, and thus the confidence intervals of the EVaR estimate are much wider than the EVaR values obtained for closer maturities, which are based on RNDs extracted using around 25 contracts.

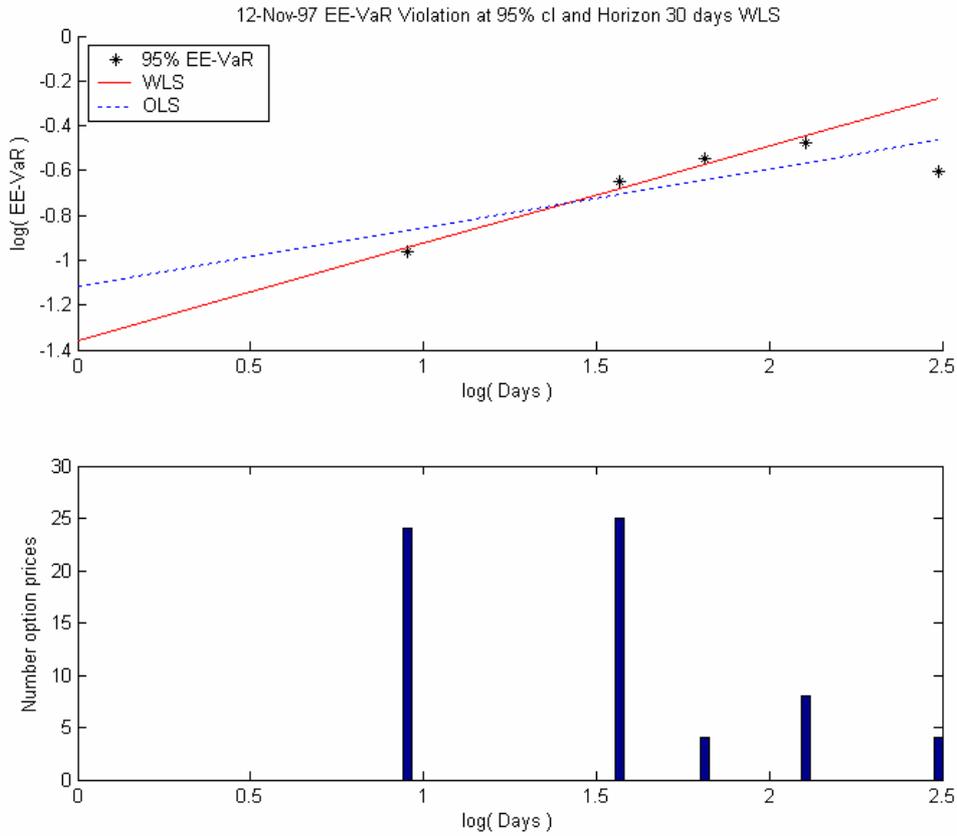


Figure A.1: Example of linear regression using OLS vs. WLS for a day when there are some maturities with very few option prices available.

One method to solve this issue is to use a *Weighted Linear Squares* (WLS) regression, using the number of option prices available at each maturity to weight the EVaR values. The Weighted RMSE is given by:

$$WRMSE = \frac{1}{N} \sqrt{\sum_{i=T_1}^{T_N} w_i (y_i - \hat{y}_i)^2} \quad \text{with} \quad \sum_{i=T_1}^{T_N} w_i = 1, \quad (\text{A.1})$$

where the weights w are calculated as the percentage of option prices in each maturity

$$w_i = \frac{\text{Number of prices at maturity } i}{\text{Total number of prices}} \quad (\text{A.2})$$

The weighed R^2 is given by

$$\text{Weighted } R^2 = 1 - \frac{\sum_{i=T_1}^{T_N} w_i (y_i - \hat{y}_i)^2}{\sum_{i=T_1}^{T_N} w_i (y_i - \bar{y}_i)^2} \quad (\text{A.3})$$

Table A.1 below shows the average weighted R^2 at each confidence level for the GEV model. Note how the fitting performance increases with confidence level, while is lowest at 96.8% for the lowest quantile of 70%.

| Confidence level | 70% | 80% | 90% | 95% | 99% |
|-------------------------|------------|------------|------------|------------|------------|
| OLS R^2 | 87.9% | 97.9% | 98.8% | 98.7% | 97.9% |
| WLS R^2 | 96.8% | 99.4% | 99.6% | 99.5% | 99.3% |

Table A.1: Average weighted R^2 at each quantile

The linear fit of the 3 parametric models is shown in the table A.2 below. The GEV has the best linear fit, while the non-parametric method has the worst linear fit.

| RND method | WLS R^2 |
|----------------------|-----------------------------|
| GEV | 99.7% |
| Mixture lognormals | 98.7% |
| Black-Scholes | 98.5% |
| Non-parametric (BoE) | 69.95 |

Table A.2: Average WLS R^2 of the linear regression in the quantile space for the 95% confidence level over 1733 days (1997-2004)