

**Evaluation of Power EWMA VaR Models for Bank
Portfolio Conservativeness, Accuracy and
Efficiency**

EFM codes : 450, 510, 520

ABSTRACT

To measure the risk inherent in their trading activities, major banks have increasingly employed value-at-risk (VaR) models. The successful implementation of VaR depends on the accurate estimation of the conditional distribution of portfolio returns. While the normality assumption is widely used to forecast VaR, the asset returns are typically found to be fat-tail distributed. Thus the VaR estimators based on the normal distribution are inefficient and lead to the true value of the risk being underestimated. Therefore, by using the power exponential distribution instead, we introduce a family of nested power EWMA estimators that are robust to fat-tailedness in the conditional distribution of returns. We also use data on the actual trading book portfolio of a large bank in Taiwan, and compare the empirical performance of various VaR models. The assessment of a model's performance is based on a range of measures that address the conservativeness, accuracy and efficiency of each model. The backtesting results demonstrate that, due to the flexibility of the power parameters of the conditional distribution, the power exponential distribution can properly capture the fat-tailedness characteristic of the asset return distributions. Most of the family of EWMA estimators that are based on the power exponential distribution outperform those VaR estimators that are based on the normal distribution, and offer an appropriate coverage of the extreme risk.

Key Words: Fat-tail, VaR, Power Exponential Distribution, Power EWMA.

1. Introduction

Increased competition has forced banks to search for more income at the expense of added risk. Over the past decade, banks have enlarged the scope of their trading activities in the highly-volatile financial markets. The rapid increase in the relative importance of market risk has also led to a number of unfavorable events, as exemplified by Orange County, Barings and LTCM, etc. Since a single bank crash could lead to a crisis engulfing the whole of the banking system, supervisory institutions have been spurred by such experiences to reconsider the guidelines related to capital requirements that were agreed upon in the 1988 Basle Capital Accord which was designed primarily to deal with credit risk and had noticeable drawbacks in that it neglected market risk. Thus, the 1996 Amendment to the Basle Accord to incorporate market risk, which remained effective in the 2000 Basle II, has required banks to maintain levels of capital adequacy to cover the market risk in their trading accounts. The capital requirements for market risk can be estimated through the use of a VaR model generated by the banks' internal risk management models. Banks whose models significantly underestimate market risks suffer penalties in that their capital requirements are adjusted upwards. Accordingly, it is important for both regulators and banks to understand the strengths and weaknesses of different modeling approaches.

Two broad types of VaR analysis are commonly employed. First, under the parametric VaR approach, the distribution of asset returns is estimated using historical data under the assumption that the distribution is a member of a given parametric class. The most common approach is to suppose that the returns are stationary, joint normal and independent over time. By using the estimated variance and covariance, VaR can be calculated with a given level of confidence. Second, the simulation approach to VaR analysis consists of using historical data that extends over a long period, in which the VaR is set equal to the percentile of the empirical distribution at the required level of confidence. As a non-parametric procedure, the latter imposes no distributional assumptions. However, a number of empirical studies show that asset returns are not normally distributed. In particular, the conditional distribution of short horizon asset returns is found to be leptokurtic, with tails that are significantly fatter than those of the normal distribution (see Mandelbrot, 1963; Fama, 1965; Baillie and de Gennaro, 1990; Jansen and de Vries, 1991; Bollerslev, Chou and Kroner, 1992; Koedijk and

Kool, 1994; Loretan and Phillips 1994; Kearns and Pagan 1997). The models that unconditionally employ the normality assumption (used in JP Morgan's RiskMetrics[®] model) will be inefficient and will understate the true value of risk if the asset returns are fat-tail distributed.

To remedy this problem, the analysis proceeds in two main directions in order to characterize the tail behavior. The first is to set up the unconditional distribution as a mixture of a normal distribution and another kind of distribution such as a normal-Poisson (Jorion, 1988), a normal-lognormal (Hsieh, 1989) or a Bernoulli-normal distribution (Vlaar and Palm, 1993), whereby the assumption of homoscedasticity is preserved, i.e. the volatility of asset returns is time-independent. The second is to use a non-normal distribution, for instance, a Student's *t*-distribution (Bollerslev, 1987; Baillie and Bollerslev, 1989; Kaiser, 1996; Beine, Laurent and Lecourt, 2002), a Laplace and a double exponential distribution (Linden, 2001) or an exponential power distribution (Varma, 1999; Guermat & Harris, 2002) to capture the fat-tailed nature of most asset returns.

A vast number of studies have addressed the weaknesses and strengths of various forms of VaR modeling (Drudi, et al., 1997; Hendricks, 1996 and Jackson, et al., 1997). To date, no single consistent measure of VaR model performance has been developed. An issue of concern to supervisors is whether the required minimum regulatory-capital calculated by the internal model of the bank can offer an appropriate coverage for its losses. Alternatively, there is the issue of the conservativeness of the model. We identify relatively conservative models as those that systematically produce higher estimates of risk in comparison with other models. With regard to accuracy, the risk manager should be concerned with whether the model's ex-post performance is compatible with the theoretically desired level. The regulatory capital-adequacy framework also provides an incentive to develop efficient models, that is, models that offer enough coverage in relation to the risk so that the supervisors' requirements can be met with the minimum amount of capital that is required to be held.

Most of these VaR studies are based on simulated portfolios, rather than the actual portfolios held by market practitioners. Our use of actual portfolios distinguishes our research from other similar studies. The main advantage of using actual book values for the predominant bank trading risks is that they ensure that the pattern of risk exposures along the yield curve and across markets is realistic and can clearly reflect

a bank's investment decisions. Randomly generated portfolios are unlikely to be representative.

In this paper, in addition to the basic VaR models, that include the variance-covariance method, historical simulation and Monte-Carlo simulation widely used by most banks, we introduce the family of nested power EWMA estimators that are based on the exponential power distribution and are robust to fat-tailedness and leptokurtosis in the conditional distribution of returns to forecast the VaR. The bank that supplied us with the data had sizeable equity, interest-, and FX-rate exposures. We then focus on three aspects of the models – conservativeness, accuracy and efficiency – and propose a range of statistics based on these criteria to compare the performances of the different models.

The paper is set up as follows. In the next section we explain the methodology, by first describing the bank's portfolio and return data, and then illustrating the tail-index estimation. Subsequently, we introduce the nested power EWMA variance estimators that have been used, and set out the parametric and simulated VaR approaches. To assess the relative performances of the different models, a range of measures based on different viewpoints is proposed. Section 3 presents the results of the empirical evaluation. Some concluding remarks are offered in Section 4.

2. Data and Methodology

2.1 Trading Book and Return Data

Our empirical evaluation employs data on the trading books¹ of a bank with significant trading exposure in Taiwan. Since the capital charges in the internal VaR model for interest rate-related instruments and equities apply only to the current market value of items in banks' trading books, the banking book assets of this particular bank are not included in our portfolio. The bank's portfolio, which is provided to us on condition of anonymity, consists of equity securities, foreign exchange² and fixed-income securities. The positions and related risk factors of the

¹ The trading book means the bank's proprietary positions in financial instruments which are intentionally held for short-term resale and/or which are taken on by the bank with the intention of benefiting in the short-term from actual and/or expected differences between their buying and selling prices, or from other price or interest-rate variations, and positions in financial instruments arising from matched principal brokering and market making, or positions taken in order to hedge other elements of the trading book. The others comprise the banking book (Basle,1996).

² The foreign exchange assets only include foreign currencies since our data bank excludes the foreign

bank's trading book portfolio are summarized in Table 1. The holdings of the bank portfolio for equities, fixed-income securities and foreign exchange are 13.72%, 83.84% and 2.44%, respectively. The risk factors are the market variables from which the value of all other instruments can be derived. For equities, foreign exchange and fixed-income securities, we deal with 73, 13 and 19 different sources of risk, respectively. There are in total 105 different sources of risk included in the bank's portfolio. For equities and foreign currencies, the risk factors are the same as their value or price. The price or value series data used comprise daily observations in relation to the stock prices, the net present value of mutual funds and the foreign exchange rates. However, the value of fixed-income securities must be derived indirectly from the pricing model and the yield to maturity obtained from the yield curves of fixed-income securities³. We employ the market risk factors data from the Taiwan Economic Journal (TEJ) databank, for which the yield curves were originally compiled by Reuters. Throughout the study we use the NT dollar as the base currency and employ the data for the period from January 1, 1998 to December 31, 2002. We obtain the continuously compounded returns by calculating the changes in the natural logarithms of each price or value series.

financial assets data.

³ Refer to Marrison (2002) pp.53-60, 89, 99.

Table1 The Bank's Portfolio and Risk Factors

Thousands of NT Dollars

Trading Book Assets	Positions	Risk Factors
Equity Position Risk	7,543,617 (13.72)	
Stocks (42)	5,031,832 (9.15)	Stock Prices
Mutual Funds (31)	2,511,785 (4.57)	Net Value of Mutual Funds
Interest Rate Risk	46,099,787 (83.84)	
Bills (668)	43,824,065 (79.70)	Different terms of Interest Rates (1 day, 10 day, 90 day, 180 day, 1 year, 2 year, 3 year, 5 year, 7 year, 8 year, 10 year, 13 year, 15 year)
Bonds (29)	2,275,722 (4.14)	
Foreign Exchange Risk		
Foreign Currencies (19)	1,342,817(2.44)	Foreign Exchange Rates
Total (770)	54,986,221(100.00)	

Note: The quantities and composition ratios of the various assets are reported in parentheses.

2.2 Tail Index Estimation

Fat tails are somewhat arbitrarily defined. Intuitively, a fat-tailed distribution is a distribution that assigns more weight to the tails (i.e. extremal events) than a normal distribution. For instance, a six-sigma event has a near zero probability in a Gaussian distribution but might have non-negligible probability in a fat-tailed distribution. The tail index is a measure of the degree of tail fatness of the underlying distribution, and the estimation involves extreme value theory (EVT) that addresses the characteristics of the tail behavior of the distribution. The most famous and commonly applied estimator of the tail index, due to its easy implementation and asymptotic unbiasedness, is that proposed by Hill (1975) as follows:

$$\xi(m) = \left[\frac{1}{m} \sum_{i=1}^m \ln(x_{n-i+1}) \right] - \ln(x_{n-m}) \quad , \quad m \geq 2 \quad (1)$$

$$= \frac{1}{m} \left[\sum_{i=1}^m \ln\left(\frac{x_{n-i+1}}{x_{n-m}}\right) \right]$$

where ξ is the Hill estimator which is the inverse of the tail index α , and m is the pre-specified number of tail observations to be included. The selection of m is crucial for obtaining unbiased estimators of the tail index, n is the sample size, and x_i is the i th increasing order statistic ($i=1,2,\dots,n$). Equation (1) illustrates that the Hill estimator measures the average of the ratios of each observed value relative to the threshold value in the predetermined tail area. The larger the average, the smaller the tail index and the greater the magnitude of the fat-tailedness.

However, there is considerable empirical evidence to show that the Hill estimator is biased in relatively small samples and limited to the cases in which a larger sample is available. To improve the Hill estimator, a recently-developed alternative approach proposed by Huisman, et al. (2001) has been especially useful for small samples. Their regression-based approach is based on an approximation of the asymptotic expected value of m :

$$E(\xi(m)) \approx \frac{1}{\alpha} - cm \quad (2)$$

where c is a constant depending on the parameters of the distribution and the sample size. If m decreases, then the bias decreases and the expectation approaches the true value ξ . The variance of the estimator increases as m decreases:

$$Var(\xi(m)) \approx \frac{1}{m\alpha^2} \quad (3)$$

The idea put forward by Huisman, et al. (2001) is to use equation (2) in a regression analysis and to regress the values of $\xi(m)$ on m as follows:

$$\xi(m) = \beta_0 + \beta_1 m + \varepsilon(m), \quad m = 1 \dots \kappa \quad (4)$$

The estimated $\hat{\beta}_0$ is an estimator of $\xi = \frac{1}{\alpha}$. Huisman, et al. propose choosing a threshold value κ equal to half the sample ($\frac{n}{2}$). Although the parameters in equation (4) can be estimated by means of ordinary least squares (OLS), equation (3) indicates

that the variance in relation to the Hill estimator is not constant for different values of m . The error term $\varepsilon(m)$ is heteroscedastic. Accordingly, they propose a weighted least squares (WLS) approach to correct for the heteroscedasticity and improve the efficiency of the estimator. We apply the modified Hill estimator to obtain the tail index estimates by using both OLS and WLS.

2.3 The Power EWMA Variance Estimator

We introduce a general power EWMA estimator proposed by Guermat and Harris (2002), which nests EWMA models that are more robust to the leptokurtosis of returns, and which would therefore be expected to be more efficient when the conditional distribution of returns is fat-tailed. The power EWMA estimator is based on the maximum likelihood estimator of the variance of the power exponential distribution (also known as the generalized error distribution, or Box-Tiao distribution). The probability density function of the power exponential distribution is given by

$$f(r, \sigma, \delta) = \left(\frac{\delta}{\varphi 2^{\frac{\delta}{\delta+1}} \Gamma(1/\delta) \sigma} \right) e^{-\frac{1}{2} \left| \frac{r}{\varphi \sigma} \right|^{\delta}} \quad (5)$$

where

$$\varphi = \left\{ \frac{2^{-\frac{2}{\delta}} \Gamma(1/\delta)}{\Gamma(3/\delta)} \right\}^{\frac{1}{2}} \quad (6)$$

and $\Gamma(\cdot)$ is the gamma function. The power exponential distribution has a variance equal to σ^2 , zero skewness and a kurtosis coefficient that depends on the value of the power parameter δ . When $\delta = 2$, the power exponential distribution reduces to the normal distribution. When $\delta > 2$, the power exponential distribution is thin-tailed and platykurtic, and when $\delta < 2$, the power exponential distribution is fat-tailed and leptokurtic. When $\delta = 1$, the power exponential distribution reduces to the Laplace distribution. The power exponential distribution with different values of the power parameter δ is shown in Figure 1. From Figure 1 we can find that, the smaller the power parameter δ is, the more fat-tailed and leptokurtic the distribution becomes.




Figure 1. Power Exponential Distribution with Different Values of the Power Parameter δ

Being conditional upon the power parameter δ , the maximum likelihood estimator of the standard deviation of the power exponential distribution is given by

$$\hat{\sigma}^{\delta} = g(\delta) \frac{1}{T} \sum_{t=1}^T |r_t|^{\delta} \quad (7)$$

where

$$g(\delta) = \delta \left[\frac{\Gamma(3/\delta)}{\Gamma(1/\delta)} \right]^{\frac{\delta}{2}} \quad (8)$$

Equation (7) is the unconditional variance estimator that is independent of past information. Guermat and Harris (2002) transform it into a conditional variance estimator and replace the unweighted average in (7) with an exponentially weighted average to yield the power EWMA estimator

$$\sigma_{t+1}^k = (1 - \lambda) g(k) \sum_{i=0}^{\infty} \lambda^i |r_{t-i}|^k \quad (9)$$

By recursive substitution, the power EWMA estimator can be rewritten as

$$\sigma_{t+1}^k = \lambda \sigma_t^k + (1 - \lambda) g(k) |r_t|^k \quad (10)$$

so that the power EWMA estimator can be seen as an infinite weighted average of past squared returns, incorporating information from all past shocks into the power parameter k of returns, but with exponentially declining weights. Alternatively, by using the fact that $r_{t+1}^2 = \sigma_{t+1}^2 + \varepsilon_{t+1}^2$ where ε_{t+1}^2 is a zero-mean random shock that is orthogonal to the time t information set, the power EWMA estimator can also be interpreted as an infinite order autoregressive model for the k th powered return. When

$k = 2$, the power EWMA estimator coincides with the standard EWMA estimator given by

$$\sigma_{t+1}^2 = \lambda \sigma_t^2 + (1 - \lambda) r_t^2 \quad (11)^4$$

The standard EWMA estimator is a special case of the generalized autoregressive conditional heteroscedasticity, or GARCH, model (Engle, 1982; Bollerslev, 1986). The GARCH(1,1) model for the conditional variance of returns is given by

$$\sigma_{t+1}^2 = \alpha_0 + \alpha_1 \sigma_t^2 + \beta_1 r_t^2 \quad (12)$$

where α_0 , α_1 and β_1 are the parameters to be estimated. When $\alpha_0 = 0$ and $\beta_1 = 1 - \alpha_1$, the GARCH model reduces to the standard EWMA estimator, and is alternatively known as Integrated GARCH or IGARCH. When $k < 2$, the power EWMA estimator is more sensitive to extreme observations and thus we may expect it to be more efficient when the conditional distribution of returns is leptokurtic. The role of the function $g(k)$ is to preserve the integrated nature of the volatility, in keeping with the standard EWMA model.

When $k = 1$, the power exponential reduces to the Laplace distribution, and the power EWMA estimator reduces to

$$\begin{aligned} \sigma_{t+1} &= (1 - \lambda) \sqrt{2} \sum_{i=0}^{\infty} \lambda^i |r_{t-i}| \\ &= \lambda \sigma_t + (1 - \lambda) \sqrt{2} |r_t| \end{aligned} \quad (13)^5$$

The Laplace distribution is commonly used in the context of robust estimation, and so the EWMA estimator given by (13) might therefore be thought of as a ‘robust’ EWMA estimator. The power EWMA estimator therefore nests the standard EWMA estimator, the robust EWMA estimator, and a continuum of estimators that lie between the two, as well as estimators that are even more sensitive to outlying observations than the standard EWMA estimator, and those that are even less

⁴ when $k = 2$, $g(k) = 2 * \left[\frac{\Gamma(3/2)}{\Gamma(1/2)} \right]^{\frac{1}{2}} = 2 * \left[\frac{\Gamma(\frac{1}{2} + 1)}{\Gamma(\frac{1}{2})} \right] = 2 * \left[\frac{\frac{1}{2} * \Gamma(1/2)}{\Gamma(1/2)} \right] = 1$

⁵ when $k = 1$, $g(k) = 1 * \left[\frac{\Gamma(3/1)}{\Gamma(1/1)} \right]^{\frac{1}{2}} = \left[\frac{\Gamma(2+1)}{\Gamma(1)} \right]^{\frac{1}{2}} = \left[\frac{2\Gamma(2)}{\Gamma(1)} \right]^{\frac{1}{2}} = 2^{\frac{1}{2}} = \sqrt{2}$

sensitive to them than the robust EWMA estimator. The power EWMA estimator described above is a special case of the NGARCH model of Higgins and Bera (1992) given by

$$\sigma_{t+1}^k = \alpha_0 + \alpha_1 \sigma_t^k + \beta_1 r_t^k \quad (14)$$

When $\alpha_0 = 0$ and $\beta_1 = (1 - \alpha_1)g(k)$, the NGARCH model reduces to the power EWMA estimator. The relationships among these nested models may be summarized in Figure 2. The parameters of the models for the bank portfolio returns are first estimated by means of the maximum likelihood approach based on the power exponential distribution, using the BHHH algorithm (Berndt, et al., 1974) with a convergence criterion of 0.00001 being applied to the function value. Next, we apply the likelihood ratio statistics⁶ to test the restrictions of the NGARCH model that are implied by the power EWMA estimators and introduce the standard EWMA based on the normal distribution (well known as RiskMetrics) as the benchmark to compare the out-of-sample performance of various power EWMA estimators.

⁶ Under regularity, the large sample of the likelihood ratio test statistic $\chi^2 = -2\ln\lambda$ is chi-squared, with degrees of freedom equal to the number of restrictions imposed, where $\lambda = L_R / L_U$, L_U denotes the likelihood function evaluated without regard to constraints, and L_R is the constrained likelihood function estimate. Since both likelihoods are positive, a restricted optimum is never superior to an unrestricted one, so that $0 < \lambda < 1$. If λ is too small, then doubt is cast on the restriction, otherwise if λ is close to 1, the null hypothesis is accepted (see Greene (2000)).

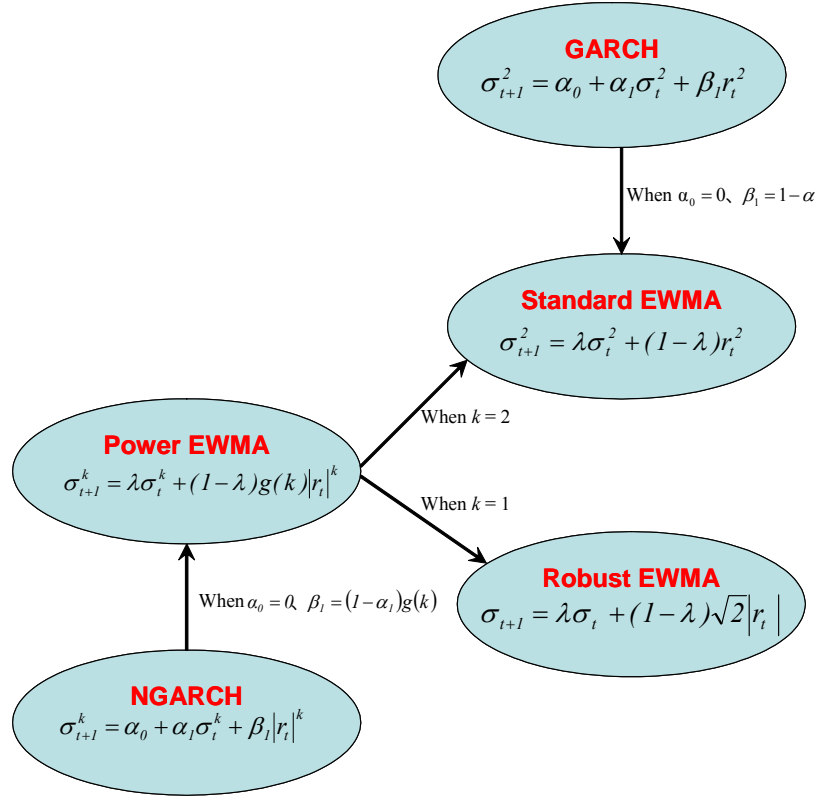


Figure 2. The Nested Power EWMA Models

2.4 The Estimation of VaR

2.4.1 Parametric Analytical Approach

We compute out-of-sample one-day VaR forecasts for the bank's portfolio, using each of the EWMA estimators as previously introduced. The VaR of the portfolio in each period t is forecast by the formula

$$VaR_{t+1} = -\delta(\alpha) \sigma_{t+1} \quad (15)$$

where σ_{t+1} is the standard deviation of the portfolio's return, r_{t+1} , that is conditional upon the time t information set, and $\delta(\alpha)$ is the α -quantile of the standardized (i.e. zero mean, unit variance) empirical power exponential distribution and α is the VaR significance level. The standardized empirical distribution is defined as the return series over the window, scaled by the estimated standard deviation for each of those days. The standard deviation estimate used to standardize the return is obtained from the EWMA model (see Hull and White, 1998).

2.4.2 Historical Simulation Approach

Historical simulation is a simple, theoretical approach that simply assumes that

market price innovations in the future are drawn from the same empirical distribution as those market price innovations generated historically. The VaR is set equal to the percentile of the observed daily return distribution at the required level of confidence. The main drawback of this historical simulation approach is that extreme percentiles are difficult to estimate precisely without a large sample of historical data.

2.4.3 Monte-Carlo Simulation Approach

In contrast to the historical simulation approach that uses simulated portfolio returns drawn directly from historical data, Monte-Carlo simulation methods interpose a stochastic process that is simulated through a computer application in order to generate a large number of possible portfolio-return outcomes. The Monte-Carlo process takes into account the events that were not observed over the historical period but may possibly occur in the future. To implement the Monte-Carlo method, the stochastic processes of the asset returns must be selected at the outset. We employ the Monte-Carlo approach based on the assumption that the asset returns are characterized by a joint normal distribution.⁷ From the hypothetical distribution, 5,000 scenarios are simulated to obtain 5,000 outcomes of the values of the bank's portfolio. Based on the distribution of the resulting changes in the portfolio values, the appropriate percentile is then determined to derive the VaR estimates. Since our approach employs the same assumptions regarding distributions as the simple variance-covariance method, it can be expected that the results will be close to those obtained from the fixed weighted variance-covariance method. All the models used to estimate the VaRs and the distributions they are based on are presented in Table 2.

A rolling window is used for the estimation of each model. Each estimated model is then used to forecast the VaRs of the portfolios with window lengths of 10, 50, 100, 250, 500, and 1,000⁸ observations for the 99% confidence levels, respectively.

⁷ Most banks operate Monte-Carlo simulation methods based on the assumption that the asset returns are normally distributed. Given the sizeable evidence that financial asset returns are nonnormal and have tails that are fatter than those in the case of normal distributions, this assumption is indeed questionable. Other multivariate distributions, for instance the *t*-distribution, which can capture the fat-tailed nature of most return series, are therefore preferred. However, more complicated models may make simulation difficult to implement (see Engel and Gizycki (1999)).

⁸ According to the guidelines of the Basel Accord, the confidence levels need to be 99% and the window lengths at least 250 business days. Our EWMA VaR estimates show that when the window lengths are greater than 100 days, the discrepancies among estimated VaRs for different window lengths are so small that they can be neglected. Consequently, a window length of 250 observations is sufficient. It is not necessary to extend the window lengths any further.

Table 2 The Models Used and their Base Distributions

	NGARCH	Power EWMA	Standard EWMA	Robust EWMA	RiskMetrics®	Historical Simulation	Monte-Carlo Simulation
Base Distribution	PED	PED	PED	PED	Normal	NONE	Normal

Note: **PED** denotes exponential power distribution.

2.5 Model Evaluation

The evaluation of VaR forecasts is not straightforward. As with the evaluation of volatility forecasting models, a direct comparison between the forecast VaR and the actual VaR cannot be made, since the latter is unobservable. A variety of evaluation methods have been proposed (see, for instance, Kupiec, 1995; Christofferson, 1998; Lopez, 1998). Up to now, no single definition of VaR model performance has been developed. To evaluate the performance of this family of models, we propose a range of statistics that address different aspects of the usefulness of VaR models to risk managers and the supervisory authorities. We focus on three aspects of the models: conservativeness, accuracy and efficiency (Engel and Gizycki, 1999).

The most important issue of concern to supervisors is whether the required minimum regulatory-capital calculated by the internal model of the bank can offer an appropriate coverage for its losses. Alternatively, the supervisors may prefer the conservative models⁹. As for the risk managers, they should be concerned with the accuracy of the model - whether the model's ex-post performance is compatible with the theoretically desired level. The regulatory capital-adequacy framework also provides an incentive to develop efficient models, that is, models that offer enough coverage in relation to the risk so that the supervisors' requirements can be met with the minimum amount of capital that is required to be held.

2.5.1 Conservativeness

Mean Relative Bias (MRB)

We follow Engel and Gizycki (1999) to identify relatively conservative models as those that systematically produce higher estimates of risk in comparison with other models. They defined the conservativeness of models in terms of the relative size of the VaR in relation to the risk assessment. The larger the VaR value was, the more

⁹ We can confirm this argument from the backtesting guidelines of the internal model method according to the Basle Accord (1996).

conservative the model became. To measure the relative size of the VaR among different models, they applied the mean relative bias developed by Hendricks (1996). The mean relative bias statistic captures the degree of the average bias of the VaR of the specific model from the all-model average. Given T time periods and N VaR models, the MRB of model i can be calculated as:

$$MRB_i = \frac{1}{T} \sum_{t=1}^T \frac{VaR_{it} - \overline{VaR}_t}{\overline{VaR}_t} \quad (16)$$

$$\text{where, } \overline{VaR}_t = \frac{1}{N} \sum_{i=1}^N VaR_{it}$$

However, the MRB measure is in terms of the relative but not absolute¹⁰ concept. If the evaluated models included are different, then we might obtain different results regarding the relative conservativeness of the models.

2.5.2 Accuracy

Different users of the VaR model will focus on different types of inaccuracies. It may be expected that supervisors will pay more attention to the underestimation of losses while the financial institutions will be more concerned with over-predictions of losses due to the capital adequacy requirement. In this study, we define accuracy as the extent to which the rate of failure of the specific model is close to the preset significance level. The four accuracy measures, namely, the binary loss function, the LR test of unconditional coverage (Kupiec, 1995), the LR test of independence and the scaling multiple to obtain coverage, are presented as follows:

a. Binary Loss Function (BLF)

The binary loss function is based on whether the actual loss is larger or smaller than the VaR estimate. Here we are simply concerned with the number of failures rather than the magnitude of the exception. If the actual loss $\Delta P_{i,t+1}$ is larger than the VaR then, it is termed an “exception” (or failure) and has a value equal to 1, with all others having a value of 0. That is

¹⁰ Based on the absolute conservativeness concept, the conservative models can be defined as the models that can provide higher VaR estimates in comparison with the theoretically desired level. In other words, conservativeness means the overestimation of the true risk of the bank portfolio.

$$L_{i,t+1} = \begin{cases} 1 & \text{if } \Delta P_{i,t+1} < VaR_{i,t} \\ 0 & \text{if } \Delta P_{i,t+1} \geq VaR_{i,t} \end{cases} \quad (17)$$

The aggregate number of failures across all dates is divided by the sample size. The BLF is obtained as the rate of failure. The closer the BLF value is to the confidence level of the model, the more accurate the model is.

b. LR Test of Unconditional Coverage (LR_{uc})

The BLF provides a point estimate of the probability of failure. In other words, the accuracy of the VaR model requires that the BLF on average be equal to one minus the prescribed confidence level of the VaR model. The model should provide the correct unconditional coverage of loss. Kupiec (1995) proposed a likelihood ratio test based on the binomial process that can be applied to determine if the rate of failure is statistically compatible with the expected level of confidence. Given the sample size T and the frequency of failure N governed by a binomial probability, the likelihood ratio statistic of the unconditional coverage hypothesis $H_0 : p = \alpha$ can be stated as

$$LR_{uc} = -2\ln[(1-p)^{T-N} p^N] + 2\ln\left[\left(1-\frac{N}{T}\right)^{T-N} \left(\frac{N}{T}\right)^N\right] \sim \chi_{1,\alpha}^2 \quad (18)$$

Under the null correct hypothesis of correct unconditional coverage, the LR_{uc} has a chi-squared distribution with one degree of freedom.

c. LR Test of Independence (LR_{ind})

If a VaR model accurately captures the conditional distribution of returns, as well as its dynamic properties such as time-varying volatility, then exceptions should be unpredictable, and hence independently distributed over time. To test the independence of the exceptions of a VaR model, Christofferson (1998) has derived an LR statistic as follows:

$$H_0 : \pi_{01} = \pi_{11} = \pi$$

$$LR_{ind} = 2(\ln L_u - \ln L_R) \sim \chi_{1,\alpha}^2 \quad (19)$$

where

$$L_U = (1 - \pi_{01})^{T_{00}} \pi_{01}^{T_{01}} (1 - \pi_{11})^{T_{10}} \pi_{11}^{T_{11}}, \quad L_R = (1 - \pi)^{T_{00}+T_{10}} \pi^{T_{01}+T_{11}},$$

$$\pi_{01} = T_{01} / (T_{00} + T_{01}) , \quad \pi_{11} = T_{11} / (T_{10} + T_{11}) , \quad \pi = (T_{01} + T_{11}) / (T_{00} + T_{01} + T_{10} + T_{11}) ,$$

and T_{ij} denotes the number of times that state i is followed by state j , where state 0 is defined as that in which the actual portfolio loss is less than the estimated VaR and state 1 is defined as that in which the actual return is greater than the estimated VaR. Under the null hypothesis that exceptions are independently distributed, the LR statistic has a chi-squared distribution with one degree of freedom.

d. Multiple to Obtain Coverage (MOC)

To highlight the magnitude of the deviation in the losses from the VaR estimate, we compare it with the multiple to obtain the coverage proposed by Engel and Gizycki (1999). The multiple equivalent, X_i , of the risk measure for model i is calculated so that

$$F_i = T_i \alpha , \quad F_i = \sum_{t=1}^{T_i} \begin{cases} 1 & \text{if } \Delta P_{i,t+1} < X_i VaR_{i,t} \\ 0 & \text{if } \Delta P_{i,t+1} \geq X_i VaR_{i,t} \end{cases} \quad (20)$$

where F_i is equivalent to the total number of failures, T_i is the sample size and α is the significance level of the model. $\Delta P_{i,t+1}$ denotes the realized profit or loss on day $t+1$. The closer the **MOC** value is to one, the more accurate is the model.

2.5.3 Efficiency

Mean Relative Scaled Bias (MRSB)

Efficiency is important since VaR measures are used by both the supervisory authorities and the internal management of financial institutions to influence investors' incentives. A more efficient VaR model provides more precise resource allocation signals to the financial institutions. Hence we address the aspect of efficiency in terms of the ability of a model to provide adequate risk coverage with a minimum capital outlay. The Mean Relative Scaled Bias (MRSB) (Engel and Gizycki (1999)) aims to evaluate which model, once the desired risk coverage level is suitably obtained, produces the smallest VaR measure. There are two steps involved in

calculating the MRSB measure. First, the scaling should be calculated by multiplying the VaR for each model by the multiple needed to obtain the 99% coverage as described in the MOC measure. Subsequently, we compare the scaled VaR measures with the average relative size in relation to the all-model average. The MRSB measure is arrived at as follows:

$$MRSB_i = \frac{1}{T} \sum_{t=1}^T \frac{X_i \cdot VaR_{i,t} - \overline{X \cdot VaR_t}}{\overline{X \cdot VaR_t}} \quad (21)$$

$$\text{where, } \overline{X \cdot VaR_t} = \frac{1}{N} \sum_{i=1}^N X_i \cdot VaR_{i,t}$$

3. Empirical Results

3.1 Descriptive Statistics

The preliminary statistics for all four asset returns¹¹ are summarized in Table 3a to 3b. As is commonly found, daily financial asset returns are not normally distributed and exhibit leptokurtic behavior. In all cases, the Jarque-Bera tests for normality are highly significant with excess kurtosis,¹² and negatively skewed except for the bond returns that are positively skewed. Moreover, bill returns exhibit the most leptokurtic distribution with kurtosis of 27.11. The standard deviation in the case of equity returns is the largest of all whereas that in relation to bill returns is the smallest, which shows that equity risk is the major market risk source of the sampled bank. The asset returns on equity securities are negatively correlated with those on interest rate instruments and foreign exchange, while the asset returns on interest rate instruments are positively correlated with those on foreign exchange as theoretically expected.

Table 3a Descriptive Statistics

	Equities	Foreign Exchange	Bills	Bonds
Mean	-0.0003	-4.53E-05	2.18E-05	-0.0001

¹¹ As stated previously, we deal with 105 different risk factors included in the bank's portfolio. The author has performed the Jarque-Bera tests for normality for all asset returns. Most of them reject the normality hypothesis significantly with excess kurtosis. According to the major components of the bank portfolio, we choose the asset returns in the case of the TSEC Capitalization Weighted Stock Index, US dollar, the interest rate on 180-day commercial paper, and the yield-to-maturity of 5-year government bonds in the secondary market, as the representative asset returns for equities, foreign exchange, bills and bonds.

¹² Excess kurtosis means the distribution has kurtosis greater than that of the normal distribution whose kurtosis is equal to 3.

Maximum	0.0637	0.0219	0.0012	0.0333
Minimum	-0.0655	-0.0205	-0.0012	-0.0267
Standard Deviation	0.0181	0.0028	0.0001	0.0041
Skewness	0.1817 (0.0000)	0.1194 (0.0000)	2.1933 (0.0000)	-0.0154 (0.0000)
Kurtosis	3.8717 (0.0000)	19.3690 (0.0000)	27.1098 (0.0000)	17.6594 (0.0000)
Jarque-Bera¹³	48.31 (0.0000)	14717.78 (0.0000)	32928.86 (0.0000)	11604.54 (0.0000)
Sample Size	1300	1318	1316	1296

Note: P-values are reported in parentheses for the skewness, kurtosis and the Jarque-Bera statistic.

Table 3b Correlation of Asset Returns Based on Risk Factor Classes

	Equity	Foreign Exchange	Interest Rate
Equity	1.0000	-0.0630	-0.0114
Foreign Exchange	-0.0630	1.0000	0.0794
Interest Rate	-0.0114	0.0794	1.0000

3.2 Estimation of Tail Index

The most important information in terms of characterizing the limiting extreme distribution of the tail is the tail index. The indexes for the right-tail, left-tail and both tails¹⁴ for the four asset returns are estimated by means of OLS and WLS. The modified Hill estimate results are shown in Table 4. From Table 4 we can find that the estimates of the normal distribution, as a benchmark, are all around 8.5. All tail index estimates, which vary between 1.98 and 7.80, are less than 8.5. This indicates that all the asset return distributions exhibit fatter tails than the normal distribution, as is commonly found in the literature, with those related to bonds having the highest degree of fat-tailedness. If these estimates differ significantly over both tails, it is inappropriate to use the estimate based on the combined information of both. However, our results show that the left-tail estimates exhibit slightly fatter tails than

¹³ The Jarque-Bera statistic has a chi-squared distribution with two degrees of freedom under the null hypothesis of normality.

¹⁴ All observations are taken in excess of their sample means. The left tail is examined by using the absolute value of all negative returns, while the right tail is examined by using the absolute value of all positive returns. To examine both tails simultaneously, we use the absolute values of all returns. (see Huisman, et al. (1998))

the right-tail estimates for the majority of asset returns.

Table 4 Tail-index Estimates

	Equities	Foreign Exchange	Bills	Bonds	Normal
Both tail_{OLS}	6.4276 (0.0000)	2.4880 (0.0000)	5.3841 (0.0000)	2.3973 (0.00000)	8.4189 (0.0000)
Both tail_{WLS}	5.8927 (0.0000)	2.3997 (0.0000)	5.2093 (0.0000)	1.9894 (0.00000)	—
Left tail_{OLS}	5.6578 (0.0000)	2.6081 (0.0000)	4.6892 (0.0129)	2.8019 (0.00000)	8.4976 (0.0000)
Left tail_{WLS}	5.8370 (0.0000)	2.3780 (0.0000)	4.8780 (0.0000)	1.9991 (0.0000)	—
Right tail_{OLS}	7.8021 (0.0000)	2.4918 (0.0000)	5.0045 (0.0000)	4.2396 (0.00000)	8.74279 (0.0000)
Right tail_{WLS}	6.4541 (0.0000)	2.7310 (0.0000)	4.1085 (0.0014)	4.2918 (0.0000)	—

Note: P-values are reported in parentheses for the estimate of the tail index $\alpha = \frac{1}{\xi}$

3.3 Estimation of Power EWMA

The NGARCH model, and the power EWMA estimator that it nests, are based on the maximum likelihood estimators of the variance of the power exponential distribution. Table 5 presents the parameter estimates of each model for the bank's portfolio return series. The first column of each panel in Table 5 gives the results for the unrestricted NGARCH model. The sum of the parameters α_1 and α_2 is 0.9722. It is close to but not equal to 1. The second column reports the results for the power EWMA estimator which imposes the restrictions $\alpha_0=0$ and $\alpha_2=1-\alpha_1$. The power parameter, k , of 1.4077 is closer to unity than to two, which suggests that the robust EWMA estimator might be expected to perform better than the standard EWMA estimator. The estimated power parameter of the conditional distribution, δ , is 1.7809 in the case of the NGARCH model, 1.7608 for the power EWMA model, 1.7507 for the standard EWMA model and 1.6526 for the robust EWMA model. That all of these are smaller

than 2 confirms again our previous findings that all the asset returns exhibit fat-tailed and leptokurtic distributions. Besides, the estimated power parameter of the conditional distribution for the family of EWMA models is in each case close to the estimated power parameter, k , of the conditional variance model, which is consistent with the results of previous studies (Nelson and Foster, 1994; Guermat and Harris, 2002). The estimated decay factor in the case of the standard EWMA model is 0.9476, and is very close to the value of 0.94 suggested by JP Morgan.

Table 5 Estimates of EWMA

	NGARCH	Power EWMA	Standard EWMA	Robust EWMA
α_0	0.0000 (0.8162)	[0]	[0]	[0]
α_1	0.9027 (0.0000)	0.9428 (0.0000)	0.9476 (0.0000)	0.9348 (0.0000)
α_2	0.0694 (0.0006)	$[1 - \alpha_1]$	$[1 - \alpha_1]$	$[1 - \alpha_1]$
k	1.4978 (0.0313)	1.4078 (0.0000)	[2]	[1]
δ	1.7809 (0.0000)	1.7608 (0.0000)	1.7507 (0.0000)	1.6526 (0.0000)
LOGL	5808.0118	5804.3378	5802.9887	5797.1850

Notes: 1. The restricted parameter values imposed on NGARCH are reported in square brackets [].

2. P-values are reported in parentheses () for the parameter estimates.

3. LOGL is the maximum value of the log likelihood function.

3.4 Restrictions Test on the Nested Model

The power EWMA estimator is nested by the NGARCH model, and therefore imposes certain restrictions on the NGARCH model. In this section, we test whether those restrictions are supported by the data. Table 6 reports the likelihood ratio tests in relation to the various restrictions. Table 6 shows that, owing to the precision with which the NGARCH parameters are estimated, the null hypothesis that the sum of the NGARCH parameters is unity can be rejected at the 1% significance level. This result

suggests that while the true data generating process is not quite integrated, the sum of the estimated parameters is close to unity and so, over a short horizon, their dynamic properties should be reasonably well described by an integrated NGARCH model, or power EWMA process.

The third and fourth rows report the results for the standard EWMA estimator and the robust EWMA estimator. These models impose restrictions on the power EWMA estimator in that $k = 2$ and $k = 1$, respectively. On the basis of the likelihood ratio tests reported in Table 6, the standard EWMA model cannot be rejected while the robust EWMA estimator can be rejected at the 1% significance level. Only the restriction $k = 2$ imposed on the EWMA models can be supported by the sample asset returns. The constrained model, the standard EWMA, might be a more appropriate one for our data.

Table 6 LR Test of Restrictions on the Nested Models

	H_0	LR Statistics
Power EWMA vs. NGARCH	$\alpha_0 = 0$ $\alpha_1 + \alpha_2 = 1$	7.3486 (0.0000)
Standard EWMA vs. Power EWMA	$k = 2$	2.6983 (0.1005)
Robust EWMA vs. Power EWMA	$k = 1$	11.6074 (0.0000)

Note: P-values are reported in parentheses for the LR estimates.

3.5 Model Evaluation

We then check the results of the bank's backtesting of its internal measurement system (i.e. by comparing value-at-risk estimates with actual profits and losses) to ensure that the model provides a reliable measure of potential losses over time. The backtesting results of our models are exhibited in Figures 3a to 3c. Differing from the other two methods – the RiskMetrics model, and the Monte-Carlo simulation approach, the historical simulation approach without the normality assumption produces a different volatility pattern of estimated VaR and has a significant 'window

effect¹⁵, in Figure 3a as expected. Due to the same normality assumption, the variance-covariance approach, i.e. the RiskMetrics model, and the Monte-Carlo simulation approach exhibit similar volatility patterns for the VaR estimates. From Figure 3b we can find that the VaR series of EWMA models based on the power exponential distribution also exhibit similar volatility patterns to each other. Among these models, the VaRs estimated by the Robust EWMA model, which are located at the bottom, provide the highest VaR estimates, whereas the VaRs estimated by the RiskMetrics model located at the top provide the lowest VaR estimates. That confirms our argument that the asset returns are typically found to be fat-tail distributed. The VaR estimators based on the normal distribution might be inefficient and lead to the underestimation of the true value of the risk.

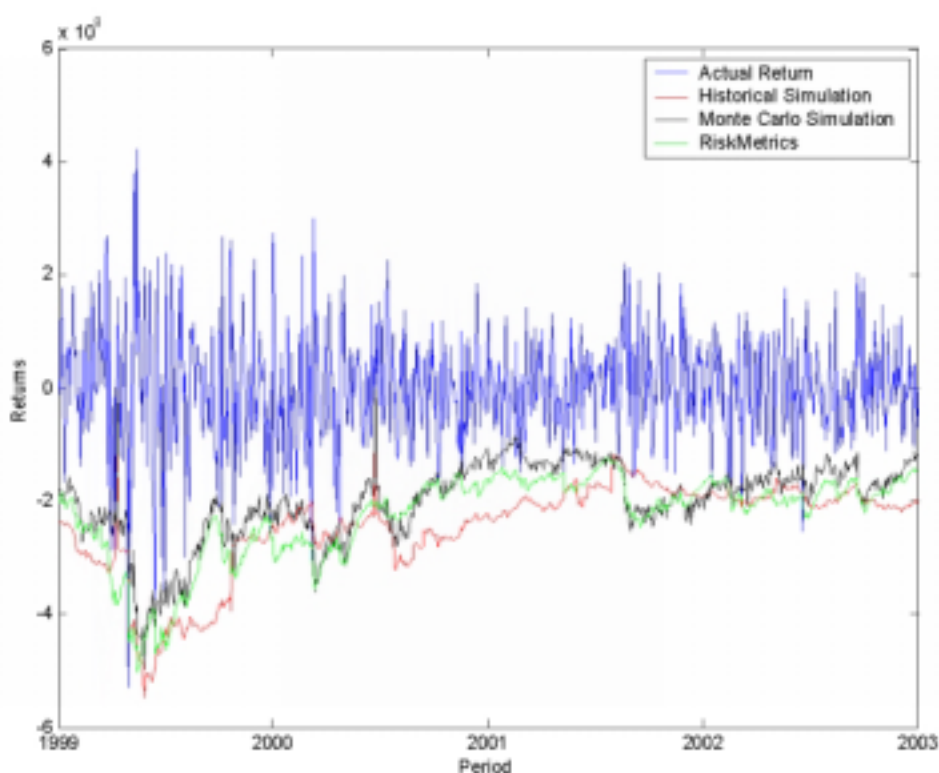


Figure 3a Backtesting of Three Common Models

¹⁵ The reported VaR suddenly increases from one day to the next when some recent, specific crisis observation moves into the rolling window for historical data and suddenly drops when it drops out of the window.

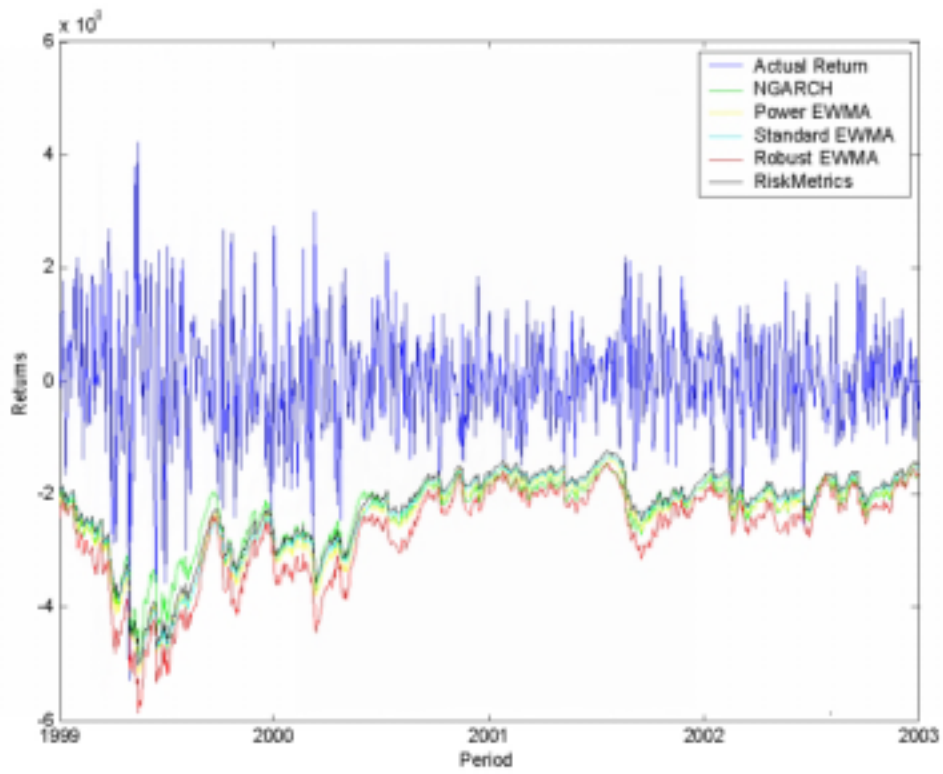


Figure 3b Backtesting of the Power EWMA Models




Figure 3c Backtesting of All Models

The estimated performance measures for each model in the case of the 99th VaR are summarized in Table 7. The mean relative bias that tends to fall between 0.1101 and -0.0899 indicates that there is little difference in the magnitude of the risk measure across the models. In terms of the concerns of the supervisory authorities, the most conservative model is the robust EWMA model, which produces the largest average VaR estimate, while the Monte-Carlo simulation is the least conservative model that produces the lowest average VaR estimate.

With regard to the accuracy of the measures, our findings are as follows. The BLFs for the robust EWMA, NGARCH, power EWMA, historical simulation and standard EWMA models are all lower than the benchmarks of 1%, indicating that these models overstate risk. On the contrary, the BLFs for the Monte-Carlo simulation and RiskMetrics[®] models are higher than 1%, indicating that these models understate risk. As theoretically expected, the RiskMetrics[®] and Monte-Carlo simulations based on the normal distribution appear to be less sensitive to the outlying observations than other models based on the exponential power distribution. The most conservative model, namely, the robust EWMA model, produces the least BLF in contrast to the least conservative model, namely, the Monte-Carlo simulation, which produces the greatest BLF. The standard EWMA model provides the BLF which is closest to the benchmark, suggesting that the standard EWMA model is the most accurate model in terms of the BLF measure.

In terms of the 99th VaR estimate, the multiples needed to obtain coverage (MOC) are all less than one except for the RiskMetrics[®] and Monte-Carlo simulation models that are both based on the normal distribution. Put another way, the RiskMetrics[®] and Monte-Carlo simulation models underestimate risk while the other models over-predict risk. That should result from the discrepancy between the fat-tailedness of the empirical return distribution and the normality assumption of the two models. The standard EWMA model needs the multiple closest to one and requires the least adjustment to provide the appropriate coverage, while also achieving the highest accuracy. The other models in terms of accuracy are the historical simulation model, the power EWMA model, the NGARCH model, the RiskMetrics[®] model, the robust EWMA model and the Monte-Carlo simulation in that order. On the whole, we can find that the MOC accuracy measures provide the same results as the BLF accuracy measures.

With the exceptions of the robust EWMA model and the Monte-Carlo simulation, the VaR measures give rise to rates of failure (i.e. BLF) close to the benchmark figure of 1%. The results of the likelihood ratio test of unconditional coverage further provide statistical evidence that only the robust EWMA and Monte-Carlo simulation models reject the null hypothesis (the number of observed exceptions is statistically consistent with the expected level of confidence). For the majority of models, the null hypothesis is accepted, which indicates that most of our models are accurate.

From Table 7, we find that the independence LR statistic across various models, LR_{ind} , is not available except in the case of the Monte-Carlo simulation. The reason for this is that once the number of times, T_{11} , in the case of an exception state followed by another exception state, is equal to 0, this would lead to both π_{11} and L_u being equal to 0. In such a situation, the independence LR statistic cannot be well-defined. That there are non-consecutive exceptions in most models results in the unavailability of an independence LR statistic.

In sum, the results of evaluating the models' accuracy highlight the fact that the standard EWMA model achieves the most accurate results in terms of different accuracy measures and most of our models are accurate. As expected, the models based on the normal distribution, RiskMetrics® and the Monte-Carlo simulation, appear to be less sensitive to the outlying observations than other models based on the exponential power distribution, and also lead to the risk of the bank portfolio being underestimated.

In comparing the MRSB across all models, we find that the historical simulation model is the most efficient in terms of providing the 99th percentile coverage, whereas the Monte-Carlo simulation method performs the least efficiently. The historical simulation method performs well, yielding good accuracy and efficient estimates. It should be noted that simple VaR model does not substantially under-perform relative to the more complex models.

As stated previously, the evaluation viewpoints of the VaR models in terms of internal risk management, capital charges for regulation and external supervision are quite different. Nevertheless, we argue that accuracy is still the most important criterion of model evaluation. The backtesting results also demonstrate that, due to the flexibility of the power parameters of the conditional distribution, the power exponential distribution can properly capture the fat-tailedness characteristic of the asset return

distributions. Most of the family of EWMA estimators based on the power exponential distribution outperform those VaR estimators that are based on the normal distribution, and offer an appropriate coverage of the extreme risk.

Table 7 The Performance Measures of VaR Models
(99% VaR window length=250)

Models	NGARCH	Power EWMA	Standard EWMA	Robust EWMA	RiskMetrics [®]	Historical Simulation	Monte-Carlo Simulation
Conservativeness							
MRB	-0.0154	-0.0035	-0.0283	0.1101	-0.0630	0.1096	-0.0899
Accuracy							
BLF	0.64%	0.64%	0.90%	0.26%	1.41%	0.75%	2.19%
MOC	0.9706	0.9804	0.9921	0.8868	1.0644	0.9832	1.2714
LR_{uc}	1.1561	1.1561	0.0838	6.1846*	1.1846	0.9852	15.6078**
LR_{ind}	N.A.	N.A.	N.A.	N.A.	N.A.	N.A.	1.7047
Efficiency							
MRSB	-0.0301	-0.0086	0.0276	-0.0010	0.0121	-0.0630	0.0435

Note: * denotes significance at the 5% level.

** denotes significance at the 1% level.

4. Conclusion

In this study, we have attempted to provide an analysis that is helpful to the user of VaR models either for risk management within a bank or for regulatory purposes to control a bank's risk exposure. The main strength of the study lies in our employment of data on the actual trading book portfolio that comprise a broad range of financial assets. The RiskMetrics[®] standard EWMA model, which is based on a normal distribution, is widely used to forecast the variance of the conditional distribution of asset returns. While it is appropriate when the asset returns are drawn from a normal distribution, there is considerable evidence to show that the distribution of most financial returns is not well approximated by a normal distribution, even conditionally. The conditional distribution of asset returns is typically found to be leptokurtic, and to have fatter tails than that of a normal distribution. We therefore introduce the power exponential distribution to construct a series of EWMA family estimators to improve on the inefficiency of the VaR estimators. By considering the different aspects of the usefulness of the VaR models to both the risk manager and the supervisory authorities,

we focus on three aspects of the model – conservativeness, accuracy and efficiency – and propose a range of statistics based on these criteria to evaluate the performance of the family models. Our concluding remarks are as follows.

From the results of the descriptive statistics, the estimated tail-indexes and the estimated power parameters of the power exponential distribution, we consistently find that the asset return distributions of the bank’s portfolio are significantly fat-tailed and leptokurtic. The estimated decay factors of the family of EWMA models are all around 0.94, as suggested by JP Morgan.

In terms of the concerns of the supervisory authorities, we find that the most conservative model is the robust EWMA model, while the Monte-Carlo simulation is the least conservative model. The results of the evaluation of model accuracy highlight the fact that the standard EWMA model achieves the most accurate results in terms of different accuracy measures and most of our models are accurate. As expected, the models based on the normal distribution, RiskMetrics[®] and the Monte-Carlo simulation, appear to be less sensitive to the outlying observations than other models based on the exponential power distribution, and then lead to underestimations of the risk of the bank portfolio. Another finding is that the historical simulation performs well, yielding good accuracy and efficient estimates. It should be noted that the simple VaR model does not substantially under-perform relative to the more complex models.

The backtesting results demonstrate that, due to the flexibility of the power parameters of the conditional distribution, the power exponential distribution can properly capture the fat-tailedness characteristic of the asset return distributions. Most of the family of EWMA estimators based on the power exponential distribution outperform those VaR estimators that are based on the normal distribution, and offer an appropriate coverage of the extreme risk.

References

- Baillie, R., and T. Bollerslev, 1989, “The Message in Daily Exchange Rates: A Conditional-Variance Tale,” *Journal of Business and Economic Statistics*, 7, 297–305.
- Baillie, R., and R. de Gennaro, 1990, “Stock Returns and Volatility,” *Journal of Financial and Quantitative Analysis*, 25, 203-14.
- Beine, M., S. Laurent, and C. Lecourt, 2002, “Accounting for Conditional Leptokurtosis and Closing Days Effects in FIGARCH Models of Daily Exchange

- Rates,” *Applied Financial Economics*, 12, 589–600.
- Berndt, E. K., B. H. Hall, R. E. Hall, and J. A. Hausman, 1974, “Estimation and Inference in Nonlinear Structural Models,” *Annals of Economic and Social Measurement*, 4, 653-65.
- Bollerslev, T., 1986, “Generalized Autoregressive Conditional Heteroskedasticity,” *Journal of Econometrics*, 31, 307-27.
- Bollerslev, T., 1987, “A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return,” *Review of Economics and Statistics*, 69, 542–547.
- Bollerslev, T., R. Chou, and K. Kroner, 1992, “ARCH Modelling in Finance,” *Journal of Econometrics*, 52, 5-59.
- Boudoukh, J., M. Richardson, and R. F. Whitelaw, 1997, “Investigation of a Class of Volatility Estimators,” *Journal of Derivatives*, 4, 63-72.
- Christoffersen, P., 1998, “Evaluating Interval Forecasts,” *International Economic Review*, 39, 841-62.
- Dowd, K., 1998, *Beyond Value-at-Risk: The New Science of Risk Management*, New York: Wiley.
- Drudi, F., A. Generale, and G. Majnoni, 1997, “Sensitivity of VaR Measures to Differing Risk Models,” Banca D’Italia Temi di Discussione del Servizio Studi, 317.
- Engel, J., and M. Gizycki, 1999, “Conservatism, Accuracy & Efficiency: Comparing Value-at-Risk Models,” *Working Paper 2*, Australian Prudential Regulation Authority.
- Engle, R., 1982, “Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation,” *Econometrica*, 50, pp. 987-1007.
- Fama, E., 1965, “The Behavior of Stock Market Prices,” *Journal of Business*, 38, 34–105.
- Greene, W. H., 2000, *Econometric Analysis*, Prentice Hall.
- Guermat, C., and R. D. F. Harris, 2002, “Robust Conditional Variance Estimation and Value at Risk,” *Journal of Risk*, 4, 25-41.
- Hendricks, D., 1996, “Evaluation of Value at Risk Models Using Historical Data,” *Federal Reserve Bank of New York Economic Policy Review*, April, 39-69.
- Higgins, M., and A. Bera., 1992, “A Class of Nonlinear ARCH Models,” *International Economic Review*, 33, 137-58.
- Hill, B., 1975, “A Simple General Approach to Inference about the Tail of a Distribution,” *Annals of Statistics*, 3, 1163-74.
- Hsieh, D., 1989, “Modeling Heteroskedasticity in Daily Foreign Exchange Rates,” *Journal of Business and Economic Statistics*, 7, 307–317.
- Huisman, R., K. G. Koedijk, C. J. M. Kool, and F. Palm, 1998, “The Fat-Tailedness of FX Returns,” *Working Paper*.
- Huisman, R., K. G. Koedijk, C. J. M. Kool, and F. Palm, 2001, “Tail-Index Estimates in Small Samples,” *Journal of Business and Economic Statistics*, 19, 208-16.
- Hull, J. and A. White, 1998, “Value at Risk when Daily Changes in Market Variables are not Normally Distributed,” *Journal of Derivatives*, 5, 9-19.
- Jackson, P., D. J. Maude, and W. Perraudin, 1997, “Bank Capital & Value at Risk,” *Journal of Derivatives*, Spring, 73-89.
- Jansen, D., and C. de Vries, 1991, “On the Frequency of Large Stock Returns: Putting Booms and Busts into Perspective,” *Review of Economics and Statistics*, 73, 18-24.
- Jorion, P., 1997, *Value at Risk: The New Benchmark for Controlling Market Risk*,

- McGraw-Hill.
- Jorion, P., 1988, "On Jump Processes in the Foreign Exchange and Stock Markets," *Review of Financial Studies*, 68, 165–176.
- Jorion, P., 2000, *Value at Risk: The New Benchmark for Managing Financial Risk*, McGraw-Hill.
- Kaiser, T., 1996, "One-Factor-GARCH Models for German Stocks – Estimation and Forecasting," *Working Paper*, Universiteit Tubingen
- Kearns, P., and Pagan, A., 1997, "Estimating the Density Tail Index for Financial Times Series," *The Review of Economics and Statistics*, 79, 171-175.
- Koedijk, K. G., and C. Kool, 1994, "Tail Estimates and the EMS Target Zone," *Review of International Economics*, 2, 153-165
- Kupiec, P. H., 1995, "Techniques for Verifying the Accuracy of Risk Measurement Models," *Journal of Derivatives*, 3, 73-84.
- Linden, M., 2001, "A Model for Stock Return Distribution," *International Journal of Finance and Economics*, 6, 159-69.
- Lopez, J., 1998, "Methods for Evaluating Value-at-Risk Estimates," *Federal Reserve Bank of New York Research Paper*, No. 9802.
- Loretan, M., and P. Phillips, 1994, "Testing the Covariance Structure of Heavy-Tailed Time Series," *Journal of Empirical Finance*, 1, 211-248.
- Marrison, C., 2002, *The Fundamentals of Risk Measurement*, McGraw-Hill.
- Mandelbrot, B., 1963, "The Variation of Certain Speculative Prices," *Journal of Business*, 36, 394–419.
- Nelson, D., and D. Foster, 1994, "Asymptotic Filtering Theory for Univariate ARCH Models," *Econometrica*, 62, 1-41.
- Varma, J. R., 1999, "Rupee-Dollar Option Pricing and Risk Measurement: Jump Processes, Changing Volatility and Kurtosis Shifts," *Journal of Foreign Exchange and International Finance*, April-June, 13 (1), 11-33.
- Vlaar, P., and F. Palm, 1993, "The Message in Weekly Exchange Rates in the European Monetary System: Mean Reversion, Conditional Heteroskedasticity and Jumps," *Journal of Business and Economic Statistics*, 11, 351–360.