Scale-Consistent Value-at-Risk

Citation for published version (APA):

Lehnert, T., & Wolff, C. C. P. (2004). Scale-Consistent Value-at-Risk. *Finance Research Letters*, 1(2), 127-134. <https://doi.org/10.1016/j.frl.2004.03.002>

Document status and date:

Published: 01/01/2004

DOI:

[10.1016/j.frl.2004.03.002](https://doi.org/10.1016/j.frl.2004.03.002)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Scale-consistent Value-at-Risk

Thorsten Lehnert^{a,*} and Christian C.P. Wolff^{a,b}

^a *Limburg Institute of Financial Economics (LIFE), Maastricht University, PO Box 616,
6200 MD Maastricht, The Netherlands*

^b *Centre for Economic Policy Research (CEPR), London, UK*

Received 11 December 2003; accepted 12 March 2004

Abstract

Returns in financial assets show consistent excess kurtosis and skewness, indicating the presence of large fluctuations not predicted by Gaussian models. In this paper we propose a generalization to the popular RiskMetrics approach to Value-at-Risk. In order to model scale-consistent Value-at-Risk (VaR), we propose a model with a time varying scale parameter and error terms that are truncated Lévy distributed. Lévy flights include a method for scaling up from a single-day volatility to a multi-day volatility. We use this rule to approximate future volatility and estimate Value-at-Risk several days ahead, and compare it to the popular approach, which is a special case of our method. Back-testing results suggest that the inclusion of more sophisticated tail properties and the data-driven scaling rule improves the performance of the VaR model significantly, for short and long time horizons. Our approach is easier to implement and is less time and computer intensive compared to Monte Carlo simulation methods.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Lévy processes; GARCH; Scale consistency; Value-at-Risk

1. Introduction

The dynamics of financial markets constitute a major challenge for financial econometricians. While financial applications involve many different time intervals, ranging from a few minutes (intraday) to a number of years, most techniques used in econometrics focus on modeling the fluctuations of price series in a single time interval. But the distribution

* Corresponding author.

E-mail address: t.lehnert@berfin.unimaas.nl (T. Lehnert).

that successfully explains daily price changes, for example, is typically unable to characterize the nature of hourly price changes. On the other hand, the statistical properties of monthly price changes are often not fully covered by a model for daily price changes. In order to describe the statistics of future prices of a financial asset, one needs *a priori* a distribution of price fluctuations for different time intervals, corresponding to different trading time horizons.

Mandelbrot (1963) first proposed the idea that price changes are distributed according to a Lévy stable law. This model was frequently criticized, because the tails are now much overestimated and the infinite variance makes it impossible to apply the Central Limit Theorem. In physical systems, so-called Lévy flights have been observed experimentally and have been used very successfully to describe, for instance, the spectral random walk of a single molecule embedded in a solid. Typically, unavoidable cutoffs in the tails of the distribution are always present. This cutoff ensures that the variance will be finite and the distribution converges to a Gaussian in the limit. To model financial prices over time the so-called truncated Lévy flight (TLF) can be constructed by the sum of independent and identically distributed random variables described by a truncated Lévy distribution (TLD). Mantegna and Stanley (1994) showed that the scale invariant behavior of a Lévy flight can be observed for short time scales using ultra high frequency data. Cont et al. (1997) showed that it breaks down for longer time scales. These observations have been explained as a structural break in terms of the truncated Lévy flight. Among others, Mantegna and Stanley (1998, 2000) and Bouchaud and Potters (2000) confirmed that the truncated Lévy distribution better captures the skewness and excess kurtosis in financial return series.

In contrast to physical systems, research in finance has shown that there is a strong non-i.i.d. clustering effect in financial data. Fortunately the class of GARCH models has been very successful in modeling the significant volatility clustering and non-i.i.d. nature of the data. More specifically, the GARCH model produces a mean reverting time dependent volatility process that “filters out” the correlations in the data and the remaining residuals are assumed to be i.i.d. We account for non-i.i.d.-ness in the data by modeling the location and scale parameter of the truncated Lévy distribution time varying. A good description of the distribution of price changes, especially in the tails, is important for risk measures like Value-at-Risk (VaR). The VaR of a particular portfolio of assets (in our study an index portfolio) is directly related to the quantile of the asset’s return distribution. In a GARCH framework, typically Monte Carlo simulation techniques are used in order to derive multi-day density forecasts (see, e.g., Bams et al., 2003). A related approach was recently proposed by McNeil and Frey (2000). They combine extreme value theory (EVT) and GARCH-modeling and confirm that the model outperforms existing methods. Mittnik and Paolella (2000), Mittnik et al. (2000) demonstrate that more general GARCH structures and skewed fat tailed distributions improve the precision of the model in-sample and out-of-sample. Both methods also rely on simulation methods to approximate the multi-day density function. In recent option pricing models, the same methods are successfully employed (see, e.g., Lehnert, 2003). In all of these studies, the techniques are experienced to be quite time and computer intensive, but there is a lack of alternative methods. We propose an extension of the well-known RiskMetrics (1996) approach and approximate the time scaling behavior of the quantiles by a scaling rule, namely the alpha-root-of-time rule of the Lévy flight, estimated from the data directly.

2. The econometric framework

The Lévy distribution with a cutoff and exponentially declining tails was introduced in the physics literature by Mantegna and Stanley (1994). Koponen (1995) developed the characteristic function (CF) of a truncated Lévy distribution. Note the misprint in the original publication, the characteristic function of a TLD random variable k should read:

$$\psi_{\text{TL}}(k, \mu, C, \alpha, \lambda, \beta) = i\mu k - C^\alpha \left\{ \frac{\lambda^\alpha - (k^2 + \lambda^2)^{\alpha/2}}{\cos(\pi\alpha/2)} \cos\left(\alpha \arctan\left(\frac{|k|}{\lambda}\right)\right) \right. \\ \left. \times \left[1 + i \operatorname{sgn}(k)\beta \tan\left(\alpha \arctan\left(\frac{|k|}{\lambda}\right)\right) \right] \right\}, \quad (1)$$

where μ is a location parameter, $C > 0$ is a scale parameter, α is the characteristic exponent determining the shape of the distribution and especially the fatness of the tails ($0 < \alpha \leq 2$, but $\alpha \neq 1$), and λ is the so-called cutoff parameter, which determines the speed of the decay and as a result the cutoff region. The parameter β ($\beta \in [-1, 1]$) determines the skewness when $\beta \neq 0$. The distribution is skewed to the right when $-1 < \beta < 0$ and skewed to the left when $0 < \beta < 1$. Accurate numerical values for the density can be calculated Fourier-transforming the CF and evaluating the integral numerically. We use Romberg integration, which allows ex-ante specification of the tolerated error and in fact a calculation of the density as precise as necessary (see Lambert and Lindsey, 1999).

In practice, comparing the distributional properties of price increments at various time intervals provides insight into the temporal dependence structure of the time series. Analyzing the time scaling behavior of financial fluctuations means comparing the increments for shorter time scales τ and for longer time scales $N\tau$. This formally corresponds to summing N random variables. In the case of the Lévy distribution, the characteristic function satisfies $N\psi_L(k) = \psi_L(N^{1/\alpha}k)$. The distribution for various time scales for stationary and independent variables is related by a convolution relation $P_{N\tau} = P_\tau \otimes P_\tau \otimes \dots \otimes P_\tau$. More generally, the distribution $P(x)$ of price changes on a time scale $N\tau$ may be obtained from that of a shorter time scale τ by a rescaling of the variable $P_L^N(x) = N^{-1/\alpha} P_L^1(N^{-1/\alpha}x)$, where $P^N(x)$ denotes a N -times convoluted distribution of $P^1(x)$.

Since we introduced a cut-off for the CF of the truncated Lévy $\psi_{\text{TL}}(k, \lambda)$, it is no longer self-similar or uni-fractal by the criteria mentioned above, but bi-fractal, the simplest version of a multi-fractal process (see Nakao, 2000). The convolution of the probability distribution can still be obtained by scaling both x and λ . The CF $\psi_{\text{TL}}(k, \lambda)$ satisfies $N\psi_{\text{TL}}(k, \lambda) = \psi_{\text{TL}}(N^{1/\alpha}k, N^{1/\alpha}\lambda)$ and the N -times convoluted probability distribution satisfies $P_{\text{TL}}^N(x, \lambda) = N^{-1/\alpha} P_{\text{TL}}^1(N^{-1/\alpha}x, N^{1/\alpha}\lambda)$. For short time scales (daily) the process behaves like a Lévy flight, but converges towards a Gaussian for longer time scales (say monthly) (see Matacz, 2000). The scaling of λ , which means that for increasing time scales the cutoff is introduced earlier in the tails, ensures that the process converges towards a Gaussian process instead of staying a Lévy flight.

Volatilities can be predicted reasonably successfully with a parametric model such as GARCH. Traditional GARCH models (with Normal- or Student- t distributed error terms) were designed to capture clustering of large and small innovations, which can be modeled

as serially correlated conditional variances when the variance exists (Bollerslev et al., 1992). The analogue of the standard deviation σ in the family of Lévy distributions is the scale parameter C . If we replace the standard deviation σ by the scale parameter C , we allow C_t to be serially correlated, which produces the volatility clustering.

In the following we are considering the single lag ($p = q = 1$) version of a very general GARCH process, which is typically sufficient in practice. The complete *augmented* GARCH (1, 1) model reads:¹

$$\begin{aligned} r_t &= \mu_t + \sigma_t \varepsilon_t, \quad \varepsilon_t \sim D(0, 1), \quad \phi_t = \alpha_0 + \gamma_{1,t-1} \phi_{t-1} + \gamma_{2,t-1}, \\ \sigma_t &= \begin{cases} |\delta \phi_t - \delta + 1|^{1/2\delta}, & \text{if } \delta \neq 0, \\ \sqrt{\exp(\phi_t - 1)}, & \text{if } \delta = 0, \end{cases} \\ \gamma_{1,t-1} &= \alpha_1 + \alpha_2 |\varepsilon_t - b|^\kappa + \alpha_3 \max(0, b - \varepsilon_t)^\kappa, \\ \gamma_{2,t-1} &= \alpha_4 \frac{|\varepsilon_t - b|^\kappa - 1}{\kappa} + \alpha_5 \frac{\max(0, b - \varepsilon_t)^\kappa - 1}{\kappa}, \end{aligned} \quad (2)$$

where the conditional location parameter μ_t can be specified additionally and the conditional scale parameter σ_t is assumed to vary over time. The parameters δ and κ allow for a Box–Cox transformation of the scale parameter and the innovation term, respectively. For example, for $\delta = 0$, $\kappa = 1$, and $\alpha_2 = \alpha_3 = b = 0$ the model reduces to the well-known EGARCH specification. In this case α_1 refers to the persistence of the volatility process and a combination of α_4 and α_5 captures the leverage effect. A value of $b \neq 0$ would additionally shift the news impact curve. Our model is not restricted to zero–mean and variance–one distributions, it also allows for location–zero and scale–one continuous distributions $D(0, 1)$. Since there is no evidence as to which GARCH specification should be used, we develop a Lagrange Multiplier specification test (LM test) based on our generalization of the augmented GARCH process of Duan (1997), which allows us to reject several specifications among the models analyzed and to derive the ‘best’ specification for each particular return series (see Duan, 1997, for details).

Under the hypothesis of conditional leptokurtosis and skewness, the conditional volatility estimate is used together with the characteristics of the truncated Lévy flight. Once the parameters of the model are estimated, we know the probability distribution of one particular time scale, say daily. The N -times convoluted probability distribution can be obtained by scaling the parameter C and the cut-off parameter λ by a α -root-of-time rule to derive the multi-day parameters. Therefore, the VaR of a position with a confidence level q and for N periods ahead can be defined as

$$\text{VaR}_{q,N} = W_0(1 - \exp(R_N^*)), \quad (3)$$

where $\exp(R_N^*)$ can be determined by setting R_N^* equal to the $(1 - q)$ th quantile of the truncated Lévy distribution with the characteristic function $\psi_{\text{TL}}(\mu_t N, N^{1/\alpha} C_t, N^{1/\alpha} \lambda, \beta)$, where μ_t and C_t are the location and scale parameter, respectively, α is the characteristic exponent of the truncated Lévy distribution, and $N^{1/\alpha}$ is the adjustment factor needed to scale the parameter. Equivalently, for the RiskMetrics EWMA method, $\exp(R_N^*)$ can be

¹ Our model contains most of the existing GARCH specifications.

determined by setting R_N^* equal to the $(1 - q)$ th quantile of the normal distribution with the characteristic function $\psi_G(\mu_t N, \sqrt{N}\sigma_t)$, where μ_t and σ_t are the mean and standard deviation, respectively, $\sqrt{N}$ is the adjustment factor needed to scale the standard deviation. The method is frequently applied in practice, but it has been shown to underestimate the downside risk consistently (see, e.g., Pownall and Koedijk, 1999).

3. Empirical results

In this study, daily closing prices for some major stock market price indices between May 1992 and April 2000 are used for the analysis. In particular, we examine the S&P500, NASDAQ, and FTSE 100 from May 4, 1992, to April 3, 2000. The total number of trading days covered by the data is 2000 (FTSE 100) and 2001 (S&P500 and NASDAQ). The data are obtained from DataStream. We used the percentage daily logarithmic change $100 * \ln(p_t / p_{t-1})$, where p_t is the price index at time t .

There is already some evidence that the inclusion of a time-varying asymmetric volatility process can capture only some of the excess kurtosis and skewness in financial data. Additionally, Hentschel (1995) showed that the differences in the conditional volatility estimates could be substantial among the various specifications. Since the choice of a particular specification will affect the conditional volatility estimate, it will also affect VaR estimates, which is of our particular interest. We considered the following restricted specifications of our augmented GARCH process for the specification test: the simple GARCH, the power-GARCH processes with either shift or rotation of the news impact curve, the NGARCH, and the EGARCH specification. The results of the Lagrange Multiplier test for robust standard errors based on White (1982) confirm that standard GARCH model can be rejected for all indexes at the 1 or 5% level (results not reported). For the S&P500 and the NASDAQ returns the test suggests that an asymmetric GARCH process with a rotation of the news impact curve (Power-GARCH) and for the FTSE 100 returns the asymmetric volatility model with a shift of the news impact curve (NGARCH) are most appropriate.² The parameter estimates for the particular GARCH specifications and the individual index portfolios are presented in Table 1.

Since we are interested in the out-of-sample performance, we back-test our model for all indices and the particular GARCH specifications over a period May 1996 until March 2000 and compare it to the RiskMetrics EWMA approach. Every day we estimate the model using the last about 1000 trading days (that means exactly one half of each sample and a moving window) and forecasted the 99% (95%) VaR 5, 10, and 20 days ahead. Table 2 reports the out-of-sample results.

The results are very promising: the VaR estimates we obtained by using our method compared to the RiskMetrics method produces on average less violations of the expected VaR for all confidence intervals and horizons. The RiskMetrics method constantly underpredicts extreme events and this leads very often to an inappropriate number of

² Among the models that cannot be rejected, we base our decision on the Schwarz Bayesian Criterion (SBC) (Schwarz, 1978).

Table 1
Parameter estimates

Index/model	Parameters										
	δ	κ	α_1	α_2	α_3	α_4	α_5	b	α	λ	β
S&P500											
Power-GARCH, truncated Lévy											
	0.880	1.760	0.927	0.008	0.072	0	0	0	1.698	0.235	0.276
	(0.033)	–	(0.015)	(0.009)	(0.020)	–	–	–	(0.105)	(0.117)	(0.133)
	(0.058)	–	(0.018)	(0.007)	(0.026)	–	–	–	(0.185)	(0.153)	(0.248)
NASDAQ											
Power-GARCH, truncated Lévy											
	0.791	1.582	0.877	0.040	0.071	0	0	0	1.814	0.089	0.918
	(0.052)	–	(0.022)	(0.015)	(0.022)	–	–	–	(0.045)	(0.041)	(0.141)
	(0.049)	–	(0.029)	(0.012)	(0.022)	–	–	–	(0.036)	(0.046)	(0.120)
FTSE 100											
NGARCH, truncated Lévy											
	1	2	0.935	0.019	0	0	0	1.116	1.906	0.164	0.419
	–	–	(0.010)	(0.004)	–	–	–	(0.274)	(0.065)	(0.157)	(0.289)
	–	–	(0.011)	(0.003)	–	–	–	(0.242)	(0.055)	(0.117)	(0.289)

Notes. The table reports parameter estimates for the truncated Lévy model and the particular GARCH specifications. We used our specification test to determine the appropriate GARCH specification for each return series. The underlying data set consists of daily observations for the period May 1992 until March 2000. The location parameter is assumed to be equal to the unconditional value. Standard errors and robust standard errors proposed by White (1982) are given within parentheses.

Table 2
Violations of the actual Value-at-Risk

	HORIZON (days)	VaR (conf. level) (%)	S&P500 (%)	NASDAQ (%)	FTSE 100 (%)
EWMA	5	99	2.6	3.3	2.9
		95	5.9	7.6	6.1
	10	99	1.8	2.7	1.1
		95	5.2	6.6	4.9
	20	99	0.9	2.3	1.3
		95	4.6	7.2	4.3
Augmented GARCH Truncated Lévy	5	99	1.5	1.2	1.3
		95	5.2	4.9	5.3
	10	99	1.2	1.3	0.8
		95	4.3	5.5	4.0
	20	99	0.7	1.4	0.9
		95	3.7	4.6	3.9

Notes. The table reports percentage violations of the actual Value-at-Risk for the alternative models during the ‘back-testing period’ (May 1996 until March 2000). Every day the particular model was calibrated and a Value-at-Risk estimate was obtained for different forecasting horizons and confidence levels. Subsequently, the actual, realized return over a certain horizon was compared with the reported Value-at-Risk. Given a confidence level of 95% (99%), the model assumes a violation rate of 5% (1%). A higher (lower) violation rate indicates that the model underestimates (overestimates) the actual Value-at-Risk.

violations. But in particular this underprediction is slightly reduced for lower confidence intervals or longer forecasting horizons. This is a well-known result for the Gaussian

distribution and the square-root-of-time scaling rule. On the other side, the GARCH models with the truncated Lévy distribution and the alpha-root-of-time scaling rule lead to an appropriate number of violations for low *and* high confidence intervals and short *and* long forecasting horizon.

In order to further evaluate our modeling approach, we compared the performance of our model with the performance of a model proposed by Mittnik and Paoletta (2000). The authors consider a GARCH process driven by a skewed Student-*t* distribution and obtain multi-period densities using Monte Carlo simulation techniques. Results suggest that the in-sample fit as well as the out-of-sample forecasting performance of our VaR model is superior to the GARCH-skewed Student-*t* model (results not reported). Further evaluation of the different forecasting techniques underlying the two models is beyond the scope of this paper.

Therefore, we can conclude that our scaling rule captures the scaling behavior of the data quite well and shows a convergence from a skewed leptokurtic distribution to a Gaussian for larger sampling intervals. This is actually the unique bi-fractal scaling behavior of the truncated Lévy flight.

4. Conclusions

In this paper we propose a generalization of the popular RiskMetrics approach to Value-at-Risk. Our approach has some advantages: using the truncated Lévy flight for the innovations of a GARCH process, we are able to capture the observed conditional tail fatness and skewness in financial returns. We propose a new scaling rule to forecast volatility. Location and scale parameter (volatility) are estimated on one time scale (daily) and the multi-day (weekly or monthly) values are derived by using the stable property of Lévy processes. The method has the implied advantage that we are able to identify the relationship between return distributions for different sampling intervals by analyzing the time series of returns on one sampling interval (say daily). As a result, based on our scaling approach, we are able to produce better forecasts compared to methods, which add up one-day forecasts to derive a multi-day forecast. Our approach is easier to implement and is less time and computer intensive compared to Monte Carlo simulation methods, which are typically used in risk management applications and in option pricing models. Additionally, scaling rules are known to improve the accuracy of VaR estimates and suffer less from ‘estimation errors’ in VaR compared to other approaches.

References

- Bams, D., Lehnert, T., Wolff, C.C.P., 2003. An evaluation framework for alternative VaR models. *Journal of International Money and Finance*. In press.
- Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance: A selective review of the theory and empirical evidence. *Journal of Econometrics* 52, 5–59.
- Bouchaud, J.-P., Potters, M., 2000. *Theory of Financial Risk*. Cambridge University Press, Cambridge, UK.
- Cont, R., Potters, M., Bouchaud, J.-P., 1997. Scaling in stock market data: stable laws and beyond. *Science & Finance Working paper*, cond-mat/9705087.

- Duan, J.D., 1997. Augmented GARCH (p, q) process and its diffusion limit. *Journal of Econometrics* 79, 97–127.
- Hentschel, L., 1995. All in the family: Nesting symmetric and asymmetric GARCH model. *Journal of Financial Economics* 39, 71–104.
- Koponen, K., 1995. Analytic approach to the problem of convergence of truncated Lévy flights towards the Gaussian stochastic process. *Physical Review E* 52, 1197–1199.
- Lambert, P., Lindsey, J.K., 1999. Analysing financial returns by using regression models based on non-symmetric stable distributions. *Journal of Applied Statistics* 48, 409–424.
- Lehnert, T., 2003. Explaining smiles: GARCH option pricing with conditional leptokurtosis and skewness. *Journal of Derivatives* 10, 27–39.
- Mandelbrot, B., 1963. The variation of certain speculative prices. *Journal of Business* 36, 394–419.
- Mantegna, R.N., Stanley, H.E., 1994. Stochastic process with ultraslow convergence to a Gaussian: the truncated Lévy flight. *Physical Review Letters* 73, 2946–2949.
- Mantegna, R.N., Stanley, H.E., 1998. Modeling of financial data: Comparison of the truncated Lévy flight and the ARCH (1) and GARCH (1,1) processes. *Physica A* 254, 77–84.
- Mantegna, R.N., Stanley, H.E., 2000. *Introduction to Econophysics*. Cambridge University Press, Cambridge, UK.
- Matacz, A., 2000. Financial modeling and option theory with the truncated Lévy process. *International Journal of Theoretical and Applied Finance* 3, 143.
- McNeil, A.J., Frey, R., 2000. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance* 7, 271–300.
- Mittnik, S., Paolella, M.S., 2000. Conditional density and Value-at-Risk prediction of Asian currency exchange rates. *Journal of Forecasting* 19, 313–333.
- Mittnik, S., Paolella, M.S., Rachev, S.T., 2000. Diagnosing and treating the fat tails in financial return data. *Journal of Empirical Finance* 7, 389–416.
- Nakao, H., 2000. Multi-scaling properties of truncated Lévy flights. Manuscript. Graduate School of Mathematical Sciences, University of Tokyo.
- Pownall, R.A.J., Koedijk, K.G., 1999. Capturing downside risk in financial markets: the case of the Asian Crisis. *Journal of International Money and Finance* 18, 853–870.
- RiskMetrics, 1996. Technical Document, fourth ed. J.P. Morgan, London.
- Schwarz, G., 1978. Estimating the dimension of a model. *Annals of Statistics* 6, 461–464.
- White, H., 1982. Maximum likelihood estimation of misspecified models. *Econometrica* 50, 1–25.