Long-Term Risk Management

Roger Kaufmann
Swiss Life
General Guisan-Quai 40
Postfach, 8022 Zürich
Switzerland
roger.kaufmann@swisslife.ch

April 28, 2005

Abstract. In this paper financial risks for long time horizons are investigated. As measures for these risks, value-at-risk and expected shortfall are considered. In a first part, questions concerning a two-week horizon are addressed. For GARCH-type processes and stochastic volatility models with jumps, methods to estimate quantiles of financial risks for two-week periods are introduced, and compared with the widely used square-root-of-time rule, which scales one-day risk measures by $\sqrt{10}$ to get ten-day risk measures.
In the second part of the paper, a framework for the measurement of one-year risks is developed. Several models for financial time series are introduced, and compared with each other. The various models are tested for their appropriateness for estimating one-year expected shortfall and value-at-risk on 95% and 99% confidence levels.

Key words and phrases. Value-at-risk, expected shortfall, long time horizons, scaling rules, stochastic volatility models.

1 Motivation

In a first part we analyse financial risks over a two-week horizon. We investigate whether converting a one-day risk estimate into a ten-day risk estimate by scaling it with the square-root-of-time formula is appropriate. The motivation for this investigation comes from the banking industry. When one searches The New Basel Accord by the Basel Committee on Banking Supervision [4] for an instruction for incorporating market risks, one is referred to an earlier publication of this committee, the Amendment to the Capital Accord to Incorporate Market Risks [3]. The following three quotes can be found in this document:

- In calculating the value-at-risk, a 99th percentile, one-tailed confidence interval is to be used.
- In calculating value-at-risk, an instantaneous price shock equivalent to a 10 day movement in prices is to be used.
- Banks may use value-at-risk numbers calculated according to shorter holding periods scaled up to ten days by the square root of time.

The key message is that market risks should be measured on a 10-day basis, and one should evaluate the one-in-a-hundred event. The Basel Committee on Banking Supervision explicitly permits banks to use the square-root-of-time rule to obtain a 10-day 99% value-at-risk out of a one-day 99% value-at-risk.

In insurance, one is even interested in the 1-year 99% value-at-risk or in the 1-year 99% expected shortfall. The estimation of such long-term risks is the subject of the second part of this paper. Let us first remind the definition of value-at-risk and expected shortfall.

Definition 1.1 The *value-at-risk* at level α of a random variable R is defined as

$$\mathrm{VaR}_\alpha(R) = -\inf\{x \in \mathbb{R} \mid \mathbb{P}[R \le x] \ge 1-\alpha\},$$

i.e. $\mathrm{VaR}_\alpha(R)$ is the negative $(1-\alpha)$-quantile of R.

For a given (daily) price process $(P_t)_{t\in\mathbb{Z}}$, alternatively to considering the returns $R_t = (P_t - P_{t-1})/P_{t-1}$, one can work with log-returns $X_t = \log(P_t/P_{t-1})$. Log-returns have the nice property that N-day log-returns $(X_t^N := \log(P_t/P_{t-N}))$ are simply the arithmetic sum of one-day log-returns:

$$\begin{aligned} X_t^N &= \log(P_t/P_{t-N}) \\ &= \log(P_t/P_{t-1}) + \log(P_{t-1}/P_{t-2}) + \cdots + \log(P_{t-(N-1)}/P_{t-N}) \\ &= X_t + X_{t-1} + \cdots + X_{t-(N-1)}. \end{aligned}$$

Since we are primarily interested in the relation between one-day and N-day value-at-risk, we use log-returns throughout this thesis. The two processes $(R_t)_{t\in\mathbb{Z}}$ and $(X_t)_{t\in\mathbb{Z}}$ are in any case related through:

$$R_t = e^{X_t} - 1, \quad \mathrm{VaR}_\alpha(R) = 1 - \exp(-\mathrm{VaR}_\alpha(X)).$$

An alternative risk measure, which has nice theoretical properties being important in many practical applications, is expected shortfall; see Artzner et al. [2] for the concept of coherent risk measures.

Definition 1.2 The *expected shortfall* at level α of R is defined as

$$\mathrm{ES}_\alpha(R) = -\mathbb{E}[R \mid R < -\mathrm{VaR}_\alpha(R)].$$

Expected shortfall is the average loss when value-at-risk is exceeded. $\mathrm{ES}_\alpha(R)$ can also be interpreted as the expected value of $S_\alpha(R) := -R \mid R < -\mathrm{VaR}_\alpha(R)$, which gives information about frequency and size of large losses. The two risk measures are illustrated in Figure 1. Note that losses are shown as positive values.

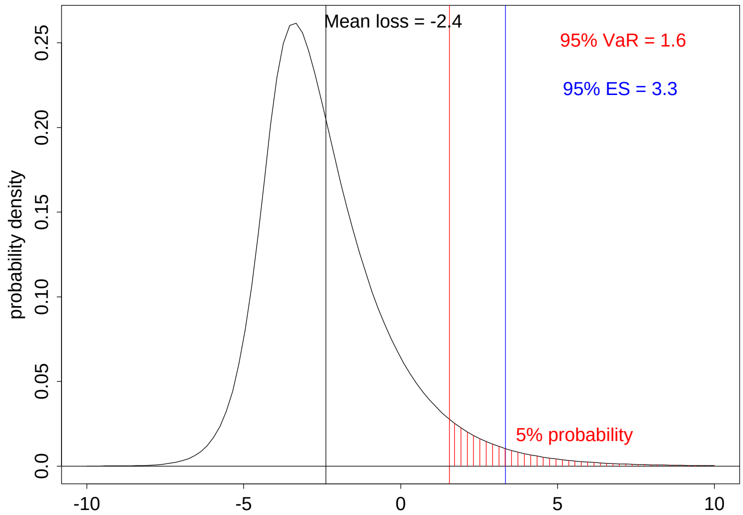

Figure 1 Loss distribution function.

In the present work, we concentrate on *unconditional* risk estimates. This choice can be motivated by the fact that it is not possible to constantly adapt risk reserves to changing market conditions.

2 10-day risks

The question whether applying the simple square-root-of-time scaling rule gives reasonable 10-day value-at-risk estimates does not have an absolute, universally valid answer, but is highly model-dependent. For our investigations we restrict to GARCH-type models and stochastic volatility models with jumps.

2.1 Scaling rule for GARCH-type models.

An important class of models for financial data are so-called GARCH-type models. We start our investigations about scaling of risks with the simplest form of such models, which is a random walk with normally distributed log-returns and no trend.

2.1.1 *Scaling under normality.*

Under the assumption of normally distributed log-returns, $X_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$, n-day log-returns are also normally distributed, that is $\sum_{t=1}^n X_t \sim \mathcal{N}(0, n\sigma^2)$. For a $\mathcal{N}(0, \widetilde{\sigma}^2)$-distributed profit X, value-at-risk can be written as $\mathrm{VaR}_\alpha(X) = \widetilde{\sigma}\, x_\alpha$, where x_α denotes the α-quantile of a standard normal distribution. Hence the square-root-of-time scaling rule

$$\mathrm{VaR}^{(n)} = \sqrt{n}\, \mathrm{VaR}^{(1)}$$

works perfectly for this model.

2.1.2 *Accounting for trends.*

When adding a constant value μ to the one-day returns, i.e. $X_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \sigma^2)$, then n-day log-returns are still normally distributed: $\sum_{t=1}^n X_t \sim \mathcal{N}(n\mu, n\sigma^2)$. Compared to the model with mean zero, the value-at-risk is decreased by the trend, which means that $\mathrm{VaR}^{(n)} + n\mu = \sqrt{n}\,(\mathrm{VaR}^{(1)} + \mu)$, i.e.

$$\mathrm{VaR}^{(n)} = \sqrt{n}\,\mathrm{VaR}^{(1)} - (n - \sqrt{n})\mu.$$

In all financial models, trends can – and should – be taken into account as shown here. Accounting for trends is very important, since the effect increases linearly with the length n of the time period.

2.1.3 *Autoregressive models.*

For an autoregressive model of order 1 with normal innovations,

$$X_t = \lambda X_{t-1} + \epsilon_t, \quad \epsilon_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2),$$

both the 1-day and the n-day log-returns are normally distributed:

$$X_t \sim \mathcal{N}\left(0, \frac{\sigma^2}{1-\lambda^2}\right) \quad \text{and} \quad \sum_{t=1}^{n} X_t \sim \mathcal{N}\left(0, \frac{\sigma^2}{(1-\lambda)^2}\left(n - 2\lambda\frac{1-\lambda^n}{1-\lambda^2}\right)\right).$$

Hence, making use of $\mathrm{VaR}_\alpha(X) = \widetilde{\sigma}\, x_\alpha$, we get

$$\mathrm{VaR}^{(n)} = \sqrt{\frac{1+\lambda}{1-\lambda}\left(n - 2\lambda\frac{1-\lambda^n}{1-\lambda^2}\right)}\mathrm{VaR}^{(1)}$$

for AR(1) models with normal innovations, which leads to the conclusion that for small values of λ, the scaled one-day value $\sqrt{n}\,\mathrm{VaR}^{(1)}$ is a good approximation of $\mathrm{VaR}^{(n)}$.

2.1.4 *Scaling for more general models.*

After having investigated these models with normal innovations, we raise the question whether scaling with $\sqrt{n}$ is still appropriate when innovations are heavier tailed. We start this analysis with a random walk with Student-t_8 distributed returns. In this model, returns are still independent and identically distributed, but the tails are heavier than normal.

In this model, the correct scaling from a 1-day to an n-day horizon depends on the value-at-risk level α and cannot be calculated analytically. Hence we examine the issue of the square-root-of-time rule from an empirical point of view. In practice, the goal is often to obtain good 10-day value-at-risk estimates based on data sets of not more than 250 daily returns. Our empirical approach consists of a comparison of some alternative 10-day value-at-risk estimators. We evaluate their performances relative to the square-root-of-time rule.

In the present situation with independent log-returns, random resampling would be the best one could do. Hence, for illustrative purposes, we restrict to this alternative 10-day risk estimator. Random resampling means that all possible convolutions of 10 returns are taken, and the quantile is evaluated based on these artificial 10-day returns.

The result of the comparison is shown in Figure 2. Every point represents the outcome of one simulation. The x-value corresponds to the $\mathrm{VaR}_{99\%}$ estimate obtained from the square-root-of-time rule, while the y-axis shows the corresponding risk estimate based on random resampling. The horizontal and the vertical line mark the true $\mathrm{VaR}_{99\%}$ for the random walk with normal innovations. The dashed line has slope 1 and goes through $(\mathrm{VaR}_{99\%}, \mathrm{VaR}_{99\%})$. The solid line goes through $(\mathrm{VaR}_{99\%}, \mathrm{VaR}_{99\%})$ as well, its slope is $S = (\sum_{i=1}^{1000} |y^i - \mathrm{VaR}_{99\%}|)/(\sum_{i=1}^{1000} |x^i - \mathrm{VaR}_{99\%}|)$. Defined like this, S is the ratio of the mean deviations of the value-at-risk estimates (y^i) and (x^i) from the true value-at-risk. The slope being larger than 1 indicates that random resampling performs slightly better than the square-root-of-time rule. Nevertheless, the scaling rule still performs reasonably well.

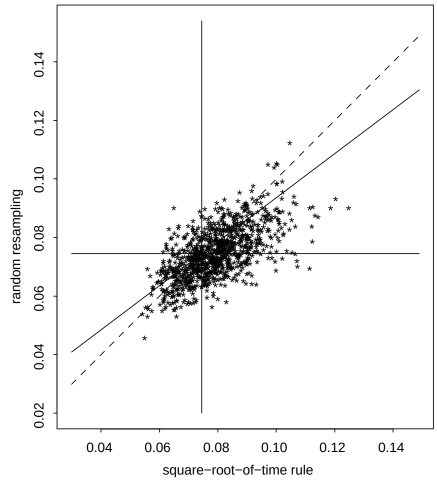

Figure 2 Comparison of two quantile estimation methods for a random walk model with Student-t_8 innovations: square-root-of-time rule vs. random resampling.

This holds true for all GARCH-types models which we investigated. We conclude that the scaling rule provides good estimates also for Student-t innovations, but other methods like random resampling might perform slightly better.

2.1.5 *AR(1)-GARCH(1,1) processes.*

A more complex process, often used for practical applications, is the GARCH(1,1) process ($\lambda = 0$ in the formula below) and its generalization, the AR(1)-GARCH(1,1) process:

$$X_t = \lambda X_{t-1} + \sigma_t \epsilon_t,$$

$$\sigma_t^2 = a_0 + a(X_{t-1} - \lambda X_{t-2})^2 + b\,\sigma_{t-1}^2,$$

$$\epsilon_t \text{ i.i.d., } \mathbb{E}[\epsilon_t] = 0,\ \mathbb{E}[\epsilon_t^2] = 1.$$

If not mentioned otherwise, we use typical parameters for this model, which are for daily financial log-return data $\lambda = 0.04$, $a_0 = 3 \cdot 10^{-6}$, $a = 0.05$, $b = 0.92$.

We concentrate on the goodness of fit of the square-root-of-time scaling rule depending on the parameter λ, which is – roughly speaking – the model's size of direct dependence (as opposed to the dependence inherent in the volatility part).

In Figure 3 the 10-day 99% value-at-risk in AR(1)-GARCH(1,1) models with normal, Student-t_8 and Student-t_4 innovations is shown for various choices of λ (keeping $\alpha_1 = 0.05$

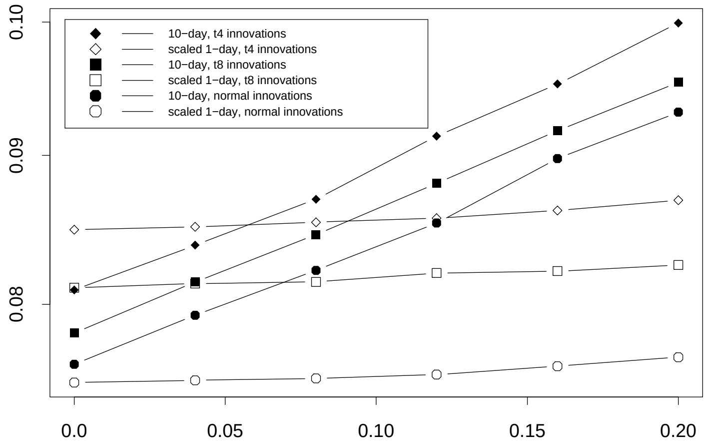

Figure 3 Simulated values for the 10-day 99% value-at-risk in AR(1)-GARCH(1,1) models: true VaR (black symbols) and VaR using the $\sqrt{10}$-rule (white symbols). Parameters: $\alpha_1 = 0.05$, $\beta_1 = 0.92$, $\lambda \in [0.00, 0.20]$.

and $\beta_1 = 0.92$ fixed). We observe that for realistic (small) values of λ the square-root-of-time scaling rule yields a very close approximation (white symbols) of the true 10-day value-at-risk (black symbols). This is good news for practical applications. In particular, for Student-t_8 innovations, for the parameters $\lambda = 0.04$, $\alpha_1 = 0.05$ and $\beta_1 = 0.92$, the estimation via the square-root-of-time rule coincides with the true 10-day value-at-risk. The reason for this is the existence of two counter-acting effects. Direct dependence (λ large) leads to an underestimation, whereas the heavy-tailedness of the one-day log-returns causes an overestimation. With the above parameters, the two effects exactly neutralise each other, which causes this perfect fit.

2.2 The importance of the confidence level α.

From the Central Limit Theorem we know that the normalised sum of n independent and identically distributed random variables with finite variance converges weakly to a standard normal distribution as n tends to infinity. In practical applications, one would typically like to approximate the sum of n independent and identically distributed random variables by a normal distribution, if n is reasonably large. Here, immediately the question arises: when is n large enough?

To find an answer to this question, we study the convolution of Student-t distributed random variables. Let $X_1, \ldots, X_n$ denote independent copies of a Student-t distributed random variable with ν degrees of freedom, expectation 0 and variance 1. Let $S := (X_1 + \cdots + X_n)/\sqrt{n}$ denote the standardised sum, and F_S the corresponding cumulative distribution function. We compare the quantiles $s_\alpha := F_S^{-1}(\alpha)$ of the sum with the quantiles $q_\alpha := \Phi^{-1}(\alpha)$ of a standard normal distribution. We first do this for $\nu = 8$ degrees of freedom. The contour plots in Figure 4 show the area where q_α is a good approximation of s_α. The x-values represent the number of convolutions n (on a logarithmic scale). In Figure 4(a), the value $1 - \alpha$ can be read off on the y-axis. The range of values for the level α goes from 0.50 (top) to $1 - 10^{-7}$ (bottom). The lines (in pairs) show the range where the approximation error $\epsilon := |\log(s_\alpha/q_\alpha)|$ equals a certain threshold. For example for the sum of $n = 8$ Student-t_8 distributed random variables, the only levels for which a normal distribution yields a very good approximation ($\epsilon \leq 0.01$) are the ones with $\alpha \in [0.897, 0.984]$ (and for symmetry reasons also $\alpha \in [0.016, 0.103]$). Allowing for an error $\epsilon \leq 0.05$, for $n = 8$ all quantiles with $\alpha \in [0.0008, 0.9992]$ can be approximated by normal quantiles. In order to read off the quantiles easier for small values of n, we plot the same lines a second time, using a linear scale for the α-values, see Figure 4(b). For the original Student-t_8 distribution ($n = 1$), asking for an error of at most $\epsilon = 0.01$, we observe that only quantiles in the range $\alpha \in [0.959, 0.971]$ (and $\alpha \in [0.029, 0.041]$) can be replaced by normal quantiles. For all levels between 0.041 and 0.959, normal quantiles exceed t_8-quantiles, while for $\alpha > 0.971$ (and $\alpha < 0.029$) replacing t_8-quantiles by normal ones leads to an underestimation (in absolute values).

Repeating this comparison for a Student-t distribution with $\nu = 4$ degrees of freedom yields the expected outcome, see Figure 5. The sum must be taken over a bigger sample (n large) in order that quantiles can be closely approximated by normal ones.

These investigations also make clear that in the limit, as $\alpha \to 1$, scaling a short-term VaR_α to a long-term risk using the square-root-of-time rule is for most situations not appropriate any more, see Brummelhuis and Guégan [5], [6].

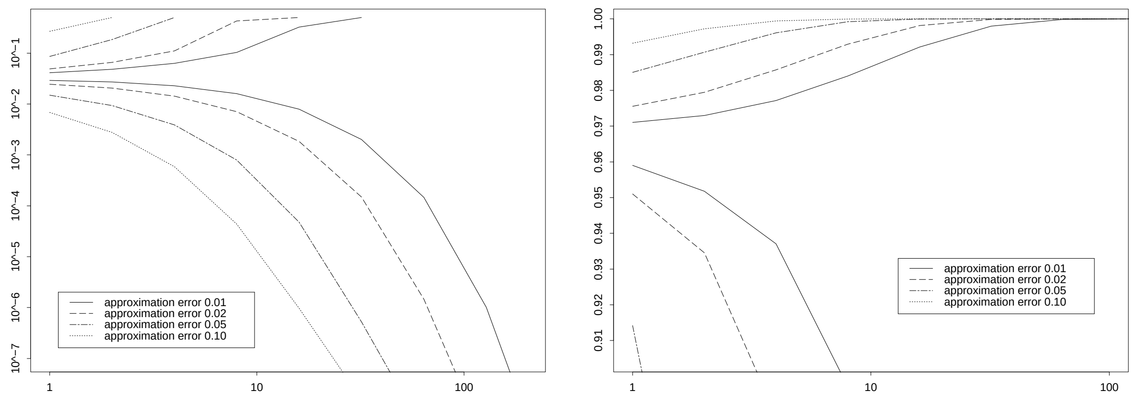

(a) logarithmic scale on the y-axis

(b) linear scale on the y-axis

Figure 4 Contour plots for the area where quantiles of a normal distribution are a good approximation for quantiles of the sum of Student-t_8 distributions. The x-axis indicates the number of convolutions. The y-axis displays one minus the level α, drawn on logarithmic scale (a), and the level α on a linear scale (b), respectively.

2.2.1 *Scaling a 1-day 95% value-at-risk to a 10-day 99% value-at-risk.*

Another problem which is related to the confidence level α is the question how to scale a 1-day 95% value-at-risk to a 10-day 99% value-at-risk. A straightforward method would be to multiply the 1-day 95% value-at-risk with the quotient $q^{\mathcal{N}}_{99\%}/q^{\mathcal{N}}_{95\%}$ (where $q^{\mathcal{N}}_{\alpha}$ denotes the α quantile of a standard normal distribution), and then scale the resulting value with the square-root-of time. But this first step – multiplying with the quotient – is in general not appropriate, as a short investigation shows.

Taking a random walk with Student-t_4 innovations as an illustrative example, we can read off from Figure 5 that using the quantile of a normal distribution as an approximation for the true α quantile yields an underestimation of about 13% for $\alpha = 99\%$, and an overestimation of more than 8% for $\alpha = 95\%$. Hence the error committed when multiplying with $q^{\mathcal{N}}_{99\%}/q^{\mathcal{N}}_{95\%}$ is more than 20%. For 10-day quantiles, the corresponding error is about 7%, composed by an underestimation of 5% and an overestimation of 2%.

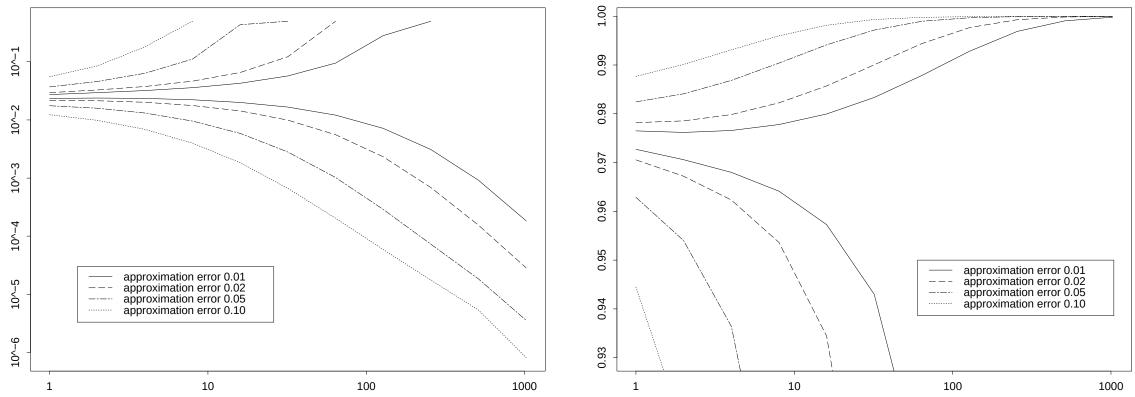

(a) logarithmic scale on the y-axis (b) linear scale on the y-axis

Figure 5 Contour plots for the area where quantiles of a normal distribution are a good approximation for quantiles of the sum of Student-t_4 distributions.

One could now argue that proceeding the other way around – first scaling the 1-day 95% value-at-risk with the square-root-of-time and then committing an error of (only) 7% by multiplying with $q^{\mathcal{N}}_{99\%}/q^{\mathcal{N}}_{95\%}$ – was not that serious. But the problem is that already the first step – multiplying a 1-day 95% value-at-risk with the square-root-of-time – produces a rather big estimation error for realistic models with dependent subsequent log-returns. While for $\alpha = 99\%$ the overestimation of quantiles caused by the square-root-of-time rule is partially compensated by the dependence in the model (which increases the 10-day quantiles), for $\alpha = 95\%$ the square-root-of-time rule causes an underestimation of quantiles, and things get even worse for dependent log-returns.

These considerations make clear that transforming a 1-day 95% value-at-risk into a 10-day 99% value-at-risk is rather delicate. One should first transform the 1-day 95% estimate appropriately into a 1-day 99% estimate, before applying the square-root-of-time scaling rule. An appropriate transformation from one quantile level to the other one requires knowledge of the tail of the one-day distribution, which corresponds to the recommendation to start directly with a 99% quantile level for daily log-returns.

2.3 Scaling rule for stochastic volatility models with jumps.

After having investigated GARCH-type models, we now focus on stochastic volatility models with jumps. We investigate estimates for the unconditional 10-day 99% value-at-risk for such models. More precisely, we are interested in finding the best estimate for the unconditional 10-day 99% value-at-risk if data for not more than 250 trading days are available.

For our investigations, we assume the following stochastic volatility model with jumps for daily log-returns $(X_t)_{t\in\mathbb{Z}}$:

$$\begin{aligned}
X_t &= a\,\sigma_t\,Z_t + b\,J_t\,\epsilon_t, \\
\sigma_t &= \sigma_{t-1}^{\phi}\,e^{c\,Y_t}, \\
\epsilon_t, Z_t, Y_t &\overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1), \\
J_t &\overset{\text{i.i.d.}}{\sim} \text{Bernoulli}(\lambda).
\end{aligned} \tag{2.1}$$

The term $a\,\sigma_t\,Z_t$ represents the stochastic volatility part, while $b\,J_t\,\epsilon_t$ is the jump term. For our analysis, we use the parameters $\lambda = 0.01$, $a = 0.01$, $b = 0.05$, $c = 0.05$ and $\phi = 0.98$. These values are typical values for financial log-return data, see for example Johannes et al. [9] and [10]. Note that the process (2.1) can also be written in the form

$$X_t = \sqrt{a^2\,\sigma_t^2 + b^2\,J_t}\;Z_t \tag{2.2}$$

with a, σ_t, b, J_t and Z_t defined as above. Typical paths of such a stochastic volatility price process and of the underlying volatility process are shown in Figures 6 and 7.

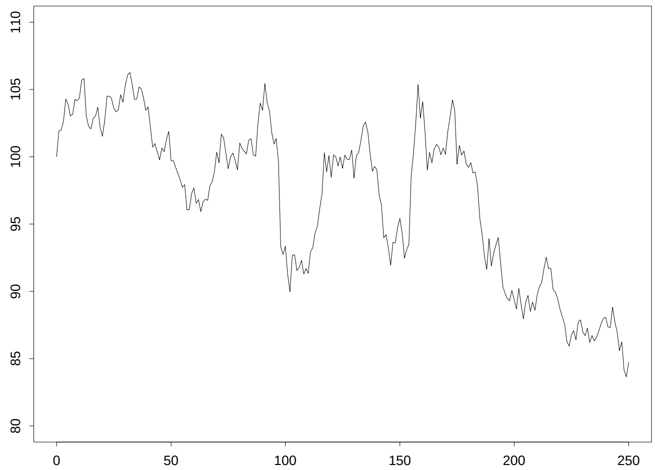

Figure 6 Typical path of the price process in a stochastic volatility model.

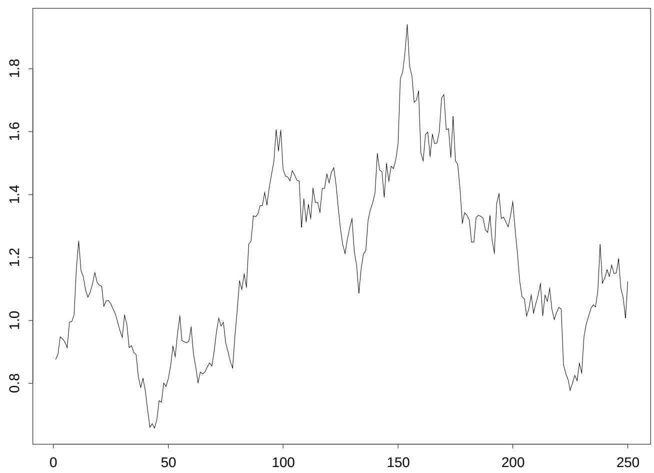

Figure 7 Typical path of the volatility process in a stochastic volatility model.

2.3.1 *Alternative estimators.*

As for the GARCH-type models in Section 2.1, we compare some alternative estimators with the simple square-root-of-time rule for the stochastic volatility model with jumps.

Details on these alternative estimators – which are an estimator using non-overlapping periods, one based on overlapping periods, random resampling, independent resampling, dependent resampling and an extreme value method – can be found in Kaufmann [11]. A graphical evaluation of this comparison in presented in Figure 8. The two methods using directly 10-day log-returns (left graphs) perform much worse than the square-root-of-time scaling rule. Also for dependent resampling (top right graph) the slope of the solid line exceeds 1, indicating a rather poor performance of this method. A comparison of the remaining three methods (random resampling, independent resampling and extreme value method) with the square-root-of-time rule shows that there is no significant difference in the performance of these methods. All four of them – including square-root-of-time scaling – are well suited for estimating the 10-day value-at-risk in the present stochastic volatility model.

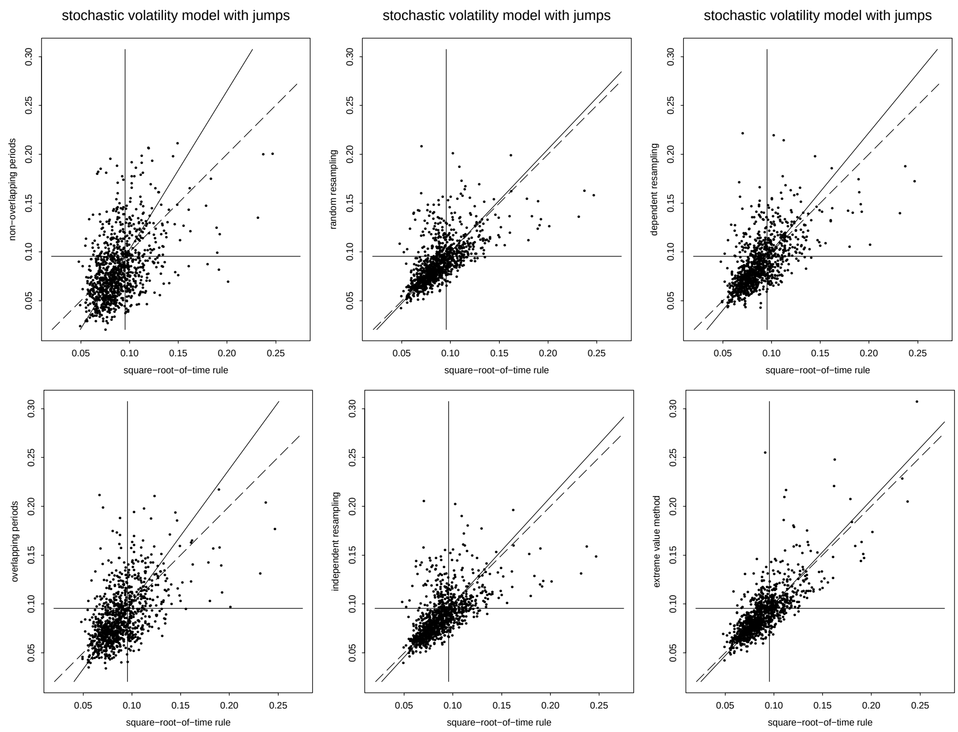

Figure 8 Stochastic volatility model: comparison of the quantile estimation methods.

2.3.2 *Sensitivity analysis.*

We conclude this section about stochastic volatility models by investigating the goodness of fit when parameters are changed. In Figure 9 the square-root-of-time rule is compared with the true 10-day value-at-risk in stochastic volatility models for various choices of λ (keeping $a = 0.01$, $b = 0.05$, $c = 0.05$ and $\phi = 0.98$ fixed). At first sight it might be surprising that the underestimation for small values of λ changes into an overestimation for $\lambda > 0.04$. The reason for this change is the fact that for low jump intensities λ, one-day returns are affected by the jump term only far out in the tail. If λ is increased, the one-day 99% value-at-risk is suddenly strongly affected. For the 10-day value-at-risk, this effect is less marked. This explains the shape of the curves in Figure 9.

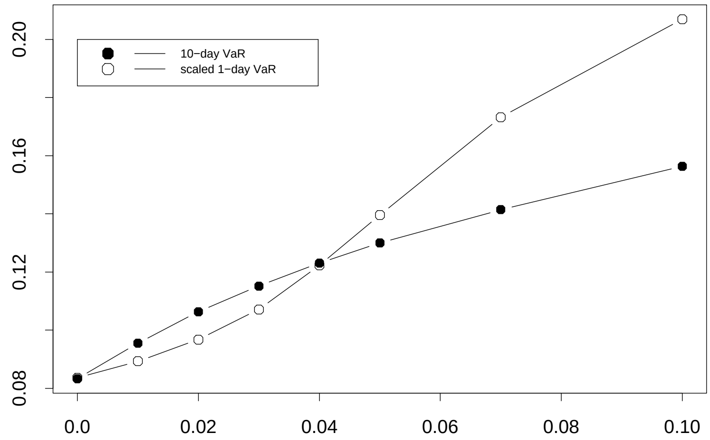

Figure 9 Simulated values for the 10-day 99% value-at-risk in stochastic volatility models: true VaR (black symbols) and VaR using the $\sqrt{10}$-rule (white symbols). Parameters: $a = 0.01$, $b = 0.05$, $c = 0.05$, $\phi = 0.98$, $\lambda \in [0.00, 0.10]$.

3 1-year risks

After concentrating on 10-day risks in the first part of this paper, in this section the evolution of risk factors for a one-year horizon is studied. We start with an overview of possible approaches that can be used to model yearly risks. We then investigate dynamical models such as random walks, AR(p) and GARCH(1,1) processes which allow for the modelling of price changes. Additionally we propose a static approach based on heavy-tailed distributions.

When modelling yearly data, one typically encounters the problem that financial time series are non-stationary. Since market conditions change over the years, it is not possible to go far back in history in order to gather data which is representative for today's situation. This is sometimes referred to as *lack of yearly returns*. Finally, properties of yearly data are different from those of daily or weekly data. They are less skewed and less heavy-tailed.

Notwithstanding this, our aim is to estimate yearly risks. One possible way to handle these inconveniences is to first fix a horizon $h < 1$ year for which data can be modelled, and to use a scaling rule for the gap between h and 1 year. This is the strategy we follow in this section, see Figure 10 for an illustration.

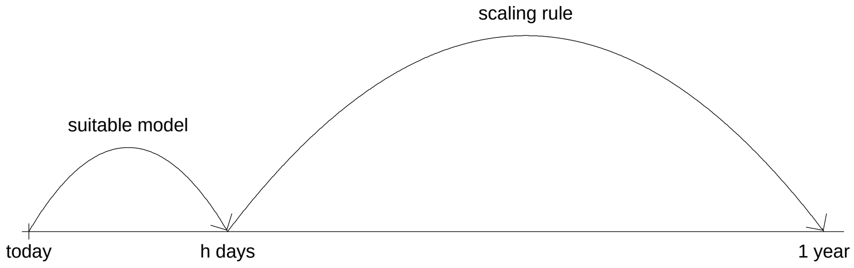

Figure 10 The two steps: first a suitable model is calibrated on a time horizon h, then the risk estimates are scaled from h to one year.

3.1 Models.

For the four models mentioned before, we implemented the above strategy, and compared their performances in estimating yearly risks. As risk measures, we took 1-year value-at-risk and 1-year expected shortfall, each on the 95% and 99% level.

We first give a quick overview of the models.

3.1.1 *Random walk with normal innovations.*

A very simple and often very useful model consists of assuming financial log-data $(s_t)_{t\in h\mathbb{N}}$ to follow a randow walk with constant trend and normal innovations:

$$s_t = s_{t-h} + X_t, \qquad X_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \sigma^2) \quad \text{for } t \in h\mathbb{N}.$$

The random variables $(X_t)_{t\in h\mathbb{N}}$ represent h-day log-returns.

For this model, the square-root-of-time rule – accounting for the trend – can be used to scale h-day risk measures to 1-year risk measures.

3.1.2 *Autoregressive processes.*

Assume (s_t) follows an AR(p) model with trend and normal innovations,

$$s_t = \sum_{i=1}^{p} a_i \, s_{t-ih} + \epsilon_t \quad \text{for } t \in h\mathbb{N},$$

where $\epsilon_t \sim \mathcal{N}(\mu_0 + \mu_1 \, t, \sigma^2)$, independent. Then the 1-year value-at-risk and expected shortfall can be calculated as a function of the parameters μ_1, σ and a_i, and the current and past values of $(s_t)_{t\in h\mathbb{N}}$.

3.1.3 *GARCH processes.*

Let (X_t) be a GARCH(1,1) process with Student-t_ν distributed innovations for h-day log-returns, i.e.

$$\begin{aligned} X_t &= \mu + \sigma_t \, \epsilon_t \quad \text{for } t \in h\mathbb{N}, \\ \sigma_t^2 &= \alpha_0 + \alpha_1 (X_{t-h} - \mu)^2 + \beta_1 \sigma_{t-h}^2, \end{aligned}$$

where $\epsilon_t \overset{\text{i.i.d.}}{\sim} t_\nu$, $\mathbb{E}[\epsilon_t] = 0$, $\mathbb{E}[\epsilon_t^2] = 1$, and the degree of freedom ν is to be estimated from data. Then the 1-year log-returns follow a so-called weak GARCH(1,1) process, see Drost and Nijman [7]. The corresponding value-at-risk and expected shortfall can be calculated as a function of the above parameters and the current and past values of $(X_t)_{t\in h\mathbb{N}}$.

3.1.4 *Random walk with heavy-tailed innovations.*

Here, the h-day log-returns $(X_t)_{t\in h\mathbb{N}}$ are assumed to have a heavy-tailed distribution, i.e.

$$\mathbb{P}[X_t < -x] = x^{-\alpha} L(x) \quad \text{as } x \to \infty,$$

where $\alpha \in \mathbb{R}^+$ and L is a slowly varying function, which means that $\lim_{x\to\infty} L(sx)/L(x) = 1$ for all $s > 0$.

Also in this case, the 1-year value-at-risk and expected shortfall can be calculated based on the parameter α and on the observed data. The details for all four models are to be found in Kaufmann [11].

3.2 Backtesting.

The suitability of these models for estimating one-year financial risks can be assessed by comparing the estimates for expected shortfall and value-at-risk with observed return data. We do this comparison for stock indices, foreign exchange rates, 10-year government bonds, and single stocks.

For backtesting the forecasted expected shortfall $\widehat{\mathrm{ES}}_{\alpha,t}$, we introduce two measures. The first measure V_1^{ES} evaluates excesses below the negative of the estimated value-at-risk $\widehat{\mathrm{VaR}}_{\alpha,t}$. This is a standard method for backtesting expected shortfall estimates. In detail we proceed as follows. Every model provides for each point of time t an estimation $\widehat{\mathrm{ES}}_{\alpha,t}$ for the one-year ahead expected shortfall $\mathrm{ES}_{\alpha,t}$. First, the difference between the observed one-year (k-period) return R_t^k and the negative of the estimation $\widehat{\mathrm{ES}}_{\alpha,t}$ is taken, and then the conditional average of these differences is calculated, conditioned on $\{R_t^k < -\widehat{\mathrm{VaR}}_{\alpha,t}\}$,

$$V_1^{\mathrm{ES}} = \frac{\sum_{t=t_0}^{t_1} \left(R_t^k - (-\widehat{\mathrm{ES}}_{\alpha,t})\right) \mathbf{1}_{\{R_t^k < -\widehat{\mathrm{VaR}}_{\alpha,t}\}}}{\sum_{t=t_0}^{t_1} \mathbf{1}_{\{R_t^k < -\widehat{\mathrm{VaR}}_{\alpha,t}\}}}.$$

A good estimation for expected shortfall leads to a low absolute value of V_1^{ES}.

This first measure is similar to the *theoretical* definition of expected shortfall. Its weakness is that it depends strongly on the value-at-risk estimates (without adequately reflecting the goodness/badness of these values), since only values which fall below the value-at-risk threshold are considered. This is possibly a fraction which is far away from $(1-\alpha)\cdot 100\%$ of the values – which is the fraction one would actually like to average over. Hence, when analysing the values of V_1^{ES}, these results should be combined with the ones given by the frequency of exceedances V^{freq} which will be described below.

In practice, one is primarily interested in the loss incurred in a one in $1/(1-\alpha)$-event, as opposed to getting information about the behaviour below a certain estimated value. Therefore we introduce a second measure V_2^{ES}, which evaluates values below the one in $1/(1-\alpha)$-event:

$$V_2^{\mathrm{ES}} = \frac{\sum_{t=t_0}^{t_1} D_t \, \mathbf{1}_{\{D_t < D^\alpha\}}}{\sum_{t=t_0}^{t_1} \mathbf{1}_{\{D_t < D^\alpha\}}},$$

where $D_t := R_t^k - (-\widehat{\mathrm{ES}}_{\alpha,t})$ and D^α denotes the empirical $(1-\alpha)$-quantile of these differences $\{D_t\}_{t_0 \le t \le t_1}$. Note that, since $\widehat{\mathrm{ES}}_{\alpha,t}$ is an estimate on a level α, we expect D_t to be negative in somewhat less than one out of $1/(1-\alpha)$ cases. A good estimation for expected shortfall again leads to a low absolute value of V_2^{ES}.

The next step is to combine the two measures V_1^{ES} and V_2^{ES}:

$$V^{\mathrm{ES}} = \frac{|V_1^{\mathrm{ES}}| + |V_2^{\mathrm{ES}}|}{2}.$$

This measure tells how well the forecasted one-year expected shortfall fits real data. It is used in our investigations to backtest the quality of the models.

We introduce one more measure that provides information about the quality of the estimators: the frequency of exceedances

$$V^{\mathrm{freq}} = \frac{1}{t_1 - t_0 + 1} \sum_{t=t_0}^{t_1} \mathbf{1}_{\{R_t^k < -\widehat{\mathrm{VaR}}_{\alpha,t}\}}.$$

This measure is used by the Basel Committee on Banking Supervision, which in order to encourage institutions to report their value-at-risk numbers, devised a system in which penalties are set depending on the frequency of violations. See The New Basel Accord by the Basel Committee on Banking Supervision [4] for a detailed description. Here, a good estimation for value-at-risk leads to a value of V^{freq} which is close to the level $1-\alpha$.

3.3 Results for 1-year risks.

We restrict to the main results here. For details we refer to Kaufmann [11].

The random walk model performs in general better than the other models under investigation. It provides satisfactory results across all classes of data and for both confidence levels investigated (95%, 99%). However, like all the other models under investigation, the risk estimates for single stocks are not as good as those for foreign exchange rates, stock indices, and 10-year bonds.

The optimal calibration horizon is about one month. Based on monthly data, the square-root-of-time rule (accounting for trends) can be applied for estimating one-year risks. An important reason for not recommending longer calibration horizons are the statistical restrictions, such as the sample size for estimating reliable model parameters and hence reliable risk measures. On the other hand, using higher frequency data (daily data for example) is not recommended either, since their properties (e.g. leptokurtosis) are clearly different from those of yearly data. Estimating a certain percentile might still be fine, but it would not be possible to estimate the whole distribution function appropriately.

In contrast to short term horizons, for a one-year period a good estimate of the trend of (log-)returns is critical when measuring risks.

4 Conclusions

In Section 2 we saw that the square-root-of-time scaling rule performs very well to scale risks from a 1 day horizon to a 10 day horizon. However, the reasons for this good performance are non-trivial. Each situation has to be investigated separately. The square-root-of-time rule should not be applied before checking its appropriateness.

For estimating 1-year risks, it can be recommended to use a random walk model with a constant trend calibrated on a time horizon of about one month, and to apply the square-root-of-time rule.

References

[1] Acerbi C. and Tasche D. (2002). *On the coherence of Expected Shortfall*, Journal of Banking and Finance 26, no. 7, 1491–1507.

[2] Artzner P., Delbaen F., Eber, J.M. and Heath D. (1999). *Coherent Measures of Risk*, Mathematical Finance 9, no. 3, 203–228.

[3] Basel Committee on Banking Supervision (1996). *Amendment to the capital accord to incorporate market risk*, BIS, Basel, Switzerland.

[4] Basel Committee on Banking Supervision (2003). *The New Basel Accord*, BIS, Basel, Switzerland.

[5] Brummelhuis R.G.M. and Guégan D. (2000). *Extreme values of conditional distributions of GARCH(1,1) processes*, prépublication 00.08, Université de Reims, Dépt. de math.

[6] Brummelhuis R.G.M. and Guégan D. (2000). *Multi-period Conditional Distribution Function for Heteroscedastic Models with Applications to VaR*, prépublication 00.13, Université de Reims, Dépt. de math.

[7] Drost F.C. and Nijman T.E. (1993). *Temporal Aggregation of GARCH Processes*, Econometrica, 61, 909–927.

[8] Embrechts P., Klüppelberg C. and Mikosch T. (1997). *Modelling Extremal Events for Insurance and Finance*, Springer-Verlag, Berlin.

[9] Johannes M.S., Kumar R. and Polson N.G. (1999). *State Dependent Jump Models: How do U.S. Equity Markets Jump?* Working Paper, available at http://www-1.gsb.columbia.edu/faculty/mjohannes/research.html.

[10] Johannes M.S., Polson N.G. and Stroud J.R. (2004). *Sequential Parameter Estimation in Stochastic Volatility Models with Jumps*, Working Paper, available at http://www-1.gsb.columbia.edu/faculty/mjohannes/research.html.

[11] Kaufmann R. (2004). *Long-Term Risk Management*, Ph.D. Thesis, ETH Zurich.

[12] McNeil A.J. and 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.