



Strategic Long-Term Financial Risks: Single Risk Factors

PAUL EMBRECHTS
ROGER KAUFMANN
PIERRE PATIE

Department of Mathematics, ETH Zurich, CH-8092 Zürich, Switzerland

embrechts@math.ethz.ch
roger.kaufmann@swisslife.ch
patie@math.ethz.ch

Abstract. The question of the measurement of strategic long-term financial risks is of considerable importance. Existing modelling instruments allow for a good measurement of market risks of trading books over relatively small time intervals. However, these approaches may have severe deficiencies if they are routinely applied to longer time periods. In this paper we give an overview on methodologies that can be used to model the evolution of risk factors over a one-year horizon. Different models are tested on financial time series data by performing backtesting on their expected shortfall predictions.

Keywords: expected shortfall, value-at-risk, scaling rules, random walk, autoregressive model, GARCH process, extreme value theory

1. Introduction

Most of the current research in asset management and market risk management focuses on short-term risk (daily, weekly). Clearly, in many cases, a long-term analysis (e.g. quarterly, yearly) is just as relevant and important. Indeed, in financial institutions (e.g. banks) risk measures are typically calculated for a one-day to two-week horizon. In insurance for instance risk exposures must be measured and managed over much longer (i.e. yearly) time spans. While for the measurement of short-term financial risks some consistent and reliable frameworks already exist, for longer time horizons, only relatively few papers can be found in the academic literature.

Existing modelling instruments such as RiskMetrics allow for a relatively good measurement of market risks of trading books. These models, however, have some severe deficiencies if they are applied to longer time periods (typically one year), as needed in the case of strategic investments of institutional investors or within insurance. A notable exception is the ForeSight Technical Document [29] of the RiskMetrics Group which focuses on forecasting beyond a two year horizon.

In this paper we develop a theoretically well-understood and empirically-founded conceptual framework for the measurement of long-term financial risk of strategic investment portfolios. The main criterion to judge the appropriateness of different models is the reliability

Research of the last two authors was supported by Credit Suisse Group, Swiss Re and UBS AG through RiskLab, Switzerland.

of expected shortfall estimates of profit and loss distributions for financial portfolios over a period of one year. Expected shortfall was chosen as risk measure because it is coherent for continuous distributions in the sense of Artzner et al. [3].

An obvious approach to model yearly financial risks is to use standard methodologies for short-term risks (i.e. for a horizon $h < 1$ year), and to apply a scaling rule for the gap between h and 1 year. It is well known that daily (or higher frequency) returns are dependent due to volatility clustering, while returns over a longer time horizon (fortnightly, monthly data) are closer to being independent. Because of the weaker dependence for lower frequency data, time aggregation rules usually perform better when starting with longer horizons. On the other hand, the serious statistical restrictions due to the small number of low frequency data should also be taken into account. There is a tradeoff between bias and variance. The pivotal point is the choice of the data frequency on which to calibrate the models.

In this paper we present our findings on the evolution of single risk factors. Sections 3 to 6 give an overview on approaches that can be used to model the risk factors. In these sections we first investigate dynamical models like random walks, $AR(p)$ and $GARCH(1,1)$ models, which allow to model price changes. Then we propose a static approach based on heavy-tailed distributions. We follow a unified framework for the presentation of the results across the different approaches. First, the model including the statistical tools used to estimate the parameters is described, then it is shown how returns can be aggregated. Finally, the methodology to compute value-at-risk and expected shortfall is given. In Section 7 the performances of the different models are compared with each other. First they are calibrated using exchange rate data, stock indices, 10-year government bonds and single stock data at different frequencies. Performing backtesting for the corresponding expected shortfall (and value-at-risk) predictions finally answers the question which model (and at which frequency) gives the most reliable estimation of one-year asset risk. Additionally, a variance analysis for the time series models is provided, and confidence intervals for expected shortfall and value-at-risk for the random walk approach are computed. The final section contains the main conclusions that can be drawn from these investigations.

In order to keep the paper relatively short, the reader is occasionally referred to a RiskLab Technical Report Kaufmann and Patie [27] for further details.

2. Definitions and notation

We introduce a minimum set of definitions and notation. Fix some real-valued continuously distributed random variable R with finite mean on some probability space $(\Omega, \mathcal{A}, \mathbb{P})$. R is considered the return of some asset. By $\mathbb{E}[\cdot]$ we denote expectation with respect to \mathbb{P} . Fix also some confidence level p .

Definition 2.1. The value-at-risk at level p of R is defined as

$$\text{VaR}^p(R) = -\inf\{x \in \mathbb{R} \mid \mathbb{P}[R \leq x] \geq p\},$$

i.e. $\text{VaR}^p(R)$ is the negative p -quantile of R .

Definition 2.2. The expected shortfall at level p of R is defined as

$$\text{ES}^p(R) = -\mathbb{E}[R \mid R < -\text{VaR}^p(R)].$$

For continuous distributions, expected shortfall is coherent, whereas value-at-risk does in general not fulfil the subadditivity property required for a risk measure to be coherent.

Definition 2.3. A measure of risk ρ is *subadditive*, if $\rho(R^a + R^b) \leq \rho(R^a) + \rho(R^b)$ for all random variables R^a, R^b .

Artzner et al. [3] explain why the subadditivity property is a natural requirement for risk measures and show that this property holds under certain restrictions on the discounted risks, as well as on the underlying probability space. Rockafellar and Uryasev [36] give an equivalent representation of the conditional expected value as the solution of an optimization problem only requiring the probability distribution of the risk to be continuous. From this minimization problem, Pflug [35] shows the convexity (including the subadditivity property) of expected shortfall. Finally, Acerbi and Tasche [1] and Rockafellar and Uryasev [37] give an alternative definition of conditional value-at-risk, which is identical to expected shortfall for continuous distributions, but stays coherent even in the case of discontinuous underlying return distributions. Moreover, this alternative definition of conditional value-at-risk (and hence for continuous distributions also expected shortfall) is consistent with second degree stochastic dominance, which is important for decision theory, see Ogryczak and Ruszczyński [34]. In Embrechts et al. [20] it is shown that value-at-risk is coherent for elliptically distributed risk factors.

Notation. Throughout the rest of the paper we use the following notation. S_t denotes the asset price at time t , $s_t = \log S_t$ is the asset log-price, h denotes the length (in days) of one period, $R_t = \frac{S_t - S_{t-h}}{S_{t-h}}$ is the one-period asset return, $r_t = s_t - s_{t-h}$ the one-period asset log-return, and r_t^k denotes the k -period (i.e. kh -day) asset log-return (think of kh days being 1 year). The mean μ of one-period asset log-returns is assumed to be constant over time, and $\bar{r}_t = r_t - \mu$ is the centered one-period asset log-return.

3. Random walks

A simple, but for many practical applications crucial starting model for h -day log-returns is the random walk model with normal innovations.

Definition 3.1. The process $(s_t)_{t \in h\mathbb{N}}$ is a random walk process with constant trend and normal innovations, if for all $t \in h\mathbb{N}$ it satisfies the equation

$$s_t = s_{t-h} + r_t, \quad \text{with } r_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \sigma^2). \quad (3.1)$$

3.1. Scaling rule

When changes in the log-price are independent and identically distributed with mean $\mu = 0$ and variance $\sigma^2 < \infty$, then the k -period log-returns $r_t^k = \sum_{i=0}^{k-1} r_{t-ih}$ have mean zero and standard deviation $\sqrt{k}\sigma$. Hence the one-period volatility can be multiplied by the square-root-of-time to calculate the k -period volatility in the i.i.d. case. Assuming additionally normality like in (3.1), the same scaling law holds for the quantiles.

It is well known that daily log-returns show neither independent nor normally distributed behaviour (see Fama [21], or Campbell et al. [8]). Lower frequency data are closer to the normality assumption which is needed to apply the square-root-of-time rule; see for instance Campbell et al. [8] and Dacorogna et al. [11]. Hence, this scaling rule should work better for models fitted to monthly or even lower frequency data.

In general, the square-root-of-time scaling rule cannot be applied directly to log-returns. First the mean μ has to be subtracted. For i.i.d. normally distributed one-period log-returns with constant mean $\mu = \mathbb{E}[r_t]$ and constant variance $\sigma^2 = \mathbb{E}[r_t^2]$, the following rule for k -period log-returns holds: $r_t^k \stackrel{d}{\sim} k\mu + \sqrt{k}(r_t - \mu)$, where $\stackrel{d}{\sim}$ denotes equality in distribution.

The one-period mean μ and the one-period standard deviation σ can be estimated via $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n r_{ih}$ and $\hat{\sigma} = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (r_{ih} - \hat{\mu})^2}$. The square-root-of-time rule leads to $\hat{\sigma}^k = \sqrt{k} \hat{\sigma}$ for the k -period variance. The trend for k periods can be estimated by $\hat{\mu}^k = k \hat{\mu}$.

3.2. Estimation of value-at-risk and expected shortfall

The estimate for a k -period (think of one year) value-at-risk at level p is

$$\widehat{\text{VaR}}^p = -(\exp(\hat{\mu}^k + \hat{\sigma}^k x^p) - 1), \quad (3.2)$$

and correspondingly for expected shortfall,

$$\widehat{\text{ES}}^p = -\left(\exp\left(\hat{\mu}^k + \frac{(\hat{\sigma}^k)^2}{2}\right) \frac{\Phi(x^p - \hat{\sigma}^k)}{p} - 1 \right), \quad (3.3)$$

where x^p is the p -quantile of a standard normal random variable, and Φ denotes the cumulative standard normal distribution function.

4. Autoregressive processes

A first model that takes into account the dependence of subsequent log-returns is the AR(p) process.

Definition 4.1. The process $(s_t)_{t \in h\mathbb{N}}$ is an AR(p) process with trend, if for all $t \in h\mathbb{N}$ it satisfies the equation

$$s_t = \sum_{i=1}^p a_i s_{t-ih} + \epsilon_t, \quad (4.1)$$

where a_i are (constant) coefficients fulfilling the stationarity condition $\sum_{i=1}^p |a_i| < 1$, $\epsilon_t \sim \mathcal{N}(\mu_0 + \mu_1 t, \sigma^2)$, $(\epsilon_t)_{t \in h\mathbb{N}}$ are independent, and σ^2 is the variance of the innovation process.

The constant term μ_0 has only a translation effect on the process $(s_t)_{t \in h\mathbb{N}}$. Hence setting the constant term to zero has no effect on log-returns $(r_t)_{t \in h\mathbb{N}}$. Equation (4.1) can be rewritten as an AR(p)-process with i.i.d. $\mathcal{N}(0, \sigma^2)$ innovations:

$$\bar{s}_t = \sum_{i=1}^p a_i \bar{s}_{t-ih} + \bar{\epsilon}_t, \quad t \in h\mathbb{N}, \quad (4.2)$$

where $\bar{s}_t = s_t - \mu t$ are the detrended log-prices, $\mu = \frac{\mu_1}{1 - \sum_{i=1}^p a_i}$, $\bar{\epsilon}_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$, and a_i, σ^2 are as in (4.1).

4.1. Scaling rule

For one-period steps, the AR(p)-process (4.2) can be used to model the detrended log-prices $(\bar{s}_t)_{t \in h\mathbb{N}}$. To calculate k -period parameters from a sample of n h -day periods, we proceed as follows:

- Subtract the linear trend from log-prices s_t : $\bar{s}_t = s_t - \hat{\mu}t$ for $t \in h\mathbb{N}$, where $\hat{\mu} = \frac{s_{nh} - s_0}{nh}$.
- Fit the AR(p)-process to the drift-free one-period log-prices $(\bar{s}_t)_{t \in h\mathbb{N}}$: maximum likelihood estimation (MLE) gives the estimates \hat{p}, \hat{a}_i ($i = 1, \dots, \hat{p}$) and $\hat{\sigma}_\epsilon$. Note that $\hat{\sigma}_\epsilon = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (\bar{\epsilon}_{ih})^2}$ is the volatility of the innovation process $(\bar{\epsilon}_t)_{t \in h\mathbb{N}}$ and not of the detrended log-prices $(\bar{s}_t)_{t \in h\mathbb{N}}$.
- At time t , forecast the k -period log-return using the relation $\hat{\mu}^k = kh\hat{\mu} + \hat{m}$, where $\hat{m} = \tilde{s}_{t+kh} - \tilde{s}_t$, and \tilde{s}_{t+kh} is defined recursively: $\tilde{s}_{t+jh} = \sum_{i=1}^{\hat{p}} \hat{a}_i \tilde{s}_{t+(j-i)h}$ ($j = 1, \dots, k$), and $\tilde{s}_u = \bar{s}_u$ for $u \leq t$.
- Forecast the k -period volatility using the scheme $\hat{\sigma}^k = \hat{\sigma}_\epsilon \sqrt{\sum_{j=0}^{k-1} \delta_j^2}$, where $\delta_0 = 1$, $\delta_j = \sum_{i=1}^j \hat{a}_i \delta_{j-i}$, $\hat{a}_i = 0$ for all $i > \hat{p}$.

Conditioned on s_t , the one-year forecast s_{t+kh} has distribution $\mathcal{N}(s_t + \mu^k, (\sigma^k)^2)$, where μ^k and σ^k can be estimated as shown in the above steps.

4.2. Estimation of value-at-risk and expected shortfall

Like in the random walk model, the k -period value-at-risk at level p can be estimated by

$$\widehat{\text{VaR}}^p = -(\exp(\hat{\mu}^k + \hat{\sigma}^k x^p) - 1), \quad (4.3)$$

and the corresponding expected shortfall by

$$\widehat{\text{ES}}^p = -\left(\exp\left(\hat{\mu}^k + \frac{(\hat{\sigma}^k)^2}{2}\right) \frac{\Phi(x^p - \hat{\sigma}^k)}{p} - 1 \right), \quad (4.4)$$

where x^p is the p -quantile of a standard normal random variable, and Φ denotes the cumulative standard normal distribution function.

5. The GARCH(1,1) model

A widely used approach to model the dependence between subsequent one-period log-returns via the volatility, is the famous family of GARCH(1,1) processes.

Definition 5.1. The process $(s_t)_{t \in h\mathbb{N}}$ is generated by a GARCH(1,1) process, if for all $t \in h\mathbb{N}$ it satisfies the equation

$$s_t = s_{t-h} + r_t, \quad \text{where } r_t = \bar{r}_t + \mu, \quad (5.1)$$

and $(\bar{r}_t)_{t \in h\mathbb{N}}$ is a GARCH(1,1) process:

$$\begin{aligned} \bar{r}_t &= \sigma_t \epsilon_t, \\ \sigma_t^2 &= \alpha_0 + \alpha_1 \bar{r}_{t-h}^2 + \beta_1 \sigma_{t-h}^2, \\ \epsilon_t & \text{ i.i.d., } \mathbb{E}[\epsilon_t] = 0, \quad \mathbb{E}[\epsilon_t^2] = 1. \end{aligned} \quad (5.2)$$

In this paper, the innovations $(\epsilon_t)_{t \in h\mathbb{N}}$ are assumed to be Student- t distributed, and the stationarity conditions $0 < \alpha_0 < \infty$, $\alpha_1 \geq 0$, $\beta_1 \geq 0$ and $\alpha_1 + \beta_1 < 1$ are assumed to be fulfilled.

This model is known to be well-suited for reproducing the heteroscedastic behaviour of the conditional volatility of financial daily returns. In practice, it is intensively used for estimating the 10-day value-at-risk; see however the critical paper Stărică [38]. One often fits the model to daily data, and then applies the square-root-of-time rule to extend the results to a longer time horizon (e.g. 10 days). Brummelhuis and Kaufmann [7] show that the square-root-of-time rule works well to scale a one-day 1% value-at-risk to a 10-day 1% value-at-risk in GARCH(1,1) processes. However, this scaling is inappropriate when applied to long term horizon conversion. Christoffersen et al. [10] discuss the square-root-of-time scaling rule, showing that in GARCH(1,1) models for large k (such as $k > 15$ days) applying this standard rule to the strongly varying short-time volatilities σ_t leads to an overestimation of (actually low) volatility fluctuations of k -period log-returns (see also McNeil and Frey [31] for a critical discussion of this scaling rule).

5.1. Scaling rule

Drost and Nijman [13] investigate the temporal aggregation of so-called weak GARCH processes. We refer to their paper for the formal definition of this version of GARCH processes. The so-called Drost-Nijman formula allows to extend short-term risk estimates to longer horizons. Suppose the one-period log-prices $(s_t)_{t \in h\mathbb{N}}$ are generated by a GARCH(1,1) process as in Definition 5.1, with symmetric innovations. Drost and Nijman [13] show that, under regularity conditions, the corresponding k -period log-prices are generated by a weak GARCH(1,1) process with

$$(\sigma_t^k)^2 = \alpha_{k,0} + \alpha_{k,1} (\bar{r}_{t-kh}^k)^2 + \beta_{k,1} (\sigma_{t-kh}^k)^2, \quad (5.3)$$

where $(\bar{r}_t^k)_{t \in kh\mathbb{N}}$ denote the centered k -period log-returns, $\alpha_{k,0} = k\alpha_0 \frac{1 - (\alpha_1 + \beta_1)^k}{1 - (\alpha_1 + \beta_1)}$, $\alpha_{k,1} = (\alpha_1 + \beta_1)^k - \beta_{k,1}$, and $|\beta_{k,1}| < 1$ is the solution of the quadratic equation $\frac{\beta_{k,1}}{1 + (\beta_{k,1})^2} = \frac{a(\alpha_1 + \beta_1)^k - b}{a(1 + (\alpha_1 + \beta_1)^{2k} - 2b)}$. Further,

$$\begin{aligned} a &= k(1 - \beta_1)^2 + 2k(k-1) \frac{(1 - \beta_1 - \alpha_1)^2(1 - \beta_1^2 - 2\alpha_1\beta_1)}{(\tilde{\kappa} - 1)(1 - (\alpha_1 + \beta_1)^2)} \\ &\quad + 4 \frac{(k-1 - k(\alpha_1 + \beta_1) + (\alpha_1 + \beta_1)^k)(\alpha_1 - \alpha_1\beta_1(\alpha_1 + \beta_1))}{1 - (\alpha_1 + \beta_1)^2} \quad \text{and} \\ b &= (\alpha_1 - \alpha_1\beta_1(\alpha_1 + \beta_1)) \frac{1 - (\alpha_1 + \beta_1)^{2k}}{1 - (\alpha_1 + \beta_1)^2}. \end{aligned}$$

Note that the Drost-Nijman scaling rule does not give exact information about the distribution of the innovations $(\epsilon_t^k)_{t \in kh\mathbb{N}}$. Only even moments (up to order four) are calculated. Indeed, all odd moments are zero because of the symmetry assumption for one-period innovations. We follow the approach of modelling k -period innovations $(\epsilon_t^k)_{t \in kh\mathbb{N}}$ with Student- t distributions with correct first five moments (i.e. two even moments). The second moment $\mathbb{E}[\epsilon_t^2]$ equals one. The fourth moment, the conditional k -period kurtosis κ^k , can be calculated via the unconditional k -period kurtosis $\tilde{\kappa}^k$:

$$\kappa^k = \frac{1 - (\alpha_1 + \beta_1)^2 + \alpha_1^2}{1 - (\alpha_1 + \beta_1)^2 + \alpha_1^2 \tilde{\kappa}^k} \tilde{\kappa}^k. \quad (5.4)$$

The unconditional kurtosis $\tilde{\kappa}^k$ of $(\bar{r}_t^k)_{t \in kh\mathbb{N}}$ is given by

$$\tilde{\kappa}^k = 3 + \frac{\tilde{\kappa} - 3}{k} + 6(\tilde{\kappa} - 1) \frac{(k-1 - k(\alpha_1 + \beta_1) + (\alpha_1 + \beta_1)^k)(\alpha_1 - \alpha_1\beta_1(\alpha_1 + \beta_1))}{k^2(1 - \alpha_1 - \beta_1)^2(1 - \beta_1^2 - 2\alpha_1\beta_1)} \quad (5.5)$$

(see Drost and Nijman [13]), where $\tilde{\kappa}$ denotes the unconditional one-period kurtosis, which can be expressed via the conditional one:

$$\tilde{\kappa} = \frac{1 - (\alpha_1 + \beta_1)^2}{1 - (\alpha_1 + \beta_1)^2 - \alpha_1^2(\kappa - 1)} \kappa. \quad (5.6)$$

Proofs of Eqs. (5.4) and (5.6) can be found in the Technical Report, Kaufmann and Patie [27].

When Student- t distributions are assumed for one-period innovations $(\epsilon_t)_{t \in h\mathbb{N}}$, the estimation procedure for the parameter ν^k of the Student- t_{ν^k} distributed k -period innovations consists of the following steps:

- Estimate the parameter ν of the Student- t_ν distributed one-period innovations.
- Calculate the conditional kurtosis κ via $\kappa = \frac{3\nu-6}{\nu-4}$.
- Calculate the unconditional kurtosis $\tilde{\kappa}$ via (5.6).
- Apply (5.5) to get the unconditional kurtosis $\tilde{\kappa}^k$ for k -period log-returns.
- Calculate the conditional kurtosis κ^k for k -period log-returns via (5.4).
- Calculate the parameter ν^k of the Student- t_{ν^k} distributed innovations $(\epsilon_t^k)_{t \in kh\mathbb{N}}$ via $\nu^k = \frac{4\kappa^k-6}{\kappa^k-3}$.

Since $\alpha_1 + \beta_1 < 1$ (and $\alpha_1 \geq 0, \beta_1 \geq 0$), the stationarity condition $0 \leq \alpha_{k,1} + \beta_{k,1} < 1$ is fulfilled for the k -period parameters, and the asymptotic behaviour $\lim_{k \rightarrow \infty} \frac{\alpha_{k,0}}{k} = \frac{\alpha_0}{1 - (\alpha_1 + \beta_1)}$ holds. From the formulas for $\alpha_{k,1}$ and $\beta_{k,1}$, it can be seen that $\alpha_{k,1} \rightarrow 0$ and $\beta_{k,1} \rightarrow 0$ as $k \rightarrow \infty$. This means that long term volatility fluctuations stay rather small. More precisely, asymptotically the volatility becomes constant, and therefore the weak GARCH(1,1) process behaves in the limit like a random walk.

To make the quasi maximum likelihood estimation (QMLE) procedure for the parameters in the one-period GARCH(1,1) process more stable (but somewhat less flexible), we simplify model (5.2) and continue working with standard normally distributed innovations for one-period log-returns:

$$\bar{r}_t = \sigma_t \epsilon_t, \quad \sigma_t^2 = \alpha_0 + \alpha_1 \bar{r}_{t-h}^2 + \beta_1 \sigma_{t-h}^2, \quad \epsilon_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1). \quad (5.7)$$

After estimating $\hat{\alpha}_0, \hat{\alpha}_1$ and $\hat{\beta}_1$, the Drost-Nijman scaling rule can be applied to calculate $\hat{\alpha}_{k,0}, \hat{\alpha}_{k,1}$ and $\hat{\beta}_{k,1}$ for the centered k -period log-returns. For estimating ν^k , we use the six steps listed above. In the first step, we take the limit $\nu \rightarrow \infty$, which corresponds to assuming standard normally distributed one-period innovations $(\epsilon_t)_{t \in h\mathbb{N}}$, and which yields $\kappa = 3$ in the second step.

An alternative approach to taking $\nu \rightarrow \infty$ would be to continue working with Student- t distributed one-period innovations, and to only assume normality during the estimation procedure (to get $\hat{\alpha}_0, \hat{\alpha}_1$ and $\hat{\beta}_1$). The parameter ν can then be estimated from the residuals. Gouriéroux [24] shows that this so-called pseudo-maximum-likelihood method yields a consistent and asymptotically normal estimator, see also Straumann [41] and Straumann and Mikosch [42].

At time t , the k -period conditional volatility $\hat{\sigma}^k = \hat{\sigma}(t, t)$ can be forecasted using a recursive relation: for $t^* = t - (n-1)kh, \dots, t - kh, t$ (where n stands for the total number of k -period log-returns) $\hat{\sigma}^2(t^*, t) = \hat{\alpha}_0 + \hat{\alpha}_1(r_{t^*}^k - \hat{\mu}^k)^2 + \hat{\beta}_1\hat{\sigma}^2(t^* - kh, t)$, starting with $\hat{\sigma}^2(t - nkh, t) = \frac{k}{nk-1} \sum_{i=0}^{nk-1} (r_{t-ih} - \hat{\mu})^2$. For the k -period mean we use $\hat{\mu}^k = k\hat{\mu}$, where $\hat{\mu}$ stands for the QMLE for the mean one-period log-return up to time t .

5.2. Estimation of value-at-risk and expected shortfall

Since the k -period innovations are fitted with Student- t distributions, the k -period (one-year) value-at-risk at level p can be estimated by

$$\widehat{\text{VaR}}^p = -(\exp(\hat{\mu}^k + \hat{\sigma}^k x_{\hat{v}^k}^p) - 1), \quad (5.8)$$

and the corresponding expected shortfall by

$$\widehat{\text{ES}}^p = -\left(\frac{1}{p} \int_0^p \exp(\hat{\mu}^k + \hat{\sigma}^k x_{\hat{v}^k}^q) dq - 1\right), \quad (5.9)$$

where $x_{\hat{v}^k}^p$ denotes the p -quantile of a Student- $t_{\hat{v}^k}$ distributed random variable with mean zero and variance one. The integral in (5.9) can be evaluated numerically.

As shown in Kaufmann and Patie [27], stationarity conditions are often violated for parameters estimated on a one-month (or longer) horizon. Hence, in the present paper, GARCH(1,1) models are only calibrated on shorter horizons.

6. Heavy-tailed distributions

Extreme Value Theory (EVT) provides procedures for estimating extreme quantiles. It has attracted much interest from the finance industry lately. For an exhaustive description of EVT see the monograph of Embrechts et al. [19]. Applications to risk management are summarised in Embrechts [18].

Definition 6.1. Log-returns $r_t = s_t - s_{t-h}$ of the process $(s_t)_{t \in h\mathbb{N}}$ are said to have a heavy-tailed distribution, if for all $t \in h\mathbb{N}$ they satisfy

$$\mathbb{P}[r_t < -x] = x^{-\alpha} L(x) \quad \text{as } x \rightarrow \infty, \quad (6.1)$$

where $\alpha \in \mathbb{R}^+$ and L is a slowly varying function, i.e. $\lim_{x \rightarrow \infty} \frac{L(tx)}{L(x)} = 1$ for all $t > 0$.

For heavy-tailed distributions (6.1), the m th moment is infinite for all $m > \alpha$.

6.1. Scaling rule

A natural estimator for the tail index α is the so-called Hill estimator

$$\hat{\alpha}_{l,n}^{-1} = \frac{1}{l} \sum_{i=1}^l \log \left(\frac{r_{(i)}}{r_{(l)}} \right), \quad (6.2)$$

where n is the sample size and $r_{(l)}$ is the l th order statistic (i.e. $r_{(1)} \leq r_{(2)} \leq \dots \leq r_{(n)}$), taken as a threshold. Other estimators, as discussed in Embrechts et al. [19], can also be used.

From (6.1), estimates for the p -quantile $\hat{x}_{l,n}^p$ can be derived. On the one hand, putting $x = -r_{(l)}$ leads to $\frac{l}{n} = \mathbb{P}[r < r_{(l)}] = (-r_{(l)})^{-\alpha} L(-r_{(l)})$. On the other hand, setting $x = -x_{l,n}^p$ yields

$$p = \mathbb{P}[r < x_{l,n}^p] = (-x_{l,n}^p)^{-\alpha} L(-x_{l,n}^p). \quad (6.3)$$

Combining these equations gives the estimate

$$\hat{x}_{l,n}^p = r_{(l)} \left(\frac{l}{np} \right)^{1/\hat{\alpha}_{l,n}}. \quad (6.4)$$

Equation (6.4) yields an easy way to calculate the p -quantile once the tail (shape) parameter α is estimated. As the Hill estimator (6.2) is based upon the lowest order statistics, only a proportionally small subsample of the whole data set is considered. Consequently, a large data set to start with is desirable. Therefore, rather high frequency data should be used to estimate α , and to calculate $\hat{x}_{l,n}^p$. Then, a theoretical scaling rule can be applied to get the quantile estimate at the required frequency.

For one-period log-returns which are i.i.d., and which satisfy (6.1), by the subexponentiality property (see e.g. Feller [23, VIII.8, p.271] and more in detail Embrechts et al. [19]), the k -period log-returns $(r_t^k)_{t \in k\mathbb{N}}$ fulfil

$$\mathbb{P}[r_t^k < -x] = k x^{-\alpha} L(x)(1 + o(1)), \quad \text{as } x \rightarrow \infty. \quad (6.5)$$

This result supplements the central limit theorem by providing information concerning the tails. It describes the self-additivity in the tails of heavy-tailed distributions. The implication of this result for portfolio analysis has been discussed in the specific case of non-normal stable distributions by Fama and Miller [22]. In that case $\alpha < 2$ holds, and the variance is therefore infinite. Here we focus on the finite variance case. Dacorogna et al. [11] show the following asymptotic result, which can be derived from (6.3) and (6.5):

$$x^{k,p} \sim k^{1/\alpha} x^p, \quad \text{as } p \rightarrow 0, \quad (6.6)$$

where $x^{k,p}$ is the p -quantile of the k -period log-returns. This means that for a constant probability p , increasing the time horizon by a factor k increases the p -quantile for the

heavy-tailed model by a factor $k^{1/\alpha}$. For financial log-returns, the estimated tail index α usually belongs to the interval (2, 5) (see Dacorogna et al. [11] and Straetmans [40] for empirical studies on foreign exchange rate and stock index data). For $\alpha > 2$, the factor $k^{1/\alpha}$ is smaller than for the normal model, where the value-at-risk is increased by the factor $k^{1/2}$. Hence the square-root-of-time rule—used very often in practice to scale quantiles—leads to an overestimation of value-at-risk. More details about this time aggregation issue for heavy-tailed distributions can be found in Kaufmann and Patie [27], Section 7.2.3.

In comparison with the normal model, the probability of an extreme one-period loss is higher for the heavy-tailed model. However, as we saw before, the multiplication factor used to obtain the multi-period value-at-risk is smaller for heavy-tailed log-returns ($k^{1/\alpha}$, $\alpha \in (2, 5)$) than for normal log-returns ($k^{1/2}$). Based on this, the value of an estimated extreme k -period quantile may still be larger for the normal model than for the EVT model, if k is chosen large enough (as shown by simulation in Dacorogna et al. [11]).

6.2. Estimation of value-at-risk and expected shortfall

Estimates for the k -period value-at-risk at level p can be derived by applying the quantile transformation (6.6) to one-period values fulfilling (6.1):

$$\widehat{\text{VaR}}^p = -\left(\exp\left(\left(\frac{k l_{n,p}}{n p}\right)^{1/\hat{\alpha}_{l,n,p,n}} r_{(l_{n,p})}\right) - 1\right), \quad (6.7)$$

and for the corresponding one-year expected shortfall

$$\widehat{\text{ES}}^p = -\left(\frac{1}{p} \int_0^p \exp\left(\left(\frac{k l_{n,p}}{n q}\right)^{1/\hat{\alpha}_{l,n,p,n}} r_{(l_{n,p})}\right) dq - 1\right), \quad (6.8)$$

where $l_{n,p} = \lfloor n(p + 0.045 + 0.005h) \rfloor$, $\hat{\alpha}_{l,n}^{-1} = \frac{1}{l} \sum_{i=1}^l \log(r_{(i)}/r_{(l)})$ is the Hill estimator, and $r_{(l)}$ denotes the l th order statistic of one-period log-returns. The integral in the estimation of $\widehat{\text{ES}}^p$ can be evaluated numerically. The choice made for $l_{n,p}$ works well in practice; it turns out that the results are rather insensitive to the exact choice of $l_{n,p}$.

7. Model comparison

In this section, the suitability of the models for estimating one-year financial risks is compared. As explained before, we first fix a horizon $h < 1$ year, on which the various models are calibrated. For the gap between h days and one year (kh days), we use the corresponding scaling rule. This finally gives the values for one-year value-at-risk and expected shortfall.

Modelling h -day log-returns causes a first uncertainty. Scaling h -day log-returns to one-year log-returns produces a second uncertainty. The optimal horizon h for a chosen model is the one leading to the minimal total uncertainty, in our case measured by the quality of

the prediction for expected shortfall. In Section 7.1.1 the measures used for backtesting expected shortfall estimates are described. In order to find eventually the best model among the ones investigated, the backtesting measures are evaluated and compared for all the models for several intermediate horizons h .

7.1. Backtesting

7.1.1. Description of the backtesting measures. For backtesting the forecasted expected shortfall $\widehat{\text{ES}}_t^p$, we introduce two measures. The first measure V_1^{ES} evaluates excesses below the negative of the estimated value-at-risk $\widehat{\text{VaR}}_t^p$. 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 $\widehat{\text{ES}}_t^p$ for the one-year ahead expected shortfall ES_t^p . Now the difference between the observed one-year (k -period) return R_t^k and the negative of the estimation $\widehat{\text{ES}}_t^p$ is taken, and then the conditional average of these differences is calculated, conditioned on $\{R_t^k < -\widehat{\text{VaR}}_t^p\}$,

$$V_1^{\text{ES}} = \frac{\sum_{t=t_0}^{t_1} (R_t^k - (-\widehat{\text{ES}}_t^p)) \mathbf{1}_{\{R_t^k < -\widehat{\text{VaR}}_t^p\}}}{\sum_{t=t_0}^{t_1} \mathbf{1}_{\{R_t^k < -\widehat{\text{VaR}}_t^p\}}}.$$

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

This first measure sticks very closely to the theoretical definition of expected shortfall. Its weakness is that it depends strongly on the VaR estimates (without adequately reflecting the goodness/badness of these values), since only values which fall below the VaR threshold are considered. This is possibly a fraction which is far away from 1% 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 $1/p$ -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 $1/p$ -event”:

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

where $D_t := R_t^k - (-\widehat{\text{ES}}_t^p)$ and D^p denotes the empirical p -quantile of $\{D_t\}_{t_0 \leq t \leq t_1}$. Note that, since $\widehat{\text{ES}}_t^p$ is an estimate on a level p , we expect D_t to be negative in less than one out of $1/p$ cases. A good estimation for expected shortfall will again lead 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} :

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

This measure—which tells how well the forecasted one-year expected shortfall fits real data—will be used in this paper 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

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

This measure is used by the Basel Committee on Banking Supervision, which in order to encourage institutions to report their true value-at-risk numbers, devised a system in which penalties are set depending on the frequency of violations. A good estimation for value-at-risk will lead to a value of V^{freq} which is close to the level p .

7.1.2. Backtesting results. We start this section by describing the set-up for backtesting expected shortfall. For each data set, the models (random walk, $\text{AR}(p)$ process, $\text{GARCH}(1,1)$ process, heavy-tailed distributions) are calibrated at different frequencies (one day, one week, one month, three months, one year). For the $\text{GARCH}(1,1)$ model, horizons from one month upwards are missing since stationarity conditions are violated for the estimated parameters. For heavy-tailed distributions there are not enough quarterly and yearly data to estimate the tail index with the Hill and indeed any other EVT based estimator.

In a second step, one-year value-at-risk and expected shortfall are estimated as described in Sections 3.2–6.2. Finally, the backtesting measures presented in Section 7.1.1 are evaluated. Comparing the results across all models and all calibration horizons then tells which model yields the most reliable estimation for one-year asset risk.

In a first study, we use foreign exchange rates, stock indices, and 10-year bonds from five different markets. In a second part, we concentrate on single stocks. For all data sets under investigation, the well-known stylized facts of financial log-returns can be observed: leptokurtosis (which decreases with increasing time horizon), skewness (which is persistent for all time horizons), and significant positive autocorrelations; see Kaufmann and Patie [27] for details.

(a) Foreign exchange rates, stock indices and 10-year bonds

We first use the following data sets, obtained from Datastream: foreign exchange rates (DEM/CHF, GBP/CHF, USD/CHF, JPY/CHF), stock indices (SMI, DAX, FTSE, S&P, NIKKEI), 10-year government bonds (CH, DE, UK, US, JP). For each of these data sets, when trying to backtest the models, one encounters the problem that the number of yearly data is rather small to estimate model parameters and proceeding the backtesting in a significant way. This problem can be handled by aggregating the backtesting results for several data sets. For each of the four foreign exchange rates, 16 years of data are available (1985–2000, $n = 4173$ daily log-returns). For both stock indices and bonds we have five samples each containing 11 years of data (1990–2000, $n = 2869$ daily log-returns). The backtesting is carried out on each sample independently, and

then the results are aggregated within each of the three types of data. In detail, for each level p , each model, each intermediate horizon h and each data sample, we proceed as follows:

- Estimate the yearly forecasted expected shortfall \widehat{ES}_t^p on a window containing half of the data. $l = \lfloor n/2h \rfloor$ non-overlapping h -day log-returns are used for this estimation.
- Compare these estimates with the observed one-year returns R_t^k .
- Move the window by one, then repeat steps 1 and 2 till the end of the data set.
- Finally, for each of the three pooled samples (foreign exchange rates, stock indices and bonds), the risk measures can be evaluated as described in 7.1.1.

The results for the three types of data for $p = 1\%$ and $p = 5\%$ are presented in Tables 1, 2 and 3. Since risk managers are mainly interested in rare event cases, our main focus will be on the outcomes for the 1% expected shortfall. Results for the 5% level will be used to confirm the reliability and the flexibility of each approach.

For foreign exchange rates (Table 1) all four models seem to perform rather well for appropriate choices of the calibration horizon h . The best results are obtained when the models are calibrated with monthly ($h = 22$) or quarterly ($h = 65$) data. This behaviour was expected, since most scaling rules (for random walks, AR(p) processes, heavy-tailed distributions) require some independence structure between returns and, in some cases, the normality assumption for innovations. Low frequency ($h \geq 22$) returns are closer to fulfil this hypothesis than daily or weekly returns. Concerning the comparison across models, it can be observed that the random walk and the heavy-tailed distributions perform better than the other models. However, one should point out that the heavy-tailed approach is not reliable at frequencies other than the monthly one. While the random walk model and the heavy-tailed distributions slightly overestimate the risk measures, the AR(p) model underestimates them, which can be explained by the frequent underestimation of variance for this approach. We will come back to this issue in Section 7.3.

For stock indices (Table 2) the same lines of analysis can be followed. Here, the random walk model with $h \leq 65$ days and the GARCH(1,1) model calibrated on daily data outperform the other models. Finally, for 10-year government bonds (Table 3), the random walk models calibrated on horizons $h \leq 65$ days provide the most reliable forecasts for one-year 1% expected shortfall, and they are at the same time reliable on the 5% level.

Summarising the results for the three types of data, we can say that random walk processes with a constant trend and normal innovations, calibrated on horizons $h \leq 65$ days are a good choice. They clearly outperform the other models under investigation. However, these backtesting results do not provide direct information about the variability of the risk measures. For the random walk models with normal innovations, we will investigate this issue in Section 7.4 by calculating confidence intervals for expected shortfall and value-at-risk.

(b) German stock data

After backtesting foreign exchange rates, stock indices and government bonds, we now do the same for single stocks. For this study we use a data set of 22 single stocks which are part of the German stock index DAX.

Like in (a), there do not exist enough stationary returns per stock in order to perform the backtesting in a straightforward way. In the case of German stock data, we evaluate the backtesting measures on the pooled sample of all 22 stocks. Each sample contains 23.5 years of data ($n = 6146$ daily values). We proceed as described in (a).

The results for $p = 1\%$ and $p = 5\%$ are presented in Table 4. It can be seen that for this data set the random walk model achieves its best performance when working directly with yearly data. For a GARCH(1,1) model a one week horizon seems to be a good choice. The results for AR(p) models do not seem to depend much on the calibration horizon. For the heavy-tailed distributions the optimal calibration horizon for $p = 1\%$ is one month. However, this choice gives very unsatisfactory results for $p = 5\%$. Moreover—as will be discussed in Section 7.2—for appropriate models the backtesting results on a 5% level should indicate a slight overestimation of the risk. The results in Table 4 show the opposite behaviour for heavy-tailed distributions applied to monthly data.

Comparing the four models with each other, and considering the shortcomings of heavy-tailed distributions, GARCH(1,1) models applied to weekly data provide the best forecasts, measured by their suitability to predict one-year 1% expected shortfall (and considering the corresponding 5% prediction). Also a random walk calibrated on yearly data seems to give reasonable forecasts for these German stock data.

7.2. Further investigations for a simulated random walk

The results of the backtesting showed that random walk models with a constant trend and normal innovations are a good choice for estimating 1% expected shortfall for a one-year period. In order to get a better understanding of these backtesting results, we repeat the same analysis for a time series where the distribution of the log-returns is known. We analyse the properties of a simulated random walk with normal innovations. We proceed as follows. We simulate one-period log-returns

$$\tilde{r}_t^{(j)} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\tilde{\mu}^{(j)}, \tilde{\sigma}^{(j)}), \quad j = 1, \dots, 5, \quad t = h, 2h, \dots, nh, \quad (7.1)$$

where $\tilde{\mu}^{(j)} = \frac{1}{n} \sum_{i=1}^n r_{ih}^{(j)}$ and $\tilde{\sigma}^{(j)} = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (r_{ih}^{(j)} - \tilde{\mu}^{(j)})^2}$ are mean and standard deviation estimated for the one-period (h -day) log-returns $r_t^{(j)}$ of the stock indices SMI ($j = 1$), DAX ($j = 2$), FTSE ($j = 3$), S&P ($j = 4$) and NIKKEI ($j = 5$).

We proceed as described in Section 3 to get estimates for the risk measures. The backtesting is done as explained in Section 7.1. Since the results vary from simulation to simulation, we repeat the whole procedure several times. The results are displayed in Table 5, together with a recapitulation of the corresponding results for stock index data (see Table 2). Note

that now for these simulated processes (7.1) it is known that the correct model (random walk with normal innovations) is fitted.

We observe that calibration horizons $h \leq 65$ days outperform the results for $h = 261$ days by far. A second observation, which we are going to discuss in more detail, is the fact that expected shortfall and value-at-risk tend to be underestimated. In particular, this clearly holds for a level of 5%. This is caused by the statistical properties of the backtesting measures. To make this clear, we examine the frequency of exceedances V^{freq} . When evaluating this measure, the portion of returns below $-\widehat{\text{VaR}}^p$ is calculated. Let Q be an estimator of a quantile of log-returns. Since the cumulative distribution function $\Phi(q)$ of a normal random variable Q is convex for $q < 0$, unbiasedness of Q (i.e. $\mathbb{E}[Q] = q$) leads to $\mathbb{E}[\Phi(Q)] \geq \Phi(q)$. This means that the probability of a log-return being below $-\widehat{\text{VaR}}^p$ tends to being overestimated. As an example, we consider the situation where the mean of a standard normal distribution is overestimated by 0.5 in the first case and underestimated by 0.5 in the second case. On average these two scenarios lead to an overestimation of the number of values which lie below the real value-at-risk of $\frac{1}{2}(\Phi(\Phi^{-1}(5\%) + 0.5) + \Phi(\Phi^{-1}(5\%) - 0.5)) - 5\% \approx 2.11\%$ for the 5% VaR, and $\frac{1}{2}(\Phi(\Phi^{-1}(1\%) + 0.5) + \Phi(\Phi^{-1}(1\%) - 0.5)) - 1\% \approx 0.81\%$ for the 1% VaR. The uncertainty in the estimated standard deviation has a similar effect. This even enlarges the above differences. Since $\widehat{\text{ES}}^p$ can be written as $\frac{1}{p} \int_0^p \widehat{\text{VaR}}^q dq$, this fact also holds for expected shortfall. This explains the indications for an underestimation of expected shortfall and value-at-risk in Table 5.

The backtesting results in Section 7.1.2 have shown that random walks with normal innovations and a constant drift give quite good results for the data sets under investigation. To be critical, we can ask ourselves whether replacing normal innovations by heavier tailed ones, e.g. Student- t_ν with $\nu = 4$, would lead to even better results. The backtesting results answer in the negative to this question. With normal innovations, for most of the 20 cases (four data sets, $h = 1, 5, 22, 65, 261$), $\text{VaR}^{5\%}$ is clearly underestimated ($V^{\text{freq},5\%} > 5\%$), while $\text{VaR}^{1\%}$ is overestimated for three of the four types of data, see Tables 1–4. This over- and underestimation is approximately of the same order as for the simulated time series in Table 5. Replacing normal log-returns by heavier tailed ones would even worsen the slight overestimation for the 1% level. This is due to the fact that replacing a normal distribution by a Student- t distribution with the same mean and the same 5% quantile increases the size of the estimated 1% quantile, which is the opposite effect to the one we would like to have. This suggests that working with normal innovations in the random walk model (as done in our study) should be preferred to assuming Student- t distributed innovations.

7.3. Variance analysis

7.3.1. Description. To learn more about the statistical properties of the models in use, we analyse the variance of forecasted yearly log-returns. This analysis consists of the following steps:

- For each data set, each model and each intermediate time horizon the model parameters for h -day (one-period) log-returns are estimated.

- Based on these estimations, the one-year (k -period) variance (and the standard deviation as its square root) is calculated.
- To get some information about the precision of these standard deviations, 1000 periods of 16 years (for foreign exchange rates) and 1000 periods of 11 years (for stock indices and 10-year government bonds) are simulated. These simulations are based on the parameters estimated in the first step. For each of the 1000 simulations the yearly standard deviation is estimated. Finally, based on these 1000 values, the 95% confidence intervals for standard deviation are calculated.

For the exact proceeding within these three steps for random walk, $AR(p)$ and GARCH(1,1) processes, we refer to Kaufmann and Patie [27]. Note that only the confidence intervals are determined via simulation, whereas the point estimate for standard deviation is the theoretical one calculated from the parameters estimated in Step 1.

7.3.2. Results. The yearly standard deviations $\tilde{\sigma}^k$, and the 95% confidence intervals are estimated for random walk, $AR(p)$ and GARCH(1,1) models, calibrated on foreign exchange rates, stock indices and 10-year government bonds for different horizons. In figures 1–3, these values (point estimates and confidence intervals) for calibration horizons varying from one day to one year are plotted.

As before, for GARCH(1,1) models only horizons up to one week are used, since for longer horizons stationarity conditions are violated for some of the estimated parameters. We observe that also for calibration horizons of one day and one week, estimated variances for GARCH(1,1) models vary significantly. Some of the 95% confidence intervals are huge.

For $AR(p)$ models, variances tend to be underestimated. The confidence intervals are very asymmetric around the point estimates the simulations are based on. For several data sets the confidence interval for the standard deviation even does not contain the value $\tilde{\sigma}^k$, which is the true value of the standard deviation belonging to the parameters underlying the simulations. This means that the proceeding in $AR(p)$ models provides unsatisfactory results for estimating the standard deviation of such processes.

Only for random walk models we observe a satisfying behaviour. For most data sets these variance estimates are fairly stable. As one would expect, the confidence intervals are increasing with increasing calibration horizon (i.e. using fewer data).

This variance investigation confirms the conclusion of Section 7.1.2, where a random walk model with a calibration horizon of at most three months was proposed. Considering the stylised fact that daily (log) returns in a portfolio show dependencies which are not present in longer horizon returns, using an intermediate horizon between one week and one month seems to be most appropriate for modelling one-year returns.

7.4. Confidence intervals for risk estimates in random walk models

In random walk models with normal innovations, confidence intervals for one-year expected shortfall can be calculated based on h -day observations. The explicit formula for these

confidence intervals (and also for the ones for value-at-risk) can be found in Kaufmann and Patie [27].

7.4.1. Confidence intervals for the Swiss Market Index. Assuming a random walk model, we calculate 95% confidence intervals for $\text{VaR}^{1\%}$ and $\text{ES}^{1\%}$ of the Swiss market index SMI. The estimations are performed using 11 years of data (January 1, 1990 to December 31, 2000). This gives one-year forecasts for December 31, 2001. The percentage losses, using daily, weekly, monthly, quarterly and yearly data, are shown graphically in figure 4. The full and the dotted lines represent 95% confidence intervals for 1% expected shortfall and 1% value-at-risk, respectively. The dots visualize the corresponding point estimates. To elucidate what this percentage loss means for the absolute value of SMI, we transform these one-year point estimates and confidence intervals into predictions for December 31, 2001, see figure 5 (value of the SMI at the end of 2000: 8135.2; at the end of 2001: 6417.8). For example for $h = 22$ days (1 month), the point estimate for the 1% value-at-risk gives a value of 0.288. This means that with a probability of 1%, the SMI will lose 28.8% of its value or more during the year 2001. Transformed to the value at the end of the year 2001, this gives a point estimate of 5795. Hence with probability 1%, the SMI will be below 5795 at the end of this one-year period. The 95% confidence interval around the percentual loss goes from 0.191 to 0.388, which corresponds to a 95% probability that the true 1% value-at-risk lies between these bounds (which corresponds to the interval [5017, 6584] for the SMI at the end of 2001). Values for expected shortfall can be read off in a similar way.

The most evident observation in figures 4 and 5 is the fact that the confidence intervals get larger as the intermediate time horizon h increases (i.e. as the number of values used for the prediction decreases). Furthermore, the point estimate for $h = 1$ year seems to differ from the ones for shorter intermediate horizons. This suggests that shorter horizons should be preferred. To figure out whether this clear result is just a coincidence, or whether there is a clear trend as h increases, we evaluate the point estimates for all intermediate horizons h between one day and one year, and plot them in figure 6. This plot does not confirm such a trend. But it does confirm the increase in variation as h increases. This suggests once more that a relatively small h ($h \leq 3$ months) should be chosen for estimating one-year asset risks.

7.4.2. Confidence intervals for a simulated random walk. Like for the SMI data, 95% confidence intervals of $\text{VaR}^{1\%}$ and $\text{ES}^{1\%}$ can also be calculated for a simulated random walk. First, from SMI log-returns (11 years of data), mean and variance are estimated. Based on these values, a random walk (with normal innovations) of length 11 years is simulated. Finally, one-year value-at-risk and expected shortfall and the corresponding confidence intervals are estimated based on daily, weekly, monthly, quarterly and yearly data. These risk estimates are shown graphically in figure 7. The full and the dotted lines represent 95% confidence intervals for 1% expected shortfall and 1% value-at-risk, respectively. The dots visualize the corresponding point estimates. The slight overestimation compared to the true values ($\text{ES}^{1\%} = 27.35\%$, full horizontal line; $\text{VaR}^{1\%} = 23.09\%$, dotted horizontal line) stems from

the fact that the average of the simulated daily log-returns is slightly smaller than the mean estimated from SMI data. At first glance, the estimates for simulated data seem to vary less than for the SMI (figure 4). To analyse this further, we evaluate the point estimates for all calibration horizons h from one day to one year. Figure 8 shows that for simulated random walk data these estimates do in fact not behave significantly more stable than for SMI data. (For comparison: 10% in figure 8 corresponds to 813.52 units in figure 6.) Like before, the variation of the point estimates around the true value increases as the intermediate horizon h is increased. This confirms again that for a reliable prediction of one-year risks, h should not be chosen too large.

8. Conclusion

This paper is concerned with the estimation of measures of market risk over a long term horizons of one year, with emphasis on the expected shortfall. We present and test dynamical models like random walk, GARCH(1,1) and AR(p) processes. We also propose a static approach based on heavy-tailed distributions. An important motivation for choosing these models is the existence of time aggregation rules. Since the main difficulty on measuring long term risks is the lack of such low frequency (e.g. one year) data, we calibrate the models on higher frequency data and then apply scaling rules to get risk estimates at a one-year horizon. We compare the models by applying backtesting methods. The outcome of these investigations can be summarised as follows:

- As opposed to short term horizons, for a one-year period a good estimate of the trend of (log-)returns is critical when measuring risks.
- In general, the best frequencies for calibration are the intermediate ones, like one month. The statistical restrictions, like the sample size for estimating the models parameters and the confidence interval of the risk estimates, play an important role in this choice. A reason why using relatively high frequency data (daily data for example) is not perfect is the fact that scaling rules are based on certain assumptions which are not fulfilled by such data.
- The random walk model performs better on average than other models. It provides satisfactory results across all classes of data and for both confidence levels investigated (1%, 5%). 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.

Based on these results, we recommend 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 for estimating the one-year 1% expected shortfall.

Appendix A: Tables

Table 1. Backtesting results for the levels $p = 1\%$ and $p = 5\%$, using foreign exchange rate data. Measures for one-year value-at-risk and expected shortfall are evaluated (The values in Tables 1–5 are percentages).

Model	Freq days	ES ¹			VaR ¹ V ^{freq}	ES ⁵			VaR ⁵ V ^{freq}
		V ^{ES}	V ₁ ^{ES}	V ₂ ^{ES}		V ^{ES}	V ₁ ^{ES}	V ₂ ^{ES}	
Optimal		0	0	0	1	0	0	0	5
Random walk	1	1.4	1.1	1.7	0.4	1.0	0.8	-1.3	8.6
	5	1.1	0.8	1.4	0.5	1.1	1.0	-1.3	8.7
	22	1.0	0.7	1.3	0.5	1.2	1.2	-1.2	8.1
	65	0.9	0.5	1.4	0.5	1.1	1.1	-1.1	7.1
	261	0.8	-0.4	-1.1	1.6	1.7	0.1	-3.3	9.4
GARCH(1,1)	1	N/A	N/A	3.9	0.0	1.9	3.7	0.0	7.3
	5	2.2	1.9	2.6	0.3	1.5	2.0	-1.1	10.1
AR(p)	1	2.3	-0.7	-3.9	6.6	3.9	-1.1	-6.8	19.5
	5	2.4	-0.9	-4.0	5.9	3.9	-0.9	-6.9	19.1
	22	2.3	-0.8	-3.8	5.9	3.8	-0.8	-6.8	18.8
	65	2.1	-0.7	-3.4	5.8	3.6	-0.8	-6.4	18.3
	261	5.2	-1.9	-8.6	10.7	6.9	-2.3	-11.4	23.4
Heavy-tailed distribution	1	5.1	-1.2	-8.9	15.2	9.4	-3.1	-15.8	41.4
	5	3.7	0.2	-7.1	9.3	7.8	-1.8	-13.8	35.8
	22	0.8	1.6	0.1	1.4	4.8	0.1	-9.4	19.8

Table 2. Backtesting results for the levels $p = 1\%$ and $p = 5\%$, using stock index data. Measures for one-year value-at-risk and expected shortfall are evaluated.

Model	Freq days	ES ¹			VaR ¹ V ^{freq}	ES ⁵			VaR ⁵ V ^{freq}
		V ^{ES}	V ₁ ^{ES}	V ₂ ^{ES}		V ^{ES}	V ₁ ^{ES}	V ₂ ^{ES}	
Optimal		0	0	0	1	0	0	0	5
Random walk	1	0.8	0.3	1.3	0.8	3.5	0.0	-7.0	8.3
	5	1.2	0.5	1.9	0.7	3.2	0.0	-6.4	8.1
	22	0.7	0.2	1.1	0.8	3.7	-0.2	-7.1	8.3
	65	1.3	-1.2	-1.3	1.0	4.7	-0.8	-8.5	8.6
	261	10.5	-6.0	-15.0	2.5	11.0	-4.0	-18.0	9.2
GARCH(1,1)	1	0.6	0.2	-1.1	1.3	5.4	-0.3	-10.5	8.8
	5	3.7	2.6	4.9	0.5	3.1	1.6	-4.6	7.6
AR(p)	1	6.3	-3.2	-9.3	3.3	8.5	-3.5	-13.5	10.6
	5	6.4	-3.2	-9.6	3.4	8.5	-3.4	-13.7	11.0
	22	7.1	-3.8	-10.4	3.2	8.8	-3.2	-14.4	11.2
	65	8.8	-4.1	-13.4	3.8	9.9	-3.3	-16.6	12.3
	261	13.7	-6.4	-21.0	12.4	14.7	-6.8	-22.6	20.9
Heavy-tailed distribution	1	3.0	4.1	1.9	2.0	2.5	1.5	3.5	3.6
	5	2.4	1.8	2.9	0.8	4.5	2.4	6.7	2.3
	22	1.7	-0.5	2.9	0.5	8.4	3.0	13.8	0.9

Table 3. Backtesting results for the levels $p = 1\%$ and $p = 5\%$, using 10-year government bond data. Measures for one-year value-at-risk and expected shortfall are evaluated.

Model	Freq days	ES ¹			VaR ¹ V^{freq}	ES ⁵			VaR ⁵ V^{freq}
		V^{ES}	V_1^{ES}	V_2^{ES}		V^{ES}	V_1^{ES}	V_2^{ES}	
Optimal		0	0	0	1	0	0	0	5
Random walk	1	1.0	0.3	1.8	0.6	1.9	0.5	-3.3	7.3
	5	1.8	0.5	3.2	0.4	1.6	0.8	-2.4	7.1
	22	2.4	-0.4	4.4	0.2	1.2	1.3	-1.1	6.7
	65	3.6	1.1	6.1	0.1	1.0	1.6	0.5	5.6
	261	4.1	-4.7	-3.4	0.9	5.2	-2.8	-7.7	6.0
GARCH(1,1)	1	6.1	-1.8	10.4	0.0	4.8	4.4	5.2	4.4
	5	10.7	11.8	9.5	0.1	2.3	2.6	2.0	5.1
AR(p)	1	5.9	-2.4	-9.3	3.4	7.8	-2.1	-13.5	12.4
	5	5.8	-2.6	-9.0	3.0	7.6	-1.9	-13.2	11.9
	22	5.7	-2.8	-8.5	2.8	7.4	-1.9	-12.8	11.9
	65	6.9	-3.6	-10.2	3.4	8.2	-2.3	-14.2	13.6
	261	13.1	-4.9	-21.2	13.5	15.0	-5.9	-24.1	23.8
Heavy-tailed distribution	1	11.4	-2.2	-20.5	35.0	9.4	-3.1	-15.8	41.4
	5	8.4	-1.2	-15.6	25.2	7.8	-1.8	-13.8	35.8
	22	7.6	-1.4	-13.7	11.7	4.8	0.1	-9.4	19.8

Table 4. Backtesting results for the levels $p = 1\%$ and $p = 5\%$ for stocks of the DAX. Measures for one-year value-at-risk and expected shortfall are evaluated.

Model	Freq days	ES ¹			VaR ¹ V^{freq}	ES ⁵			VaR ⁵ V^{freq}
		V^{ES}	V_1^{ES}	V_2^{ES}		V^{ES}	V_1^{ES}	V_2^{ES}	
Optimal		0	0	0	1	0	0	0	5
Random walk	1	10.3	-5.2	-15.4	3.3	6.8	-4.2	-9.4	8.3
	5	9.6	-4.7	-14.6	3.3	6.5	-4.0	-9.1	8.3
	22	8.3	-4.1	-12.5	2.9	5.8	-3.7	-7.9	7.7
	65	9.0	-4.5	-13.6	3.0	6.1	-4.0	-8.2	7.5
	261	4.8	-2.6	-7.0	2.1	3.3	-2.5	-4.1	6.0
GARCH(1,1)	1	6.1	-2.6	-9.5	2.3	4.2	-2.2	-6.1	7.5
	5	2.6	-1.4	-3.8	1.4	1.3	-0.2	-2.4	6.5
AR(p)	1	8.8	-3.8	-13.9	4.4	6.8	-3.3	-10.2	11.2
	5	8.6	-3.8	-13.3	4.0	6.3	-3.0	-9.6	10.8
	22	7.7	-3.6	-11.8	3.4	5.3	-2.5	-8.2	9.9
	65	7.3	-3.3	-11.2	3.5	5.2	-2.4	-8.1	10.0
	261	9.8	-4.3	-15.3	3.2	5.8	-2.8	-8.9	8.9
Heavy-tailed distribution	1	6.1	1.3	-10.8	5.2	2.8	2.0	-3.5	10.1
	5	6.2	-0.3	-12.1	4.7	2.6	0.5	-4.7	8.9
	22	1.7	2.6	0.8	1.1	7.2	5.7	8.6	3.0

Table 5. Backtesting results for a random walk model for the levels $p = 1\%$ and $p = 5\%$, using real stock indices and simulated random walk data, respectively. Measures for one-year value-at-risk and expected shortfall are evaluated.

Data	Freq days	ES ¹			VaR ¹ V^{freq}	ES ⁵			VaR ⁵ V^{freq}
		V^{ES}	V_1^{ES}	V_2^{ES}		V^{ES}	V_1^{ES}	V_2^{ES}	
Optimal		0	0	0	1	0	0	0	5
Stock indices	1	0.8	0.3	1.3	0.8	3.5	0.0	-7.0	8.3
	5	1.2	0.5	1.9	0.7	3.2	0.0	-6.4	8.1
	22	0.7	0.2	1.1	0.8	3.7	-0.2	-7.1	8.3
	65	1.3	-1.2	-1.3	1.0	4.7	-0.8	-8.5	8.6
	261	10.5	-6.0	-15.0	2.5	11.0	-4.0	-18.0	9.2
Simulation 1	1	1.3	0.9	-1.7	2.9	5.6	-1.1	-10.0	12.7
	5	1.4	0.4	-2.4	3.0	6.0	-1.3	-10.7	13.1
	22	2.1	-0.3	-3.9	3.2	6.9	-1.8	-11.9	14.0
	65	2.4	-0.2	-4.5	3.9	7.1	-2.0	-12.3	15.0
	261	9.6	-4.4	-14.8	6.3	12.7	-3.8	-21.6	18.2
Simulation 2	1	0.9	1.2	0.6	1.2	4.4	0.4	-8.5	9.4
	5	1.8	1.9	1.7	1.1	4.1	0.6	-7.6	9.1
	22	0.4	0.7	0.1	1.2	4.6	0.1	-9.0	9.2
	65	2.5	-1.1	-3.8	1.9	6.6	-0.9	-12.3	11.1
	261	12.0	-5.2	-18.8	7.1	14.9	-5.3	-24.6	16.5
Simulation 3	1	0.7	1.2	0.2	1.5	3.6	-0.1	-7.0	10.5
	5	0.8	0.8	-0.8	1.8	4.1	-0.4	-7.8	10.8
	22	0.6	1.0	0.3	1.3	3.5	0.2	-6.9	10.8
	65	2.7	2.3	3.0	0.7	2.5	0.6	-4.4	10.2
	261	7.1	-4.0	-10.2	4.3	8.4	-2.7	-14.0	13.8

Appendix B: Graphs

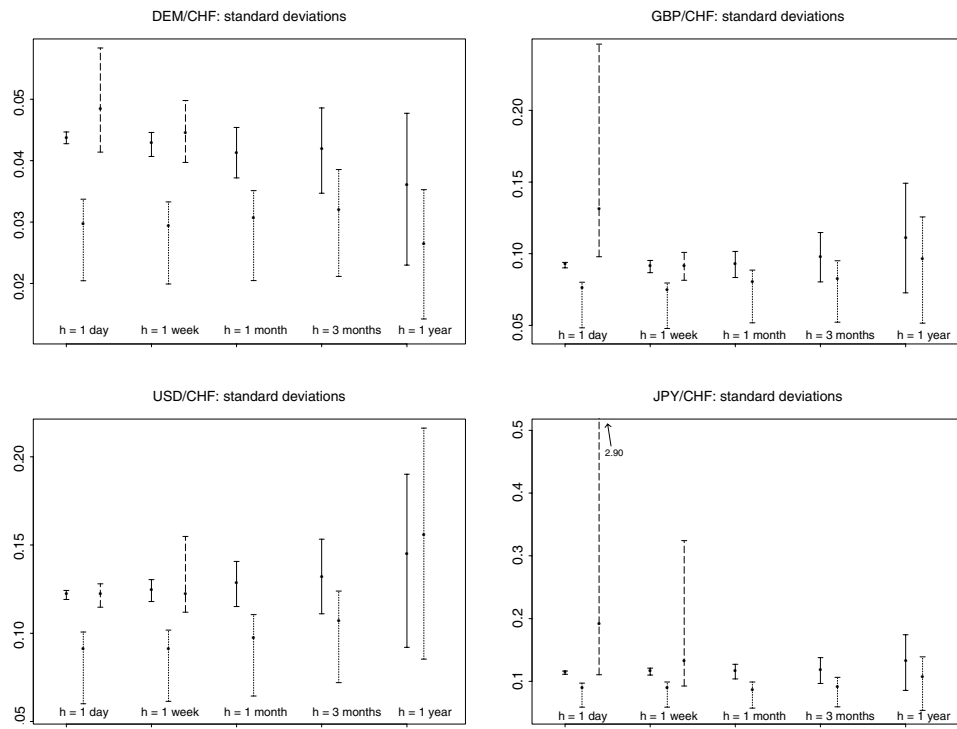


Figure 1. Foreign exchange rates: estimates for standard deviations (including confidence intervals constructed via simulation), using random walk models (full lines), AR(p) models (dotted lines) and GARCH(1,1) models (dashed lines).

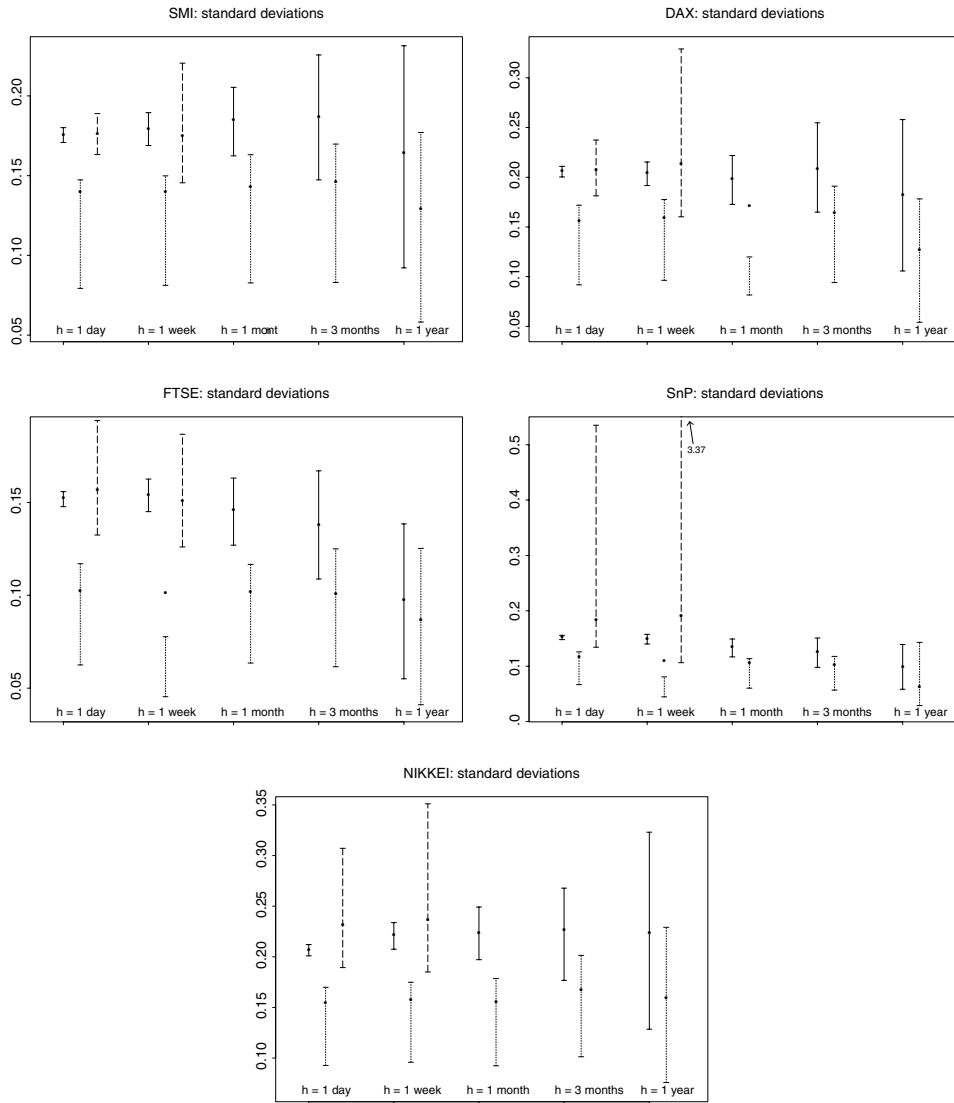


Figure 2. Stock indices: estimates for standard deviations (including confidence intervals constructed via simulation), using random walk models (full lines), $AR(p)$ models (dotted lines) and GARCH(1,1) models (dashed lines).




Figure 3. 10-year government bonds: estimates for standard deviations (including confidence intervals constructed via simulation), using random walk models (full lines), AR(p) models (dotted lines) and GARCH(1,1) models (dashed lines).




Figure 4. Point estimates and 95% confidence intervals for one-year 1% expected shortfall and 1% value-at-risk for the SMI (percentage loss), based on different calibration horizons h . Underlying model: random walk with normal innovations.




Figure 5. Point estimates and 95% confidence intervals for one-year 1% expected shortfall and 1% value-at-risk for the SMI, transformed in such a way that the values at the end of the year 2001 can be read off. The estimates are based on different calibration horizons h . Underlying model: random walk with normal innovations.




Figure 6. Point estimates for one-year 1% expected shortfall for the SMI, transformed in such a way that the values at the end of the year 2001 can be read off. The estimates are based on calibration horizons h from one day to one year. Underlying model: random walk with normal innovations.




Figure 7. Point estimates and 95% confidence intervals for one-year 1% expected shortfall and 1% value-at-risk (percentage loss) for a simulated random walk, based on five different calibration horizons h . Underlying model: random walk with normal innovations. The true values (horizontal lines) are 27.35% for expected shortfall and 23.09% for value-at-risk.




Figure 8. Point estimates for one-year 1% expected shortfall (percentage loss) for a simulated random walk, based on calibration horizons from one day to one year. Underlying model: random walk with normal innovations. The true value (horizontal line) is 27.35%.

Acknowledgment

The authors take pleasure in thanking Niklaus Bühlmann, John Hancock, Uwe Schmock and Fabio Trojani for fruitful discussions, Peter Blum for providing the data, and Filip Lindskog for careful reading.

References

1. C. Acerbi and D. Tasche, "On the coherence of expected shortfall," *Journal of Banking and Finance*, vol. 26, no. 7, pp. 1491–1507, 2002.
2. A. Ang and J. Chen "Asymmetric correlations on equity portfolios," *Journal of Financial Economics*, vol. 63, no. 3, pp. 443–494, 2002.
3. P. Artzner, F. Delbaen, J.M. Eber, and D. Heath, "Coherent measures of risk," *Mathematical Finance*, vol. 9, no. 3, pp. 203–228, 1999.
4. O.E. Barndorff-Nielsen, "Exponentially decreasing distributions for the logarithm of particle size," *Proceedings of the Royal Society London A*, vol. 353, pp. 401–419, 1977.
5. O.E. Barndorff-Nielsen, "Processes of normal inverse gaussian type," *Finance and Stochastics*, vol. 2, pp. 41–68, 1998.
6. M. Borkovec, "Large fluctuations in finance models," PhD Thesis, Munich University of Technology, 1999.
7. R. Brummelhuis and R. Kaufmann, "Estimating 10-day value-at-risk," Working Paper, ETH Zurich, 2003.
8. J.Y. Campbell, A.W. Lo, and A.C. MacKinlay, *The Econometrics of Financial Markets*, Princeton University Press, 1997.
9. P.F. Christoffersen, "Evaluating interval forecasts," *International Economic Review*, vol. 39, pp. 841–862, 1998.
10. P.F. Christoffersen, F.X. Diebold, and T. Schuermann, "Horizon problems and extreme events in financial risk management," *Economic Policy Review*, Federal Reserve Bank of New York, October 1998, pp. 109–118.

11. M.M. Dacorogna, U.A. Müller, O.V. Pictet, and C.G. de Vries, "Extremal forex returns in extremely large data sets," *Extremes*, vol. 4, pp. 105–127, 1999.
12. F.X. Diebold and A. Inoue, "Long memory and regime switching," *Journal of Econometrics*, vol. 105, pp. 131–159, 2001.
13. F.C. Drost and T.E. Nijman, "Temporal aggregation of GARCH processes," *Econometrica*, vol. 61, pp. 909–927, 1993.
14. D. Duffie and J. Pan, "An overview of value at risk," *The Journal of Derivatives*, Spring, pp. 7–49, 1997.
15. E. Eberlein, "Application of generalized hyperbolic Lévy motions to finance," in *Lévy Processes: Theory and Applications*, O.E. Barndorff-Nielsen, T. Mikosch, and S. Resnick (Eds.), Birkhäuser, 1999, pp. 319–337.
16. E. Eberlein and U. Keller, "Hyperbolic distributions in finance," *Bernoulli*, vol. 1, pp. 281–299, 1995.
17. E. Eberlein and K. Prause, "The generalized hyperbolic model: Financial derivatives and risk measures," in *Mathematical Finance—Bachelier Congress 2000*, H. Geman, D. Madan, S. Pliska, and T. Vorst (Eds.), Springer, 2002, pp. 245–267.
18. P. Embrechts (Ed.), *Extremes and Integrated Risk Management*, Risk Waters Group, London, 2000.
19. P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling Extremal Events for Insurance and Finance*, Springer, New York, 1997.
20. P. Embrechts, A.J. McNeil, and D. Straumann, "Correlation and dependence in risk management: Properties and pitfalls," in *Risk Management: Value at Risk and Beyond*, M.A.H. Dempster (Ed.), Cambridge University Press: Cambridge, 2002, pp. 176–223.
21. E.T. Fama, "The behavior of stock market prices," *Journal of Business*, vol. 38, pp. 34–105, 1965.
22. E.T. Fama and M.H. Miller, *The Theory of Finance*, Dryden Press, 1972.
23. W. Feller, *An Introduction to Probability Theory and its Applications*, Volume II, 2nd edition, John Wiley: New York, 1971.
24. C. Gouriéroux, "ARCH-models and financial applications," *Springer Series in Statistics*: Springer, New York, 1997.
25. J. Hancock, N. Bühlmann, and B. Schmolck, "Results of IRIS project: Long time-horizon VaR," *Internal Swiss Re Report*, 2000.
26. J.C. Hull, *Options, Futures and Other Derivatives*, 4th edition, Prentice-Hall, 1999.
27. R. Kaufmann and P. Patie, "Strategic long-term financial risks—The one dimensional case," *RiskLab Report*, ETH Zurich, 2003. Available at <http://www.risklab.ch/Papers.html#SLTFR>.
28. J. Kim, A.M. Malz, and J. Mina, *LongRun Technical Document*, RiskMetrics Group: New York, 1999.
29. J. Kim and J. Mina, *ForeSight Technical Document*, RiskMetrics Group: New York, 2000.
30. J.A. Lopez, "Regulatory evaluation of value-at-risk models," *Federal Reserve Bank of New York, Research Paper FRBNY—No. 9710*, 1998.
31. A.J. McNeil and R. Frey, "Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach," *Journal of Empirical Finance*, vol. 7, pp. 271–300, 2000.
32. T. Mikosch and C. Stărică, "Changes of structure in financial time series and the GARCH model," *Revstat Statistical Journal*, vol. 2, no. 1, pp. 41–73, 2004.
33. Morgan Guaranty Trust Company, *RiskMetrics Technical Document*, 4th edition, 1996.
34. W. Ogryczak and A. Ruszczyński, "Dual stochastic dominance and related mean–risk models," *SIAM Journal on Optimization*, vol. 13, pp. 60–78, 2002.
35. G.C. Pflug, "Some remarks on the value-at-risk and the conditional value-at-risk," in *Probabilistic Constrained Optimization: Methodology and Applications*, Kluwer Academic Publishers: Dordrecht, Netherlands, 2000, pp. 272–281.
36. R.T. Rockafellar and S.P. Uryasev, "Optimization of conditional value-at-risk," *The Journal of Risk*, vol. 2, no. 3, pp. 21–42, 2000.
37. R.T. Rockafellar and S.P. Uryasev, "Conditional value-at-risk for general loss distributions," *Journal of Banking and Finance*, vol. 26, no. 7, pp. 1443–1471, 2001.
38. C. Stărică, "Is GARCH(1,1) as good a model as the accolades of the Nobel prize would imply?," Preprint, Chalmers University of Technology, Gothenburg, 2003.
39. E. Stern, "Specific risk for long-term horizons," *RiskMetrics Journal*, vol. 3, no. 2, pp. 3–31, 2002.
40. S. Straetmans, "Extreme financial returns and their comovements," *PhD Thesis*, Erasmus University Rotterdam, 1998.

41. D. Straumann, "Estimation in conditionally heteroscedastic time series models," PhD Thesis, University of Copenhagen, 2003.
42. D. Straumann and T. Mikosch, "Quasi-maximum-likelihood estimation in heteroscedastic time series: A stochastic recurrence equations approach," Preprint 184, Laboratory of Actuarial Mathematics, University of Copenhagen, 2003.
43. O. Strub, "On the normality of long-term financial log-returns," Master Thesis, ETH Zurich, 2000.
44. C.G. de Vries, "Second order tail effects," in *Statistics and Finance: An Interface*, W.S. Chan, W.K. Li, and H. Tong (Eds.), Imperial College Press, 2000, pp. 153–165.