Modelling of Long-Term Risk

Roger Kaufmann
Swiss Life
roger.kaufmann@swisslife.ch

15th International AFIR Colloquium

6-9 September 2005, Zurich

Contents

A. Basel II

B. Scaling of Risks

C. One-Year Risks

D. Conclusions

A. Basel II

Amendment to the Capital Accord to Incorporate Market Risks (Basel Committee on Banking Supervision, 1996):

- "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."

Basel II (cont.)

- Market risk: 10-day value-at-risk, 99%

  Standard: 1-day value-at-risk, 95%

- Insurance: 1-year value-at-risk, 99%

  1-year expected shortfall, 99%

VaR in Visual Terms

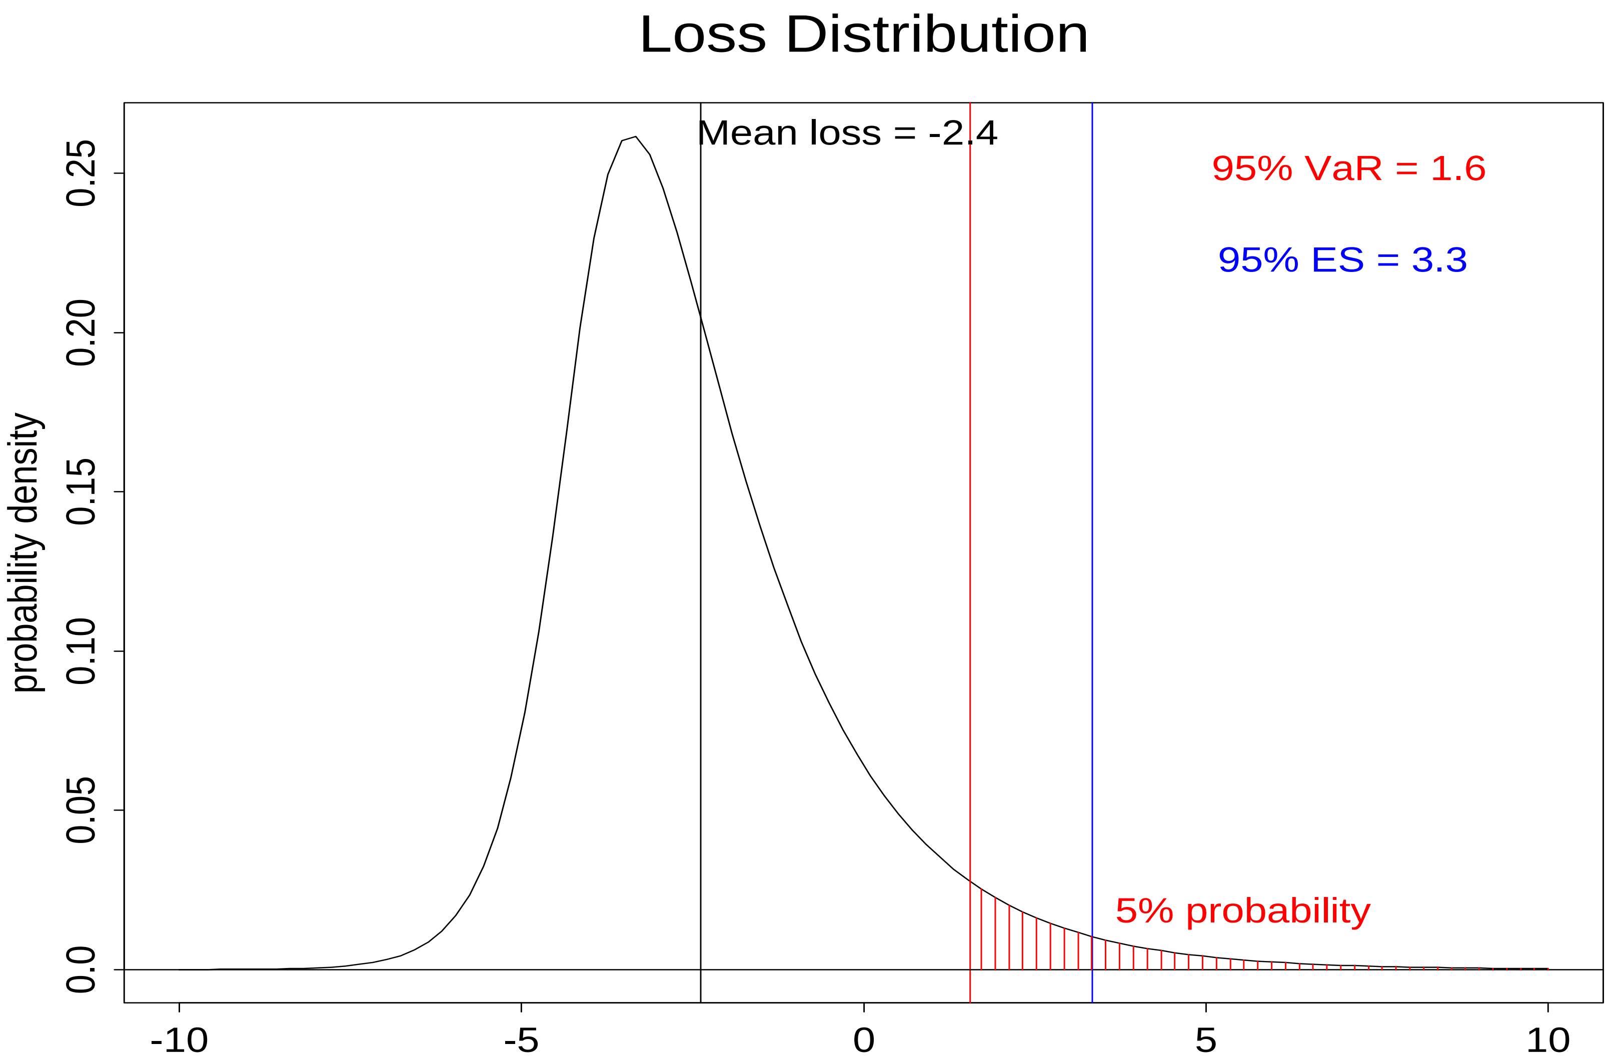

B. Scaling

Question 1: how to get a 10-day VaR (or 1-year VaR)?

Solution in the praxis: scale the 1-day VaR by $\sqrt{10}$ (or $\sqrt{250}$).

Question 2: how good is scaling?

$\longrightarrow$ model dependent!

Scaling under Normality

Under the assumption

$$X_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2),$$

n-day log-returns are normally distributed as well:

$$\sum_{i=1}^{n} X_i \sim \mathcal{N}(0, n\sigma^2).$$

For a $\mathcal{N}(0, \widetilde{\sigma}^2)$-distributed profit X, $\mathrm{VaR}_p(X) = \widetilde{\sigma}\, x_p$, where x_p denotes the p-quantile of a standard normal distribution. Hence

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

AR(1)-GARCH(1,1) Processes

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

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

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

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

(typical parameters: $\lambda = 0.04$, $a_0 = 3 \cdot 10^{-6}$, $a = 0.05$, $b = 0.92$)

Scaling for AR(1)-GARCH(1,1) Processes

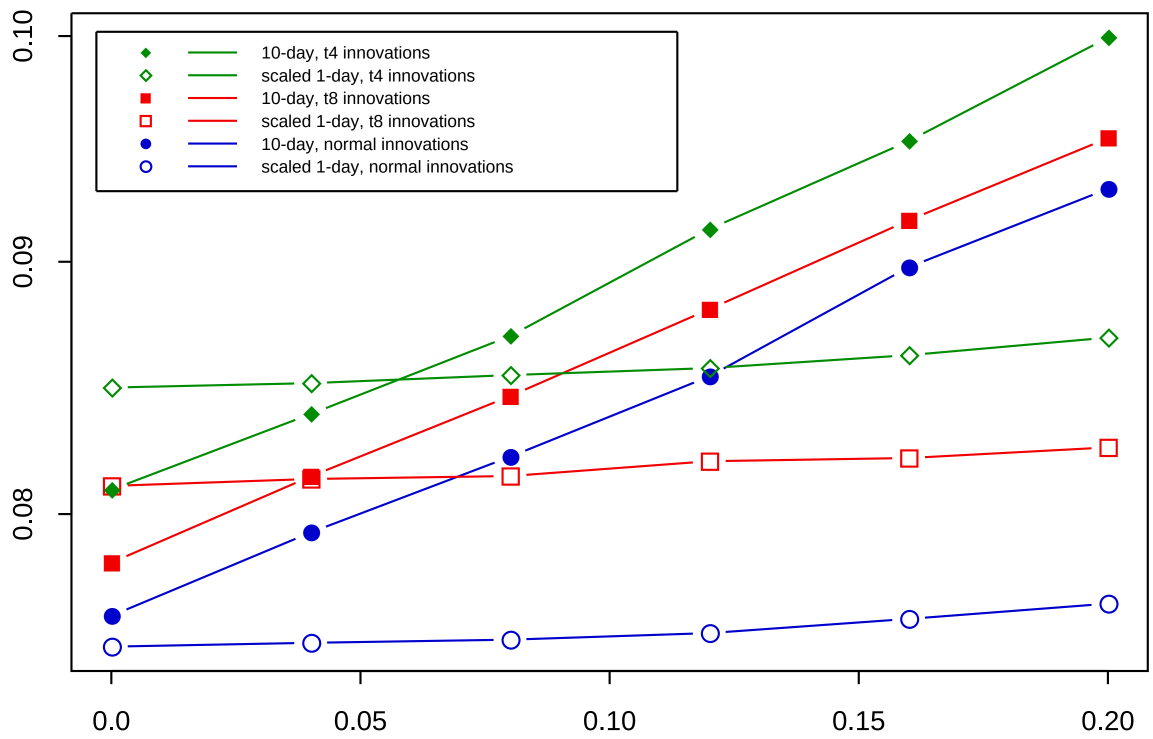

Goodness of fit of the scaling rule, depending on different values of λ (x axis) for different distributions of the innovations ϵ_t.

For typical parameters ($\lambda = 0.04$, $\epsilon_t \sim t_8$), the fit is almost perfect.

C. One-Year Risks

Problems when modelling yearly data:

- Non-stationarity of data.
- Lack of yearly returns.
- Properties of yearly data are different from those of daily data.

How to Estimate Yearly Risks

- Fix a horizon $h < 1$ year, for which data can be modelled.
- Use a scaling rule for the gap between h and 1 year.

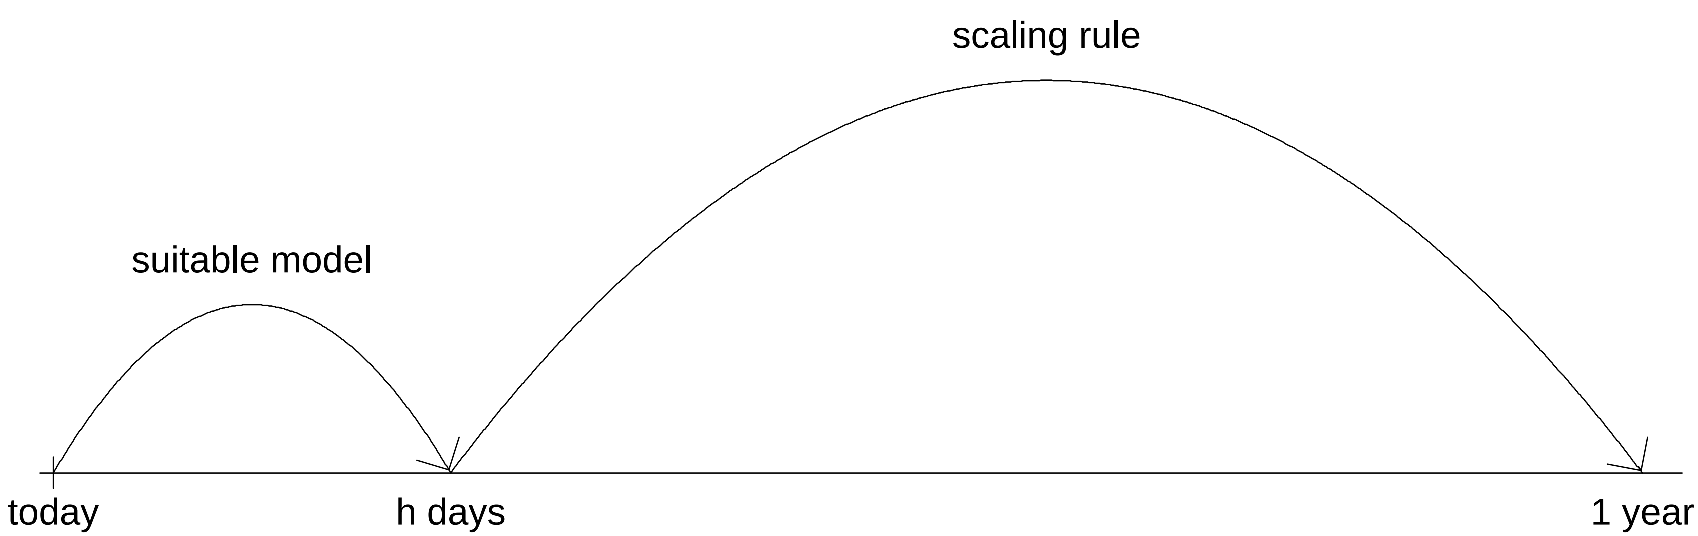

Models for One-Year Risks

- Random Walks
- Autoregressive Processes
- GARCH(1,1) Processes
- Heavy-tailed Distributions

Backtesting

The suitability of these models for estimating one-year financial risks can be assessed by backtesting estimated value-at-risk and expected shortfall using observed return data for

- stock indices,
- foreign exchange rates,
- 10-year government bonds,
- single stocks.

Conclusions for One-Year Forecasts

- In general, the random walk model performs 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 these data, the square-root-of-time rule (accounting for trends) can be applied for estimating one-year risks.

D. Conclusions

- The square-root-of-time scaling rule performs very well to scale risks from a short horizon (1 day) to a longer one (10 days, 1 year).

- The reasons for this good performance are non-trivial. Each situation has to be investigated individually. The square-root-of-time rule should not be applied before checking its appropriateness.

- In the limit, as $\alpha \to 1$, scaling a short-term VaR_α to a long-term VaR_α using the square-root-of-time rule is in most situations not appropriate any more.