Reliable Quantification and Efficient Estimation of Credit Risk

Jörn Dunkel
University of Cambridge

Stefan Weber*
Leibniz Universität Hannover

August 23, 2011

Abstract

The recent crisis in the global financial markets requires a critical review of current regulatory practice. Corporate governance is a key issue since improper due diligence and myopic incentive schemes for employees contributed largely to the credit bubble and the ensuing recession. A better oversight by regulatory authorities and institutional changes of financial firms pose political and economic challenges in their own right. In addition, substantial efforts are required to devise efficient quantitative methods that allow one to more reliably measure financial risks in the future. These tools should be able to detect extreme scenarios that are very unlikely to occur but whose impact may be dramatic as illustrated by the recent liquidity crisis of Lehman Brothers, Merrill Lynch, Fannie Mae, Freddy Mac, AIG, and others. We report here a novel Monte-Carlo (MC) approach for the direct and efficient computation of an important class of financial portfolio risk measures, known as *Shortfall Risk* (SR). Unlike the current industry standard *Value-at-Risk* (VaR), convex risk measures such as SR are sensitive to the tails of loss distributions and provide financial institutions with incentives to diversify properly.

Portfolio models.– Risk management in practice involves two complementary tasks: the construction of accurate portfolio models (Cre 1997, Gupton, Finger & Bhatia 1997, Gordy 2000, Frey & McNeil 2003, McNeil, Frey & Embrechts 2005, Frey, Popp & Weber 2008), and the reliable quantification of the downside risk for these models (Artzner, Delbaen, Eber & Heath 1999, Föllmer & Schied 2011, Frey & McNeil 2002, Tasche 2002, Weber 2006). The first task covers both the design and the calibration of models to available data. In domains where data are scarce, models need to be extrapolated based on an understanding of the underlying economic mechanisms. The second task, the definition of well-defined benchmarks, is crucial since applied risk management and financial regulation require simple summary statistics which correctly reflect the loss exposure (Artzner et al. 1999, Frey & McNeil 2002, Tasche 2002).

A broad class of credit portfolio models (Cre 1997, Gupton et al. 1997) specify the total loss $L \geq 0$ over a fixed period (e.g., a day, month or year) by

$$L = \sum_{i=1}^{m} v_i D_i. \tag{1}$$

Here, m is the number of portfolio positions (obligors), v_i the partial monetary loss that occurs if the obligor i defaults within this period, and D_i is the random default variable taking

Acknowledgement: The authors would like to thank Thomas Knispel for helpful remarks.

*Corresponding author: Leibniz Universität Hannover, Institut für Mathematische Stochastik, Welfengarten 1, 30167 Hannover, Germany, email sweber@stochastik.uni-hannover.de.

values 0 ('no default') or 1 ('default'). Realistic models take into account that the default risk of different positions may be interdependent (Cre 1997, Gupton et al. 1997, Gordy 2000, Frey & McNeil 2003). Underlying mechanisms include observable and hidden economic risk factors, global feedback effects and local interactions (Embrechts, McNeil & Straumann 2002). A pragmatic and popular way for modeling dependences uses a factor structure (Cre 1997, Gupton et al. 1997, Frey & McNeil 2003, Kang & Shahabuddin 2005). Within this approach default indicators are constructed as binary functions $D_i = \Theta(\boldsymbol{a}_i \cdot \boldsymbol{Z} - x_i) \in \{0, 1\}$ with Θ denoting the Heaviside function $[\Theta(z) := 0, z \leq 0;\ \Theta(z) := 1, z > 0]$ and fixed threshold parameters $(x_1, \ldots, x_m)$. The random vector $\boldsymbol{Z} = (Z_1, \ldots, Z_d)$ comprises common or individual risk factors whose joint distribution is specified as an ingredient of the model. Potential dependences between portfolio positions are encoded through coupling parameters $\boldsymbol{a}_i = (a_{ij})_{j=1,\ldots,d}$, which have to be deduced from historical data. For realistic models, it is usually impossible to analytically evaluate the full loss distribution $\mathbb{P}[L \leq x]$, and numerical simulations must be employed. A naive computer experiment would first sample the random numbers (Z_j) under the model probability measure $\mathbb{P}$ and subsequently calculate D_i and L. By repeating this procedure several times, one can estimate specific values of the loss distribution function, the mean loss $\mathbb{E}[L]$, the variance, or other relevant quantities. Of particular interest with regard to risk estimation are quantities which characterize extreme events that cause large losses (Artzner et al. 1999, Glasserman, Heidelberger & Shahabuddin 2002, Föllmer & Schied 2011, Giesecke, Schmidt & Weber 2008). The recent market turmoil has clearly demonstrated that such scenarios may have serious global consequences for the stability of the financial system as well as the real economy, but they typically occur with very low probability; i.e., reliable predictions require advanced MC simulation techniques (Glasserman 2004). The novel method reported here allows for an efficient estimation and sensible characterization of big-loss scenarios through convex risk measures. The approach is generically applicable whenever the loss variable L can be sampled from a given set of rules similar to those outlined above.

Risk measures.– The theoretical foundations for the systematic measurement of financial risks were laid almost a decade ago (Artzner et al. 1999, Föllmer & Schied 2002); the numerical implementation of well-defined risk quantification schemes is, however, still a work-in-progress (Glasserman et al. 2002, Kang & Shahabuddin 2005). Risk measures, as specified in Eqs. (2) or (4) below, define the monetary amount s_* that should be available to insure against potentially large losses. The value s_* is called 'capital requirement' or 'economic capital' and depends on both the underlying portfolio model and the adopted risk measure. A major responsibility of regulatory authorities consists in identifying appropriate standards for risk measurement that prevent improper management of financial risks. Below, we describe an efficient MC method for estimating the risk measure *Shortfall Risk* (SR). Unlike the current industry standard of risk assessment *Value-at-Risk* (VaR) (Jorion 2000, Glasserman et al. 2002), SR encourages diversification and is well-suited for characterizing rare big-loss scenarios. The severe deficiencies of VaR become evident upon analyzing its definition: For a fixed loss level $\lambda \in (0, 1)$, VaR is defined by

$$\begin{aligned} \mathrm{VaR}_\lambda &:= \inf\{s \in \mathbb{R} \mid \mathbb{P}[L > s] \leq \lambda\} \\ &= \inf\{s \in \mathbb{R} \mid \mathbb{E}[\Theta(L - s)] \leq \lambda\}. \end{aligned} \tag{2}$$

Representing a quantile of the loss distribution, VaR provides the threshold value that is exceeded by the loss L only with a small probability λ, but it ignores the shape of the loss distribution beyond the threshold. Very large losses are systematically underestimated by VaR. Consider

e.g. a portfolio with loss distribution

$$L = \begin{cases} 0, & \text{with probability } 99.9\,\% \text{ (no loss)}, \\ \$10^{10} & \text{with probability } \ 0.1\ \% \text{ (big loss)}. \end{cases} \tag{3}$$

Adopting the customary value $\lambda = 0.01$, one finds in this case $\mathrm{VaR}_\lambda = 0$, i.e., according to this risk measure, the portfolio does not require any economic capital although there exists a considerable chance of losing billions of dollars.

The severe deficiencies of VaR can be fixed by replacing the Θ-function in Eq. (2) with a convex, increasing loss function $\ell \geq 0$, which leads to the definition of SR, see e.g. Chapter 4.9 in Föllmer & Schied (2011):

$$\mathrm{SR}_\lambda := \inf\{s \in \mathbb{R} \mid \mathbb{E}[\ell(L-s)] \leq \lambda\} \tag{4}$$

where now $\lambda > 0$. Typical examples are exponential or (piecewise) polynomial loss functions,

$$\ell_\beta(y) \ = \ \exp(y/\beta), \qquad \ell_{\alpha,\eta}(y) = \eta^{-1}(y/\alpha)^\eta\, \Theta(y), \tag{5}$$

with scale parameters $\alpha, \beta > 0$ and $\eta \geq 1$. The function ℓ determines how strongly large losses are penalized. In the case of example (3), exponential SR with $\lambda = 0.01$ and $\beta = 2$ demands a capital requirement $s_* = \mathrm{SR}^\beta_\lambda(L) \approx \$10^9$, reflecting the actual size of potentially large losses. In contrast to VaR, SR risk measures provide a flexible tool for regulatory authorities to devise good risk measurement schemes.

Shortfall-Risk & importance sampling.– Equation (4) implies that SR is equal to the unique root s_* of the function

$$g(s) := \mathbb{E}[\ell(L-s) - \lambda] \tag{6}$$

(see e.g. Chapter 4.9 in Föllmer & Schied (2011)).

For realistic portfolio models, the functional value g at a given argument s can only be estimated numerically. A naive algorithm would sample n random variables L_k according to the rules of the model, cf. Eq. (1), and compute the simple estimator $\hat{g}_n(s) = n^{-1}\sum_{k=1}^n \hat{G}(L_k, s)$ where $\hat{G}(L_k, s) = \ell(L_k - s) - \lambda$. This procedure is often inefficient for practically relevant loss distributions, since the variance of $\hat{g}_n(s)$ can be large. Improved estimates can be obtained by importance sampling (IS), defined as follows: Assume L is governed by the probability density $p(x)$, abbreviated by $L \sim p$. For another, possibly s-dependent probability density $x \mapsto f_s(x)$, we may rewrite [1]

$$g(s) = -\lambda + \int \mathrm{d}x\, f_s(x)\, \frac{p(x)}{f_s(x)}\, \ell(x-s). \tag{7}$$

Consequently, $g_n(s) = n^{-1}\sum_{k=1}^n G(L_k, s)$ with

$$G(L_k, s) := -\lambda + \frac{p(L_k)}{f_s(L_k)}\, \ell(L_k - s), \quad L_k \sim f \tag{8}$$

is another estimator for $g(s)$. Compared with the naive estimator $\hat{g}_n$, the variance of g_n can be substantially reduced, if the IS density f_s is chosen appropriately (Dunkel & Weber 2007).

[1] $f_s(x)$ is assumed to be non-zero if $p(x)\ell(x-s) > 0$.

Hence, to estimate SR one could try to combine IS with conventional root finding schemes, e.g., by defining a recursive sequence $s_j = R[s_{j-1}, \dots, s_1; g(s_{j-1}), \dots]$ using the secant method (Press, Vetterling & Flannery 2002). However, this approach suffers from drawbacks: Firstly, accurate estimates $g_n(s_j)$ of $g(s_j)$ at each point of the sequence $\{s_j\}$ are required which can be computationally expensive. Secondly, cross-averaging of errors for different values of s is not exploited. The algorithm below resolves these problems and yields a direct estimate of the SR value s_* by combining importance sampling with a stochastic root-finding scheme (Ruppert 1988, Ruppert 1991, Polyak & Juditsky 1992).

Stochastic root-finding algorithm.– We focus here only on those aspects that are relevant for the practical implementation; a theoretical analysis will be given elsewhere (Dunkel & Weber 2010). The proposed algorithm consists of the following steps:

1. Choose a fixed interval $[a, b]$ that contains the root s_*. Fix an initial value $s_1 \in [a, b]$, and constants $\gamma \in (\frac{1}{2}, 1]$ and $c > 0$.

2. Sample L_n from the IS density f_{s_n} and calculate

$$s_{n+1} = \Pi \left\{ s_n + \frac{c}{n^\gamma}\, G(L_n, s_n) \right\}, \tag{9}$$

   where Π denotes a projection on the interval $[a, b]$, i.e., $\Pi\{x\} := a$ if $x < a$, $\Pi\{x\} := x$ if $x \in [a, b]$, and $\Pi\{x\} := b$ if $x > b$.

The sequence s_n defined by (9) converges to the SR value s_* as $n \to \infty$. More precisely, one can prove that, if c is chosen large enough, so that $c > [-2g'(s_*)]^{-1}$, then the distribution of the rescaled quantity

$$S_n := \sqrt{n^\gamma}\,(s_n - s_*) \tag{10}$$

converges to a Gaussian normal distribution $\mathcal{N}(\mu_*, \Sigma_*^2)$ with mean $\mu_* = 0$ and constant variance

$$\Sigma_*^2 = c^2\,\sigma^2(s_*) \begin{cases} [2c\,g'(s_*)]^{-1}, & \gamma \in (\frac{1}{2}, 1), \\ [2c\,g'(s_*) + 1]^{-1}, & \gamma = 1, \end{cases} \tag{11}$$

where $\sigma^2(s)$ is the variance of the random variable $G(L, s)$ defined in (8). Equation (10) shows that γ determines the rate of convergence of the algorithm to s_*. The asymptotic variance in (11) can be improved by applying IS techniques that reduce the variance of $\sigma^2(s)$. This feature is particularly important when dealing with realistic loss distributions. Numerical values for the *a priori* unknown quantities $\sigma^2(s_*)$ and $g'(s_*)$ can be obtained from previously stored simulation data $\{(s_i, L_i, p(L_i), f_{s_i}(L_i))\}$ by using the numerically obtained root s_* to evaluate the estimators

$$\sigma_n^2(s_*) \;=\; \frac{1}{\rho n} \sum_{i=n(1-\rho)}^{n} G(L_i, s_i)^2, \qquad \rho \in (0, 1), \tag{12}$$

$$g'_{n,\epsilon}(s_*) \;=\; \frac{1}{\epsilon n} \sum_{i=1}^{n} \left[\frac{p(L_i)}{f_{s_i}(L_i)} \ell(L_i - (s_* + \epsilon)) - \lambda \right] \tag{13}$$

for a sufficiently small $\epsilon > 0$. Estimates of $\sigma^2(s_*)$ and $g'(s_*)$ can be used for the construction of confidence intervals for s_*.

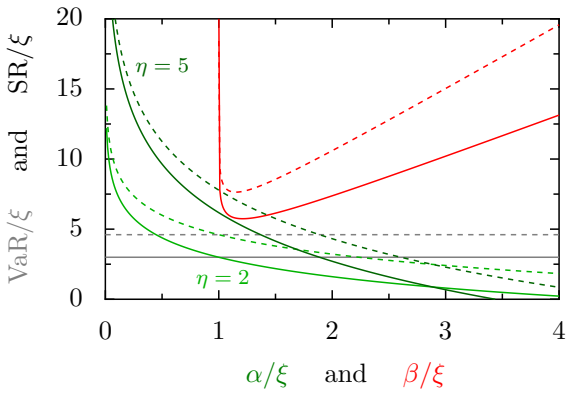

Figure 1: Comparison of risk measures for a light-tailed exponential loss distribution (15): VaR_λ (grey), exponential $\mathrm{SR}_\lambda^\beta$ (red), and polynomial $\mathrm{SR}_\lambda^{\alpha,\eta}$ (green) in units of the mean loss ξ for levels $\lambda = 0.05$ (solid) and $\lambda = 0.01$ (dashed) plotted as functions of the rescaled parameters β/ξ and α/ξ, respectively.

Variance reduction is not only important to decrease the asymptotic variance in (11), but also for improving the finite sample properties of the algorithm. S_n shows quasi-Gaussian behavior for much smaller values of n if IS is applied. At the same time, the estimators in (12) and (13) perform considerably better. The optimal choice of the constant c, which minimizes the variance in (11), is not known a priori. In practice, the optimal asymptotic variance can thus hardly be achieved. A solution to this problem is to average the estimator s_n given by (9) over the last $\rho \times n$ sampling steps, i.e., to return the estimator

$$\bar{s}_n = \frac{1}{\rho n} \sum_{i=n(1-\rho)}^{n} s_i, \qquad \rho \in (0,1). \tag{14}$$

In this case, one can show that, for $\gamma \in (\frac{1}{2}, 1)$ and $c > [-2g'(s_*)]^{-1}$, the distribution of the rescaled quantity $\bar{S}_n := \sqrt{\rho n}\,(\bar{s}_n - s_*)$ converges to the Gaussian distribution $\mathcal{N}(0, \sigma^2(s_*)/[g'(s_*)]^2)$ as $n \to \infty$. Apart from a factor $1/\rho$, the asymptotic variance then corresponds to the optimal choice for c in (11) in the case of an optimal convergence rate $\gamma = 1$.

Applications.– Due to the generic definition of the sequences s_n and $\bar{s}_n$, the above scheme is applicable to a wide range of portfolio models, and can be combined with various model-specific variance reduction techniques (Glasserman 2004). To explicitly demonstrate the efficiency of the algorithms and to further illustrate the advantages of SR compared with VaR, we study two generic, stylized scenarios: a light-tailed exponential loss distribution with density

$$p(x) := \mathrm{d}\mathbb{P}[L < x]/\mathrm{d}x = \xi^{-1} \exp(-x/\xi)\Theta(x), \tag{15}$$

and a heavy-tailed power law distribution with density

$$p_\kappa(x) = \frac{(\kappa-1)\,[(\kappa-2)\xi]^{\kappa-1}}{[x+(\kappa-2)\xi]^\kappa}\,\Theta(x), \qquad \kappa > 2, \tag{16}$$

where $\xi > 0$, respectively. In both cases, the mean loss is given by $\mathbb{E}[L] = \xi$, but ruinous losses are more likely to occur under the power law distribution (16).

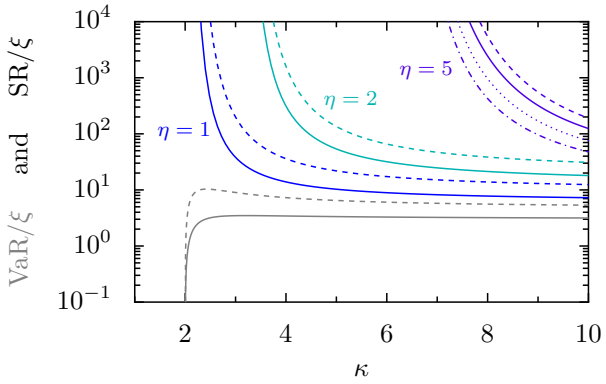

Figure 2: VaR_λ (grey) and polynomial $\mathrm{SR}_\lambda^{\alpha,\eta}$ (colored) for the heavy-tailed distribution (16) plotted as a function of the exponent κ. Solid (dashed) lines correspond to levels $\lambda = 0.05$ (0.01), with $\alpha = 0.5$ in the case of SR. For SR with $\eta = 5$ (violet), additional curves with $\alpha = 1.0, \lambda = 0.05$ (dash-dotted) and $\alpha = 1.0, \lambda = 0.01$ (dotted) are shown. In the heavy-tail limit $\kappa \to 2$, VaR tends to zero and, thus, becomes inadequate for defining securities in this regime.

For exponentially distributed losses and loss functions (5), the risk measures VaR and SR can be calculated analytically as

$$\begin{aligned} \mathrm{VaR}_\lambda &= \xi \log(\lambda^{-1}), \\ \mathrm{SR}_\lambda^{\beta} &= \beta \log[\lambda^{-1}(1-\xi/\beta)^{-1}], \\ \mathrm{SR}_\lambda^{\alpha,\eta} &= \xi \log[\lambda^{-1}(\xi/\alpha)^{\eta}\Gamma(\eta)], \end{aligned} \tag{17}$$

with Γ denoting the Gamma function. Finite positive SR values are obtained for $\beta > \xi$ and $\alpha < \xi[\Gamma(\eta)/\lambda]^{1/\eta}$. Figure 1 compares the three risk measures (17) for two values for λ. We plot the risk measures in units of ξ as a function of the normalized scale parameters α/ξ and β/ξ. For the exponential loss distribution, the probability of large losses increases with its mean value ξ. Figure 1 illustrates not only the dependence on α or β, respectively, but also how the risk measures behave as a function of the mean loss ξ. While VaR (grey) is proportional to ξ, polynomial SR (green) grows more than proportionally with ξ. Exponential SR (red), on the other hand, increases for small ξ less than proportionally, but diverges rapidly as ξ approaches the parameter β. These specific characteristics must be taken into account by regulatory authorities and risk managers in order to devise and implement reasonable policies.

In the case of the heavy-tail distribution (16) exponential SR diverges, but VaR and polynomial SR with $1 \le \eta < \kappa - 1$ remain finite, yielding

$$\begin{aligned} \mathrm{VaR}_\lambda &= (\kappa-2)\left(\lambda^{-1/(\kappa-1)} - 1\right)\xi, \\ \mathrm{SR}_\lambda^{\alpha,\eta} &= (2-\kappa)\xi + \left\{\frac{[(\kappa-2)\xi]^{\kappa-1}}{\lambda\,\alpha^\eta\, C(\eta,\kappa)}\right\}^{1/(\kappa-1-\eta)}, \end{aligned} \tag{18}$$

where $C(\eta,\kappa) = \Gamma(\kappa-1)/[\Gamma(\eta)\,\Gamma(\kappa-1-\eta)]$. As evident from Eqs. (18) and Fig. 2, VaR (grey) vanishes in the heavy-tail limit $\kappa \to 2$, even though the tail risk is increased for smaller values of κ. By contrast, SR (colored) provides a reasonable risk measure for the whole parameter range.

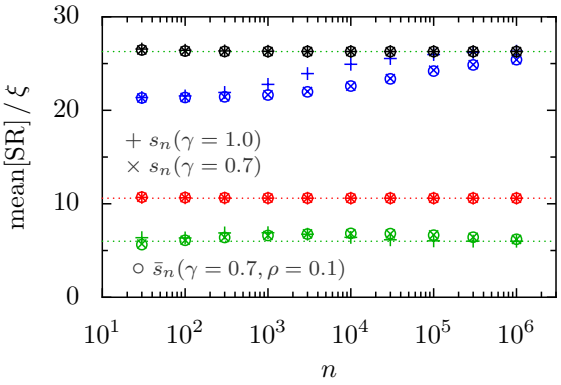

Figure 3: Numerical SR estimates s_n and $\bar{s}_n$ as obtained from $N = 10^4$ simulation runs using $\lambda = 0.01$; the corresponding variances are depicted in Fig. 4. Red/green symbols: Exponential/polynomial SR for a light-tailed exponential loss distribution (15), using $c = 500$ and direct sampling. The estimators converge rapidly to the exact theoretical value (dotted, cf. Fig. 1) for exponential $\mathrm{SR}_\lambda^\beta$ ($\beta = 2\xi$; red), while the convergence is considerably slower for polynomial $\mathrm{SR}_\lambda^{\alpha,\eta}$ ($\alpha = 0.5\xi, \eta = 2$; green). Blue/black symbols: Polynomial $\mathrm{SR}_\lambda^{\alpha,\eta}$ estimated for the heavy-tail power-law distribution (16), using $c = 10^3$ and parameters $\alpha = 0.5\xi, \eta = 1, \kappa = 4$, cf. Fig. 2. Compared with direct sampling (blue), the importance sampling estimators (black) converge much faster.

We can use the analytic expressions (17) and (18) to verify the convergence behavior of the proposed algorithm. Figures 3 and 4 depict numerical results obtained from $N = 10^4$ sample runs for a fixed loss level $\lambda = 0.01$ and different values of n and γ (colors correspond to those in Figs. 1, 2). The diagrams show the sample mean values and variances of data sets $\{s_n^{(1)}, \ldots, s_n^{(N)}\}$ and $\{\bar{s}_n^{(1)}, \ldots, \bar{s}_n^{(N)}\}$, respectively. For each run (k) the initial value $s_1^{(k)}$ was randomly chosen from the search interval $[a, b] = [s_* - 5, s_* + 5]$ where s_* is the exact analytical value. One readily observes that in all examples the estimators converge to the exact values (dotted lines in Fig. 3). The convergence speed, however, depends on the underlying loss distribution as well as on the loss function. Generally, SR estimates based on s_n or $\bar{s}_n$ can be considered reliable when the variance decreases with $n^{-\gamma}$ or n^{-1}, respectively, in accordance with Gaussian asymptotics for the rescaled quantities $S_n = \sqrt{n^\gamma}(s_n - s_*)$ and $\bar{S}_n = \sqrt{n}(\bar{s}_n - s_*)$. As evident from both Fig. 3 and 4, for the light-tailed exponential distribution $p(x) = \xi^{-1} \exp(-x/\xi)\Theta(x)$ the exponential SR estimators (red) converge very rapidly, while for polynomial SR (green) the convergence is slower but still acceptable even without importance sampling (i.e., if L_n is directly sampled from p).

By contrast, and not surprisingly, for the heavy-tail distribution (16) direct sampling (blue) of L_n from p_κ results in poor convergence behavior. In such cases, variance reduction techniques like IS (black) can significantly improve the performance. As guidance for future implementations, we outline the IS procedure in more detail: Instead of sampling losses from the original distribution p_κ, we consider the following 'shifted' power law density

$$f_{\nu,s}(x) = \frac{(\nu - 1)\,(\zeta + s)^{\nu-1}}{(x + \zeta)^\nu}\,\Theta(x - s), \tag{19}$$

where $\zeta > -s$ and $1 < \nu < 2(\kappa - \eta) - 1 =: \nu_+$. The latter condition ensures finiteness of the second moment. We have to determine ν and ζ such that sampling from $f_{\nu,s}$ yields a better convergence to the correct SR value s_*. To this end, we note that the likelihood ratio

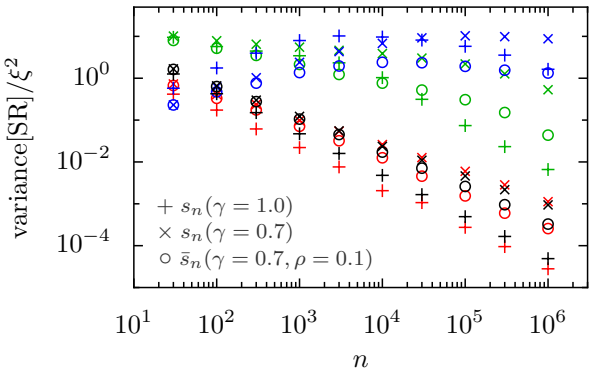

Figure 4: Sample variances of the SR estimates from Fig. 3, using the same colors/symbols. Estimates can be considered as reliable when the variance of s_n (+/×) or $\bar{s}_n$ (○) decreases with $n^{-\gamma}$ or n^{-1}, respectively. For the heavy-tail distribution (16), importance sampling (black) is much more efficient than direct sampling (blue).

$h(x) := p_\kappa(x)/f_{\nu,s}(x)$ takes its maximum at $x_- = s$, representing the effective lower integral boundary in Eq. (7), if

$$\zeta \geq [(\kappa - 2)\nu\xi - (\kappa - \nu)s]/\kappa. \tag{20}$$

We fulfill this condition by fixing $\zeta = \nu(\xi + s)$. One then finds that variance reduction, corresponding to $h(x) < 1$ for $x \geq s$, is achieved if

$$\nu > \nu_*(s) := -\frac{2\kappa_1 s\,(\kappa_2\xi)^\kappa + \kappa_2\xi\,(s + \kappa_2\xi)^\kappa}{\kappa_1\xi\,(\kappa_2\xi)^\kappa - \kappa_2\xi\,(s + \kappa_2\xi)^\kappa}, \tag{21}$$

where $\kappa_2 = \kappa - 2, \kappa_1 = \kappa - 1$, and $\nu_* \to 1$ for $s \to \infty$. Accordingly, we sample $L_n \sim p_\kappa$ if $\nu_*(s_n) > \nu_+$, and $L_n \sim f_{\nu,s_n}$ with $\nu = 0.5[\nu_*(s_n) + \nu_+]$ if $\nu_*(s_n) < \nu_+$. Intuitively, by sampling from f_{ν,s_n} losses beyond s_n become more likely, while simultaneously suppressing the tail if $\nu_+ > \kappa$. As evident from Fig. 3 and 4, both aspects contribute to a vastly improved convergence. Most importantly, however, this strategy can be extended to more general models without much difficulty, e.g., by combining the stochastic root finding scheme with standard variance reduction techniques (Glasserman 2004) for the factor variables $\boldsymbol{Z}$ in Eq. (1).

Summary.– Financial risk measures have been studied systematically for almost a decade (Artzner et al. 1999, Gordy 2000, Föllmer & Schied 2002, Weber 2006, McNeil et al. 2005). The financial industry, however, is still almost exclusively relying on the deficient risk measure Value-at-Risk (Glasserman et al. 2002, Jorion 2000), or even less sophisticated methodologies. The recent financial turmoil leaves little doubt about the importance of adequate risk quantification schemes. The above discussion clarifies how well-defined, tail-sensitive shortfall risk measures can be efficiently evaluated by combining stochastic root-approximation algorithms with variance reduction techniques. These tools can provide a basis for more sensible risk management policies and, thus, help to prevent future crises.

References

Artzner, P., F. Delbaen, J.-M. Eber & D. Heath (1999), 'Coherent measures of risk', *Mathematical Finance* **9**(3), 203–228.

Cre (1997), *CreditRisk$^+$: A CreditRisk Management Framework*.

Dunkel, J. & S. Weber (2007), Efficient Monte Carlo Methods for Convex Risk Measures in Portfolio Credit Risk Models, *in* S. G.Henderson, B.Biller, M.-H.Hsieh, J.Shortle, J. D.Tew & R. R.Barton, eds, 'Proceedings of the 2007 Winter Simulation Conference', IEEE (Piscataway, NJ), Washington, D.C., pp. 958–966.

Dunkel, J. & S. Weber (2010), 'Stochastic root finding and efficient estimation of convex risk measures', *Operations Research* **58**(5), 1505–1521.

Embrechts, P., A. J. McNeil & D. Straumann (2002), Correlation and dependency in risk management: Properties and pitfalls, *in* M.Dempster, ed., 'Risk Management: Value at Risk and Beyond', Cambridge University Press, pp. 176–223.

Föllmer, H. & A. Schied (2002), Robust representation of convex measures of risk, *in* 'Advances in Finance and Stochastics. Essays in Honour of Dieter Sondermann', Springer-Verlag, pp. 39–56.

Föllmer, H. & A. Schied (2011), *Stochastic Finance - An Introduction in Discrete Time*, 3. edn, Walter de Gruyter.

Frey, R. & A. J. McNeil (2002), 'VaR und expected shortfall in portfolios of dependent credit risks: conceptual and practical insights', *Journal of Banking and Finance* **26**, 1317–1334.

Frey, R. & A. J. McNeil (2003), 'Dependant defaults in models of portfolio credit risk', *Journal of Risk* **6**(1), 59–92.

Frey, R., M. Popp & S. Weber (2008), 'An approximation for credit portfolio losses', *The Journal of Credit Risk* **4**(1), 3–20.

Giesecke, K., T. Schmidt & S. Weber (2008), 'Measuring the risk of large losses', *Journal of Investment Management* **6**(4), 1–15.

Glasserman, P. (2004), *Monte Carlo Methods in Financial Engineering*, number 53 *in* 'Applications of Mathematics', Springer, New York.

Glasserman, P., P. Heidelberger & P. Shahabuddin (2002), 'Portfolio value-at-risk with heavy-tailed risk factors', *Mathematical Finance* **12**(3), 239–269.

Gordy, M. (2000), 'A comparative anatomy of credit risk models', *Journal of Banking and Finance* **24**, 119–149.

Gupton, C., C. Finger & M. Bhatia (1997), *CreditMetrics Technical Document*, J. P. Morgan & Co., New York. `www.riskmetrics.com`.

Jorion, P. (2000), *Value at Risk*, 2. edn, McGraw-Hill Companies.

Kang, W. & P. Shahabuddin (2005), Fast simulation for multifactor portfolio credit risk in the t-copula model, *in* M. E.Kuhl, N. M.Steiger, F. B.Armstrong & J. A.Joines, eds, 'Proceedings of the 2005 Winter Simulation Conference', INFORMS, pp. 1859–1868.

McNeil, A. J., R. Frey & P. Embrechts (2005), *Quantitative Risk Management: Concepts, Techniques and Tools*, Princeton Series in Finance, Princeton University Press, Princeton, NJ.

Polyak, B. T. & A. B. Juditsky (1992), 'Acceleration of stochastic approximation by averaging', *SIAM Journal on Control and Optimization* **30**, 838–855.

Press, W. H., W. T. Vetterling & B. P. Flannery (2002), *Numerical recipes in C++: The art of scientific computing*, 2 edn, Cambridge University Press, Cambridge.

Ruppert, D. (1988), 'Efficient estimators from a slowly convergent Robbins-Monro procedure', *ORIE Technical Report 781, Cornell University* .

Ruppert, D. (1991), Stochastic approximation, *in* B.Gosh & P.Sen, eds, 'Handbook of Sequential Analysis', Marcel Dekker, New York, pp. 503–529.

Tasche, D. (2002), 'Expected shortfall and beyond', *Journal of Banking and Finance* **26**(7), 1519–1533.

Weber, S. (2006), 'Distribution-invariant risk measures, information, and dynamic consistency', *Mathematical Finance* **16**(2), 419–442.