Modeling Market Risk

by Marius Bochniak

This Chapter Covers

The definition of market risk.
General approaches to the measurement of market risk.
Regulatory requirements concerning market risk.
The definition of value-at-risk (VaR).
Historical and Monte Carlo simulation of VaR.
The scaling of VaR between different time horizons.
A comparison of historical simulation and Monte Carlo methods.

Introduction

Every financial institution with a portfolio exposed to market risk should have a model in place which is designed to measure that risk. Such a model allows one to control and to limit the market risk taken by each desk or by the traders and to charge each portfolio position a cost of capital required to cover its market risk. Success in meeting these objectives serves the interests of the stakeholders in the firm.

The measures of market risk employed by financial institutions are usually based on the mathematical concept of value-at-risk (VaR). Roughly speaking, VaR is defined as the maximum loss of a portfolio over some target period that will be not exceeded with a specified probability, or confidence level. Various techniques have been developed to estimate the VaR of trading portfolios. Some, like the variance–covariance approach, are outdated and rarely used in practice as they are incompatible with regulatory requirements to capture the nonlinear behavior of derivative instruments and they ignore event risk. The two main approaches currently used in financial institutions are historical simulation and Monte Carlo methods.

In the present overview we briefly describe the main ideas of market risk modeling and present the two main approaches to such modeling in more detail. As the simple forms of both historical simulation and the Monte Carlo have some undesirable properties, we show how both approaches can be improved. Finally, we compare the benefits and the drawbacks of the different approaches.

Market Risk

According to the Basel II framework, the banks must mark-to-market their trading books at least daily, which means that they must revalue all trading book positions at readily available close-out prices in orderly transactions that are sourced independently (Basel Committee on Banking Supervision (BCBS), 2006). Market risk is the risk that the market value of positions may change in the future. It is of little consequence to investors who purchase financial instruments with the intention of holding them until maturity or for long time periods.

The market value of trading book positions depends on:

risk factors that are observable on the market, such as share and commodity prices, interest rates, credit spreads and foreign exchange rates;
additional pricing parameters like volatilities and correlations that are not directly observable and which must be implied from the market prices of financial instruments.

All pricing parameters contribute to the market risk of positions, but the set of pricing parameters that must be captured in a particular market risk model depends on the regulatory requirements.

The same framework sets capital requirements to cover potential losses resulting from market risk in the trading book. These capital requirements are formulated in terms of several risk measures which capture different types of market risk.

How to Measure Market Risk

Due to the stochastic nature of financial markets, it is obviously not possible to predict exactly the future value of a portfolio. However, knowing the past behavior of the risk factors that drive the market value of the portfolio, it is possible to generate a stochastic set of possible scenarios for the future value of the portfolio. The idea is as follows:

we describe the past behavior of underlying risk factors by means of probability distributions;
we generate stochastic forecasts of the future behavior of the risk factors using their probability distributions;
we revalue the portfolio for each forecast of the risk factors.

In this way we obtain a probability distribution of the possible future values of the portfolio. The risk measures can now be defined as special characteristics of this distribution.

The risk measure used in the Basel II capital requirements is the value-at-risk, or VaR, which is defined in the following way. Let us denote by

P_t

the value of the portfolio at time t and by

P_{t+h}

the value at the end of the forecast period [t, t + h]. Then the loss distribution at time t for the forecast horizon h is defined by
Unknown element imgmain

Here we take the time value of money into account, i.e.

$B_{t,h}$

is a discounting factor such that

$B_{t,h}P_{t+h}$

is the value of

P_{t+h}

at time t.

VaR is defined as a threshold value such that the probability that loss on the portfolio over the given time horizon will exceed this value is the given probability level, 1 – α, i.e. value-at-risk is the negative of the α - quantile of the loss distribution
Unknown element imgmain

Note that in this notation losses correspond to negative values and profits correspond to positive values of

$X_{t,h}$

. Furthermore VaR is positive.

Regulatory Requirements

The Basel II framework defines several market risk measures based on the value-at-risk concept, which differ in:

•the set of risk factors to be simulated;

•the capital horizon, *h*;

•the confidence level, α.

The original Basel II framework from June 2004 required banks to capture by means of a VaR model the market risk due to changes of all pricing parameters that were deemed relevant. The confidence level for market risk VaR was set to 99% and the capital horizon was 10 days. In the case that a risk factor was incorporated in a pricing model but not in the VaR model, the bank had to justify the omission to the satisfaction of its supervisor.

As a consequence of the 2007–08 financial crisis, in July 2009 the Basel Committee revised the Basel II market risk framework and introduced several supplementary risk measures to capture market risks that were not properly taken into account by the market risk VaR: these are the so-called stressed VaR, the incremental risk charge, and the comprehensive risk measure.

The stressed VaR is intended to replicate a value-at-risk calculation that would be generated on the bank's current portfolio if the relevant market factors were experiencing a period of stress. It should be based on the 10-day, 99th percentile, one-tailed confidence interval value-at-risk measure of the current portfolio, with model inputs calibrated to historical data from a continuous 12-month period of significant financial stress relevant to the bank's portfolio (BCBS, 2009a). In this way, the stressed VaR addresses some shortcomings of traditional stress tests that do not assign probabilities to stress scenarios.

The incremental risk charge applies to all credit spread-sensitive financial instruments with the exception of securitizations and resecuritizations. It is defined as the 99.9% quantile of the loss distribution due to defaults and rating migrations over the capital horizon of one year. Thus, the incremental risk charge is a risk measure with credit risk and market risk components (BCBS, 2009b).

The comprehensive risk measure, which applies to the correlation trading portfolio only, combines the incremental risk charge with the market risk VaR. It captures not only incremental default and migration risks, but all price risks. The confidence level and capital horizon are the same as for the incremental risk charge (BCBS, 2009a).

In the following we describe modeling techniques for the market risk VaR.

Quality of Market Risk VaR Models

The quality of a market risk VaR model is usually investigated by looking at the VaR exceptions—i.e. days when portfolio losses exceed VaR estimates. The VaR exceptions should have the following two properties.

The number of VaR exceptions should be consistent with the confidence level. A VaR model with too few exceptions overestimates the risk. If the number of exceptions is too large, the risk will be underestimated. The VaR exceptions should be independent and uniformly distributed in time.

In order to decide whether the number of exceptions is reasonable or not—i.e. whether the market risk VaR model is correct—some statistical analysis is required.

The most widely used statistical test for the consistency of the number of VaR exceptions with a given confidence level is the proportion of failures (POF) test suggested by Kupiec (1995). Under the null hypothesis of Kupiec's test that the VaR model is correct, the number of exceptions follows the binomial distribution. Hence, the only information required to implement a POF test is the number of observations, the number of exceptions, and the confidence level, as shown in Table 1.

Table 1. 95% confidence regions for the POF test. (Source: Nieppola, 2009)

VaR confidence level	Nonrejection region for number, N, of VaR exceptions		
255 days	510 days	1,000 days	
99%	*N* < 7	1 < *N* < 11	4 < *N* < 17
97.5%	2 < *N* < 12	6 < *N* < 21	15 < *N* < 36
95%	6 < *N* < 21	16 < *N* < 36	37 < *N* < 65
92.5%	11 < *N* < 28	27 < *N* < 51	59 < *N* < 92

90%	16 < *N* < 36	38 < *N* < 65	81 < *N* < 120

Historical Simulation

Historical simulation is a nonparametric approach to the estimation of VaR that does not require any assumption about the probability distribution of changes of risk factors. Instead, the future values of risk factors are obtained by applying the observed historical changes of risk factors to their current values. This can be formalized in the following way. We assume that for each risk factor its current value and the value on M previous days is available (here 0 denotes the current day):
Unknown element imgmain

Then we can define M scenarios,

$$\tilde{X}_h^i, i = 1, \cdots, M$$

, for the value

$$X_h$$

of the risk factor after time h by applying the historical changes (returns),

$$R_h^i = X_{-(i-1)h} / X_{-ih}, i = 1, \cdots, M$$

, to the current value,

$$X_0$$

, of the risk factor:
Unknown element imgmain

The historical simulation consists of the following steps.

For each risk factor, convert the time series of its M + 1 values into a series of M daily returns.
For each risk factor, compute M possible future values (scenarios) by applying the daily historical returns to the current values of the risk factors.
For each of the M scenarios, reprice the portfolio and determine the corresponding profit-and-loss (the difference between the value of the portfolio under the new scenario and the current value of the portfolio). This results in M possible profit-and-loss values for the portfolio.
Order the profit-and-loss scenarios in ascending order and determine the required risk measure. For example, if the risk measure is the standard 99% VaR and M = 300, then VaR is the negative of the third smallest value; if M = 500, then VaR is the negative of the fifth smallest value. For M = 200, some interpolation between the second and the third smallest values will be required.

Sometimes, instead of historical returns, the VaR model employs historical differences,

$$X_{-(i-1)h} - X_{-ih}, i = 1, \cdots, M,$$

between the values of a risk factor. This leads to historical simulation with additive scenarios:
Unknown element imgmain

Historical simulation is probably the most widely used VaR method. Finger (2006) remarks that this popularity is due to the following views of many risk managers about historical simulation:

it is easy to explain;
it gives insight into what could go wrong;
it is “assumption-free”;
it is conservative;
it captures fat tails.

As the data for several banks about the number of VaR exceptions in 2008 show (see, for example, Lazaregue-Bazard (2010)), the last two views are clearly wrong. Banks that employed historical simulation with equally weighted input data reported up to 50 VaR exceptions in 2008. One of the reasons for the failure of these VaR models is that the historical simulation is not really assumption-free. For example, it implicitly assumes that the historical scenarios (historical returns in Equation (1) or historical differences in Equation (2)) are stationary, i.e. that their volatility does not vary with time.

Furthermore, VaR models based on historical simulation encounter different practical difficulties.

The time series used for the generation of historical scenarios must not have any gaps. If there are gaps, they must be filled or appropriate proxies must be used. However, the introduction of proxies leads to artificially high correlations that are not present in the historical data.
Historical simulation based on additive scenarios may result in negative values for the risk factors which are nonsensical for CDS spreads, equity prices, etc. Flooring such scenarios at zero changes the distribution of the historical scenarios and the correlation structure between the risk factors.

Historical Simulation with Volatility Adjustment (Filtered Historical Simulation)

The problem of time-varying volatility in the historical simulation can be dealt with by adapting the volatility of the historical returns. However, this requires some parameterization of the volatility evolution, resulting in a semiparametric model.

The basic idea is to remove the volatility from the observed sample in such a way that the resulting adjusted historical returns have time-independent volatility and then to scale the volatility of the adjusted returns to the current volatility of the risk factor.

If we denote by

σ_i

the volatility of the returns

$R_h^i = \frac{X_{-(i-1)h}}{X_{-ih}}$

at time t = –(I – 1)h, then the volatility-adjusted returns are defined by

R_h^i / σ_i

. The resulting scenarios are
Unknown element imgmain

The volatility adjustment solves the problem of nonstationary input data and improves significantly the performance of VaR models based on historical simulation. In fact, both of the bank VaR models that performed best in 2008 (Lehman Brothers and Goldman Sachs) were based on historical simulation with four years of input data and volatility updating (Lazaregue-Bazard, 2010).

Monte Carlo as an Extension of Historical Simulation Equation (1) for the scenarios of a risk factor can be written in the following equivalent way:

Equation (1) for the scenarios of a risk factor can be written in the following equivalent way:
Unknown element imgmain

The idea of Monte Carlo VaR simulation is to fit the log-returns

$\ln\left(\frac{X_{-(i-1)h}}{X_{-ih}}\right), i = 1, \cdots, M,$

to some parametric distribution and to replace the historical log-returns in the above equation by samples for the probability distribution. The choice of the probability distribution has a big impact on the properties and performance of the VaR model. Most banks model the log-returns of the risk factors by a multivariate Gaussian distribution, which has the disadvantage that scenarios larger than three times the standard deviation will be very unlikely, excluding from the simulation large, rare events which occur in the market more frequently than predicted by the Gaussian distribution. Therefore some banks use multivariate heavy-tailed distributions like the Student's t or the normal inverse Gaussian distribution.

The Monte Carlo simulation consists of the following steps.

For each risk factor, convert the time series of its M + 1 values into a series of M daily log-returns.
For each risk factor, fit the log-returns to the normal distribution

$$N(0,\sigma^2)$$

.
Generate for each risk factor N samples,

$$\xi_i, i = 1, \cdots, N,$$

, from the corresponding probability distribution

$$N(0,\sigma^2)$$

.
For each risk factor, compute N possible future values (scenarios) by

$$\tilde{X}_h^i = X_0 \exp^{\xi_i}$$

For each of the N scenarios, reprice the portfolio and determine the corresponding profit-and-loss (difference between the value of the portfolio under the new scenario and the current value of the portfolio). This results in N possible profit-and-loss values for the portfolio.
Order the profit-and-loss scenarios in ascending order and determine the required risk measure.

If the market value of the portfolio depends on more than one risk factor, the market risk model must capture the dependence of the risk factors, which is usually described by means of linear correlations. We assume that the generated scenarios should have the same correlation (covariance) structure as the input data.

Let us suppose that the value of the portfolio depends on K different risk factors, and let us denote by

$$X^k = (X^k_{-Mh}, X^k_{-(M-1)h}, \cdots, X^k_{-h}, X^k_0), k = 1, \cdots, K,$$

the historical values of the kth risk factor, by

$$R^k$$

the corresponding vector of historical log-returns, and by # the corresponding covariance matrix, i.e.:
Unknown element imgmain

If we generate for each risk factor N normally distributed random numbers
Unknown element imgmain

then the covariance of the random vectors

$$\xi^k = \left(\xi_1^k, \cdots \xi_N^k\right), k = 1, \cdots, K$$

, should coincide with #.

The generation of normally distributed random numbers with a given covariance matrix is usually based on the Cholesky decomposition of the covariance matrix.

Let

$$\zeta = (\zeta^1, \cdots \zeta^K)$$

be a vector of pairwise independent N(0, 1) normally distributed random numbers and let # be a given semipositive definite covariance matrix with Cholesky decomposition

$$\Sigma = QQ^T$$

. Then the random numbers

$$\xi = (\xi^1, \cdots, \xi^K)$$

defined by

$$\xi = Q\zeta$$

are normally distributed with covariance matrix #.

Monte Carlo with Linear Regression

The Monte Carlo approach presented in the previous section has two large drawbacks. If the number of risk factors is very large, the Cholesky decomposition of the covariance matrix may be very time-consuming. Moreover, the rank of the covariance matrix is min{K, M}. Thus, if the number of risk factors is larger than the length of the historical time series, i.e. K > M, the covariance matrix will have rank M and will be singular. Consequently, there can be at most M risk factors which can have pairwise linearly independent scenarios. The same is true for historical simulation: the number of linearly independent risk factors is limited by the length of the historical time series.

Both drawbacks can be resolved by regressing all risk factors against a smaller set of suitable risk factors (regression factors). In practice, risk factors used as regression factors are commonly used industry sector and region indices: equity indices, CDS spread indices, etc.

Let us denote by

$$R_h^i = \ln\left(\frac{X_{-(i-1)h}}{X_{-ih}}\right)$$

the log-return of a risk factor and by

$$R_h = (R_h^1, R_h^2, \cdots, R_h^M)$$

the corresponding time series of historical log-returns. Suppose that Rh is regressed linearly against a regression factor

$$I_h = (I_h^1, I_h^2, \cdots, I_h^M)$$

, i.e. that there exist a parameter β and an error term

$$\varepsilon \sim N(0, \sigma^2)$$

independent of

$$I_h$$

such that
Unknown element imgmain

This leads to a new Monte Carlo scheme with steps 3 and 4 of the sequence given in the previous section replaced by the following.

3a. For each regression factor

I_h

, generate N samples

$\xi_i, i = 1, \cdots, N,$

from the corresponding probability distribution

$N(0, \sigma^2(I_h))$

.

3b. For each risk factor, generate N samples

$\varepsilon_i, i = 1, \cdots, N$

, from the probability distribution

$N(0, \sigma^2(\varepsilon))$

of the corresponding error term.

4. For each risk factor, compute N possible future values (scenarios) by

$\tilde{X}_h^i = X_0 \exp^{\beta\xi_i + \varepsilon_i}$

Regressing the log-returns of risk factors against a smaller set of industry sector and region factors has the following two benefits.

Instead of generating correlated random numbers directly for all risk factors, we generate only correlated random numbers for the much smaller set of regression factors. Thus the covariance matrix will have a much smaller size.
Since the error terms of all risk factors are by definition pairwise-independent, the number of risk factors which receive linearly independent scenarios is limited only by the number of scenarios and not by the length of the time series.

General Monte Carlo

The Monte Carlo scheme presented in the previous section but one (as an extension of historical simulation) was formulated in a way which makes it clear that historical simulation is in fact a Monte Carlo method with scenarios generated in a particular way. However, both historical simulation and the Monte Carlo approach presented in the previous sections are applicable in practice only to short time horizons as they require historical market data for the past Mh days. For example, in order to estimate 10-day VaR with historical simulation based on 250 scenarios, the model would require historical market data for the past 2,500 days—i.e. 10 years.

Furthermore, both VaR approaches presented above have the following serious drawbacks.

They lead to implausible scenarios for risk factors over longer time horizons.
They are not applicable to all types of risk factors. For example, the values of implied correlations are limited to the (0, 1) interval, but the approaches described above cannot guarantee that the simulated values for the risk factors will be less than one.

Luckily, it is possible to easily get rid of these restrictions and to widen the applicability of Monte Carlo VaR methods. The starting-point is the observation that the simulation Equation (4) written as
Unknown element imgmain

where

$\xi \sim N(0,1)$

is the exact solution of the stochastic differential equation
Unknown element imgmain

Thus, the idea is to choose for each risk factor a stochastic differential equation
Unknown element imgmain

with appropriately chosen coefficients μ and σ, which describe its time evolution, and to replace Equation (4) for this risk factor by the exact solution of Equation (6) or by the Euler discretization of (6):
Unknown element imgmain

where

$\xi_i \sim N(0,1)$

.

The coefficients of Equation (6) can be determined from the time series of daily values of the risk factor, even in the case that the time horizon is longer than one day. Consequently, longer time horizons do not require longer historical time series.

Simulation of Base Correlations

Base correlations are of fundamental importance in pricing models for collateralized debt obligations (CDOs). As their values are limited to the (0, 1) interval and they exhibit strong mean-reverting behavior, their time evolution may be modeled by means of so-called Jacobi processes, which have the form:
Unknown element imgmain

Under some weak conditions, the

$\sqrt{X_t(1-X_t)}$

term in the stochastic differential equation prevents the base correlations from becoming negative or larger than one.

Monte Carlo simulation of risk factors based on special stochastic differential equations should be used with great care. The model may be misspecified due to a bad choice of the form of the stochastic differential equation and the assumption of a Gaussian Brownian motion, Wt, may be inappropriate in some applications.

The parameter estimation for a stochastic differential equation may pose considerable practical difficulties. A good example is Jacobi processes, where there is no closed-form expression for the log-likelihood function and one has to resort to approximate maximum-likelihood techniques. Moreover, the lack of long enough time series may prevent stable parameter estimation.

Another difficult point with general Monte Carlo methods is the modeling of the dependence structure between different risk factors. If the scenarios of each risk factor are driven by a Brownian motion Wt, the dependence structure between the risk factors is captured by the correlation structure between the Brownian motions of the risk factors. This implies a Gaussian copula, which contradicts the empirical finding that large market movements are usually more strongly correlated than “ordinary” ones. However, banks refrain from using more realistic dependence models due to the heavy estimation problems that are encountered in the case of copulas with tail dependence.

Scaling of VaR Between Different Risk Horizons

As we have seen before, the calculation of one-day market risk VaR by historical simulation with 250 scenarios would require one year of historical data, but the calculation of 10-day VaR—which is required by Basel II—would require 10 years of historical data, much more than is available in most banks. Therefore the banks usually prefer to calculate the one-day market risk VaR and then scale it up to the horizon of 10 days.

Under the assumption that log-returns of the portfolio are i.i.d. and that

$N(0,\sigma^2)$

is normally distributed, we have

$VaR_{h,\alpha} = \Phi^{-1}(\alpha)\sigma\sqrt{h}$

, where # denotes the standard normal distribution function. In particular, for

$\alpha = 0.99$

one has

$VaR_{h,0.99} = 2.32635\sigma\sqrt{h}$

.

Thus, for two different capital horizons

h_1

and

h_2

:

Unknown element imgmain

Consequently

Unknown element imgmain

Unfortunately, the portfolio log-returns are rarely i.i.d. and normally distributed, so applying the square root rule to obtain the 10-day market risk VaR may lead to a serious misstatement of the risk. For this reason, in July 2009 the Basel Committee included in its “Revisions to the Basel II market risk framework” the requirement that banks must demonstrate that the use of the square root of time to scale from a one-day holding period to a 10-day holding period is appropriate and does not underestimate risk (BCBS, 2009a).

See Provinzionatou, Markose, and Menkens (2005) and Degen and Embrechts (2011) for some theoretical and empirical studies of scaling rules for non-Gaussian distributions.

Which Method Is the Best?

Monte Carlo is the only VaR methodology that is applicable to long time horizons or to high confidence levels like 99.9%. The question as to which method for measuring one-day 99% VaR is the best—Monte Carlo or historical simulation—has been a topic of discussion for a long time, but opinion is still divided. This is somewhat surprising as the performance of banks’ market risk VaR models during the 2007–08 financial crisis shows a clear failure of VaR models based on historical simulation without volatility adjustment. The VaR models that performed best were, according to Lazaregue-Bazard (2010), models based on historical simulation with long observation horizons (four years) and volatility adjustment. This is not surprising as the filtered historical simulation seems to be very popular in academia due to the mathematical properties of this method. Unfortunately, most banks do not have sufficient quality-assured historical market data to be able to implement it.

The overview of VaR model performance presented by Lazaregue-Bazard (2010) contains only a limited amount of information about banks that employ Monte Carlo methods, but the few data that are available give an indication of performance between historical simulation with and without volatility adjustment.

The bad performance of VaR models based on historical simulation without volatility updating is due to some serious drawbacks of this approach.

The VaR estimates react very slowly to changes in the market regime, especially if longer observation periods (in excess of one year) are used.

The number of scenarios is limited by the length of the historical time series used in the simulation, leading to inaccurate VaR estimates.
The number of linearly independent risk factors is limited by the number of scenarios.
With short observation horizons (one to two years), rare events are not properly captured. For example, if the observation period contains events that happen once in N years, their frequency will be overestimated. If no such events are present in the history, they will be not be simulated at all.

All these issues can be solved to a large extent by choosing a longer observation period and updating the volatility of the scenarios (filtered historical simulation).

VaR models based on Monte Carlo have no limitations caused by the length of the historical time series and the number of Monte Carlo scenarios can be chosen to be arbitrarily large. Rare events can be captured by replacing the normal distribution by an appropriate probability distribution with heavy tails, such as Student's t or the normal inverse Gaussian.

The main difficulty of Monte Carlo-based VaR models is the modeling of the dependence structure of the risk factors. In the case of historical simulation, the scenarios have by definition the same dependence structure as the historical market data. This is an issue for models with short observation horizons as the historical data may capture only a very limited set of correlated market movements. VaR models based on Monte Carlo must make some assumption about the dependence structure of risk factors. Most such models describe the dependence by means of linear correlations (Gaussian copula). This assumption contradicts empirical findings as the market data for most risk factors exhibit tail dependence that cannot be captured by simple correlation models.

Summary and Further Steps

Two main approaches are currently used to estimate market risk VaR: historical simulation and Monte Carlo. Although VaR models based on simple historical simulation performed very badly during the recent financial crisis, they continue to be widely used by most banks. This shows that the properties of VaR models are not well understood by many market risk quants and market risk managers. This is partly due to the lack of appropriate literature describing advanced market risk modeling techniques. The books available on the market, such as McNeil, Frey, and Embrechts (2005), Jorion (2007), and Alexander (2008), are somewhat academic and are useful to people who want to learn about the main modeling techniques, but they mainly consider elementary methods that are rarely used in practice, or modeling techniques that are too difficult to calibrate. In particular, none of them is devoted to the new risk measures introduced in 2009 that combine credit and market risk—the incremental risk charge and the comprehensive risk measure. Therefore, for advanced modeling techniques applicable to those risk measures we refer readers to the very few available original technical papers, in particular Grundke (2005) and Wilkens, Brunac, and Chorniy (2011).

More Info

Books

Alexander, Carol. *Market Risk Analysis: Value-at-Risk Models v4*. Chichester, UK: Wiley, 2008.

Jorion, Philippe. *Value at Risk: The New Benchmark for Managing Financial Risk. 3rd ed.* New York: McGraw-Hill, 2007.

McNeil, Alexander J., Rüdiger Frey, and Paul Embrechts. *Quantitative Risk Management: Concepts, Techniques, and Tools.* Princeton, NJ: Princeton University Press, 2005.

Articles

Basel Committee on Banking Supervision. "International convergence of capital measurement and capital standards." Revised framework, comprehensive version. Bank for International Settlements, Basel, June 2006. Online at: www.bis.org/publ/bcbs128.pdf

Basel Committee on Banking Supervision. "Revisions to the Basel II market risk framework." Bank for International Settlements, Basel, July 2009a. Online at: www.bis.org/publ/bcbs158.pdf

Basel Committee on Banking Supervision. “Guidelines for computing capital for incremental risk in the trading book—Final version.” Bank for International Settlements, Basel, July 2009b. Online at: www.bis.org/publ/bcbs159.pdf

Degen, Matthias, and Paul Embrechts. “Scaling of high-quantile estimators.” *Journal of Applied Probability* 48:4 (2011): 968–983. Online at: www.math.ethz.ch/%7Edegen/DE_scaling_2009.pdf

Finger, C. “How historical simulation made me lazy.” *RiskMetrics Research Monthly* April 2006. Online at: tinyurl.com/74r4tab

Grundke, Peter. “Risk measurement with integrated market and credit portfolio models.” Journal of Risk 7:3 (Spring, 2005): 63–94.

Kupiec, Paul H. “Techniques for verifying the accuracy of risk measurement models.” *Journal of Derivatives* 3:2 (Winter, 1995): 73–84. Online at: dx.doi.org/10.3905/jod.1995.407942

Lazaregue-Bazard, Céline. “Exceptions to the rule.” Risk 23:1 (2010): 106–108. Online at: tinyurl.com/7skytmh

Nieppola, Olli. “Backtesting value-at-risk models.” Master’s thesis, Helsinki School of Economics, 2009. Online at: tinyurl.com/7tandnj

Provinzionatou, Vikentia, Sheri Markose, and Olaf Menkens. “Empirical scaling rules for value-at-risk (VaR).” Preprint, Department of Economics, University of Essex, April 2005. Online at: tinyurl.com/6n7umy3

Wilkens, Sascha, Jean-Baptiste Brunac, and Vladimir Chorniy. “IRC and CRM: Modelling approaches for new market risk measures.” SSRN (April 2011). Online at: dx.doi.org/10.2139/ssrn.1818042

See Also

Best Practice

- Measuring Country Risk
- Risk Management Revisited
- Measuring Company Exposure to Country Risk

To see this article on-line, please visit
http://www.qfinance.com/financial-risk-management-best-practice/modeling-market-risk?full