

McKinsey Working Papers on Risk, Number 32



Managing market risk: Today and tomorrow

Amit Mehta
Max Neukirchen
Sonja Pfetsch
Thomas Poppensieker

May 2012

© Copyright 2012 McKinsey & Company

Contents

Managing market risk: Today and tomorrow

Introduction	1
Modeling market risk	3
VAR-iations on a theme	3
Economic capital	6
<i>Current modeling practices</i>	7
Two complements to VAR	8
Implications for IT, the steering framework, and governance	9
Better risk aggregation	9
A simplified steering framework	11
Improved governance	11
Appendix: A note on the performance of historical simulations	14
Some advanced approaches	15
Comparing the techniques	16

McKinsey Working Papers on Risk presents McKinsey's best current thinking on risk and risk management. The papers represent a broad range of views, both sector-specific and cross-cutting, and are intended to encourage discussion internally and externally. Working papers may be republished through other internal or external channels. Please address correspondence to the managing editor, Rob McNish (rob_mcnish@mckinsey.com).

Introduction

Market risk refers to the risk of losses in the bank's trading book due to changes in equity prices, interest rates, credit spreads, foreign-exchange rates, commodity prices, and other indicators whose values are set in a public market. To manage market risk, banks deploy a number of highly sophisticated mathematical and statistical techniques. Chief among these is value-at-risk (VAR) analysis, which over the past 15 years has become established as the industry and regulatory standard in measuring market risk.

The demands placed on VAR and other similar techniques have grown tremendously, driven by new products such as correlation trading, multi-asset options, power-reverse dual currency swaps, swaps whose notional value amortizes unpredictably, and dozens of other such innovations. To keep up, the tools have evolved. For example, the number of risk factors required to price the trading book at a global institution has now grown to several thousand, and sometimes as many as 10,000. Valuation models have become increasingly complex. And most banks are now in the process of integrating new stress-testing analytics that can anticipate a broad spectrum of macroeconomic changes.

Despite these accomplishments, VAR and other risk models have continually come up short. The 1998 crisis at Long Term Capital Management demonstrated the limitations of risk modeling. In the violent market upheavals of 2007–08, many banks reported more than 30 days when losses exceeded VAR, a span in which 3 to 5 such days would be the norm. In 2011, just before the European sovereign crisis got under way, many banks' risk models treated eurozone government bonds as virtually risk free.

Indeed, the perceived limitations of VAR are bringing the industry under severe scrutiny. As one pundit puts it: "A decade rarely passes without a market event that some respected economist claims, with a straight face, to be a perfect storm, a 10-sigma event, or a catastrophe so fantastically improbable that it should not have been expected to occur in the entire history of the universe from the big bang onward."¹

In the wake of the most recent troubles, critics have noted VAR's reliance on normal market distributions and its fundamental assumption that positions can be readily liquidated. Regulators have attempted to compensate for some of these limitations, notably through Basel II.5, a comprehensive upgrade to the market-risk framework that took effect in December 2011. Some new elements in the framework, such as a requirement to calculate stressed VAR, are driving risk-weighted assets (RWAs) higher and boosting capital requirements by a factor of two to three.² (Basel III will bump the stakes even higher, notably through the introduction of the credit-valuation adjustment (CVA), which measures the market risk in OTC derivatives from counterparty credit spreads.)

The imposition of higher capital requirements may make the financial system safer, but from a modeling perspective this is a fairly blunt instrument. The ongoing refinements in stress testing are a welcome complement to the main work on VAR, but almost all banks would agree that risk models need more work. Banks are curious about the design choices entailed in simulation and valuation; they are probing for the right balance between sophistication and accuracy, on the one hand, and simplicity, transparency, and speed on the other. Having high-quality market data turns out to be just as critical as the models themselves, but many banks are uncertain about where to draw the line between acceptable and unacceptable levels of quality.

¹ William Poundstone, *Fortune's Formula: The Untold Story of the Scientific Betting System That Beat the Casinos and Wall Street*, New York: Hill and Wang, 2005, p. 296.

² For more on Basel II.5, Basel III, and the effects on capital, see McKinsey's Working Papers on Risk, especially "Day of reckoning? New regulation and its impact on capital markets businesses," Number 29, October 2011; "Mastering ICAAP: Achieving excellence in the new world of scarce capital," Number 27, May 2011; and "Basel III and European banking: Its impact, how banks might respond, and the challenges of implementation," Number 26, November 2010 (www.mckinsey.com).

And questions about market-risk governance have also arisen. Weak connections between the front office and the risk group are suspected by many to contribute to the often substantial differences between calculated and actual risk. Further, risk and finance teams often come up with conflicting versions of the bank's exposures, and find it difficult to explain why.

Charting a way forward

To help banks develop answers to these and other questions, McKinsey and Solum Financial Partners have conducted detailed interviews with practitioners at 13 large banks from Europe and North America, including 9 of the 15 biggest European banks (measured by assets), 2 of the top 10 North American banks, and 2 regional banks); additionally we have drawn on information supplied by 9 other large banks as part of their Pillar 3 requirements under Basel II.³

In this paper we will present the findings from this research as part of a broader discussion of current practices. The paper is divided into two main sections: modeling market risk and implications for risk IT and governance. We focus on some recent enhancements to market-risk modeling and on efforts that banks can take to make market-risk models more applicable. In addition, we propose a new analysis to untangle the pros and cons of different variations to the historical-simulation approach (see the appendix).

Banks should use the findings from this paper to challenge and validate their market-risk practices, and in so doing deepen their knowledge of the bank's risks. In an era when, in the course of a single day, banks routinely swap huge amounts of risk, and their trading activities lock up significant portions of their risk capital, even modest improvements in risk awareness are well worth the effort.

³ Solum has also conducted a similar survey on the credit-valuation adjustment (CVA); see "Counterparty credit risk: The impact of CVA, Basel III, funding and central clearing," www.solum-financial.com.

Modeling market risk

What do banks want from the market-risk management group? Primarily, they want to understand their market-risk profile, including both short-term profit-and-loss (P&L) volatilities and long-term economic risk. They want to know how much risk they have accumulated and how the total compares with the bank’s stated risk appetite. And they want the group to develop and win regulatory approval of a fair treatment of RWAs, allowing the bank to get maximum efficiency out of its capital.

These needs are supported by risk models. But while the requirements for market-risk modeling are quite consistent among banks, actual practices vary substantially. Below we highlight many of the differences in VAR usage; review current practices in the economic capital model; and finally, take a look at the ways banks use stressed VAR and stress testing to probe the fat tails that VAR may not address.

VAR-iations on a theme

Banks employ a cluster of tools to define and measure market risk and to allocate capital. VAR, of course, is at the center of the model. (See “Current modeling practices” on p. 7 for the basics on how banks use this tool and economic capital models.)

The essential choices in VAR design are the approach used to generate simulation scenarios (Monte Carlo versus historical simulation) and the valuation approach (full revaluation versus sensitivities). In the following sections, we explore the choices banks are making today, followed by a brief discussion of the growing importance of understanding VAR and the individual risks it comprises.

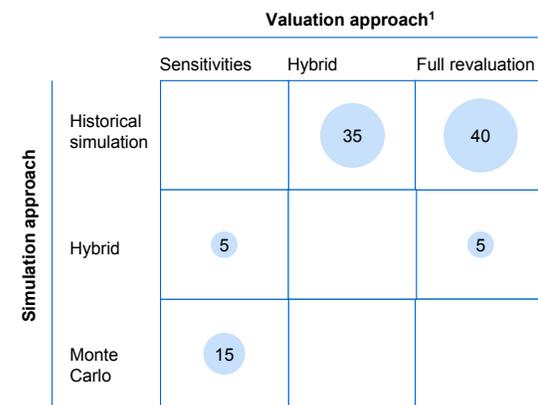
Monte Carlo versus historical simulation

The Monte Carlo method is widely considered the better theoretical approach to simulation of risk. Its chief advantage is that it provides a more comprehensive picture of potential risks embedded in the “tail” of the distribution. Moreover, it allows the bank to modify individual risk factors and correlation assumptions with some precision, making it a quite flexible approach. Proponents also argue for its greater consistency and synergies with other trading-book modeling approaches, such as the expected-potential-exposure (EPE) approach used for counterparty risk modeling.

But Monte Carlo, which typically requires about 10,000 simulations per risk factor, places a burden of complexity on the bank. Especially when used in combination with full revaluation, discussed later, the result is often a computational bottleneck that leads to much longer reaction times compared with the easier but less accurate historical simulation. In addition, many complain that it is a “black box,” which is not easily understood by either the businesses or management. As a result, only about 15 percent of banks surveyed use it as their main approach (Exhibit 1).⁴ (Some others use Monte Carlo techniques in limited circumstances, for example, on some complex portfolios or specific risks.) This is a shift; as recently

Exhibit 1 Banks use various approaches to simulation and valuation.

Market-risk practices at 18 financial institutions, 2011, %



¹ Banks are deemed to use the sensitivities approach if they use it exclusively, hybrid if they use it at least 30 percent of the time, and full revaluation if less than 30 percent. Source: McKinsey Market Risk Survey and Benchmarking 2011

⁴ All figures from the survey have been rounded, to avoid a suggestion of precision that the small sample size does not warrant.

as 2002, 20 percent of banks were using Monte Carlo, and that figure was rising.⁵ The tide now seems to be flowing the other way; three of nine banks in our sample either recently made the change away from Monte Carlo, or are about to.

They are joining the majority (75 percent) that use historical simulation (10 percent use a hybrid approach). To be sure, banks recognize the theoretical limitations of this approach: its projections are directly derived from the distribution of past occurrences and may be irrelevant or even unhelpful if the future is statistically different from the past. But because of its greater simplicity—it requires far fewer simulations, its calculations are transparent,⁶ and it requires no additional assumptions regarding market-factor distribution shapes—most banks now favor the historical approach.

Historical-simulation models differ primarily in the span of time they include and the relative weights they assign to the immediate and more distant past. These differences become critical in periods of stress. The longer the look-back period, the more conservative the model; but the model also becomes less reactive to recent events. Put another way, the shorter the look back, the more likely it is that the model will provide early-warning signals—though these more reactive models also create faster change in front-office risk limits, making compliance more challenging.

The survey found three types of historical simulation, based on the length of the look-back period and the use of the weighting of data points over time (Exhibit 2):

- One year, equal weighting (40 percent of banks using historical simulation).** Regulations require an effective one-year minimum for historical time series, in which the weighted average time lag of the individual observations cannot be less than six months; all banks that use the minimum also have to use the equal-weighting approach.⁷

Exhibit 2 Most banks use equal weighting and look back for one year.

VAR¹ historical-simulation practices at 18 financial institutions, %

		Year				
		1	2	3	4	5
Equal weighting		40	20	5	10	5
Time weighting			10	5		

¹ Value at risk.

Note: Numbers may not add up to 100 due to rounding.

Source: McKinsey Market Risk Survey and Benchmarking 2011

- Christophe Perignon and Daniel Smith, “The level and quality of value-at-risk disclosure by commercial banks,” *Journal of Banking & Finance*, February 2010.
- It should be noted that the transparency associated with historical simulation is not universally admired. Aaron Brown, the chief risk officer at hedge fund AQR Capital Management, says that value-at-risk (VAR) systems can “give results that are as hard to explain as market-price changes. In fact, their value is that they can surprise you.” He laments that at some institutions, “the pressure for rationality, control, auditability, and explainability has killed the value of VAR.”
- “Enhancements to the Basel II framework,” Basel Committee on Banking Supervision, July 2009.

- **Multiyear, equal weighting (45 percent).** Time horizons range between two and five years. This approach is more conservative and stable than the one-year approach but is also less reactive. While this conservatism would have been appropriate before the crisis, it might be less important now, as Basel II.5 has added stressed VAR, another conservative calculation.
- **Multiyear, with weighting (20 percent).** The lower weight assigned to the distant past makes it easier to stretch the look-back period beyond one year. The advantage is that recent data are generally agreed to be more relevant for near-term projections and count for more in this approach.

A few leading banks have also begun to experiment with “volatility scaling,” in which historical periods that exhibit volatility similar to the present are given greater weight in the model. Our analysis, detailed in the appendix, suggests that volatility scaling significantly increases the responsiveness of VAR calculations to changes in the loss distribution. However, that increased responsiveness comes at the cost of a less stable VAR if volatility scaling is used in its purest form.

Full revaluation versus sensitivities

VAR and other tools operate in one of two ways:

- **Full revaluation.** The positions whose risk is being measured are fully revalued, in all their complex detail, typically by using the front-office pricing models. (In some cases, such as in times of market stress, the full revaluation is difficult to run; many banks instead do an approximation for purposes of risk management and daily P&L.)
- **Sensitivities.** The positions are mathematically reduced to a few critical characteristics of their market behavior—the “Greeks” made familiar by option-valuation techniques (such as the “Delta-Gamma approach,” a Taylor expansion).

Because both approaches produce nearly identical results for simple products, most banks opt to use sensitivities where appropriate. For complex products, full revaluation provides a more accurate calculation of risk; despite the additional cost in time and effort, most believe that full revaluation is a must for these products.

To calculate VAR across the entire portfolio, approximations or sensitivities are often deployed; the computing power needed to reprice a typical trading portfolio is so enormous that banks sometimes cannot do it in a timely manner, particularly if Monte Carlo techniques are used on some parts of the portfolio. Across all banks, the survey found that average VAR run time ranges between 2 and 15 hours; in stressed environments, it can take much longer.

Explaining VAR’s results

Regulators are increasingly asking banks to produce a statement of their risk and returns in which the VAR perspective offered by the risk team agrees with the P&L perspective produced by the finance group. Misalignments in data and reporting hierarchies or book definitions can cause major differences between VAR and P&L, leading to a burdensome reconciliation process.

Accordingly, the “decomposition” of VAR into its component risk drivers is a high priority for many banks. Attribution of the variation with P&L to various risk factors (changes in rates, equity prices, and so on) can then be tested for correspondence with the bank’s ingoing hypothesis on risk sensitivities. A good decomposition can also better ensure that a given trader’s positions genuinely reflect declared strategies.

For banks that use a sensitivities-based approach, developing this breakdown is straightforward enough, but some higher-order risks are neglected. Broadly speaking, full revaluation will lend itself to a more thorough and thoughtful decomposition of VAR.

Banks also check VAR and their understanding of their risk positions through “back testing.” Regulators require back testing only “on the 99th percentile.” However, some leading banks undertake full back testing of the distributions that are lower and higher than the stated confidence level, in the belief that this better captures the performance of their VAR model.

Economic capital

Banks’ economic capital models for market risk are designed to capture the potential for more extreme events. Typically, economic capital models seek to identify the maximum loss within a one-year time horizon and assume a confidence interval of 99.95 to 99.97 percent. This will yield events so rare that they might happen only 1 year in every 2,000; ; a confidence level of 99.97 will yield events that happen only 1 year in every 3,300. By comparison, VAR at 99.0 percent captures 1-day-in-100 events.

Eighty-five percent of survey participants do not use a discrete economic capital model; instead, they use VAR and shareholder value at risk to calculate economic capital requirements for market risk. The remainder, 8 percent, uses an independent model, building mainly on stress testing.

There are three main design choices: (1) the confidence interval; (2) how to consider liquidity horizons; and (3) how best to account for fat tails.

1. The range of **confidence intervals** employed lies between 99.91 and 99.99 percent; banks with significant capital markets activity tend to use 99.98 percent (Exhibit 3). Historically, the choice of confidence interval was dependent on the bank’s risk appetite and on a specific target the bank had for its rating, but the survey suggests this relationship no longer holds.

Exhibit 3 Banks use a range of confidence intervals for economic capital models.

Economic capital-modeling practices at 17 financial institutions, %



¹ Some public information of additional regional banks has been included for comparison.
Note: Numbers may not add up to 100 due to rounding.

Source: McKinsey Market Risk Survey and Benchmarking 2011

2. **Liquidity horizons** establish the effective holding period by making assumptions on how easily assets can be liquidated. The most accurate method is to apply them at the level of individual risk factors. But this requires a detailed modeling of risk factors. In practice, many banks in our sample choose to simply scale up VAR and stressed VAR, using a single liquidity horizon across asset classes.
3. To capture severe periods of stress as required for economic capital, 14 percent of banks use particular **fat-tail assumptions** (that is, in addition to best-fit or other distribution assumptions), typically a predefined fat-tail add-on.

Current modeling practices

In the survey, banks provided a good deal of information about how they operate their models:

- Fifty percent of banks run at least two different iterations of VAR, one for regulatory reporting, with parameters set appropriately to regulatory guidelines, and others for internal management, with parameters adjusted for various modeling and testing needs. All calculate VAR on a one-day holding period.
- Sixty-two percent of banks use the calculated 1-day VAR to extrapolate the 10-day VAR required by regulators. They do this predominantly by multiplying by the square root of 10, either assuming normally distributed returns or a non-normal distribution. At the time of the survey, in late 2011, 40 percent of banks separately calculated a 10-day VAR for regulatory purposes.
- Banks that calculate VAR for internal use do so with shorter holding periods and sometimes even a lower confidence interval; 60 percent calculate internal VAR at the 99 percent confidence interval required for regulatory VAR.
- Some banks use relative returns exclusively; yet most banks take a hybrid approach, using an historical analysis to determine whether relative or absolute returns best fit a given asset class. For example, absolute returns are thought by many banks and regulators to be more appropriate for credit spreads.
- To determine the return at the required percentile of their chosen confidence interval, 75 percent of banks simply count down the values in the distribution until they reach the corresponding one (for example, in a two-year time horizon, the fifth data point would represent the return at the 99th percentile). But this may or may not represent the extent of the potential loss; in the tail, calculations are quite unstable. To cast a broader net and make a more conservative estimate, 25 percent of banks use averaging approaches or scale up from a lower quantile.
- Most banks calculate economic capital on a daily (30 percent) or weekly (40 percent) basis, actively using it for risk steering and definition of limits in accordance with the risk appetite. Ten percent calculate economic capital quarterly, mainly for regulatory purposes.

Two complements to VAR

Stressed VAR has long complemented VAR for use in market-risk management, and is now explicitly required by Basel II.5. The approach seeks to alter VAR measurement, focusing on high volatility periods to better capture tail or stress events. Some banks find it a useful halfway house between VAR and economic capital calculations. VAR looks only for 1-in-100 events, while economic capital identifies 1-in-2,000 events. Stressed VAR falls in between; it uses scenario analysis to look for rare events but remains a short-term measure (using a low quantile and a 1- to 10-day horizon).

However, stressed VAR has its complications. The degree of conservatism it introduces may well have unintended consequences, by discouraging banks from making the kinds of improvements to their models that would boost the accuracy of market-risk calculations but would also increase the overall market-risk RWA. And the high volatility of the data often requires robust capacity for full revaluation, which for many banks means work to improve their pricing models.

Banks are increasingly emphasizing a second technique as an additional complement to VAR: stress testing against coherent and internally consistent scenarios, systematically generated from historical stresses (for example, a repeat of the Asia crisis) as well as hypothetical emerging risks, such as a eurozone breakup, further downgrade of US sovereign debt, or currency devaluations. Moreover, banks now seek to coordinate stress testing of market risk with tests of credit and liquidity risk, particularly in regulatory stress tests conducted across the entire bank, and they test stresses on the balance sheet as well as P&L.

However, stress testing is no panacea: it is nonstatistical; it is limited by the imagination of scenario planners; and it is prone to recency bias, as yesterday's market events are invariably top of mind. Finally, some say it is nothing more than a form of proprietary bet, since the bank's returns will suffer if the scenario that unfolds differs from the one against which the bank has chosen to plan its risk management.

While banks largely use stress testing as a way of limiting losses or setting capital requirements, more nimble organizations on the buy side use them as dress rehearsals for crisis management. On both buy and sell sides, stress testing has become more prominent since 2008; both groups find that this different and differently hamstrung risk-management tool complements VAR.

Implications for IT, the steering framework, and governance

The theoretical approaches to modeling become concrete in the IT systems that support market risk, and in the output of these systems, which bank leaders use to steer the business. Any improvements in market risk modeling must be supported by corresponding changes in risk IT. Risk governance too is affected; for VAR models and other tools to be reliable, the ways they are managed and operated must be thoughtfully designed.

Risk IT and governance are subject to some of the same forces bedeviling risk modeling. The complexity of today's trading portfolios and pricing models presents obvious challenges to systems and infrastructure. An avalanche of new products and businesses during the run-up to the crisis has resulted in an exceptionally complex risk IT architecture at most banks, characterized by lots of inefficient and hard-to-maintain bespoke systems. Banks increasingly struggle with low-quality data, which throws sand in the gears of models such as VAR, EPE, economic capital, and stress-testing systems. And new rules on market risk put forward in reaction to the crisis (such as the securitization framework, stressed VAR, and new reporting requirements) have further strained organizations, not least by multiplying the number of models and systems that constitute the framework that leaders use to steer the bank.

A recent McKinsey research project, conducted jointly with the Institute of International Finance,⁸ along with the market-risk survey, identified three strategic thrusts for banks to manage these complexities:

1. Better alignment of risk calculations through stronger integration of the risk IT architecture and higher data quality
2. A simplified steering framework (risk and economic capital models and the limit systems they drive)
3. Improved risk governance, especially new product and trade-approval processes

Better alignment of risk calculations

Historically, many banks set up separate IT systems and data storage for each trading business. The risk group and the finance function also had their own models and data. Each setup was optimized to meet the specific needs of its owner: while the front-office teams prized high flexibility and finely calibrated pricing models to facilitate innovation in quickly changing markets, the finance function and risk group were focused on meeting regulatory, accounting, and internal standards. As business complexity increased, these separate systems agreed less and less often, and the data flows among them became more and more convoluted. At many banks today, aggregating and verifying market risk across the bank in real time has become a significant challenge.

Banks have been looking for ways to unite these disparate systems. As a first principle, they recognize that improvements to data design and storage will radiate throughout the business system and empower the critical task of risk aggregation. One step that most banks are considering is the creation of a common data model (a standard description of how the bank stores and works with data), to be used throughout the bank, wherever possible. If every business within a firm creates its data in the same way, all the downstream tasks become much simpler.

Banks are also weighing the benefits of consolidated data warehouses, another challenging but worthwhile aspiration. A single trusted source of data, sometimes called a "golden source," can vastly simplify the risk IT architecture and make application development and maintenance more straightforward. To ensure consistency and transparency of market

⁸ In 2011, McKinsey partnered with the Institute of International Finance to research current practices in risk IT. More than 40 banks participated in extensive surveys, interviews, and working-group sessions. The analysis and recommendations are detailed in "Risk IT and operations: Strengthening capabilities" (www.mckinsey.com and www.iif.com).

data among geographies and businesses, most banks surveyed use central data repositories for VAR computation. To get timely aggregation of risk, 50 percent use a system of individual cutoff times by trading or booking location (for example, end-of-day New York or London). Others use a single global cutoff time; these banks tend to have most of their business in one time zone.

Some banks are going another route, by creating a common data model and then implementing that model across several warehouses. For these banks, a single golden source is too difficult or costly to create, and so they focus on building a robust data “fabric” to support the needs of different users (for example, the business, risk, or finance groups) for reference data. Others are pursuing a “metadata” repository, a kind of master data layer that allows users to navigate reference data effectively.

Data quality

Of course, risk aggregation is undermined if banks’ data are unreliable. Banks buy external market data (price information, volatilities, and so on) to meet their risk systems’ needs. Where the data are unavailable banks usually approximate them in some way, either by simplifying the calculations in VAR and other tools or by identifying appropriate proxies and substitutes.

These gaps have become more numerous recently, in two ways. First, full revaluation models, increasingly used for complex products, require a higher number of risk factors than sensitivity simulation. Second, banks that use very long time series in their historical simulation—some use up to five years—often have gaps that must be filled. (Stressed VAR can require even longer time horizons, with similar challenges.)

To fill gaps in data (in cases of newly introduced products, market illiquidity, and so on), banks take two approaches: 60 percent simply “roll” data forward—that is, they replace the missing value with the preceding value. The other 40 percent now apply a more sophisticated approach, using algorithms to backfill missing data.

Issues of data quality are more important to banks than the sheer volume of information and the resulting computational challenge. In our survey, market data and the “brittleness” of risk models (the tendency to break down when low-quality data are used, as data values exceed their defined limits) were identified as core challenges by most banks. The burden of computation resulting from enormous volumes was deemed slightly less relevant, as at many banks infrastructure has already been expanded.

Banks are working to assure the quality of the data they buy and the data they generate internally. With respect to external data, 40 percent of banks surveyed source this directly from front-office platforms (that is, from external vendors such as Bloomberg and Reuters, or from exchanges). Fifty-eight percent have a separate market-data repository for use in the back office, typically managed by the risk group. A very few banks use teams outside the center to collect market data.

In all cases, leading banks have systematic, periodic, and largely automated validation processes to monitor data quality. Many banks are also trying to improve data quality and integrity at the source, by appropriately incentivizing front-office staff. One, for example, applies capital-usage penalties to business units that have too many missing data fields; it does this by applying artificially conservative values to missing fields. Some banks have established owners of every data domain, such as interest-rate market-risk data, ensuring a single point of accountability for all data-related issues.

Another measure that banks are starting to explore is the creation of a dedicated function to manage data quality. A dedicated team can manage routine data aggregation and also fulfill specific requests from regulators. Its tasks include managing market-data infrastructure, and the “golden source” where one exists, to improve the daily P&L/VAR

reconciliation process, and the embedded market-data verification process; delivering comprehensive aggregated exposure reports for all risk types for each counterparty; and producing regular reports on limits and exposures by subcategory.

A simplified steering framework

The complexity of the market-risk steering framework (including VAR and other risk models, the economic capital model, and the risk-limit system) has increased significantly in recent years. Risk models have proliferated and the limit systems have become more detailed, driven by increasing regulatory requirements, higher internal risk standards, and businesses' expanding needs. Managing the steering framework has become a challenging task, and its output for traders, finance professionals, and risk managers is too often contradictory or inconsistent.

Banks have tried to minimize these complexities in several ways. First, many are in the process of consolidating their VAR models, VAR-derived limit structures, and economic capital models, often by putting in place integrated models for use by the front office, risk, and finance. In addition, they have harmonized assumptions and parameters as much as possible.

In another step, banks are seeking greater alignment between VAR and models for other risk types (for example, CVA models for counterparty credit risk). In some cases, banks have built integrated credit-risk and market-risk models, typically by aligning the assumptions on which the model is based and agreeing on the output and use of the models.

Finally, as discussed in the previous section, banks and regulators alike are looking for close alignment between VAR models and P&L. Misalignments caused by different hierarchies or book definitions can cause major differences in P&L and often lead to a burdensome reconciliation process. Deviations need to be investigated in the market-risk report, particularly if they lead to limit breaches. Most banks surveyed said that their VAR calculation is aligned with P&L; that is, they use the P&L as the base case and seek to explain anomalies in VAR. Some banks have a special team for the task of alignment.

While these moves have accomplished much, more can be done. But simplifying the steering framework requires a high degree of coordination among different parts of the bank and a flexible IT architecture that can accommodate the integrated framework effectively.

Improved governance

Better risk aggregation and a simplified steering framework can only be achieved if they are properly embedded in the market-risk governance structure. We surveyed banks on the ways they assign roles and responsibilities for each of the important activities in market-risk management:

- **Setting the risk appetite and deriving the limit structure.** In most cases, the risk group drives the process of setting the risk appetite for the trading operation and breaking it out into limits for individual desks. This is a highly iterative process that seeks to align top-down management targets with bottom-up business initiatives. For the system to work, the parameters chosen to express the risk appetite should accurately reflect the risks' underlying sources of revenue; if not, problems arise. For instance, the omission of liquidity risk and some basis risks were exposed in the crisis as fundamental flaws in the limit structure. Hence, limits on VAR and shareholder value at risk (SVAR) are typically supplemented by limits on single risk-factor sensitivities and sometimes on notional amounts.

- **Limit monitoring and treatment of limit breaches.** Limit monitoring is a responsibility of the control functions and is typically entrusted to the risk group. Limit breaches are typically addressed directly or brought to joint decision-making bodies (for example, teams from risk and trading). Limit systems must be managed closely to investigate the root cause of limit breaches—which may be only technical violations caused by data errors, or legitimate violations caused by excessive positions—and to enforce prompt and sensible corrections, for instance, rebooking or data cleanup for technical breaches, and escalation processes to request limit adjustments, enforce risk reduction, or decide on sanctions for actual breaches.
- **Data management.** This is typically a back-office function and can be well managed by either the risk group or the front office, or even an associated operations unit. Ninety percent of the banks we surveyed use a centralized model to provide market data to businesses. About one-third of these are not purely centralized but use additional data from distributed “stores” (such as trading desks) to supplement the centrally available data.
- **Pricing models and valuations.** These are best owned by the front-office analytics team. When businesses develop their own pricing models, they tend to be stronger and more reflective of the business’s activities; moreover, they can often be calculated faster. We found that 90 percent of banks that use full revaluation use the same pricing models and IT architecture as the front office. However, the task of reconciliation and aggregation of a number of disparate pricing models becomes hopelessly complex as the inconsistencies between them mount. While the front office should own the activity, an independent pricing valuation should be performed by risk at least monthly; many banks do it daily. Either way, the risk group must sign off on the final valuation and must regularly align its risk factors with finance.⁹
- **Simulations, risk aggregation, and risk decomposition.** These are almost always owned by risk. Risk decomposition—the task of breaking down variations between VAR results and P&L by risk factor and asset class—should be done from the center, by finance or risk. In some cases, while most of the function is owned by risk, sensitivities can be calculated by the front office.
- **Risk reporting.** This is universally seen as a function best managed by the independent risk group. Nevertheless, close alignment with the front office is required, as desk heads and business leaders have the ultimate frontline responsibility for managing market risk. This alignment should focus less on a quantitative portfolio view (of “Greeks,” VAR, SVAR, stress results, and limit breaches) and more on an overarching view that includes a synthesized evaluation of market-risk outlook by business unit, a prioritized list of significant concerns, and the placement of

9 “International convergence of capital measurement and capital standards: A revised framework,” Basel Committee on Banking Supervision, June 2004.

current risk levels in appropriate competitive, historical, and macroeconomic contexts. This overarching view is essential to inform management discussions that seek alignment between portfolios and holistic risk appetite.

Over the last three to five years and in response to the credit crisis, many banks have strengthened the role of the risk function in the governance structure. Close collaboration of all stakeholders, clear accountabilities, and equal seats at the table remain the most important success factors in market-risk governance.

□□□

Banks have recognized the shortfalls in their management of market risk and are working hard to compensate for them. An understanding of the current state of the art, as outlined in this paper, can highlight their strengths, from which they can gain confidence, and their weaknesses, which may be areas that require further investment.

Appendix: A note on the performance of historical simulations

In the period leading up to the crisis, three of the best-performing banks in terms of number of outliers¹⁰ were JP Morgan, Goldman Sachs, and (believe it or not) Lehman Brothers. JP Morgan did not then (and still does not) use any form of weighting but made use of a short time horizon, while Goldman and Lehman used a significantly longer time period coupled with a time-weighting factor.¹¹ It appears that each of these models was reactive to a changing market environment.

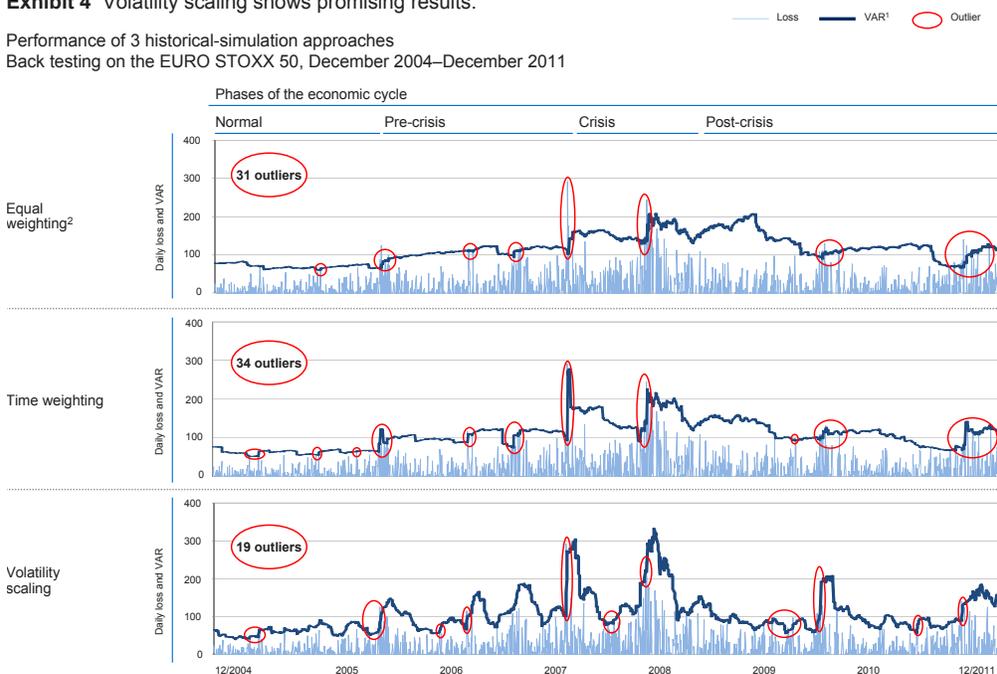
Obviously, however, a reactive VAR did not prevent Lehman's failure, suggesting once again that VAR has substantial scope for improvement. To better understand the true drivers of success in VAR modeling, we undertook an analysis of current approaches to historical simulation, which is the dominant mode of simulation. Seventy-two percent of banks use these techniques. Of these, 85 percent apply equal weighting to the chosen time span (of which about half use a one-year time span and half use a two- to five-year time span). Only 15 percent apply some sort of weighting.

Equal weighting often overstates VAR (and associated RWA requirements) because the longer the time horizon, the higher the VAR. It is a conservative and stable approach. On the other hand, it is not very reactive to jumps in volatility as markets enter periods of turbulence, and in periods of market stress, historical simulation will report a high number of outliers (Exhibit 4).

Time weighting can reduce the number of outliers as well as improve stability and accuracy, but at some cost, as these models can produce false signals and can produce rapid changes in risk limits that make compliance challenging.

Exhibit 4 Volatility scaling shows promising results.

Performance of 3 historical-simulation approaches
Back testing on the EURO STOXX 50, December 2004–December 2011



1 Value at risk.

2 One-year historical time horizon.

Source: Thomson Reuters; McKinsey Market Risk Survey and Benchmarking 2011

¹⁰ As we discuss later, outliers are times when daily profit and loss exceeds the previous day's value at risk.

¹¹ *Risk*, January 2010; annual reports.

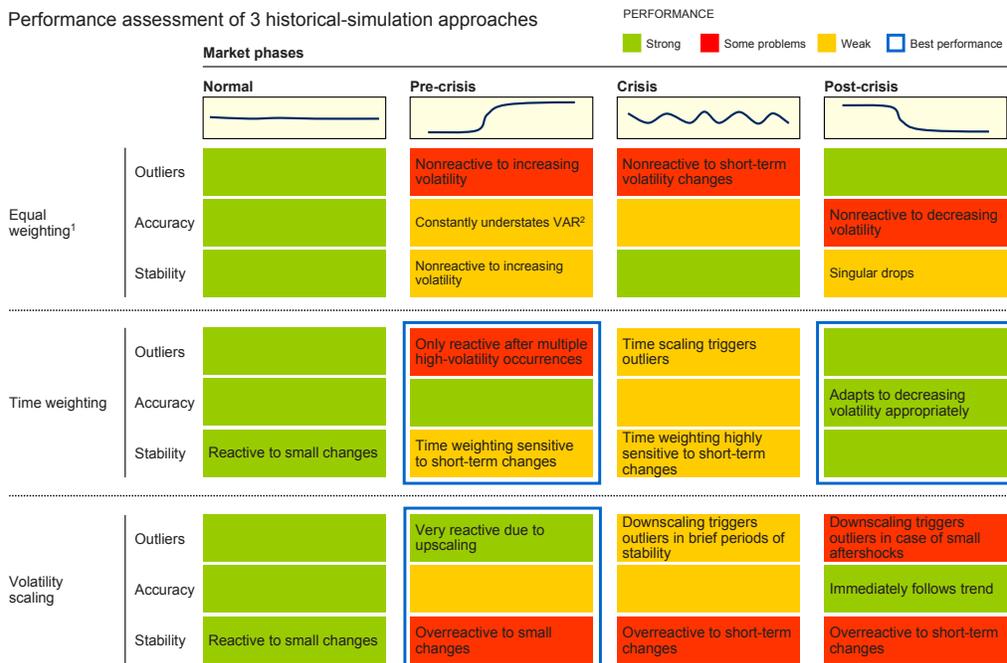
In light of these shortcomings, some banks are experimenting with alternative approaches in which the current VAR is adjusted to match the volatility of similar periods in the past. In the following, we compare the accuracy of three approaches being used by banks across a range of market conditions. The analysis has been conducted on about 20 risk drivers in a number of asset classes, but for the sake of clarity it has been based in part on some pragmatic assumptions and should therefore be considered only as a basis for discussion and further analysis.

Some advanced approaches

About 50 percent of the participants in our sample use a multiyear time horizon for their historical-simulation data set, exceeding the minimum regulatory requirement of using just one year. This extended horizon allows for the inclusion of more market events and thus increases the range of shocks included in the model, but it also reduces the accuracy and reactivity of the model, as very old data points continue to drive VAR. For the purpose of this analysis, each approach was modeled using a one-year time horizon.

The equal-weighting and time-weighting approaches are explained on pp. 4–5. In the new volatility-scaling approach,¹² returns are adjusted by a factor based on the ratio of the market volatility at the time of the observed data point to the current market volatility. The relevant data point within the time series is selected in the usual way, according to a percentile cutoff, but this return is normalized to current market volatility. For example, if the market was less volatile 100 days ago than it is now, the return from that day would be scaled up to its current market equivalent. Likewise, returns would be adjusted downward if current market conditions were less volatile than in the past. Because returns are scaled both up and down, we call this version of volatility scaling unlimited. Some banks use a limited variation of volatility scaling, in which returns are scaled up but never down.

Exhibit 5 Models' performance varies substantially across the cycle.



¹ One-year historical time horizon.
² Value at risk.
 Source: McKinsey Market Risk Survey and Benchmarking 2011

¹² Also known as GARCH model; a key difference is that our prototype uses some simplified parameters, for example implied volatilities.

Banks can also combine the approaches to fine-tune the responsiveness of the model. For example, combining limited volatility scaling and time weighting increases responsiveness and adds an element of conservatism in periods of decreasing volatility.

Comparing the techniques

We analyzed the performance of VAR models on their back-testing results as measured by accuracy, stability, and outliers, using four market conditions: normal, pre-crisis, crisis, and post-crisis. A desirable model will score high on accuracy and stability with a low (or expected) number of outliers in all phases of the economic cycle.

We define the evaluation criteria in this way:

- **Accuracy:**¹³ the absolute difference between VAR and actual daily P&L. An accurate model will be highly reactive, in the sense that it will rise and fall in a way that corresponds to daily fluctuations in the P&L. As a result, it will have high information content; management will be able to see changes in market conditions reflected quickly. Excess RWAs will be avoided, as VAR reduces rapidly when volatility declines.
- **Stability:**¹⁴ the change in VAR from day to day. A stable model will show minimal fluctuations in VAR and will not be prone to surprising leaps in VAR when risk positions change only slightly. A more stable model will usually be better understood and accepted by risk takers and management. A stable model will avoid sudden drops in VAR when data points fall out of the time series and will not be overly reactive to small, short-term changes in market conditions.
- **Outliers:** the number of times in a given period that the daily P&L exceeds the previous day's VAR. Regulators will alter the RWA multiplier applied to the VAR calculation based on the frequency with which these outliers occur, and so minimizing the number of outliers produced by a model can help to keep capital requirements lower. On the other hand, a very conservative model may produce few outliers but overstate the VAR, which would also be inefficient from a capital perspective.

We tested performance in four distinct market phases: normal (that is, average volatility), pre-crisis (moving from a phase of normal to high volatility), crisis (a period of sustained high volatility), and post-crisis (moving from a phase of high volatility back to normal).

Because VAR must accommodate all asset classes, we tested each technique on several risk drivers, including shares of listed companies on the EURO STOXX 50; the 10-year German Bund; and a currency pair, the euro/dollar exchange rate. Exhibit 4 illustrates the result for the EURO STOXX 50 sample; the results on other asset classes look much the same.

It is clear that applying some sort of weighting, whether volatility based or time based, significantly improves performance, at least by measures such as accuracy and number of outliers. Exhibit 5 shows the combined results of the tests of each asset class.

¹³ Defined as the absolute difference between one-day value at risk (VAR) and actual daily profit and loss, in relation to the long-term average VAR of the equal-weighting approach; this calculation is then averaged across the time period under consideration.

¹⁴ Defined as the absolute one-day change in VAR in relation to the long-term average VAR of the equal-weighting approach; this calculation is then averaged across the time period under consideration.

We see a few key findings:

- When **equal weighting** is applied, the model is highly stable but produces an unacceptably high number of outliers before and during a crisis scenario. We also note that accuracy levels are poor when entering the post-crisis phase; VAR remains high because it continues to be driven in part by the losses in the highest volatility period. Some portion of the weakness of this approach is due to the one-year time horizon used in the analysis; as noted, many banks use a multiyear horizon.
- When **time weighting** is introduced, we notice a marked increase in accuracy throughout the cycle due to the greater weight placed on the most recent market moves. The number of outliers, however, remains high as we move into the pre-crisis period, as many days of high volatility are required to trigger a reaction in VAR.
- When **volatility scaling** is applied, we see the accuracy of the model increase significantly as the model reacts immediately to changes in market volatility. This reduces the number of outliers in the pre-crisis phase but has the effect of reducing VAR very quickly, perhaps too quickly, post-crisis. Stability is compromised in most scenarios due to the quick reactions to changes in market conditions.

Assessing the techniques' effect on capital requirements

The choice of VAR methodology is critical to effective risk management—the objective function of all the techniques we assessed. As a side benefit, better VAR models can also lower the bank's capital requirements. These are defined by regulators as a function of the average VAR and stressed VAR over the past 60 days and the number of outliers during a 250-day period. Thus for banks that use some form of weighting, the choice of VAR model has a twofold impact (weighting and number of outliers) on this calculation. (Banks that do not use weighting can only optimize the multiplier. Use of weighting also allows the bank to optimize VAR; a combination of the two can often produce a lower capital requirement.)

A very conservative model will produce a high VAR and a low multiplier, while an aggressive model does the opposite. In making the modeling decision, the bank should try to find a methodology that occupies an ideal spot in the trade-off between VAR returns and the multiplier, in order to minimize the capital requirement.

Amit Mehta is an associate principal in McKinsey's New York office, where **Max Neukirchen** is a principal. **Sonja Pfetsch** is an associate principal in the Düsseldorf office. **Thomas Poppensieker** is a director in the Munich office.

The authors gratefully acknowledge the contributions of colleagues Diana Kapsa, Bary Pradelski, and Christian Raubenheimer, and the editorial support of Mark Staples.

Contact for distribution: Francine Martin
Phone: +1 (514) 939-6940
E-mail: Francine_Martin@mckinsey.com

McKinsey Working Papers on Risk

- 1. The risk revolution**
Kevin Buehler, Andrew Freeman, and Ron Hulme
- 2. Making risk management a value-added function in the boardroom**
Gunnar Pritsch and André Brodeur
- 3. Incorporating risk and flexibility in manufacturing footprint decisions**
Martin Pergler, Eric Lamarre, and Gregory Vainberg
- 4. Liquidity: Managing an undervalued resource in banking after the crisis of 2007–08**
Alberto Alvarez, Claudio Fabiani, Andrew Freeman, Matthias Hauser, Thomas Poppensieker, and Anthony Santomero
- 5. Turning risk management into a true competitive advantage: Lessons from the recent crisis**
Gunnar Pritsch, Andrew Freeman, and Uwe Stegemann
- 6. Probabilistic modeling as an exploratory decision-making tool**
Martin Pergler and Andrew Freeman
- 7. Option games: Filling the hole in the valuation toolkit for strategic investment**
Nelson Ferreira, Jayanti Kar, and Lenos Trigeorgis
- 8. Shaping strategy in a highly uncertain macroeconomic environment**
Natalie Davis, Stephan Görner, and Ezra Greenberg
- 9. Upgrading your risk assessment for uncertain times**
Martin Pergler and Eric Lamarre
- 10. Responding to the variable annuity crisis**
Dinesh Chopra, Onur Erzan, Guillaume de Gantes, Leo Grepin, and Chad Slawner
- 11. Best practices for estimating credit economic capital**
Tobias Baer, Venkata Krishna Kishore, and Akbar N. Sheriff
- 12. Bad banks: Finding the right exit from the financial crisis**
Luca Martini, Uwe Stegemann, Eckart Windhagen, Matthias Heuser, Sebastian Schneider, Thomas Poppensieker, Martin Fest, and Gabriel Brennan
- 13. Developing a post-crisis funding strategy for banks**
Arno Gerken, Matthias Heuser, and Thomas Kuhnt
- 14. The National Credit Bureau: A key enabler of financial infrastructure and lending in developing economies**
Tobias Baer, Massimo Carassinu, Andrea Del Miglio, Claudio Fabiani, and Edoardo Ginevra
- 15. Capital ratios and financial distress: Lessons from the crisis**
Kevin Buehler, Christopher Mazingo, and Hamid Samandari
- 16. Taking control of organizational risk culture**
Eric Lamarre, Cindy Levy, and James Twining
- 17. After black swans and red ink: How institutional investors can rethink risk management**
Leo Grepin, Jonathan Tétrault, and Greg Vainberg
- 18. A board perspective on enterprise risk management**
André Brodeur, Kevin Buehler, Michael Patsalos-Fox, and Martin Pergler
- 19. Variable annuities in Europe after the crisis: Blockbuster or niche product?**
Lukas Junker and Sirius Ramezani
- 20. Getting to grips with counterparty risk**
Nils Beier, Holger Harreis, Thomas Poppensieker, Dirk Sojka, and Mario Thaten

EDITORIAL BOARD

Rob McNish
Managing Editor
Director
Washington, DC
rob_mcnish@mckinsey.com

Lorenz Jüngling
Principal
Düsseldorf

Maria del Mar Martínez Márquez
Principal
Madrid

Martin Pergler
Senior Expert
Montréal

Sebastian Schneider
Principal
Munich

Andrew Sellgren
Principal
Washington, DC

Mark Staples
Senior Editor
New York

Dennis Swinford
Senior Editor
Seattle

McKinsey Working Papers on Risk

21. **Credit underwriting after the crisis**
Daniel Becker, Holger Harreis, Stefano E. Manzonetto, Marco Piccitto, and Michal Skalsky
22. **Top-down ERM: A pragmatic approach to manage risk from the C-suite**
André Brodeur and Martin Pergler
23. **Getting risk ownership right**
Arno Gerken, Nils Hoffmann, Andreas Kremer, Uwe Stegemann, and Gabriele Vigo
24. **The use of economic capital in performance management for banks: A perspective**
Tobias Baer, Amit Mehta, and Hamid Samandari
25. **Assessing and addressing the implications of new financial regulations for the US banking industry**
Del Anderson, Kevin Buehler, Rob Ceske, Benjamin Ellis, Hamid Samandari, and Greg Wilson
26. **Basel III and European banking: Its impact, how banks might respond, and the challenges of implementation**
Philipp Härle, Erik Lüders, Theo Papanides, Sonja Pfetsch, Thomas Poppensieker, and Uwe Stegemann
27. **Mastering ICAAP: Achieving excellence in the new world of scarce capital**
Sonja Pfetsch, Thomas Poppensieker, Sebastian Schneider, Diana Serova
28. **Strengthening risk management in the US public sector**
Stephan Braig, Biniam Gebre, Andrew Sellgren
29. **Day of reckoning? New regulation and its impact on capital markets businesses**
Markus Böhme, Daniele Chiarella, Philipp Härle, Max Neukirchen, Thomas Poppensieker, and Anke Raufuss
30. **New credit-risk models for the unbanked**
Tobias Baer, Tony Goland, and Robert Schiff
31. **Good riddance: Excellence in managing wind-down portfolios**
Sameer Aggarwal, Keiichi Aritomo, Gabriel Brenna, Joyce Clark, Frank Guse and Philipp Härle
32. **Managing market risk: Today and Tomorrow**
Amit Mehta, Max Neukirchen, Sonja Pfetsch, and Thomas Poppensieker

