

**STANDARD
& POOR'S**

Enterprise Risk Management For Financial Institutions

Rating Criteria And Best Practices

GOVERNANCE

Policies
Risk Tolerance
Conflict of Interests
Disclosure

STRATEGIC/BUSINESS

Competition
Product Innovation
Management Incentives

MARKET

Trading
ALM
Investing

CREDIT

Underwriting
Counterparty
Portfolio Management

OPERATIONAL

Compliance
Legal
Reputation
Audit
Fraud
Accounting
Fiduciary
Back Office
Technology
Personnel



STANDARD
&POOR'S

Enterprise Risk
Management For
Financial Institutions:
Rating Criteria And Best Practices

November 2005

Published by Standard & Poor's, a Division of The McGraw-Hill Companies, Inc. Executive offices: 1221 Avenue of the Americas, New York, NY 10020. Editorial offices: 55 Water Street, New York, NY 10041. Subscriber services: (1) 212-438-7280. Copyright 2005 by The McGraw-Hill Companies, Inc. Reproduction in whole or in part prohibited except by permission. All rights reserved. Information has been obtained by Standard & Poor's from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Standard & Poor's or others, Standard & Poor's does not guarantee the accuracy, adequacy, or completeness of any information and is not responsible for any errors or omissions or the result obtained from the use of such information. Ratings are statements of opinion, not statements of fact or recommendations to buy, hold, or sell any securities.

Standard & Poor's uses billing and contact data collected from subscribers for billing and order fulfillment purposes, and occasionally to inform subscribers about products or services from Standard & Poor's, our parent, The McGraw-Hill Companies, and reputable third parties that may be of interest to them. All subscriber billing and contact data collected is stored in a secure database in the U.S. and access is limited to authorized persons. If you would prefer not to have your information used as outlined in this notice, if you wish to review your information for accuracy, or for more information on our privacy practices, please call us at (1) 212-438-7280 or write us at: privacy@standardandpoors.com. For more information about The McGraw-Hill Companies Privacy Policy please visit www.mcgraw-hill.com/privacy.html.

Analytic services provided by Standard & Poor's Ratings Services ("Ratings Services") are the result of separate activities designed to preserve the independence and objectivity of ratings opinions. The credit ratings and observations contained herein are solely statements of opinion and not statements of fact or recommendations to purchase, hold, or sell any securities or make any other investment decisions. Accordingly, any user of the information contained herein should not rely on any credit rating or other opinion contained herein in making any investment decision. Ratings are based on information received by Ratings Services. Other divisions of Standard & Poor's may have information that is not available to Ratings Services. Standard & Poor's has established policies and procedures to maintain the confidentiality of non-public information received during the ratings process.

Ratings Services receives compensation for its ratings. Such compensation is normally paid either by the issuers of such securities or third parties participating in marketing the securities. While Standard & Poor's reserves the right to disseminate the rating, it receives no payment for doing so, except for subscriptions to its publications. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.

Permissions: To reprint, translate, or quote Standard & Poor's publications, contact: Client Services, 55 Water Street, New York, NY 10041; (1) 212-438-9823; or by e-mail to: research_request@standardandpoors.com.

Letter To Readers

Standard & Poor's Ratings Services is pleased to present *Enterprise Risk Management For Financial Institutions: Rating Criteria And Best Practices*.

This guide presents the latest ratings criteria for assessing the trading risk management practices of financial institutions, as well as a broad look at current best practices within financial institutions with respect to Enterprise Risk Management.

The industry faces greater challenges in assessing risks within this dynamic and evolving market structure. Regulatory and accounting practices are rapidly altering their code of requirements to "best practices." The principles of best practices of risk management are not viewed that differently around the world. There is, for example, a common language, as evidenced by Basel II (SOX, IAS, and more recently the SEC as well), which is essentially an attempt to codify some of those principles and link them into regulatory practices and capital assessment. The extent to which banks have adopted best practices in managing risk is a key benchmark in Standard & Poor's ERM initiative.

Our approach focuses on the effectiveness of policies, infrastructure, and methodologies (PIM) for controlling risks. Historically, most serious losses in trading operations could be traced to generally not one but a series of weaknesses in policies, infrastructure, and methodologies. Our aim, then, is to evaluate the robustness of the risk control practices.

In focusing on PIM, we are essentially focusing on the governance of the trading floor risk management function, its consistency with corporate strategies and risk appetite, and the way in which management assures itself that it has the methodologies and infrastructure in place to monitor and control the risks it undertakes.

To be effective and to add value, the risk assessment process must be approached from an enterprise-wide perspective because the risk assessment process is subject to the same evolutionary dynamics as the risk measuring and monitoring aspects. By adding discipline and analytical rigor to the process of assessing the risks of financial institutions, Standard & Poor's reaffirms its ongoing commitment to providing leadership and guidance to the market during times of significant change.



Jayan Dhru
*Managing Director & Head
of North American Financial Institutions Group*

Table of Contents

The PIM Approach To Assessing The TRM Practices Of Financial Institutions	7
Credit FAQ: Assessing Trading Risk Management Practices Of Financial Institutions	10
FI Criteria: Assessing Trading Risk Management Practices Of Financial Institutions	13
In Pursuit Of Best Practices For Enterprise Risk Management	20
Chasing Their Tails: Banks Look Beyond Value-At-Risk	25
U.S. Financial Institutions Rethink Trading Risk In Volatile Markets	32

The PIM Approach To Assessing The TRM Practices Of Financial Institutions

Standard & Poor's Ratings Services' approach to assessing the Trading Risk Management (TRM) practices of financial institutions focuses on the robustness of policies, infrastructure, and methodologies (an approach we call PIM) that govern the trading operations of financial institutions. At a holistic level, Standard & Poor's believes that a thorough assessment of the TRM practices within financial institutions must span multiple domains, including market risk, credit risk, operational risk, business risk, reputation risk, audit, and governance. We take our cue from the fact that historically, serious losses in trading operations have been traced to a series of weaknesses around issues related to policies, infrastructure, and methodologies.

We have developed a series of questions intended to delve into the policies, processes, and controls for identifying, measuring, and monitoring risks of the trading operations within financial institutions. Most of the issues are not amenable to quick, factual answers, so the focus of the questions is to elicit insights into the risk culture of the firm, and how that affects its choices on policies, processes, and controls.

The Policy Dimension of PIM

Within the policy dimensions of PIM, we review and assess several key attributes. First is the institutions' philosophy toward risk. In assessing this aspect, Standard & Poor's looks at the process by which institutions define their risk appetite and the dynamics by which this is tied in to the business strategy of the firm. There must be a level of consistency in the risk culture of the firm, its strategy, and its appetite for trading risk. A disjoint between them would signal a lack of awareness by senior management or a reach for profit by a badly controlled business unit. We are also interested to see whether risk management initiatives get championed at a senior level within the organization or are an outgrowth of a specific unit or trading arm. This attribute also examines the extent of senior management's involvement in setting the risk appetite for the firm and monitoring those risks.

Also of importance to Standard & Poor's is the organizational structure and authority of the risk management function. Organizational and reporting structures provide a clue as to where the authority to govern the risk measurement and monitoring process resides, and who is empowered with the ownership of the risks posed to the firm. The concept of a strong and independent risk management function, one that provides a check to balance the business considerations that drives trading executives, is not an easy one. It is not a simple matter of checking for independent reporting lines from the business unit. At times, a risk culture that engenders an environment in which business managers question each others' risk-taking decisions, as opposed to operating in silos with a narrow focus on their own risks, can also provide an effective independent check.

The second attribute of this dimension evaluates the institution's risk tolerance and senior management's awareness and control of the risks. Standard & Poor's examines the policies and the process by which trading limits are set for each risk category. We seek to understand the different risk measures

Prodyot Samanta
New York
(1) 212-438-2009

Tanya Azarchs
New York
(1) 212-438-7365

Jesus Martinez
Madrid
(34) 91-389-6941

against which limits are set and how limit breaches are handled and tracked. Who is the final arbiter for approving limits and who is authorized to approve excesses? Another important issue related to risk tolerance is the new product approval process and the limit-setting policies that govern it.

The extent to which senior management is aware of the risks facing the institution and the process by which they control the risks globally are also important. The practices in place for controlling the risks in remote offices are of particular interest. Do these offices have independent risk management and back-office functions?

The third key aspect of the policy dimension of PIM is the degree of risk communication and disclosure within the firm. In assessing this attribute, Standard & Poor's examines the quality and sophistication of disseminating risk information that is practiced by the institution. The key elements that the institution focuses on in its discussions with the board and senior management, and the nature of the questions asked of risk management, are indications of the quality of communication. Also important are the depth, clarity, and frequency of risk reports the institution uses in the normal course of business.

The Infrastructure Dimension of PIM

There are three key attributes of the infrastructure. The first focuses on the institution's risk architecture and data quality. The integration between the data warehouses and risk engines utilized in the firm's TRM function is important in assessing whether gaps could exist that might be exploited by rogue traders. The institution's awareness of the potential factors that could cause a system failure or disruption of business is also reviewed. A well-articulated disaster recovery process and business continuity plan in the event of a massive system failure or terrorist attack is critical to the institution. The source and integrity of the transaction and market data utilized by the institution in its TRM process is also important.

The second key aspect of this dimension of PIM relates to mid- and back-office operations. In reviewing this attribute, Standard & Poor's looks into the quality of the operational risk controls and processes that are in place within the trading room. We assess the extent to which mid- and back offices are capable of accurately capturing complex transactions and providing timely confirms. Back offices have been the source of trading losses in the past, when overloaded work forces allowed failures in trade processing to occur. This is also reviewed in the context of remote offices.

Auditors can also play an important role in detecting how well institutions have adhered to policies and procedures. Thus, we examine the audit reports and look into the institution's timely resolution of audit issues surrounding the risk management framework.

The third element of the infrastructure dimension reviewed by Standard & Poor's is the quality of risk education and training. Specifically, we examine the background and educational qualifications of the risk personnel and the number of years of experience within risk management. Training programs that are recommended or are in place for those involved with the trading risk management function, as well as the percentage of the trading risk management budget that is devoted to training programs, are looked into.

The Methodology Dimension of PIM

Standard & Poor's reviews three key aspects of the methodology dimension of PIM. The first relates to the risk-measurement tools employed by the institution. The issue is whether an institution can identify and quantify its trading risk exposures and whether it can incorporate the effects of diversification across risk categories and spot concentrations of correlated risks.

Value at Risk (VaR) is a widely used measure of risk for the trading book within financial institutions. In assessing the risk measurement tools used by financial institutions, Standard & Poor's reviews the approach the institution takes to calculating VaR. We look at how an institution uses VaR: its importance in limit-setting for the institution, its use in defining risk appetite, and for communicating risk within the firm. VaR is also generally not the only measure, nor even the principal risk measure used by firms, so we look at how other measures may be used.

Scenario analyses are a critical aspect of any risk management process, and Standard & Poor's reviews the financial institution's approach to constructing scenarios and stress tests. Of interest is the use of scenario analysis and its interpretation in conjunction with the other measures of risk reported by the firm.

Credit risk within the trading portfolio is an important consideration. Standard & Poor's reviews the approach that the financial institutions adopt for calculating the potential credit exposure profile of their derivatives portfolio: the process by which collateral and netting are captured and factored into the exposure calculations and the techniques by which risks are aggregated.

The second aspect of the methodology dimension addresses the process for model vetting and back testing. How does the institution insure against model risk? Standard & Poor's reviews the institution's approach to model approval and model vetting and the frequency with which models are reviewed. In assessing the model vetting aspects employed by the financial institution, Standard & Poor's reviews the process that is in place for establishing and verifying the inputs to the pricing models, the assumptions made, and the methods employed for verification. Standard & Poor's also examines the institutions' approach and methodologies established for back testing VaR.

The third aspect of the methodology dimension of PIM reviewed by Standard & Poor's is the financial institution's approach to capital attribution and management of the trading portfolio. Standard & Poor's reviews the methodology adopted by the financial institution in attributing capital to the trading operations. Is VaR of the trading book considered as an input for allocating capital by the institution? Also of importance in the assessment are the metrics used by the institution for assessing the performance of its trading operations.

Standard & Poor's PIM approach to assessing the TRM practices of financial institutions provides for an integrated view of risk, as it allows for the TRM practices within financial institutions to be assessed across multiple business lines, portfolios, and products. In addition, it allows for scalability within an institution. If an institution enters into a new business or geographic location, or if it adopts new methodologies or processes, the PIM approach can still be applied consistently.

Standard & Poor's envisions that the PIM approach to assessing the TRM practices at financial institutions will provide us with a deeper awareness of the ownership of and accountability for risks.

Credit FAQ: Assessing Trading Risk Management Practices Of Financial Institutions

By Michael Shari

Prodyot Samanta
New York
(1) 212-438-2009
Tanya Azarchs
New York
(1) 212-438-7365

In August 2004, Standard & Poor's Ratings Services expanded its review of the trading risk management (TRM) practices of global banks and other financial institutions that have significant trading operations as an integral part of the overall credit rating process for these issuers. (See "Credit Policy Update: S&P Expands Review of Trading Risk Management Within Financial Institutions," published on RatingsDirect Aug. 3, 2004.)

At a holistic level, Standard & Poor's believes that a thorough assessment of the TRM practices within financial institutions must span multiple domains, including market risk, credit risk, operational risk, business risk, reputation risk, audit, and governance. Historically, serious losses in trading operations have been traced to a series of weaknesses around issues related to policies, infrastructure, and methodologies.

Standard & Poor's has developed a series of questions intended to delve into the policies, processes, and controls for identifying, measuring, and monitoring risks of the trading operations within financial institutions. Since quick, factual answers are hard to come by in most cases, the questions are intended to elicit insights into the risk culture of a given firm and its role in choosing policies, processes, and controls.

Since October 2004, Standard & Poor's has used the PIM (policies, infrastructure, and methodology) approach to review the TRM practices of about 25 global financial institutions with annual trading revenues in excess of \$100 million. Standard & Poor's had expected to observe some institutions showing best practices across the board. Instead, Standard & Poor's has found pockets of good practices at some institutions, but no concentration of best practices at any single institution. Weak links in risk management tend to get masked during economic good times but can lead to downgrades in bad times, as some companies will be worse positioned than others because of their less robust risk management practices.

Frequently Asked Questions

What exactly is the PIM approach, and what is it intended to accomplish?

PIM stands for policies, infrastructure, and methodology. It's part of Standard & Poor's enhanced enterprise risk management (ERM) initiative. It's intended to provide insight into the quality of risk management practices within financial institutions.

How is PIM different from the previous approach to the analysis of risk management that Standard & Poor's has undertaken?

Assessing risk management from an enterprise perspective is something that Standard & Poor's has always done. This initiative is part of our culture to continually enhance our internal process, tighten it, make it more robust, put in more structure, and delve deeper. It's just good business practice. In the past, assessing the risk management practices was done at a more general or aggregate level.

Now we're getting more specific and going deeper into certain aspects so as to get a firmer understanding of the robustness of an institution's risk management practices.

Could you elaborate on how PIM enhances the surveillance process?

In assessing the TRM practices of financial institutions, we look at the three broad dimensions of policy, infrastructure, and methodology. Within each dimension, we have a number of attributes that are assessed.

What's in the policy dimension?

We look at the structure and stature of the risk function, the way the risk function sets the trading risk tolerance or risk appetite, the robustness of the structure for the approval of new transactions, the extent to which risk management is involved in setting and monitoring limits, and the extent of external and internal disclosure and awareness of risk.

What's in the infrastructure?

We look at the quality of the risk architecture that the firm has adopted and how that dovetails with its legacy systems. We also look at the reporting structure and quality of the back-office operations. We prefer the back office to report independently of the business. We also assess the degree to which transactions are automated, and the approach adopted by the firm for data recovery and business continuity planning in the event of a significant mishap or disaster.

What about methodology?

We evaluate the various risk measures that the institution utilizes in assessing its risks. For example, we look at the extent to which the firm uses measures other than value-at-risk (VAR) to assess its risk profile, the framework and quality of its stress testing, and the approach to disclosing risk measures other than that required by regulation to its shareholders. We also evaluate the extent to which valuation models exist for all transactions and how frequently these models are back-tested and vetted for changing market conditions. We look into the extent to which the risk function is involved in the process of validating models and setting reserves for model risk. If the models get validated and vetted by the front office or by the business where they get originated, then that's not really vetting. A good practice requires that the model validation group be independent of the business. That's where the risk function adds value to the firm.

How does Standard & Poor's use PIM?

Internally, an ERM Committee consisting of senior analysts uses the PIM approach to opine on the robustness of an institution's risk management function. The primary analyst for that institution then incorporates the ERM committee's assessment into the overall credit rating for that institution. There is no stand-alone rating for PIM.

If an institution falls short of expectations or exceeds them, would it warrant a downgrade or an upgrade?

Certainly. If we find that an institution is significantly weak in certain attributes, then it definitely would have an impact on the overall credit score, which implies that it could either lead to a downgrade or that an upgrade is less likely. On the other hand, if we find that a given institution turns out to be exceptionally strong in terms of aggregate risk management practices, then it clearly adds weight to the overall credit rating.

What is the range of practices that Standard & Poor's has observed thus far?

We started out with the expectation that some of the 25 institutions we looked at in North America, Europe, and Australia would have best practices. But we didn't find that. Instead, we found one or two institutions that come very close to having good practices across the board, but the majority of institutions have pockets of good practices.

We have observed a range of practices along each of the dimensions of PIM. In methodology, for example, we found that the way institutions measure VaR ranges widely in terms of the assumptions

used in the model to estimate VaR. Standard & Poor's believes that VaR is only one measure of aggregate risk, and indicates to the shareholder nothing about the tail risk that the institution is exposed to. It is important for the investor community to understand the limitations of VaR. (See "Chasing Their Tails: Banks Look Beyond VaR," published on RatingsDirect, July 12, 2005.) We emphasize the importance of comprehensive and well-designed stress tests and scenario analyses as a complement to VAR.

What are the worst practices that Standard & Poor's has observed?

We saw one large complex financial institution that had no trader limits and an extremely shoddy new product approval process. That's considered really bad. While this may not lead to an overall credit downgrade, given that the overall credit rating comprises a number of other factors, this will clearly weigh in as a very strong factor against recommending an upgrade. We also saw some institutions that were exposed to high model risk and a lax product approval process. This kind of scenario has the potential for causing possible problems and will get accounted for in the overall credit rating of the institution.

Is it Standard & Poor's goal to lead financial institutions to disclose more than is required by the SEC?

We wouldn't say it's a goal. But if an institution claims to be exercising best practices within risk management, then Standard & Poor's believes that it is important that they disclose more than is required by regulation, as that itself is good practice. Still, we're not penalizing them for not doing so. We are setting the bar for what best practices ought to be. We're saying, "If you do have best practices, we would expect that you look at risks from different angles, and report on that to your shareholders so that they are aware and confident that their management is taking care of their company."

What more should a financial institution disclose?

It would be nice to see companies that claim to have best practices talk about the framework and the analysis behind the design of some of their stress tests and scenario analyses, and provide additional reporting that might help shareholders understand where the risks lie. However, some firms may believe that the benefits of such disclosure may not be sufficient to warrant the additional computational costs, the drain on IT resources, and possible strategic disadvantages.

Is the PIM approach focused only on financial institutions, or is Standard & Poor's considering it more broadly?

PIM is a framework and has a wide applicability. It's a matrix that cuts across risk types and across industries and sectors. It can be used to assess the quality of the market risk function as well as the credit or operational risk functions. It can also be used across different industries and sectors, such as banking, asset management, insurance, utilities, or even airlines. PIM was developed to have a much broader applicability than just banking.

The articles, "Credit Policy Update: S&P Expands Review of Trading Risk Management Within Financial Institutions," published Aug. 3, 2004, and "Chasing Their Tails: Banks Look Beyond VaR," published July 12, 2005, can be found on RatingsDirect, Standard & Poor's Web-based credit analysis system. PDF files of these articles can also be found on Standard & Poor's Web site, www.standardandpoors.com. Click "Credit Ratings," then "News & Analysis" from the left navigation column, and scroll down to search for the articles.

FI Criteria: Assessing Trading Risk Management Practices Of Financial Institutions

The criteria for assessing the trading risk management (TRM) practices of financial institutions is outlined in detail in the following sections. These criteria will be applied to institutions where trading is a line of business. To assess trading risk management practices, Standard & Poor's Ratings Services deconstructs the analysis into three principal components: policies, infrastructure, and methodology (the PIM approach).

This framework—and the attributes that are employed to assess the risk management practices of financial institutions with trading operations as described later in this report—represents what Standard & Poor's believes to be “best practices,” which are not necessarily those practices that are widely applied in the industry.

The assessments based on these criteria form an integral part of Standard & Poor's overall credit ratings. For institutions where trading is a significant line of business, the application of these criteria will be of greater importance in the overall credit analysis.

Overview of the Analytical Framework

For the Policy dimension, Standard & Poor's analysts look at four key variables; stature of risk management, risk appetite, risk control process, and risk disclosure. In assessing the stature of the risk function, we evaluate the role and structure of risk management (RM) and the overall quality of the risk function. For the risk appetite, we assess the process by which the risk tolerance is established qualitatively and quantitatively and the robustness of the new product approval process. The risk control process opines on the established policies, the limit-setting process, and the limit-monitoring policies. For risk disclosure, we assess the quality of both internal disclosure and external risk disclosure.

For the Infrastructure component, Standard & Poor's assesses the quality of two primary attributes; risk architecture and back-office operations. For risk architecture, we look at the degree to which the risk systems are integrated, the data recovery process, and the quality of the business continuity planning strategy. In the case of back-office operations, we evaluate the structure of the back-office operations, the quality of the personnel employed, and the integrity of the data sources.

The Methodology component assesses the quality of the valuation techniques employed and the robustness of the model vetting process. The valuation techniques evaluate the process employed by the institution to determine the relevant pricing models for all transactions and the methodology adopted to assess counterparty credit risk in the trading book. Standard & Poor's also assesses the various risk metrics, including the value-at-risk (VaR) methodology, stress tests, sensitivity analysis, and other risk measures employed to assess and measure risk in the traded portfolio. The model vetting process evaluates the quality of the process for validating pricing models and the frequency with which they are reviewed. In addition, the process and frequency for back testing the VaR models as well as the pricing models are evaluated.

Prodyot Samanta
New York
(1) 212-438-2009

Tanya Azarchs
New York
(1) 212-438-7365

Jesus Martinez
Madrid
(34) 91-389-6941

Nick Hill
London
(44) 20-7176-7216

Policies

The stature of risk management

Role and structure of risk management. The risk management function within the institution must play a role that involves more than just monitoring and administration. The chief risk officer (CRO) should be an equal partner with the head of the business, interacting continuously with all units, discussing their risks, and ensuring that it is aligned with the stated tolerance, budget and, strategy of the institution. The risk management function should be proactive and constantly act in the interest of the institution's shareholders, carefully weighing the risk-reward trade-off and being a key decision making unit in the budget and planning process for the institution. There should be sufficient risk managers within the independent risk management group to ensure that constant dialogue and communication of the risks with the business and senior management takes place.

The authority to cut positions and/or halt the trading from a specific desk in the event of material risk-taking should reside with the CRO. In circumstances where there is a stalemate in the decision between the business and the risk management function, an escalation process to the most senior levels of management must exist.

Structurally, risk management must be a separate unit independent of the business, with the CRO preferably reporting to the CEO. This structure should be the one adopted by the smaller and remote offices as well. Standard & Poor's believes that a structure where the CRO reports to the CFO of the institution is indicative of a weaker risk function and stature. For the risk function to be effective and respected, the CRO must report independently of the business to the most senior management of the institution and be an integral part of the strategy and budget-making process. While this structure is necessary to insure independence, it is not a sufficient condition. More important are the intangible aspects that determine the respect earned by risk management.

Quality of the risk management function. Senior risk managers should hold advanced degrees that are relevant to the trading functions with significant business and trading experience. The modeling group or quantitative risk management unit within the risk management function must have advanced quantitative degrees with sufficient relevant business/trading experience.

Compensation of risk management function. At senior levels, the compensation of the risk function should be at least at par with comparable ranks in the business unit. While Standard & Poor's views the relatively high turnover of the risk management function at junior levels as a normal phenomenon for all institutions, a high turnover of the risk management function at senior levels is viewed as a weakness of the practice. However, Standard & Poor's does recognize that there would be some movement to and from Risk Management to the trading floor.

Risk tolerance

Establishing the risk appetite. In establishing the risk tolerance, the risk management function and the business must collectively consider appropriate business opportunities and the associated risks. The risk tolerance should be expressed at a holistic level in qualitative terms that clearly identify how it fits in with the tactical and strategic objectives of the institution. The qualitative expression of the risk appetite should be translated into day-to-day tractable quantitative measures as well. This level of risk should be expressed quantitatively in terms of how much the firm is "prepared to lose" from market, credit, or operational events over a one-year time frame. Taken together, the quantitative and qualitative expressions of the risk appetite should serve as a "moral compass" to guide day-to-day decisions.

In developing a qualitative expression of the risk appetite, the institution should consider the impact of the stated appetite on earnings, volatility of revenues, capital, and reputation. Market risk tolerances should be quantitatively expressed in terms of VaR limits and other measures such as expected shortfall, stress limits, stop-loss limits, and intra-day limits in times of excess market volatility. Illiquid products or positions should have additional limits that are clearly specified with adequate guidelines and clearly articulated disciplinary action that can ensue in the event of noncompliance.

New product approval process. An established new product approval (NPA) committee must exist that oversees the approval of new products in addition to new business initiatives. The NPA process should be clearly documented in the risk management policy handbook and should clearly articulate the steps required for approval.

There should be a clear definition of the types of transactions/business initiatives that are required to pass through the NPA process. For example, this may include products that have not been previously offered to the firm's counterparties or clients; variations on existing products; and existing products to be sold in a new market/region.

The CRO must be a member of the NPA committee and should be granted the authority to sign off on all transactions. The NPA committee should include senior individuals from all the decision-support areas, namely Operations, Legal, Tax, Audit, Accounting, Risk Systems, Risk Analytics, Compliance, and Documentation, as well as representatives of the business units. Each member must have veto power. Decisions to approve a transaction should consider all risks, not just the financial risks associated with the transaction alone.

No proposed transaction should be expected to trade without a thorough model validation and vetting by Risk Management and a clear understanding of the risk profile of the transaction and its implications on the overall risk appetite of the firm and its reputation. Should a proposed transaction require a new model or a modification/extension to an existing model, the product or trading manager should provide a proposed model to be reviewed by Risk Management. The proposed model must comply with the stated policies of the Model Validation Group and be detailed in providing the relevant theoretical basis, the proposed implementation, all systems requirements, booking, and any other simplifying assumptions.

There should be a specified time period, of no longer than six months after the transaction is initially approved for trading, after which it is reviewed for consistency. All transactions that go through the NPA process should be documented and tracked through secure databases and electronic media.

Risk control process

Risk management policies. In consultation with the business, all policies pertaining to risk appetite, risk control, model vetting, back-testing, risk infrastructure, and risk disclosure should be established and clearly documented by Group Risk Management indicating roles, responsibilities, and penalties for noncompliance and flagrant policy breaches. Standard & Poor's prefers that policies be reviewed periodically, but no longer than 12 to 15 months, to incorporate evolving market and business environments. All policies must be recommended by Group Risk Management and approved by the Board Risk Committee.

Limits setting. The limits framework for the firm should balance risk and reward in conjunction with the financial targets. In collaboration with the business, Risk Management should recommend a firm-wide limit for the aggregate risk appetite of the firm that would generally be approved by the Board. In addition, limits should be established for VaR, single transactions, illiquid transactions, and concentrations for asset classes, industries/sectors, and single names.

Risk Management must assign firm-wide limits and should have the authority to assign limits down to the divisions, businesses, and desks after sufficient dialogue with the business heads and other senior management. While Standard & Poor's recognizes that Risk Management would be micromanaging if it delegated limits down to the trader level, clear policies should exist that require the Business Unit management to assign and monitor limits for traders. Desk heads must assign limits to individual traders after adequate dialogue with Risk Management.

Limits should be established with appropriate triggers (as early warning signals) determined so as to ensure sufficient dialogue between the business units and Risk Management. Limits should be assigned at various levels within the entity down through to the trader or the lowest layer in the limit structure. Aggregate level limits should be assigned down to the first (or lowest) aggregate level node in the limit structure (in most cases the desk level), while sensitivity, stress, and stop-loss limits should be assigned down to the trader level as well.

A variety of limits, including, not only VaR limits but also Stress limits, Sensitivity limits, Concentration limits, and Stop-loss limits, should be employed to ensure that there is sufficient control of the risks at both the aggregate and disaggregate levels. Intra-day limits should be set as an additional line of defense to ensure against the above-mentioned limit types falling short of controlling the appropriate risks, especially during times of excess market turbulence and volatility.

Limits monitoring. Standard and Poor's prefers that there be a philosophy of zero tolerance toward ignoring limits. Traders breaching limits should be penalized through compensation adjustments or employment termination. Risk Management should be the final approver of limit exceptions after consideration of recommendations from the chief market/credit/operational risk officers. In approving limit exceptions, the market/credit/operational risk officers would have consulted with the business on the impact of the limit breach and the tenor and conditions of potential limit exceptions. Those exceptions should have been granted in all cases before the limit is exceeded, not retroactively. In the event of the firm-wide limit being breached, the CRO must immediately notify the CFO and the Risk Committee. All breaches and exception approvals must be documented and tracked by the Risk Committee.

Authority to grant limit exceptions must reside with Risk Management after sufficient dialogue with the Business Units. If a limit is breached, the market risk manager will discuss this with the traders involved, the Business Heads, and the CRO. The market risk managers should challenge the extent of market risk being taken and escalate to Group Risk Management as deemed necessary. Risk Management could then instruct the business to do one of three things:

- Reduce the risk profile back within the limit;
- Allow the excess to remain for an agreed-upon but short time after discussing the trading strategy with the trader; or
- Agree to revise the limit after careful consideration of the impact on the current risk appetite of the firm and the business strategy.

All limit breaches/exceptions should be communicated electronically to both Risk Management and the Business through daily reports describing exposures, limits, and exceptions.

The tenor of temporary limit exceptions should be no longer than 10 business days and in all cases must be approved by Risk Management after careful consideration and dialogue with the Business. Only in special cases, after sufficient dialogue with the Business, should Risk Management or Group Risk Management approve an extension of an existing temporary limit for an additional period of no longer than five business days.

In smaller/remote offices from which trading activity is conducted, there should be at least one senior Risk Management representation at all times to escalate issues or concerns as deemed necessary. Periodic visits to remote trading locations by senior Group Risk Management personnel should be conducted.

Risk disclosure

Quality of internal communications. Through weekly, monthly, and quarterly meetings with Risk Management and the Business, senior management should be made aware of their institution's exposure to both financial and nonfinancial risks. It is the responsibility of Risk Management to ensure that senior management at all times be in a position to clearly articulate to their shareholders the impact of adverse market and economic conditions on their institution's earnings and capital.

Risk Management should articulate to senior management all risks through clear and high quality internal reporting. Standard and Poor's prefers that reports contain qualitative and quantitative descriptions of the limit usage, risk concentrations and exceptions, and, where appropriate, the impact on earnings and capital. Reports to senior management should also include a concise description of the non-financial risks to which the firm is exposed. Minutes of the Risk Committee meetings with senior management should be documented and clearly highlight steps agreed upon to address any key issues or concerns discussed and outstanding.

Internal audit of the risk function should be conducted on a periodic basis, and issues deemed critical must be addressed immediately. Audit reports must be concise and should identify all issues, with time lines and agreed-upon resolution of critical issues clearly stated.

Quality of external disclosure. In the interest of their shareholders and as a discipline toward best risk management practices, institutions must provide external risk disclosure that goes beyond the minimum regulatory requirements. Institutions should provide in their external disclosure a clear qualitative and quantitative articulation of the firm-wide risk appetite and limit usage and indicate the potential impact of “tail risk” on earnings, capital, and return on equity, as well as the concentration of financial and non-financial risk exposures. This would also include expressions of how the institution has hedged against adverse market and economic environments.

Reporting of quantitative measures of risk should include VaR, expected shortfall, volatility of trading revenues, key stress tests versus limits, and qualitative expressions of the extent of “tail risk” the institution is exposed to.

Risk Infrastructure

Risk technology

The use of technology should be well dovetailed with the risk management process of institutions. Institutions must have a well-designed schematic of the technical risk architecture. This would include a clear depiction of databases, valuation engines, and reporting layers with the interconnections among these three broad tiers. Institutions should also move toward having a significant portion of transactions that have straight-through processing (STP).

Institutions should have a high quality data recovery processes (DRP). These must be tested and maintained frequently. A clear business continuity plan (BCP) must be designed and in place. The BCP must be discussed and communicated to all key individuals and fully tested at least twice a year.

Back-office operations

Standard & Poor's prefers that all back-office (B/O) operations report independently of the business with the roles and responsibilities of each function with regard to a transaction from its inception to settlement and through to its maturity, clearly articulated and documented by the Risk Management function of the institution. The quality of the B/O operations should be tracked and reported to Risk Management through the accurate capture of transactions, the integrity of the data, the production of timely confirms, and the production of accurate risk reports and profit and loss statements.

Risk Management should periodically identify, track, and report key risk indications for the B/O operations such as staff adequacy, employee satisfaction/turnover, failed settlements, trading volumes, inaccurate documentation, fraud, and outstanding reconciliations. Risk Management should have robust thresholds established for each risk indicator and assess the potential loss in the event that a breach in the threshold affects the institution's capital, earnings, and reputation.

In addition, there should be clear policy regarding security access to pricing models and the transaction processes to prevent tampering. Historically, such tampering has exacerbated most of the severe losses. Standard & Poor's prefers, therefore, that the B/O operations be under the control of Risk Management or Operations and outside of the Business Unit.

Methodology

Valuation techniques

Pricing models. Pricing models must exist for all transactions. The Model Validation/Quantitative Research Group within the Risk Management function should review all pricing models periodically with a gap of no longer than six to eight months for models associated with structured and complex products. All pricing models should have a sound theoretical financial-economics framework with all underlying assumptions clearly documented. Clear calibration routines, with well documented mathematical and software implementation specifications should exist for all pricing models. Documentation should state the permissible range of parameter values and the model's sensitivity to unobservable parameters, as well as the model's limitations. There should be appropriate reserves for model risk. Documentation must include the date when the model was last reviewed and signed off on by Risk Management.

Volatility for all products with optionality should be calculated using implied or stochastic volatility measures. Risk Management should specify the underlying assumptions and accuracy of the methodology used for estimating volatility skews, especially for FX products with significant non-linearity or optionality.

A potential future exposure (PFE) measure should be used in calculating counterparty credit exposure for derivatives in the trading book. It is preferable that the PFE measure be simulated under stochastic assumptions rather than an add-on measure based on residual tenor and historical volatility of the specific transaction. In simulating the PFE, careful consideration should be given to the assumptions made for the probability distribution of the underlying drivers of the credit exposure and the associated correlations/volatilities. If credit derivatives are traded (including CDO/CDO-squared transactions), either to hedge the banking book or to gain exposure, Risk Management should clearly specify how that gets integrated with the credit exposure calculations in the trading book.

Risk metrics

VaR. Standard & Poor's considers VaR calculations to be good discipline for robust risk management practices and believes that the "spillover" effects of the risk systems required to run these models are beneficial in many respects. However, given the wide range of possible assumptions made by institutions in the calculation of historical VaR, Standard & Poor's favors other risk measures (such as those specified below) to complement VaR.

VaR methodologies other than parametric (or variance-covariance) VaR must be considered as the appropriate technique for estimating VaR. If institutions calculated a historically simulated VaR or a Monte Carlo VaR, the assumptions about the risk factor correlations/volatilities and all model parameter estimates should be clearly specified. While recognizing that a full revaluation of the trading book for a VaR calculation can be time consuming and systems intensive, Standard & Poor's prefers institutions to perform a full revaluation of their trading book in calculating a daily VaR, especially if the book has significant optionality.

Daily VaR calculation should be run at varying levels of granularity within the trading book and aggregated across all trading locations globally. Globally aggregated VaR reports must be generated and analyzed by Risk Management no later than T+1.

Stress testing. As a criteria for robust risk management practices, extensive stress testing should be conducted by institutions to better understand their risk profile under abnormal market conditions/events and where VaR analysis provides little or limited information of “tail risk.” Traders/Desk Heads and Risk Managers should discuss and clearly identify the vulnerabilities of the portfolio. Construction of stress scenarios must then be undertaken in conjunction with macroeconomic analysis, historical events, and hypothetical possibilities. Correlations across risk factors should be carefully analyzed and accounted for in the scenario construction to assess implied “knock-on” effects. Stress tests should be run at varying levels of granularity depending on the concentrations and vulnerability of the portfolio.

Stress tests should be run on a daily basis with the resulting impact on firm-wide risk appetite, capital, and earnings reported and analyzed by Risk Management. Scenarios must be revised every six to nine months to incorporate evolving and changing market conditions, especially during periods of prolonged uncertainty and high market volatility.

Standard & Poor’s recognizes the fact that the ability for an institution to conduct “what-if” or “one-off” scenario analysis on an ad-hoc basis for the entire trading portfolio or for a specific book or transaction may be limited by the extent to which the institution’s risk systems are integrated. However, Standard & Poor’s views this as an important attribute for robust risk management practices and looks for institutions to be in a position to execute “what-if” or “one-off” scenario analysis on an ad-hoc or as-needed basis.

Sensitivity analysis. In addition to stress testing, sensitivity analysis should be conducted by institutions to assess portfolio impact and concentrations of changes to specific risk factors. Risk managers within the line of business should ensure that such measures are reported to Group Risk Management on a daily basis.

Other risk measures. In the interest of shareholders and robust risk management practices, Group Risk Management must go beyond regulatory compliance requirements and calculate other risk measures such as conditional expected shortfall to assess “tail risk” in the traded portfolio.

Model vetting process

Validating pricing models. The Model Validation/Quantitative Analysis Group within the Group Risk Management function must be responsible for vetting all pricing models independently of the Business. It is the responsibility of Risk Management to clearly specify the mathematical logic, theoretical assumptions, and parameter estimates underlying each of the valuation models. Risk Management must document all appropriate data sources for the inputs to the models and the relevant procedures for estimating model parameters. Sufficient validation and stressing of the models should be undertaken by the analytics group to ensure the accuracy of the models and conditions under which the models return invalid/nonstationary results.

Back testing. Institutions must back test valuation models and the VaR estimates. The back testing of VaR should be done using the synthetic/hypothetical or “clean” P/L. Back testing of VaR should be conducted at varying levels of granularity.

All pricing models should be back tested on an ongoing basis for predictive accuracy, and the results from this back testing should be documented and reported on a periodic basis. Comparison of the model’s output to other relevant model results and/or to market values should be performed and reported by Risk Management.

In Pursuit Of Best Practices For Enterprise Risk Management

Tanya Azarchs
New York
(1) 212-438-7365
Prodyot Samanta
New York
(1) 212-438-2009

During in-depth discussions with leading risk managers, Standard & Poor's Ratings Services noted differences of opinion on several fundamental precepts of enterprise risk management in general and risk management of trading operations in particular. In particular, the following three areas showed a range of opinion:

- The role modern portfolio management measurements methods should play in risk management practices;
- What "independence" of risk managers should actually mean; and
- The extent of the power of the risk management function.

Standard & Poor's has drawn some conclusions from these discussions about best practices.

Measurements

The field of risk management has evolved in several significant ways during the past 20 years. The most visible—and perhaps most seductive—change has been in risk measurement. High-powered computing has made possible sophisticated modeling of market risks, based on the mathematics of portfolio theory and statistics known as Value-at-Risk (VaR) models. That modeling capability is now being grafted onto the areas of credit and operational risk measurement.

The seeming simplicity and precision of the answers one can get using these measurement techniques is very attractive to managers and regulators. Meanwhile, some risk managers, even at major financial firms, remain unconvinced of VaR's benefits. These managers maintain models merely to placate regulators, and they may not invest in updating and enhancing them. Developing a robust model or even implementing an off-the-shelf model entails careful thought, massive amounts of data feeds, continual testing, and constant refinement. Those who do not invest effort and money do not get a robust model, or even a useful one for risk diagnostic purposes. Their models won't produce results that are comparable to other models using seemingly similar techniques.

Whether or not they find value in VaR models, firms rely more on older sensitivity measures—DV01, nationals, gamma, and vega—for day-to-day risk management of individual trader positions. The belief is that such measures are more sensitive in picking up the risks of specific instruments than are the blunter measures of VaR. Firms also maintain very sophisticated models for the purpose of pricing the instruments they sell, because an incorrect price opens the firm to the possibility of being picked off by other dealers or customers and getting the profit-and-loss statement wrong.

Nevertheless, the VaR models appear to be a favorite tool of risk managers who see their main function as ensuring that the firm does not find itself with concentrated exposures to any risks. In contrast to the more granular types of risk limits traders have—notional amounts, sensitivity measures, and Greek measures, for example—VaR alone provides a common language of risk across all asset classes and remains more useful for analyzing aggregate exposures for complex portfolios. In addition, VaR helps a firm understand its exposure to certain scenarios or stress tests.

Most firms perform some sort of scenario or stress tests. The differences lie in how the stresses are developed—whether they represent only some historical worst cases or some hypothetical ones tailored to expose the firm’s special vulnerabilities as well—and how sophisticated the models are in capturing the correlated effects of a shock to a specific market. For some firms, the stressed VaR becomes the firm-wide basis for limit-setting and capital allocation. That is good if the stressed scenarios were thoughtfully elaborated. If not, however, the stressed VaR becomes merely a very high limit that will never be breached and therefore will never trigger hard discussions about risk exposure.

Most seasoned risk managers understand that VaR models are merely to be used as diagnostic tools; they do not provide precise or scientific predictions of worst-case losses for the firm. Assumption-laden and data-dependent, VaR is best suited to depicting the recent past and not a future that could always suffer a paradigm shift or temporary discontinuity. Stress tests are more realistic than are daily VaR calculations for imagining a worst case and for setting capital. Daily VaR is more useful as a way of managing day-to-day trading exposures under normal market exposures.

We place emphasis on having robust models not because of their ability to produce precise measures of risk, but because of the systems requirements to operate them. These systems are very useful in providing a comprehensive aggregated view of risk positions in a framework that makes it possible to analyze the risk positions and their correlations. VaR models also are the only way in which correlations can be viewed across multiple instruments.

Risk Function Independence

Given that risk measurement is not the be-all and end-all of managing risk, we place a great deal more emphasis on risk governance. Risk management is about the policies, the built-in incentives, and the communication of risk. It is about how a firm defines and enforces its risk culture.

It would seem obvious that the risk management function, whether it pertains to credit risk or market risk management, should be independent. But what does independence really mean? In terms of organizational structure, risk managers and back offices should not report to the front office—to traders or heads of markets. On that there is little disagreement; however, beyond that there is little consistency.

Market risk managers often report to the CFO or CEO of the investment banking or the capital markets division of a universal bank. We believe this does not represent sufficient separation of the risk function from those who are charged with revenue production. In the purist model, we see market risk managers with reporting lines that lead to the credit risk officer, who in turn reports to the CEO of the entire organization. This kind of centralized model seems preferable, but it is insufficient in itself to ensure the effectiveness of the risk management function. The danger of this model is that it could result in too much independence. Separating risk management too much from the business functions could result in too little interaction with and understanding of the trading business to facilitate truly informed decision making.

More important than the reporting lines to insuring independence is the quality of the people in the risk management function. They must have a stature within the organization that warrants the respect of the traders. That is a matter of personality, education, and experience, not of official designation. If risk managers understand the traders’ businesses, can engage in meaningful, constructive dialogue with the traders, and can make forceful and sensible arguments, they can gain the respect of the traders. Otherwise, traders will always prevail, no matter how strict the enforcement approach proves to be.

Reporting Structure And Power

What emerges, then, is the inherent tension between the requirement for independence and another critical tenant of good risk management: the need to foster communication within the firm between the risk takers, risk managers, and the full hierarchy of senior management. It is a question of the delicate balance between the need to control conflicts of interest and the need for a business partnership between the risk managers and the risk takers, for the risk takers must ultimately own the risk.

We have seen a variety of practices between the two extremes, with arguments supporting each model. The argument for the distributed model of organization, where the risk managers report to the business head or CFO of the business unit, rests on the belief that the interaction between risk management and the business is best fostered by risk management being embedded in the business. We have seen variations of this theme, where there is a relatively small corporate-level risk management function reporting to the CFO, and risk managers who face off with the business unit reporting to the CFO of that business unit, with only a dotted line to the central risk management function.

The argument for the purist model, with centralized risk management reporting to the CEO, is that it elevates the stature of risk management within the firm. Most importantly, risk will have a direct voice in strategic decisions of the firm, increasing the likelihood that the risk dimension will become integral to the process. It may also be easier to get funding for the risk function if the request is not made through the voice of the CFO. Knowing that risk management has direct access to the ear of the CEO could also exert a subtle influence on traders' attention to the opinions of risk management. True respect for risk management, however, must be earned in other ways that are more important than reporting structures. The key issue is the competencies of the people in the risk function, the quality of the dialogue they can have with the traders, and the value of the insights they can provide into the trading business itself.

While the industry is generally evolving toward a more centralized model, many firms still stop short of having risk management report to the CEO rather than the CFO, and some have switched back and forth on the issue. The distributed model seems to be more prevalent in the small, more specialized firms. It is the large, universal banks that have moved to the centralized model with a formal hierarchy that may start at the top with a chief risk officer, with those in charge of the various risk areas (credit operations, market risk, etc.) reporting in to that person. The issue is perhaps a function of the complexity of the business. For brokers, whose trading operations lie at the heart of their business, the CFOs tend to be very close to the business; a risk function that reports to the CFO renders the CFO as the de facto chief risk officer, similar to the more distributed models that have the risk function reporting to the CFO of the investment bank or capital markets division. For banks that also have retail and commercial banking operations, a more centralized function is needed.

The Risk Partnership

Perhaps more important than the reporting lines are the subtler issues of the nature of the business partnership between risk and business units. Risk management must have the respect of the business units. Otherwise it will be ineffective, a sort of policeman or "ivory tower quant" group. The competencies of risk management must be on a par with those of the traders, which generally means that risk management has to be prepared to pay up for talent. A culture of constant contact and communication is also necessary. It is really a tricky issue of collaboration between the risk function and the traders. The business unit must see value in the risk function. It is, after all, in the best interests of the business to stay within the desired risk tolerance of the firm. The risk function can certainly help the business analyze its risks. It can provide a fresh and objective view of the proposed trades, one that should be valued by the traders. Ideally, traders should be willing to accept the risk function's opinion not because it has the authority to enforce that opinion, but because of the force of the argument it presents.

If the risk function is sometimes weak, it's because the role is merely one of calculating the risk equivalents and generating reports. In that case, risk management becomes a reporting and control function, which is not desirable. It reduces risk management to the role of a policeman, who is only there to tell traders when they have done wrong. We have also seen organizations in which reporting and control reported to the business unit, but the risk management unit was an independent, centralized function, using the reports and analysis generated by the control team. The business unit can want to control the analysis and reporting function because it is useful to management. While that arrangement lacks the attribute of independence, it could work if the interests of risk management and the business unit are allied. Nevertheless, it indicates a lack of respect for the risk function if the business feels the need to control the reporting.

On the question of limits-setting, all firms have some formal limits structure to govern market risk takers, not unlike the limits structures governing credit risk-taking. They vary in how involved the risk management function is in setting those limits at each level of the limits hierarchy—the overall firm-wide limit, the business unit-level limits, the desk limits, and the trader limits. The industry norm is for risk management to participate in the process down to the desk level. In large firms, which can have many traders to a desk, the responsibility for cascading the limits down to the trader level is the responsibility of the desk head. In some firms, however, individual traders have no limits at all! In some firms, where a centralized risk management function is newer and where the business units control more of the risk process, risk cascades the firm-wide limit only down to the business unit level. In others, with a purist form of centralized risk management, risk cascades down to the individual trader. The industry norm, however, strikes a balance between giving risk management control over risk concentrations and making risk management so involved in individual trades that they risk becoming like traders or business managers themselves. Just as the cascading process for limits is important, so is the methodology used to arrive at the limits. How much input does the business have in the process versus risk management? We believe risk management should drive the process, though it should do so in collaboration with the business units, and with a deep understanding of the business needs and prospects, the risk appetite, and the budget goals of the firm.

There is also a variation in limit-setting philosophy. Some like to set limits high, at a level that expresses a stressed case level of potential losses that is deemed tolerable to the firm. Risk management then typically would not interfere unless the limits were in danger of being breached. Others like to set limits low, so that there would be frequent requests for temporary excesses, on the theory that this process triggers communication and keeps risk managers informed. Communication would seem to be a good thing, but if excesses are granted on the basis of the business sense of the trade and its risk-reward trade-off, does frequent granting of excesses turn risk managers into business managers themselves? Is it not better to let traders operate freely, subject only to broad exposure constraints?

There is also industry variation on how important intraday limits are. Some desks would normally experience a high daily turnover of positions (currencies, government bonds, cash equities), others less. Most firms believe the desk heads would expect traders to stay within their limits during the entire day, and that desk heads would have a feel for whether that was true based on real-time position-tracking systems. But only a couple of firms have an ability to monitor intraday limits in real time. Merely having a set of limits is not a sufficient condition for defining the risk appetite of a firm. As carefully as the limits may be defined and attuned to the budget and the realities of the business environment, they do not capture the essence of what the risk culture should be. In fact, they could backfire, in that traders could game them—try to maximize the amount of trades they can do while staying within their risk limits. Something else is needed to provide a moral compass, which everyone in the firm can use to make myriads of daily decisions about appropriateness beyond the question of whether a trade is within limits. “No surprises” is one attempt at such a standard but is too vague and trite, because it could apply to any firm. A clear qualitative description of the risk appetite should include reference to the desired risk profile of the firm, its strategic and its budget goals. A holistic perspective on risk appetite would also include a crisp quantification in terms of stress limits or stop-loss limits and its impact on the firm’s capital and earnings.

One important piece of the intricate risk management puzzle is the back office, the processing function. Not only is it important that the back office cannot be subverted-incented to allow the front office to manipulate the trade processing in ways that alter the profit and loss statement; it is also important that it be adequately resourced to handle large volumes of trades. Controls in the back office have played a large role in many of the most visible trading debacles. Yet risk managers frequently are unfamiliar with that area. The back office still is seen as a service center, frequently reporting to the business unit, though not to the front office traders within it. Less frequently, it reports to the COO or CEO of the group. It seems to us that this critical process of the risk management framework should be under the purview of risk management. For one thing, it may have better access to resources if it reported to a risk management framework that itself had a high stature within the firm. As it is, back-office capacity is frequently outrun by volumes. The caveat is that at present, risk management does

not generally have the skill set or the experience. That skill could be acquired over time, and would further the more holistic view of risk in an enterprise-wide risk framework.

As different as the risk structures at various large firms are, it is sometimes difficult to say that different structures may work for different types of firms. In particular, the smaller, more focused types of trading houses have thrived for years with a relatively decentralized type of structure, with less formal documentation and procedures and less in the way of checks and balances from independent bodies. They rely instead on excellent communication flow and a strong understanding of the firm's risk appetite on the part of all business managers. The problem is that as a firm grows, it may become too complex to work in this fashion. The point at which that happens will not necessarily be apparent ex ante. Thus, our position is that the following attributes of best practices are important:

- A robust internal information system is needed that permits sophisticated analysis of portfolio risks and stress testing;
- Risk limits need to be assigned by the risk function down to the desk level; and
- A holistic and well-articulated risk tolerance statement is an important aspect of developing a strong risk culture.

The risk function needs to strive to attain a high stature within the firm. A centralized reporting framework reporting to the CEO is preferable but insufficient to ensure the independence of the risk function. An emphasis on the quality of personnel who can work effectively in a partnership with the business units is a critical aspect of a strong risk function.

Chasing Their Tails: Banks Look Beyond Value-At-Risk

In August 2004, Standard & Poor's Ratings Services expanded its review of the trading risk management (TRM) practices of global banks and other financial institutions that have significant trading operations as an integral part of its overall credit rating process for these issuers. (See "Credit Policy Update: S&P Expands Review of Trading Risk Management Within Financial Institutions," Aug. 3, 2004 on RatingsDirect.) Using its policy, infrastructure, and methodology (PIM) approach, Standard & Poor's assesses the impact of the quality of an institution's TRM practices on its overall creditworthiness. This article provides further insight into the methodology dimension of PIM and states our evaluation of Value-at-Risk (VaR) as an aggregate measure of market risk.

VaR is the most commonly reported aggregate measure of market risk by banks and other financial institutions. VaR has been around for more than a decade now, but the understanding of its benefits and limitations remains spotty within the investor community. Clearly, VaR is a significant and useful step forward; it can be applied to any financial instrument, portfolio, or risk factor and can be used as a measure of the relative riskiness among subportfolios. To that extent, Standard & Poor's prefers that an institution have a disciplined approach to calculating VaR, as the spillover effects of this process in terms of infrastructure and data requirements are beneficial to the overall risk awareness of the institution.

However, VaR has some severe limitations that, if not properly appreciated, can lull a company into a false sense of security. For instance, VaR lacks the criteria to provide a consistent measure for comparing the relative risk appetite across institutions, as the assumptions used by firms in calculating VaR can be vastly different and have varying degrees of precision. In addition, as a stand-alone measure VaR ignores the extent of tail risk that an institution is exposed to, especially under abnormal market conditions, and falls short of satisfying a key mathematical property required of a robust measure of risk (1, 2). For these reasons, Standard & Poor's and other market analysts believe that VaR should be interpreted with caution in evaluating market risk and should ideally be used in conjunction with other risk measures. The additional measures that Standard & Poor's analysts review in conjunction with VaR are explained later in this article.

Why VaR Is Not A Consistent Measure For Comparing The Risk Appetite Across Financial Institutions

It is well understood that VaR as a measure of market risk is a function of several underlying assumptions. Hence, to look at the reported trading VaR numbers of two institutions and comment on their relative risk appetite is not meaningful.

Most financial institutions use one of three measures for estimating VaR: parametric (or variance-covariance) VaR, Monte Carlo VaR, and historical VaR. Not surprisingly, Standard & Poor's has found through its TRM surveillance sessions that most institutions do not use parametric VaR or Monte Carlo VaR as an aggregate measure of risk. One must remember that parametric VaR assumes normality of the profit and loss distribution and requires an estimate of the volatilities (variances) and correlations (covariances) of the risk factors in the portfolio. Then, for a given confidence

Prodyot Samanta
New York
(1) 212-438-2009

Tanya Azarchs
New York
(1) 212-438-7365

Nick Hill
London
(44) 20-7176-7216

level and time horizon, the parametric VaR of the portfolio is a multiple of the portfolio's standard deviation, which is derived from the estimated matrix of variances and covariances. Parametric VaR is easy to compute, but in addition to the assumption of normality of returns, it has several limitations. The most important of these is its inability to capture the risk arising from the nonlinearity of the positions in the portfolio, or gamma risk. Nonlinearity is present in fixed-income positions (such as a bond) and in those positions with optionality. In addition, this approach requires an estimation of the volatilities and correlations of the risk factors. Obtaining a "true" estimate of the future volatilities and correlations (or variance-covariance (VCV) matrix) of the risk factors is nontrivial.

A second approach to estimating VaR is employing a Monte Carlo simulation technique. Under this approach, a future probability distribution for the relevant risk factors is assumed and randomly simulated several times over a given time horizon. The portfolio is re-evaluated under each simulation and a histogram of the profit and loss is obtained, from which the VaR is inferred. However, the Monte Carlo methodology is also dependent on the future VCV matrix and the distributional assumptions of the risk factors. In addition, the systems requirements to run a full-fledged Monte Carlo VaR engine are quite demanding, especially if the portfolio has a number of complex options and mortgages. Hence, regulators have resorted to the more practical and easily understood historical method of calculating VaR.

Briefly, historical VaR calculates a time series of historical changes in the relevant risk factors over a specified time period, applies those changes to the current levels of the risk factors, and then re-evaluates the portfolio to obtain a histogram of the profit and loss, from which the VaR is inferred. This approach assumes that the historical distribution will hold (or is a good proxy for the "true" distribution of returns) for the period over which VaR is being calculated. Also, as this approach is non-parametric, it does not require an estimate of the variances and covariances as these are already reflected in the historical time series. This approach is clearly simple to understand and, given the availability of historical data, is easy to compute.

Standard & Poor's has observed that the industry tends to use a range of values as inputs to an historical VaR computation. The key inputs are the length of the historical look-back period for the risk factors, the choice of the confidence interval, and the method for revaluing the portfolio given the set of historical scenarios. For all these key inputs, the range of values can vary significantly across institutions, making a meaningful comparison of the VaR estimate very difficult.

For example, Standard & Poor's has found that the length of the look-back period for the time series of risk factors used by a number of financial institutions could vary from as short as 250 days (the equivalent of one trading year) to more than 1,000 days (approximately four years). Institutions that weigh recent market volatility higher tend to use a shorter time series, while those that aim at having a relatively more stable VaR over time tend to use a longer time series for the risk factors.

There is no general consensus on whether the level of the confidence interval should be 99%, 98%, or 95%. The higher the confidence interval, the higher might be the VaR estimate, depending on the shape of the tail of the profit and loss distribution.

In addition, the methods employed for revaluing the positions over the set of historical scenarios vary across institutions. At one end of the spectrum are institutions that use an approximation (such as a delta or delta-gamma approach) to revalue the positions. This is analogous to approximating the value of a bond using either duration only or duration and convexity (to capture the curvature or non-linearity in the price-yield relationship of the bond). Also, one must remember that approximating the value of a bond by duration only can lead to a significant underestimation in value depending on the size of the yield change.

At the other end of the valuation spectrum are institutions that employ a full revaluation. To continue with the bond example, this is equivalent to present-valuing the cash flows from the bond at an appropriate discount rate. Between these approaches are some institutions that use a combination of these methods; an approximation for some products (usually the ones that are complex with a fair degree of non-linearity, and hence the ones that are most likely to produce the valuation errors) and a full revaluation for others. While the approximation techniques speed up the computational time for the VaR calculations, they do have the drawback of producing potential valuation errors depending on the degree to which the non-linearity inherent in the positions is accurately captured.

Given the range of possible assumptions behind estimating VaR, it is not always possible to use the reported VaR estimates as a common denominator for comparing the risk appetite across financial institutions. Further, VaR is only one aggregate measure of risk. VaR does not capture *tail risk*; it says nothing about what the worst case loss might be. Hence, it is critical for investors and analysts to *look beyond VaR* and request that institutions report on alternative risk measures as well.

Why VaR Is Not A Robust Measure Of Risk

Philippe Jorion and Nassim Taleb have debated the pros and cons of VaR as an appropriate measure of risk (5,6). It is well known that if markets suddenly experience severe turbulence, VaR will inevitably underestimate the risk and will only capture the sudden jump in volatility a few days down the road. In addition, the VaR for complex portfolios may be computationally challenging and may not clearly identify the risks inherent in the portfolio. For example, if one portfolio contains a single long position in an option and another a short futures position, VaR will indicate the same level of risk, even though the potential loss in the short futures position is arbitrarily high!

As highlighted above, calculating VaR can be quite challenging, as it depends on a number of assumptions employed by the user. In addition, from a systems perspective for large portfolios, the aggregate computation cannot be split into smaller subcomputations, as VaR is *not additive* by position or by risk factor. In other words, if a portfolio is split into smaller subportfolios, the VaR of the entire portfolio is not the sum of VaRs of the individual subportfolios. Further, if a portfolio is split by the relevant risk factors, such as foreign exchange (FX), equities (EQ), interest rates (IR), or commodities (COM), then the VaR of the entire portfolio is *not* the sum of the VaRs of the individual risk factors. That is, $FX(VaR) + EQ(VaR) + IR(VaR) + COM(VaR)$ is *not equal* to $Portfolio(VaR)$.

One important reason why VaR is not a robust risk measure is that it does not always provide an accurate sense of the degree of diversification within a portfolio. This occurs because VaR lacks an important mathematical property referred to by Philippe Artzner et al. (3, 4) as *subadditivity*. A risk measure such as VaR is said to be subadditive if the VaR of the portfolio is less than or equal to the sum of the VaRs of its individual components. [Mathematically, if a portfolio is split into two subportfolios X and Y, then subadditivity implies that $VaR(X+Y) \leq VaR(X) + VaR(Y)$.]

Intuitively, subadditivity is associated with the notion of risk reduction through diversification. If there is no diversification benefit in a portfolio (i.e., all components were perfectly correlated), then one would expect the VaR of the portfolio to be as much as the VaR of the individual components. On the other hand, if there were diversification, one would expect the VaR of the portfolio to be less than the VaR of the individual components. The risk in a diversified portfolio should be less than the risk in a portfolio that is not diversified. However, VaR does not guarantee this property, unless one assumes a Gaussian space (or that asset returns are normally distributed). Hence, the extent of diversification reported by firms based on a VaR measure can be misleading and incorrect! (*See the Appendix of this article for an example that demonstrates how VaR might violate this property.*) Artzner et al. have also demonstrated that because VaR lacks this property of being subadditive, it is not an appropriate metric for allocating capital or assessing risk-adjusted performance.

Expected Shortfall

A risk measure that looks beyond VaR and into the tail of the distribution is expected shortfall (ES). While a 95% VaR is the *minimum* potential loss of the 5% worst-case scenarios on a portfolio over a given time horizon, ES on the other hand is the mean loss of the 5% worst-case scenarios on a portfolio over a given time horizon. ES looks beyond the quantile on the left tail of the distribution that is used to compute VaR. Hence, it is also referred to in the literature as “mean excess loss” or “tail conditional expectation.” What differentiates this measure from VaR is that it captures tail risk.

More generally, ES is the average of the $(1 - X\%)$ worst-case outcomes stipulated by the calculated $X\%$ VaR estimate. For example, if a 99% VaR is calculated using 1,000 days of historical data (or outcomes), the ES is the average of the (1-99%) or 0.01 worst-case outcomes. In this case, since there are 1,000 outcomes, the ES is the average of the 10 (= $0.01 \times 1,000$) worst outcomes.

In addition to capturing tail risk, ES satisfies the property of subadditivity in that the ES of a diversified portfolio is always less than the ES of the portfolio's individual components. This makes this risk measure more robust than VaR in expressing diversification benefits. An important characteristic of ES is that it is aware of the *shape of the conditional distribution* of the worst-case scenarios beyond a specified quantile, while VaR (by definition) ignores this altogether and hence underestimates the potential risk. (See the Appendix for an example of why ES is a "coherent" and robust risk measure, and to that extent more appropriate for capital allocation decisions.)

There is very little additional work, either from a computational or from a systems perspective, that is required of an institution to calculate ES. It is not surprising that some institutions currently look at this measure in addition to their VaR. As a discipline leading to a high quality risk management practice, it makes sense for all institutions to adopt ES within their risk assessment framework. In addition, because of ES' simplicity, it is equally important for analysts and investors to request this information of institutions when assessing their market risk appetite and the robustness of their risk management practices.

Stress Testing

Stress testing is an important risk management tool, and the quality and framework adopted by institutions around stress testing is an important criteria used by Standard & Poor's analysts in assessing the risk management practices of financial institutions. Broadly speaking, stress testing assesses the potential impact of historical or hypothetical events or specific movements in risk factors (sensitivity-type analysis) on a given portfolio. All institutions conduct some form of stress testing.

Stress testing evaluates the risks that VaR and ES may not capture, such as events with a very low probability of occurrence but with a significant impact for the firm. The primary objective of stress testing is to achieve an understanding of the risk profile of the firm that goes well beyond what can be described by VaR and ES. A fair amount of thought coupled with macroeconomic analysis goes into designing meaningful stress tests, whether they are historical, hypothetical, or sensitivity-type analysis. In addition, correlations and implied effect among risk factors need to be carefully assessed when designing these tests. Hence, the quality and framework around stress testing adopted by institutions is what adds value to the risk management process and complements measures such as VaR and ES.

For example, some of the most commonly run historical stress tests are the events of "October 1987," "Bond Markets of 1994," "Asian Crisis of 1997," "Long-Term Capital Management/Russian Crisis of 1998," and "Sept. 11, 2001." However, when an institution stress tests its portfolio against "Black Monday," what exactly does that mean? Do they apply the drop in equity markets as experienced on Oct. 19, 1987? Or do they use a time window around that date? Do they also apply the implied correlations of the bond markets as well? These questions, and similar ones, are important issues, as the outcomes will vary depending on how the stress test is designed. Designing meaningful stress tests is critical and, when they are performed correctly, they constitute a powerful tool that complements standard risk measures such as VaR and provides the institution and its shareholders with a deeper understanding of the firm's risk profile.

In conclusion, it should be clear for the reasons and examples cited above that VaR is only one measure of aggregate risk and that it is important to look beyond VaR when measuring market risk. ES and well-designed stress tests are excellent complements to VaR, as they capture tail risk and provide for a more meaningful assessment of the institution's risk profile.

Notes

- [1] Carlo Acerbi, Dirk Tasche, “Expected Shortfall: A Natural Coherent Alternative to Value-at-Risk,” (Milan, AbaxBank, 2001).
- [2] Carlo Acerbi, Claudio Nardio, Carlo Sirtori, “Expected Shortfall as a Tool for Financial Risk Management” (Milan, AbaxBank, 2001).
- [3] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, David Heath, “Thinking Coherently,” RISK, November 1997, 68-71.
- [4] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, David Heath, “Coherent Measures of Risk,” Mathematical Finance, July 1999, 203-228.
- [5] Philippe Jorion, “Value at Risk: The New Benchmark for Managing Financial Risk,” (New York, McGraw-Hill, 2000).
- [6] Nassim Taleb, “Fooled by Randomness: The Hidden Role of Chance in the Markets and in Life” (New York, Thomson Texere, 2001).

Appendix

Artzner et al. have identified a risk measure to be “coherent” if it possesses certain mathematical properties or axioms.

One important property for coherence is subadditivity. To see how VaR may violate the subadditivity property, consider the following example described by Acerbi and Tasche. Suppose there are two bonds, A and B, with non-overlapping probabilities of default. Each has two different default states with recovery values of 70 and 90 and associated probabilities of 3% and 2%, respectively. In all other scenarios, the bonds redeem at 100. Table 1 below summarizes this information under five possible scenarios. For simplicity of exposition, assume that the initial value (or current market value) of the bonds and the portfolio of the two bonds is the expected value (EV) of the payoff given the assumed probability distribution. Taking the sum of the last three columns in Table 1, this can be calculated as being 98.9 for Bond A and Bond B and 197.8 for the portfolio of the two bonds.

Scenarios	Probability	Bond A	Bond B	Portfolio (A+B)	EV of Bond A	EV of Bond B	EV of Portfolio (A+B)
1	0.03	70	100	170	2.1	3	5.1
2	0.02	90	100	190	1.8	2	3.8
3	0.03	100	70	170	3	2.1	5.1
4	0.02	100	90	190	2	1.8	3.8
5	0.9	100	100	200	90	90	180

EV—Expected value.

Table 2 below shows the profit and loss (P/L) for each bond and the portfolio under each scenario. From this table, it is easy to calculate the 95% VaR or 5% worst-case losses.

Scenarios	Probability	P/L Bond A	P/L Bond B	P/L for Portfolio (A+B)
1	0.03	(28.9)	1.1	(27.8)
2	0.02	(8.9)	1.1	(7.8)
3	0.03	1.1	(28.9)	(27.8)
4	0.02	1.1	(8.9)	(7.8)
5	0.9	1.1	1.1	2.2

The 95% VaR for each bond is 8.9 and 27.8 for the portfolio. Clearly, in this example VaR violates the property of subadditivity. (The portfolio VaR turns out to be greater than the sum of the individual VaRs for each bond). Based on VaR, diversification between these two bonds would be discouraged! Notice also that VaR understates the risk; it is the “best” of the worst-case scenarios and totally ignores tail risk or the risk beyond the percentile used to compute VaR. It seems natural, then, that investors should look at other measures beyond VaR in assessing the potential level of market risk within an institution.

To see why ES is a coherent (or robust) measure of risk in terms of satisfying the axiom of subadditivity, we can refer back to the bond example in Table 2. In looking at the 5% left tail of the distribution of P/L for Bond A, Bond B, and the portfolio of the two bonds, it can be seen that the ES for Bond A and Bond B is 18.9 while the ES for the portfolio of the two bonds is 27.8.

The first thing to note is that ES satisfies the axiom of subadditivity in that the ES of the portfolio is less than the sum of the ES of the individual bonds. In addition, one should note how VaR might underestimate the potential risk on the individual bonds. According to the 95% VaR measure (or 5% worst-case losses), the potential loss on the individual bonds is only 8.9, while ES estimates the loss at 18.9!

U.S. Financial Institutions Rethink Trading Risk In Volatile Markets

By Robert McNatt

Tanya Azarchs
New York
(1) 212-438-7365
Prodyot Samanta
New York
(1) 212-438-2009
Marc Eiger
New York
(1) 212-438-1280

Given the recent volatility in the capital markets—at a time when some financial institutions are making riskier bets than ever in their proprietary trading—Standard & Poor’s Ratings Services is looking ever more carefully at how banks, brokerages, and other financial firms manage trading risk. While Standard & Poor’s has long recognized the usefulness of Value at Risk models to help determine how much might be at stake, it is also clear that numerical indicators can only show so much. Since the spectacular collapse of Long-Term Capital Management in 1998, many financial institutions have been paying more attention to nonstatistical risk measures. Standard & Poor’s increasingly has incorporated such considerations into its analyses to draw a fuller picture of the risks financial institutions are undertaking at their trading desks.

Evaluating Risk Beyond the Numbers

Assessing risk in nonquantitative areas is neither easy nor simple. Nevertheless, there are three areas to look at when doing so. First, any Standard & Poor’s analysis will consider a company’s risk management structure and its philosophy of risk: How does the firm establish what risks it will allow its traders to take, and how are those risks aligned with its business strategy? In addition, what training and education do the company’s risk managers have? Next, our analyses will look at the infrastructure involved—and by this we mean more than just the operational integrity of the back office, though, of course, that is part of it. Badly managed back-office operations have caused serious losses for some institutions. More broadly, this part of the analysis entails seeking a sense of how well the institution has integrated its systems with the decision-making process. And finally, the quantitative measures of risk management receive a hard look. They include pricing models, historical and hypothetical stress scenarios, and back-testing processes used.

One risk management initiative that has been apparent at many trading operations is a move toward centralization of the risk management structure, separating it from business units. In addition, some firms, particularly those that evolved from the close-knit partnership cultures, rely on their cohesive culture but can lack the formal processes, procedures, and controls that help more centralized organizations avoid disaster at the trading desks. Is one structure inherently better than the other?

At the Standard & Poor’s Global Bank Conference in New York City on Nov. 4, 2004, Lesley Daniels Webster, executive vice president and head of market risk management at JP Morgan Chase & Co., noted that the ideal may lie somewhere between the two extremes. Without doubt, traders need established policies and procedures. But rules should be revisited periodically and some flexibility should be extended, especially when one goal is achieving a consistent agenda among different trading cultures at the same institution. “My own preference,” said Webster, “is to find a middle ground.”

Trading Limits and Flexibility

Financial institutions can also have very different policies about what a trader's limits should be—a key way to look at an institution's risk profile. At some places, limits are not meant to be breached, and when they are, consequences can be severe. At others, limits are breached frequently, but excesses are granted liberally, the theory being that requests for excesses spur discussion within the firm and raise issues that might not otherwise surface.

Again, Ms. Webster suggests that the thing to look for might be a balance between the two extremes. Clearly, traders need to have limits on the positions that they can take. But the trick is to establish sufficient limits while not micromanaging on the trading floor. Part of that means setting limits that are aggregate in nature and also have a reasonable number of limits that are relevant at the desk level and down to the trading floor. And all of them, of course, should be in line with the firm's overall risk appetite. Smart risk managers, then, will look at trading performance from three important perspectives: that of the individual trader, the firm, and the bondholders. Indeed, it is not at all unreasonable for those managers to consider whether heavy losses would affect those who hold the institution's debt, much as a Standard & Poor's analyst would.

Ultimately, the truth is that risk management shows its effectiveness best when things are at their worst. When markets are less volatile and there are no surprises, it is more difficult to evaluate how well an institution's risk management policies are. Unfortunately, it is often only when unpleasant surprises arise that the effectiveness of risk management policy becomes most apparent.

Key Contacts

Tanya Azarchs
212-438-7365
tanya_azarchs@standardandpoors.com

Prodyot Samanta
212-438-2009
prodyot_samanta@standardandpoors.com

Nick Hill
(44) 20-7176-7216
nick_hill@standardandpoors.com

Jesus Martinez
(34) 91-389-6941
jesus_martinez@standardandpoors.com