The Ten Great Challenges of Risk Management

Carl Batlin* and Barry Schachter**

March, 2000

*Carl Batlin has been the head of Quantitative Research at Chase Manhattan Bank and its predecessor banks, Chemical and Manufacturers Hanover, since 1986. He also served on the ISDA task forces that responded to the B.I.S. proposals on market risk capital in 1993 –1996.

**Barry Schachter is the head of Enterprise-wide Risk Management at Caxton Corporation. He also maintains the All About Value-at-Risk website, http://www.GloriaMundi.org, and he can be contacted at bns@gloriamundi.org

The Ten Great Challenges of Risk Management

Risk management has been practised as a professional discipline in financial institutions that maintain derivative dealerships for at least a decade. The fundamental goal of this function is to improve the quality of business decision-making at all levels of the firm and thereby to increase shareholder wealth. It executes this responsibility through activities that clarify the firm's exposure to all forms of risk to its future earnings and analyses them on a sound economic foundation. One decision that emerges from this analysis is the amount of capital that the institution should hold to absorb future losses that can occur in the course of its trading activities. Since trading income is uncertain, though, an important risk management activity focuses on the risk measurement problem – estimating how large future losses could be. Complementary activities involve the monitoring and enforcement of risk-based trading limits, thereby facilitating the risk-adjusted performance evaluation of individual trading desks. These two aspects of the risk management role turn the capital allocation decision into an optimisation problem: too small an amount exposes the firm to excessive levels of risk, but too large an amount raises funding costs and reduces profitability.

The level of sophistication with which the risk management function is performed has advanced significantly in recent years. Of course, there are still important problems whose solutions would improve the function's effectiveness even more. While there are undoubtedly differing views about these, we have identified what we believe are the ten most important issues facing risk managers, and, in doing so, we hope to characterise the current state of the discipline. These issues fall into three categories: risk management applications, risk types, and risk measurement implementation issues.

The Applications of Risk Management

1. Firm-wide Risk Measurement vs. Business-specific Capital Allocation

The risk management function places great emphasis on trying to derive the probability distribution for the earnings of the whole firm. Risk managers are specifically interested in the extreme left-hand tail of this distribution as a predictor of the largest ex post loss that the firm could experience. Then, to provide confidence of firm-wide solvency, the institution would hold capital in excess of that amount. On the other hand, this total amount of capital must also be allocated among the individual trading desks so that their performance can be evaluated on a risk-adjusted basis.

Because of imperfect correlations among desks' income streams, the magnitude of potential losses for the whole institution is less than its sum for the individual trading businesses, creating the following methodological dilemma: if risk managers estimate the firm's capital adequacy requirement first and then disaggregate the result for allocation to individual businesses (the top-down approach), they are vulnerable to resource misallocations among desks. The alternative of estimating capital requirements at the individual business level first and then aggregating the results with the aid of a set of estimated correlations (the bottom-up approach), on the other hand, is susceptible to the risk of overcapitalising the firm as a whole. While many firms accept the inconsistency of relying on both approaches, the current best practice in the derivatives industry favors the bottom-up approach, in part because of a tradition of measuring the various types of risk (e.g., credit, market, operational) independently and recognising the different relative importance of these risk types among businesses. The development of a unified methodology which preserved this level of specialisation without sacrificing accuracy at the macro level would be a significant step forward.

2. Regulatory vs. Managerial Capital

Regulatory agencies impose capital requirements on financial institutions as part of their oversight responsibility to ensure the safety and soundness of the financial system. The global nature of the derivatives business led the Bank of International Settlements to sponsor an effort to ensure consistency across countries, first for credit risk capital in the late 1980's and then for market risk capital in the mid-1990's, with local regulatory bodies responsible for enforcement. Where the use of formulas are mandated by the regulators to compute the required amount of regulatory capital, they tend to be crude mechanisms that correspond only loosely to the economic calculations that risk managers make for internal managerial purposes. Even where internal models are allowed, as with market risk capital, the regulators still specify the parameter values that must be used in those calculations and an add-on factor for specific issuer risk, thus causing a divergence from the internal risk measurement process. The result is that in most institutions, significant resources are expended to produce two sets of capital calculations – one to measure the firm's level of risk, and one to satisfy the regulators. The risk management community can save considerable resources if it will be able to articulate a measurement methodology that satisfies both purposes with minimal incremental cost.

3. The Extension to Non-financial Firms

The success of the risk management role in financial institutions has been due in large part to its imposition of economic standards in businesses that otherwise would be evaluated by accounting measures. This was a feasible accomplishment because of the economic conventions of the trading business, especially with regard to the measurement of income on a mark-to-market basis and to the active management of risk exposures through hedging strategies. During the 1990's, some well-publicised financial losses at non-financial firms raised the issue of whether the methodology of risk management could be transplanted to those environments effectively. Much has been written on the adaptation of trading-oriented risk measurement methods to the corporate environment or, within financial services firms, to traditional banking businesses. In both cases, for instance, the time horizon to measure risk is much longer than in trading businesses. It remains an open question whether these adaptations are robust enough to permit comparable measurement of risk across trading and non-trading activities.

Risk Measurement by Type of Risk

4. Market Risk

The measurement of market risk – the potential losses due to adverse changes in the market prices of financial assets, including interest rates, foreign exchange rates, equity prices, and commodity prices – probably reflects the most advanced methodology for risk measurement. Its main tool is a position's or portfolio's Value-at-Risk (VaR), or the loss in market value that could occur over a given time interval at a specified level of confidence. There are two main approaches for measuring VaR, which differ in their characterisations of the probability distribution of earnings. Those which depend on a specific sample of historical data are more flexible in representing the empirical distribution which occurred, but their ability to predict future losses might be limited if the true distribution is not stationary. On the other hand, those methods that parameterise the distribution with the assumption that returns follow a Gaussian process could be flawed if the underlying distribution were actually non-normal. The normal distribution's association of a given confidence level with a particular size movement makes it computationally convenient to use for loss prediction, but empirical evidence seems to suggest that actual return distributions are leptokurtotic (i.e., they display fatter tails than the normal distribution).

Risk managers often try to proxy fat tails, or the unusually high probability of very large changes in financial market prices, by supplementing VaR calculations with the use of stress testing. By substituting, for example, the largest historical financial price movement ever recorded within a given time interval in place of a conservative VaR calculation, they hope to better represent statistical confidence of at least the 1% level in the analysis. In addition, they argue that stress testing can incorporate the risk of an unusually turbulent market environment not captured in VaR calculations. The crudeness of this strategy, though, commits their firms to holding a much higher level of market risk capital than should be required to maintain that level of confidence most of the time, thus reducing the efficiency of those trading businesses. The development of alternative approaches to modelling the probability distributions of the underlying financial prices more accurately, especially with an eye toward developing a better understanding of the phenomena responsible for observed leptokurtosis, is the subject of much risk management research activity.

5. Credit Risk

The oldest efforts to measure the risk of derivative products are in the area of credit risk. The observation that the replacement cost of a portfolio in the event of a counterparty's default is the mark-to-market value of the in-the-money contracts with that counterparty, rather than the notional principal amount, led to the tracking of gross and net mark-to-market exposure as the measures of counterparty credit risk. The additional observation that, unlike in the case of loans, this exposure from derivative contracts changes over time led to the allocation and monitoring of credit lines on the basis of potential credit exposure over the life of these contracts. Most methodologies to measure potential exposure, as well as the required level of credit risk capital, resembled VaR calculations and tended to overstate the actual level of credit line utilisation due to their inability to adjust to new market information after the inception of the transactions. With the development of the credit derivative market in the mid-1990's, a practical method for the active management of derivative credit exposure became available. The methodology for pricing credit derivatives, incorporating the probability of default as well as the conditional cost of default, also presents the possibility of an efficient alternative to potential exposure for the measurement of credit risk. To date, this approach is not used extensively, but such a development would be a significant boost to the management of derivative credit risk.

6. Liquidity Risk

The assumption underlying both regulatory and internal calculations of market risk is that the firm's financial assets are traded in competitive and liquid markets. For circumstances in which this assumption is not warranted due to the size of the firm's positions being large relative to the total size of the market, a special case of the overall market risk created by this imperfect degree of liquidity needs to be evaluated. In general, illiquidity is manifested by the difficulty of exiting positions quickly without depressing existing market prices. It therefore increases the time requirement of such an operation, exposes the firm to the cost of market fluctuations in the interim, and is reflected in higher transaction costs in the form of the bid/offer spread. While a separate methodology beyond VaR and stress testing does not need to be developed to evaluate this risk, modifications are needed to measure it accurately and consistently for aggregation across markets of different liquidity characteristics. The main examples of these are the lengthening of the time horizon in VaR calculations and the consideration of transaction costs in computing potential trading losses. Research that reduced the arbitrariness with which these adjustments were made would be welcome.

7. Operational Risk

Operational risk, which generically includes such risks as legal (i.e., contracts in the portfolio may not be enforceable) and model (i.e., the analytical models upon which the firm depends for its valuation exercises may contain errors), refers to the possibility that the firm will be unable to process its portfolio of completed transactions. A recent ISDA survey found that the industry standard for measuring and managing this risk was significantly less advanced and more subjective than that of either market or credit risk. This is not surprising, when one considers that operational risk is fundamentally different from these other forms of risk. While all types of risk can cause large financial losses, operational risk is not readily related to changes in the market prices of financial assets. Because it is more idiosyncratic, firms do not view it as hedgeable in the way market and credit risk are. While many firms allocate operational risk capital among businesses according to internal audit scores, there is currently no consensus on an economically more meaningful framework to measure and control this important source of risk.

Implementation Issues in the Measurement of Risk

8. Delta vs. Full Revaluation Calculation of VaR

The traditional, computationally efficient VaR calculation for an individual derivative contract is the product of three elements: its current value, its delta or sensitivity to an adverse marginal change in the market risk factor (e.g., interest rate or equity price) on which it depends, and the amount of change in that risk factor that can occur at the stated confidence level. While this calculation is commonly used to assess the risk of swap contracts and other two-sided derivatives and has the advantages of ease of implementation and speed of computation, it can be particularly misleading in two situations: First, when the projected change in the risk factor is large, the convexity or gamma effect causes the original delta to understate the VaR. Second, when the derivative contract has sensitivity to more than one risk factor, such as in the case of option contracts, this method offers no guidance on how to proceed. Some alternatives have attempted to modify the basic method by adding in the effects of other risk factors (e.g., gamma, vega, or the effect of a change in market volatility, and theta, or the effect of the passage of time), but the most complete approach is the full revaluation method, which recalculates the value of each derivative position using the relevant pricing model and estimates of the potential movements in each risk factor. The outstanding issues to be addressed for this method to be effective are a means of estimating the future correlation structure among risk factors to ensure its accuracy, and a technological solution to the slow computation time associated with it. In the absence of the latter, the choice between delta-based methods and the full revaluation method becomes a trade-off between accuracy and computation time, and, as yet, there is no consensus on how to determine the optimality of that trade-off.

9. Volatility Estimation

The volatilities of market risk factors play a central role in market and credit risk measurement. The volatility or standard deviation of financial market prices provides the basis for determining the appropriate amount of price change to incorporate into the VaR calculation, especially in those approaches that depend on distributional assumptions, such as the normal distribution. In addition, for options, the level of volatility is an important determinant of the current market value of each contract. Despite the important role of this parameter, though, risk measurement methodologies rarely take into account the complexities of volatility estimation. While it is quite common for risk managers to compute a rolling set of volatilities from a fixed sample of historical data on financial market prices, updated at some specified regular interval, such a procedure usually ignores the economic reasons that dictate the appropriate data management choices for different types of contracts and market conditions. In addition, the choice of predictors between historical data, implied volatility data from traded option prices, and forecasts of volatility using advanced statistical techniques such as GARCH is largely an unexplored area in the practice of risk management. Finally, incorporation of the dependence of volatility estimates on the tenor of the derivative contracts and the level of the principal underlying risk factor price (i.e. the volatility curve and smile) could have a significant impact on the accuracy of the risk measurement exercise, although risk managers rarely take them into account.

10. Correlation Estimation

Correlation assumptions link individual VaR calculations to the measurement of the risk of a portfolio. Their two main uses are in the aggregation of trading desk risk allocations to derive a firm-wide risk estimate, and the risk measurement of contracts that are sensitive to multiple risk factors (e.g., options). The task of estimating correlation values is subject to many of the same issues that affect the estimation of volatilities, with perhaps the most important difference being that the greater depth of vanilla option markets relative to those for exotic options make it more feasible to rely on implied volatility data than on implied correlation data. Unfortunately, the estimation of correlations from historical data is not only impacted by the same difficulties as volatility estimation (e.g., instability over time), it can also be invalidated by a unique problem -- the presence of nonsynchronous data. If the recorded values of financial variables did not really occur concurrently, correlation calculations based on that data will not accurately represent the probability of co-movement. Methods to correct this problem still need to be developed, and as with most risk measurement issues, risk managers will continue to backtest the resulting risk estimates to monitor their adequacy.