IMPROVING COUNTERPARTY RISK MANAGEMENT PRACTICES

Appendix A:
Risk Measurement, Liquidity Risk and Leverage Estimation

Appendix B:
Counterparty Credit Exposure and Risk Estimation

Appendix C:
Model Regulatory Report Formats and Definitions

Appendix D:
Glossary of terms used in report

Counterparty Risk Management Policy Group
June 1999

Appendix A:
Risk Measurement, Liquidity Risk and Leverage Estimation

Introduction

Many commentators have characterized leverage as a contributor to firm instability and as a source of systemic risk. Indeed, the President's Working Group has concluded that the central public policy issue raised by the LTCM episode is how to constrain leverage more effectively.

Leverage exists whenever an institution is exposed to changes in the value of an asset over time, without having first disbursed cash equal to the value of that asset at the beginning of the period. As the main Policy Group report suggests, funds and their creditors should not focus strictly on the nominal level of leverage, however measured, but rather on how leverage amplifies market risk, funding risk and asset liquidity risk. This Appendix offers an analysis of the relationships among these risks, first describing various measures and identifying their respective strengths and weaknesses. The intent is not so much to prescribe better technical measures of leverage, but to suggest a better framework for analysis.

Traditionally, leverage measures have related a notional or gross exposure to equity. This treatment helps to measure the degree to which a change in the value of a portfolio would affect the value of equity (Net Asset Value), but does nothing to illuminate the probability of change occurring, or the likely magnitude of change in portfolio value. By contrast, risk measures are intended to estimate potential adverse change based on the specific characteristics of the portfolio. Two portfolios of like size may have quite different risk. For a given portfolio or strategy, higher leverage implies higher risk. However, evaluating the risk of the portfolio is a necessary first step because a leveraged portfolio of low risk assets may have less aggregate risk than an unleveraged portfolio of high-risk assets. The framework for risk analysis which follows attempts to reflect the underlying risk of the positions in a portfolio, and the economic and funding structure of the portfolio as a whole.

The analysis takes as a starting point the observation that, broadly speaking, there are two ways in which highly leveraged financial institutions fail:

1. They become insolvent – that is, their liabilities exceed assets ("capital insolvency"). While historically some highly leveraged institutions– for example, savings and loans in the early 1980s – continued to operate with mark-to-market negative equity, today most leveraged institutions would find it difficult, if not impossible, to continue in business if their net asset value approached, let alone dipped below, zero. We define risk measures which attempt to estimate the potential risk of NAV becoming negative as measures of *leverage*.

2. They run out of cash and are unable to raise new funds, even though, on an economic basis, they still have positive capital. This is the cause of financial distress far more frequently than actually becoming capital insolvent, although ultimately the reluctance of credit providers to extend more financing may often be traced to a fear of impending capital insolvency. Most of the hedge funds which experienced distress during the fall of 1998 did so because they were unable to meet margin calls in a timely fashion, even though their mark-to-market NAV appeared to be well above zero. We define measures which attempt to estimate the potential of an institution running out of cash as measures of *funding liquidity*. Because most highly leveraged institutions obtain much of their financing on a mark-to-market basis, the greater the size of a portfolio of assets relative to an institution's funding sources, the greater its

funding liquidity risk (a given percentage change in the price of the assets will deplete the funding faster). Broadly speaking, funding sources scale with an institution's capital, so increased leverage amplifies funding liquidity risk.

A more leveraged portfolio may accumulate larger positions for a given amount of capital. These positions, if they need to be liquidated, may take longer to wind down, or may cause greater market impact during liquidation. Moreover, the presence of greater leverage makes it more likely that such a liquidation will have to occur, as the institution either approaches capital insolvency or has to meet margin calls in an adverse market environment. We refer to the risk that the liquidation value of assets may differ significantly from their current mark to market value as *asset liquidity risk.*

Because some of the leverage and funding liquidity measurement frameworks we describe are based on value at risk (VAR) and stress tests, the following section describes the evolution of market risk measures and some of their shortcomings. Next, the section on liquidity discusses how the interaction of leverage with asset liquidity can make simplistic uses of VAR and stress tests break down; it then suggests some funding liquidity risk measures. The final section describes a series of leverage measures, shows how traditional balance sheet measures fail to reflect the true risk of insolvency, and suggests improved leverage measures.

Risk Measurement

Prior to the advent of risk quantification, the most commonly used figure for measuring the size of a loss contingency which might be sustained by a financial institution was the total assets of the firm. This is not to say that everyone thought all assets were equally risky; it was simply that the tools for distinguishing risk were not well established. The Basel Accord in 1988 promulgated a set of the risk distinctions: three in terms of counterparty quality (0% for sovereigns, 1.6% for OECD banks, and 8% for all other counterparties) and one in terms of collateral quality (4% for mortgages). The total for each asset class was multiplied by its respective percentage, and these risk amounts were summed across the bank's entire loan portfolio. The resultant "risk weighted assets" number quickly supplanted "total assets" as a more meaningful measure of the risk of a bank, although US banks are still held to a simple leverage ratio requirement as one test of capital adequacy. The Basel Committee also took great care to insist that the percentages were not indicative of default probability, potential asset deterioration or any other particular contingency. The Accord included capital charges for OTC derivatives, using current exposure plus add-ons (reflecting potential future exposure) as a loan-equivalent surrogate. The risk weights applied to derivative current exposure were allowed to be half of the risk weight of an unsecured loan to the same counterparty; the rationale was that only high-quality counterparties were engaged in derivatives business. In July 1994, the Accord was amended to include bilateral netting agreements, substituting net for gross replacement value in the calculation of current exposure and reducing the add-ons for potential future exposure.

The reduced credit exposure of swaps and options (current exposure averaging between 1.5% and 2.5% of notional value) was ultimately accompanied by additional charges for market risk as set forth in the Market Risk Amendment to the Accord in January 1996. This amendment split each bank's portfolio into two segments – the investment book and the trading book. The investment book continued to be capitalized under the original provisions of the Accord covering credit risk; banks were urged to measure the interest rate risk on their investment books but were not required to assign capital explicitly to it. The trading book, containing nearly all of the securities and derivatives positions, was required to be marked-to-market daily and to be capitalized for market risk at a minimum of three times the Value at Risk (VAR) calculation for a 2-week holding period, with a 99%-confidence interval. (The credit risk charge for OTC derivatives remained unchanged, and a "specific risk" charge for the credit risk of securities was introduced.) The volatilities and correlations underlying the VAR model, as well as the choice of aggregation method (variance-co variance, historical simulation, or Monte-Carlo simulation), was left to each individual firm. Models were to be validated by back-testing: if the daily P&L variation forecast by the model understated the actual P&L on more than a handful of days per year, the multiplier put to the basic VAR number would rise from three to four or even five. At present, all major international banks have implemented this regime.

Much of the discussion of market risk since this Capital Amendment was adopted has centered on stress-testing, which was a qualitative requirement in the Amendment but not included in the capital computation. In contrast to VAR, which weights each outcome by its probability and then sums these increments across a portfolio, stress-testing considers the consequences of particular outcomes, without regard to their forecast probability of occurrence. VAR's strength is in measuring the 99%-probability "boundary"; VAR is less useful for saying how great the loss will be for outlier events beyond that boundary. For example, if a firm's risk profile leaves it fatally vulnerable to a potential market shock, it is of scant comfort to predict that such a shock will happen on average "only" once per year. However, attempting to fix capital requirements at such "worst-case" figures is generally considered unrealistic, as it would lead to a severe contraction in financial intermediation. Nevertheless, banks must be aware of their potential vulnerability to market shocks and many consider placing limits on their aggregate exposure to potential stress events.

The evolution of risk measurement techniques – from balance sheet totals, to risk weights, to inclusion of off-balance sheet products, to netting, to marking-to-market, to market risk VAR, to stress-testing, to more rigorous credit risk weights, and ultimately to a VAR which captures both market and credit risk – is, in essence, the search for increasingly precise delineation of the distribution of future returns (or values) of a given portfolio. The value of any portfolio of tradable assets (including off-balance-sheet liabilities) will fluctuate over time. Understanding the range of possible variations and the probability of each is equivalent to capturing the risk profile of that portfolio. Although the full depiction for any portfolio of the distribution of future returns over a given time period is still years in the future, it is already clear that there are three fundamental components to this analysis. These three methods of examination correspond to the three basic parameters of a probability distribution: the mean, the standard deviation, and the downside tail.

Valuation. Techniques of valuation seek to uncover the *mean* of the distribution of future returns in a portfolio. Marking-to-market is an accepted tool of valuation because efficient market prices have proven to be far better predictors of future value than historical-cost or accrual-accounting figures. Marking-to-model may have similar advantages to the degree the model can be anchored against market prices. The recent development of applying counterparty default probability to the amount owed by that counterparty and then summing across counterparties to an "expected credit loss" for the entire portfolio, is an overt attempt to forecast the mean of the distribution of future credit losses. A lively debate is in progress over whether all current valuation information should be imported into the financial accounts of an institution; skeptics argue that doing so will only exacerbate market shocks when they occur. However, valuation methods will continue to evolve, even if the results are not required to be shown in the financial accounts, because discovering the mean is the first result of a successful investigation into the details of any distribution.

Value at Risk. Measuring the *standard deviation* of a distribution is the most helpful step in understanding the possible future variation of a portfolio. (Valuation deals essentially with historical risk, that which a portfolio has already endured.) Indeed, the ascendance of market risk VAR in the 1990's was due to its accuracy in assessing the ordinary changes in value of a trading portfolio. Efforts are now being made to translate VAR concepts into the realm of credit risk. These attempts to calculate the potential "unexpected losses" in a credit portfolio show much promise but face unique validation challenges. VAR is often used in management accounting to attribute "economic" or "risk" capital and the costs of that capital to a particular business activity. Supervisors rightly chose VAR as the basis for regulatory capital because a firm must be able to withstand the ordinary variation in its positions. Moreover, the best practices in the industry included using VAR measures, and supervisors recognized the advantages of harmonizing internal processes for risk management with supervisory standards. (The analogy with "initial margin" on a futures exchange, which is also a VAR-style calculation, is evident. If listed futures exchanges required initial margin large enough to cover every conceivable contingency, the initial margin would approach the notional size of the contract itself.)

Stress-Testing. The purpose of stress testing is to learn more about the *downside tail* of a return distribution. All styles of VAR rely at some point on explicit or implicit forecasts of the volatility (standard deviation) and correlation of underlying market factors. The distributions of these financial factors are rarely bell-curves; the frequent arrival of new information distinguishes the financial marketplace from truly random settings. In addition, the correlations of these factors are notoriously non-stationary. One of the benefits of normal (bell curve) distributions with stationary correlations is that once the standard deviation is computed, the entire distribution can be specified. For example, one can conclude that, in a normal distribution, a move larger than three-standard-deviations has a 1% chance of occurring. However, in less well-behaved distributions, it is very dangerous to reason from the standard deviation to any conclusion about tail probabilities. Further, because the standard deviation and all "confidence interval" analyses are probability-weighted,[1] the sizes of individual outcomes in the downside tail are hidden, because they only contribute to the calculation after being multiplied by their (small) probabilities.

In other words, VAR does not yield information about whether a particular downside event might be catastrophic for the firm in question; that can only be ascertained by examining the events individually. The challenge here is choosing a set of events for intensive analysis (i.e. stress-testing) from the nearly infinite universe of possible events. At present, firms are only beginning to share information on how they determine which events are simulated in stress testing, and no consensus has yet emerged: some simulate historical market shocks, some distill market risks to a manageable set of independent factors and shock the factors individually and in combinations, and some stress the factors underlying their major risk positions, some stress those factors which show the most current volatility in the market. As progress is made in the art and science of stress-testing, the results are increasingly being used for setting risk limits, and, in some judgmental fashion, are a factor in determining internal capital allocation. (Using stress-test results exclusively or mechanically for capital allocation would be the rough equivalent of requiring every resident of an earthquake zone to conduct daily activities as if the earthquake were occurring today; ordinary business would come to a standstill.)

The real benefit in stress-test analysis comes from studying the correlation of risks that in ordinary times appear independent, for instance, market and credit risk. This type of knowledge

---

[1] The variance (i.e. the square of the standard deviation) is mathematically equal to the sum of the probability of each outcome multiplied by the square of the difference between that outcome and the mean.