The Financial Crisis, VaR, and Banking Capital

Gregg E. Berman, Ph.D.
RiskMetrics Group

House Committee on Science & Technology
Subcommittee on Investigations and Oversight

10 September 2009

I'd like to begin by thanking the committee for this opportunity to present our thoughts on Value-at-Risk and banking capital in the context of the present financial crisis. My name is Gregg Berman and I am currently the head of the risk business at RiskMetrics Group, a provider of risk and corporate governance services to the financial community. I have been at RiskMetrics since its founding 11 years ago and in the last decade have worked with many of the world's largest financial institutions on the development of risk models, their use by hedge funds, asset managers, and banks.

Simple Roots of a Complex Crisis

My comments today start with a rather bold assertion---the current crisis was not unpredictable, unforeseeable, or unknowable. In that sense I'm not sure it should be classified as a fat-tailed event. Rather, it was caused by the coupling of two fundamental problems, namely:

1. the inability of market participants to acknowledge and prepare for the consequences of long-term trends, such as a protracted downward spiral in home prices, or a leveraging of the credit market through the use of CDS, and
2. the inability of market participants to recognize the economic exposures they had to those trends through holdings such as asset-backed securities and derivative contracts.

The fact that these issues went unchecked for many years led directly to the creation of multiple, unsustainable market bubbles, which when burst propelled us downwards into a full-blown crisis.

But if my assertion is correct and these events were foreseeable, then what does that imply about all the financial models and risk methodologies that were supposed to monitor and protect us from such a crisis? It is the answer to this question that I'd like to explore.

THE INEVITABILITY OF VALUE-AT-RISK

In the early days of risk management size was used as a primary measure of risk. After all, intuition tells us that $10,000,000 in Ford bonds should be ten times riskier than $1,000,000 in Ford bonds. But soon the market realized that $10,000,000 of Ford bonds is probably riskier than a similar

value of government bonds, but not as risky as $10,000,000 of internet start-up equity[1].

To address these issues practitioners switched from asking "how large is your position" to "how much can you lose". But there is not just one answer to that question since for any given security differing amounts can be lost with different probabilities. One can estimate these probabilities by polling traders, by building econometric models, by relying on intuition, or by using variations of history to observe relevant patterns of past losses. Each of these methods has their own benefits and weaknesses. And unless we consider only one security at a time, it will also be necessary to make estimates of how the movements in each security are related to the movements of every other security in a given portfolio.

These concepts are encapsulated by two well-known statistical terms: volatility and correlation. If one could measure the volatility and correlation of every security in a portfolio the question "how much can you lose" could be meaningfully addressed. This process is the basis of a popular risk methodology known as Value-at-Risk, or VaR.

[1] The matter is further complicated by derivative contracts that do not even have a well-defined measure of size. For example, what is the size of a contract that pays the holder $1,000 for each penny-increase in the average spread throughout September between the price of natural gas for delivery in November and the price for delivery in January? Technically the answer is zero since the holder owns no natural gas, but the risk is certainly not zero.

How VaR is Computed and How it is Used

Because security valuations are often driven by underlying market factors, such as equity prices, spreads, interest rates, or housing prices, VaR is usually calculated in a two-step process that mimics this behavior. In the first step a model for the economic exposure of each security is created that links its value to one or more of underlying market factors. In the second step future trends of these underlying factors are simulated using volatilities, correlations, and other probabilistic methods. These two steps are then combined to create a curve that plots potential profits-and-losses against the probability of occurrence. For any given curve VaR is defined to be the amount that can be lost at a specific level of probability. It is a way of describing the entire profit-and-loss curve without having to list every data point[2].

The accuracy of any VaR number depends on how well underlying markets have been simulated, and how well each security has been modeled. There unfortunately exists a tremendous variability in current practices and different financial institutions perform each step with varying levels of

[2] Exhibit 1 on page 10 shows the potential 1-day profit-and-loss distribution of selling a short-term at-the-money put on the S&P 500. Out of 5,000 trials we see that about 50 of them have losses of 250% or worse. Thus VaR is 250% with a 1% probability. Alternatively we can ask for the worst 5 out of 5,000 trials (a 0.1% probability) and observe these losses to be 400% or worse.

accuracy and diligence[3]. Deficiencies in how VaR is implemented at a particular firm should not be confused with limitations of VaR itself[4].

But indeed there are limitations. When computed according to current best practices, VaR is most applicable for estimating short-term market volatilities under "normal" market conditions. These techniques are based on over a decade of well-tested research demonstrating that in most circumstances recent market movements are indeed a good predictor of future short-term volatility. VaR models have seen tremendous success in a wide range of applications including portfolio construction, multi-asset-class aggregation, revealing unexpected bets, investor communication, the extension of margin, and general transparency.

As such, VaR has become an essential part of risk management, and when properly integrated into an overall investment process it provides an

[3] The marketplace is rife with common fallacies about VaR due to poor implementations. When VaR first became popular in the mid 1990's computing power limited how accurately instruments, especially derivatives, could be modeled. Approximations that relied on the use of so-called normal distributions (bell-shaped curves) were often required. Also, the amount of market data that could be used, and the frequency at which this data was updated, was limited by technical and mathematical challenges resulting in further approximations. However, by the early part of this decade many of these challenges were overcome and today's simulation techniques do not rely on normal distributions and are not restricted by limited data. Unfortunately many institutions with older implementations still use somewhat outdated and approximate methods that do a poor job in estimating the risk of multi-asset, derivative-heavy portfolios.

[4] One fundamental criticism of VaR is that it can be "gamed" or manipulated since one number cannot by itself represent or reveal all possible "tail-loss" events. This is easily rectified by simply asking for VaR numbers at more than one level of probability, by computing the average of all losses comprising a tail event (often called conditional VaR or expected loss), or by examining the entire distribution of estimated future losses and their corresponding probabilities.

excellent framework for deploying capital in areas that properly balance risk and reward.

VaR and Banking Capital

So why did this not foretell the current crisis? First and foremost, many institutions and market participants did not perform step one correctly---they failed to correctly model how their securities would behave under changing market conditions. This failure is one of the leading causes of current crisis[5].

The second issue is where banking capital comes in. Recall that our crisis stems from long-term trends, not short-term volatility. And as mentioned, most of today's VaR techniques are only applicable for estimating potential

[5]Many institutions and market participants did not recognize nor understand how their portfolios and strategies would be affected by a fall in housing prices or a widening of credit spreads. Regulators had even less information on these effects and almost no information on how they were linked across institutions.

It could be argued that if investors had understood the nature of the mortgage-backed products they had purchased, many would not have purchased them in the first place (which would have significantly curtailed the formation of the bubble itself). If regulators had understood how CDS contracts inherently lever the credit markets they may not have allowed their unbridled expansion. And if insurance companies understood that changes in the mark-to-market values of their derivative contracts would require the posting of collateral to their counterparties many would not have entered into those deals. None of these decisions involve predicting the future or modeling fat tails. They do involve understanding the present, spending time on the details of financial instruments, and being incented to care about their risk.

Tackling these significant shortcomings may require new regulations regarding data availability, disclosure, and the analytical capabilities of each market participant. Central oversight of the markets themselves will be needed to monitor, and sometimes even limit, actions that could trigger systemic risk and future liquidity crises.

short-term movements in well-functioning markets[6]. But it is long-term trends and non-functioning markets that are the concerns of banking capital.

Nevertheless regulators rely on VaR as the basis for many bank capital calculations[7]. And even today they continue to recast VaR-like models in order to address VaR's perceived shortcomings[8]. We propose that it is not the model that needs to be recast but rather the question that regulators want the model to address.

POLICY-BASED BANKING CAPITAL

We believe that the foundation of banking capital is rooted in the following two questions:

[6] There is nothing endemic to VaR that limits its applicability to short-term estimates or functioning markets. However, current methodologies are optimized for those conditions and this is where most parameters have been tested for proper use. Research into new models that lengthen the prediction horizon and include factors like liquidity to account for non-functioning markets is underway. As development of these methodologies progresses we may see the domain of VaR extended into more areas of risk.

[7] One technique employed to "fix" the short-term aspect of VaR models is to utilize long-term historical data as the basis for "better" future estimates. This is a very common but dangerous practice since it both invalidates any estimates of short-term volatility (preventing proper use by risk managers trying to be reactive to rapid changes to the market) and it doesn't actually provide any better estimates of long-term trends. For a complete discussion on this and other related topics see included reference by Christopher Finger (RiskMetrics Research Monthly – April 2009) and references therein (including a March 2008 report issued by the Senior Supervisors Group on their study of how risk was implemented at a variety of large banks).

[8] See included reference by Christopher Finger (RiskMetrics Research Monthly – February 2009) containing our comments on the Basel committee's proposed Incremental Risk Charge---an extension that uses VaR for additional types of capital charges.

1) What are the adverse events that consumers, banks, and the financial system as a whole, need to be protected against?

2) What is required from our banks when those events occur?

This is not the domain of VaR. On the contrary, banking capital is more like an insurance policy designed to protect against worst-case events and their consequences. Instead of having banks report probabilities of short-term losses, banks should estimate the losses they would expect to sustain under a set of adverse conditions chosen by regulators. The question of "how much can you lose" is thus changed to "how much would you lose".

The conditions that banks are tested against should depend on what types of events policy makers decide that, in the public interest, banks should be able to withstand. In this fashion models, probabilities, simulations, and predictions are left to those making ongoing risk-reward business decisions whereas the minimum levels of capital needed to ensure a bank's survival are based on how regulators implement the broader requirements of policy makers. Perhaps one bank needs to survive a hundred-year flood whereas an orderly liquidation is all that is required for a different bank. Perhaps all banks should be able to weather a further 10% downturn in housing prices, but no bank is required to survive a 50% default rate or a 40% unemployment rate---not because these are highly improbable, but because

policy makers decide that this is too onerous a burden to expect a bank to bear[9].

To summarize, we believe that key differences between the needs of risk management and banking capital suggest different solutions are required. And in doing so each field can separately develop to meet the ever-expanding array of challenges we face today.

[9] The recent stress-tests conducted on banks by the Federal Reserve is an excellent example of how policy, as opposed to probability, can help set capital requirements. This should not diminish the role of simulations and the use of models to explore possibilities and uncover unexpected relationships, but this should be a guide of what the future may bring, not a prediction of what it will (or will not) bring.

EXHIBIT 1: VAR AND THE PROFIT AND LOSS CURVE