Comments on CEBS Guidelines on the implementation, validation and assessment of Advanced Measurement (AMA) and Internal Ratings Based (IRB) Approaches Issued on January 20, 2006

February 16, 2006

OpRisk Advisory
One Stamford Plaza
263 Tresser Boulevard, 9th Floor
Stamford, CT 06901
www.opriskadvisory.com

Introduction

We have prepared this formal response to draw attention to the fact that there are several key areas where we believe the proposed CEBS Guidelines on the implementation, validation and assessment of Advanced Measurement (AMA) and Internal Ratings Based (IRB) Approaches issued on January 20, 2006 (CEBS Guidelines) may be insufficient or inappropriate, specifically:

- **The Definition of Risk:** The CEBS Guidelines do not explicitly define the term *risk*. This is problematic because much of the banking industry has unknowingly adopted a flawed definition of risk. Any Operational Risk Management (ORM) framework built on this flawed definition will add little value. It could also leave bank employees thinking that ORM is a false science, a waste of resources and nothing more than a meaningless, regulatory driven compliance exercise. This could seriously impact the success of Basel II.

- **The Definition of Expected Loss:** The lack of specificity in the proposed definition of the term *expected loss* is a not appropriate. For example, using the median in place of the mean is technically incorrect and will lead to confusion in the future. Adopting this approach could also lead to systematic mispricing, overly optimistic risk adjusted profitability estimates and bad investment decision making. It may also increase the level of systemic risk in the banking industry.

- **External Data Uses:** Historical loss data has certain valid uses and many misuses. Some uses of external loss data which are deemed to be appropriate in the CEBS Guidelines are theoretically invalid, impractical and can cause results to vary by a factor of 1000. Because most banks will do what they have an incentive to do and will not do what they don't have an incentive to do, most banks will probably not adopt higher standards unless the CEBS requires them to do so. Consequently, unless this takes place, it is unlikely that AMA modeling will rise to the level of science; and the view that LDA modeling is a process with little substance will remain the consensus opinion in the industry. This is a completely unnecessary outcome, because objective, scientific methods for using external data have been available for years and there is really no need for banks to continue to rely on unscientific methods.

- **Confidence Intervals:** We wholeheartedly support the use of confidence intervals around the model results. However, we recommend that the confidence intervals not be based just on mechanical aspects of the modeling process, but also on the best case and worst case results that could be calculated by varying any assumptions or weights based on "expert judgment." Confidence intervals based on the weights and assumptions are the best way of ensuring that data and science, not subjective judgment and manipulation, drive the results.

One of the main purposes of Basel II is to improve ORM practices. But this will only happen if banks make a genuine effort to measure and manage operational risk through legitimate means. Unless the above-mentioned issues are comprehensively addressed it is unlikely that will happen. Unless banks start seeing value in ORM they will stop investing

in their ORM programs. To some extent this has already started happening and the window of opportunity is closing.

The Definition of Risk

The term risk is used throughout the industry, however, it does not appear to be explicitly defined anywhere in the CEBS Guidelines. This is problematic because several different interpretations of this term are currently in use in the banking industry:

I. In casual conversation, the term risk is used to describe an incident, such as a fraud.
II. In formal usage, risk is a metric used to describe the uncertainty, specifically the negative variance from the mean.
III. In the audit view, risk is also a metric, but is commonly represented as the product of likelihood and impact, which is somewhat similar to the mean.

These three definitions are inconsistent with one another. In our view, Definition II is the only appropriate use of the term. To explain why we offer the following example:

Consider the following three investments and their associated risk-and-return information:

Investment A: Guaranteed return of 10%.
Investment B: 50% probability of a 0% gain; 50% probability of a 20% gain
Investment C: 50% probability of a 10% loss; 50% probability of a 30% gain

Which investment has the highest mean return?

If you sum up the probability-weighted returns, you can calculate that all three investments have the same average or expected return, which is 10%.

Which investment has the most risk?

We all recognize that investment A, because it offers a guaranteed return of 10%, has no risk. Investment B has no chance of a loss. Its worst-case outcome is a break-even position, but it offers a 50% chance of a return that is below the mean return. Therefore, investment B has some risk. Lastly, investment C, which has the largest negative variance (–10% in absolute terms and –20% from the mean return), has the most risk.

Hence we can see that risk represents the level of uncertainty with respect to an adverse consequence – not the adverse consequence itself – and the adverse consequence need not be a loss[1].

How much risk is there in each investment?

There is not enough information to answer this question. Risk cannot be measured in absolute terms without first specifying a probability level (for example, 99%). The

[1] In ORM we are only concerned with the risk of loss.

probability level, which is also referred to as a confidence level (see next section), can be used to express risk tolerance in monetary terms.

Which investment is the best investment?

There is not enough information to answer this question. It is important to recognize that risk is neither inherently good nor bad.

A risk-neutral person ignores variance. He or she evaluates investments purely on the basis of expected outcomes – irrespective of the level of uncertainty associated with these potential outcomes. Since all three investments offer the same average (expected) return of 10%, a risk-neutral person would regard all three investments to be of equal value.

A risk lover would prefer investment C. In fact, he or she would be willing to pay a premium for an investment that offers the potential for a 30% gain, which is 20% in excess of the mean return.

A risk-averse person would choose investment A because it offers the same expected return as the other investments, but with less risk – in fact, none at all. Because the majority of people and financial institutions are risk-averse, they demand higher levels of return for higher levels of risk. This explains why riskier (more volatile) investments, when priced accurately, pay higher expected returns.

In summary, risk is not a type of incident, it is a measure. It describes a level of negative variance from the mean. Only where there is certainty is there no risk[2].

Why is this so important an issue? It is because an alternate definition, specifically Definition III, has been endorsed by COSO ERM and many banks have adopted this definition. By adopting Definition III they have implemented operational risk management (ORM) programs that are not compliant with Basel II and add very little value.

For example: under this definition, a 10% likelihood and a $10,000 hypothesized impact would give you $1,000 worth of risk. However, this does not actually give you the level of risk. Instead it gives you the probability weighted (expected) damage from a single hypothetical incident.

The difference between Definition III (the traditional ORM view of risk) and Definition II (the modern ORM conception) is shown in the figure below. As can be seen, under the traditional ORM approach, high risk means high-likelihood and high-impact, not low-likelihood and high-impact. This is a problem, because a business environment characterized by such catastrophic turbulence could never exist. (So under traditional ORM, no business could ever be described as being high risk.)

[2] We generally do not describe risk in terms of positive outcomes. For example, we do not speak about the risk of a gain, we speak about the opportunity for gain.

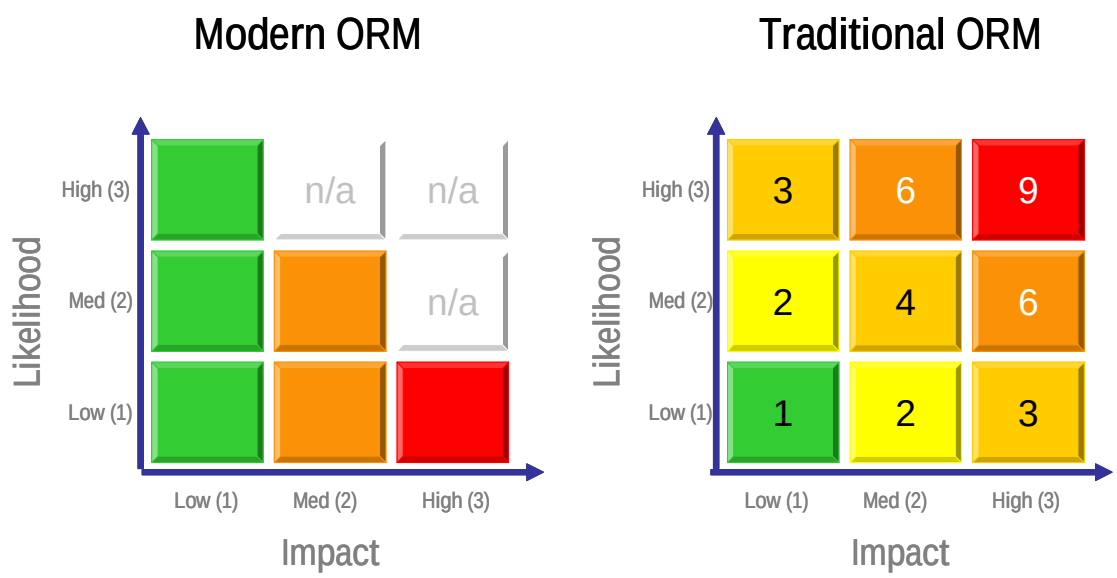

Traditional ORM risk assessment programs put managers in an awkward position. To understand why, consider this example. It is widely known that unauthorized trading is a very significant "risk". But, because unauthorized trading is driven by large, infrequent losses, the natural place to classify this risk is in the bottom right corner of the chart shown in the figure above, representing a low-likelihood and high-impact event. However, one can also see that answering correctly produces the wrong results: a score of 3 out of 9, which represents low to moderate risk. In order for unauthorized trading to be classified as a high risk, respondents must answer incorrectly and falsely classify unauthorized trading as high-likelihood and high-impact.

This flawed question places respondents on the horns of a dilemma. Should they tell the truth, or answer untruthfully to ensure the results are consistent with reality? Some may opt for the former; others the latter. No matter what happens, managers who have gone through this sort of exercise come away believing that ORM is a false science and a complete waste of time and resources. This explains to a large extent why so many banks are reluctant to invest any money in ORM. Unless the CEBS takes prompt action there is a danger that banks will permanently begin to view ORM as a meaningless compliance exercise rather that an opportunity to improve management practices.

Likelihood-impact analysis was developed by the accounting profession to identify issues – control weaknesses, not risks – in a firm's business processes. The goal was to identify the issues that could prevent a business from meeting its stated objectives[3]. And, the logical method for assessing potential damage was likelihood and impact analysis. Because this was very early in the evolutionary process, many people confused estimated damage with risk. Over time, as auditors worldwide began using this methodology, this flawed conception of risk gained broad acceptance as the standard for industry best practices.

[3] Operational risk is not the risk of a failure to meet one's business objectives. It is the risk of operational loss.

Traditional ORM uses likelihood-impact analysis to address *individual hypothetical issues/incidents*. Modern ORM uses frequency and severity distributions to evaluate risk for *general classes of events.*

Likelihood and frequency mean very different things. Again, the term likelihood is used in conjunction with an incident while the term frequency is used in association with a class of events. A frequency distribution is a probability distribution used in actuarial science. The frequency distribution shows the different probabilities (likelihoods) associated with the numbers of events that could occur during one time period. When people speak of frequency as a discrete value they are generally referring to the mean value of a frequency distribution.

Likelihood (probability) is also a component of any severity distribution. In fact, the severity distribution shows the different likelihood and impact combinations for a given class of events. In a severity distribution the higher likelihoods necessarily relate to lower impacts.

Those who don't understand the subtle differences in the meanings in these terms are generally unaware of the fact that while a high-likelihood/high-impact situation can exist, a high-likelihood/high-impact class of events cannot. The misunderstanding and misuse of these terms is a major source of the confusion in the industry.

Consider this example, suppose you are walking near the train tracks, and there is a 90% likelihood of your being hit by a train. If you estimate your value to your company at $10 million then you clearly have a high-likelihood/high-impact situation. But this situation represents a specific hypothetical scenario/incident, not a class of events. And in any case, the product of likelihood and impact (90% x $10 million = $9 million) is not the risk; instead, it is the estimated (probability-weighted) damage from the hypothesized incident. Going one step further, if the likelihood reached 100% (because 100% likelihood means certainty), the risk would become zero.

The paper on *Sound Practices for the Management and Supervision of Operational Risk* (Sound Practices Paper) published by the Basel Committee on Banking Supervision (Basel Committee) in February 2003, states unequivocally in principles four, five and six that banks must assess and monitor their operational risks and other risk-relevant information. Compliance with the principles specified in the Sound Practices Paper is mandatory for all banks – even those intending to comply only with the basic indicator approach (BIA), the minimum level of compliance under Basel II. Since the product of likelihood and impact is not risk – and, in fact is completely unrelated to risk – one must conclude that banks that use likelihood-impact analysis as a means of risk assessment cannot be found to be in compliance with the standards prescribed by the Basel Committee in the sound practices paper and therefore fail to meet the minimum requirements under Basel II, including the minimum requirements for the BIA.

Principal 10 of the Sound Practices Paper requires banks to publicly disclose such information to market participants so that they have full knowledge of the banks' risk management practices and capabilities. We recommend that the CEBS require all banks to show clear evidence that they follow the legitimate (Basel II) definition of risk. Beyond

perfunctory compliance, such evidence should pass scrutiny under the "use test" which should be validated during on-site, regulatory examinations.

The Definition of Expected Loss

The CEBS Guidelines focus on expected loss (EL) from a capital perspective. While it may be true that from a capital standpoint it does not matter how one interprets EL as long as the sum of EL and unexpected loss (UL) is equivalent to the loss value at the 99.9% level. However, there are many other issues to consider here, some with important public policy consequences.

The term EL has its roots in probability theory where the term "expected value" is defined as the arithmetic mean of a distribution. To use the term EL in a different way will lead to confusion. First of all, because the term risk represents a negative variance from the mean, defining EL as something other than the mean (such as the median) will distort the meaning of the UL (the risk). (We will then have to define risk as a negative variance from the median.) A similar misuse of the term occurred many years ago when the accounting industry defined risk as the probability of likelihood and impact. Once a flawed definition becomes the standard it is very hard to replace.

EL has important practical applications. Since the EL is the amount of money a business loses on average in a year, it is also the amount a business should budget to cover its annual cost of operational failure. Because the cost of operational failure is part of product cost, an accurate estimate of the EL is essential for proper product pricing. The EL feeds into calculations of "accounting" profitability as well as risk adjusted return (economic profit).

One of the major benefits of Basel II is that it made us realize that in operational risk, because the severity loss distributions tend to be heavily skewed, it is not possible to accurately estimate the EL without a good VaR model[4]. Consider this case, suppose you have two businesses that have the same median loss, but one has a heavily skewed distribution and has a mean loss that is twice as high. Under risk neutral pricing, if cost estimates are based on the median, both business will appear to be equally profitable. However, if product pricing were based on the mean, as it ought to be, then the latter business would be correctly observed to be less profitable.

If banks estimate the cost of operational failure based on the median loss they will be systematically underestimating their costs and over-estimating their profits and risk adjusted returns. Now one could argue that because the CEBS Guidelines do not require that banks use the median value for the EL this will not necessarily happen. In fact, one could argue that responsible organizations will continue to price products using cost estimates based on the mean. This argument ignores the many lessons we have learned

[4] One technical problem associated with calculating the EL is that sometimes the rate of decay of the tail of the theoretical distribution is so low that the distribution does not converge to zero. In such cases the mean is unbounded and, therefore, infinite. One practical solution to this problem could be to use the winsorized mean instead of the theoretical mean. The winsorized mean is calculated by specifying an upper bound on the theoretical distribution. This upper bound should not be determined arbitrarily. One reasonable approach could be to cap the loss distribution with the equity book value of the bank.

from the study of human behavior, game theory and macro-economics. Because people tend not to factor rare events into their day-to-day decision making activities, it is more than likely that at least one bank will base product cost on the median value. Other banks will have to follow suit or else they will be at a competitive disadvantage. Thus, the entire industry will soon be systematically underestimating cost and overestimating profit. This could not only lead to bad investment decision making it could also cause systemic risk in the banking industry. Since operational risk is perhaps the largest risk banks face, this could have very serious implications.

The real purpose of Basel II is to improve management practices. Redefining EL as something other than the arithmetic mean does not promote better management and, on the contrary, can lead to the opposite result.

External Data Uses

Some organizations actually "cherry-pick" losses from external data, or worse, "generate" scenario loss data, which they incorporate into their internal severity data set, to "fill in" the missing spaces, particularly in the tail region. This unscientific process has no factual basis and can cause the VAR results to vary by a factor of 1000.

External loss data is essential for operational risk modeling, but incorporating external data into the modeling process requires an objective, scientific approach. Directly combining internal and external data violates one of the fundamental precepts of operational risk modeling because loss data has meaning only in the context of the distribution from which it is drawn. A loss data point contains two integrally connected pieces of information (for severity) the loss magnitude and its relative probability with respect to the other losses in that distribution. For example if there are 10 losses at the 1,000,000 level and 100 losses at the 100,000 level, then the data tells you that the probability of a 100,000 loss is ten times higher than the probability of a 1,000,000 loss in that specific distribution. Extracting a loss data point from its original data set causes it to lose all informational value.

The CEBS Guidelines allow banks to select individual "relevant" external data for incorporation into the internal data set after "scaling" and "adjusting" for controls. This is an impossible task. In addition to the fact that it is not humanly possible to do this, it completely misstates the problem. Modeling is about analyzing data sets not manipulating data points. Individual data points carry no informational value, only data sets carry value. Even if one were to want to scale the data points, it would become necessarily to scale the entire data set. And of course, unless one knows exactly how to scale the data, it is much more likely that scaling process will do more harm than good. Generating scenario data has the same problems as extracting individual relevant external data points.

Finally, given the fact that there is no right way of manually selecting and scaling individual loss data points, any methodology that is based on such spurious techniques will necessarily degenerate into an unrestricted data manipulation exercise, where the objective is to come up with exactly the right combination of data and appropriate weights to force a result in line with the regulators expectations. This type of modeling adds absolutely no value. It makes absolutely no sense for banks to hire expensive resources if

all they are going to do is spend their time engaged in a meaningless compliance exercise. After all, what is the value in demonstrating to the regulators that is it possible to manipulate the data and assumptions in a manner that will bring about a specific result.

There are, and have been for some time, objective and scientific methods of using external data. Since these options are available, the CEBS should require banks to adopt legitimate methods that may add value, rather than follow illegitimate methods that add absolutely no value. This is very important, because many banks now believe that these spurious approaches are as good as operational risk modeling can get. Banks that are considering applying for the AMA approach may choose not to more forward because they believe this is a waste of time and resources.

Confidence Intervals

The CEBS guidelines currently recommend that banks estimate confidence intervals around their VaR results. While we strongly support the idea of confidence intervals, we believe that confidence intervals based on mechanical aspects of the model are not sufficient. The main sources of variation in results are the weights and assumptions that are part of the model hypothesis. To encourage sounder thinking in this area, the CEBS should require that, under the Pillar III requirements of Basel II, banks disclose not just both their EL and UL estimates, but also the surrounding confidence intervals. These confidence intervals should represent the minimum and maximum values that could be calculated by varying any weights and assumptions based on "expert opinion." In order to ensure these confidence intervals have been correctly estimated these limits should be stress-tested by the regulators.

Please send any comments to:

ali.samad-khan@opriskadvisory.com
OpRisk Advisory - USA (Head Office)
One Stamford Plaza
263 Tresser Boulevard, 9th Floor
Stamford, CT 06901
UNITED STATES
Telephone: +1 203 564 1990
Facsimile: +1 203 322 8364

armin.rheinbay@opriskadvisory.com
OpRisk Advisory - Switzerland
Seefeldstrasse 69
8008 Zurich
SWITZERLAND
Telephone: +41 (0) 43 488 37 69
Facsimile: +41 (0) 43 488 35 00

stephane.le-blevec@opriskadvisory.com
OpRisk Advisory - France
12-14, Rond Point des Champs Elysées
75008 Paris
FRANCE
Telephone: +33 (0) 1 53 53 16 07
Facsimile: +33 (0) 1 53 53 14 00