An LDA-Based Advanced Measurement Approach for the Measurement of Operational Risk

Ideas, Issues and Emerging Practices

Industry Technical Working Group on Operational Risk

May 29th, 2003

ABN AMRO
Banca Intesa
BNP Paribas
BMO Financial Group
Crédit Lyonnais
Citigroup
Deutsche Bank
ING
JP Morgan Chase
RBC Financial Group
Royal Bank of Scotland
San Paolo IMI
Sumitomo Mitsui BC

DRAFT DOCUMENT FOR DISCUSSION PURPOSES

I. INTRODUCTION

Objectives:

With the anticipated publishing of the new Basel Accord II by December 2003, and expectations for full implementation by the end of 2006, financial institutions around the world are increasing their focus on ensuring they are prepared. In the operational risk arena, the Accord permits an unprecedented amount of flexibility in the approach used to assess capital requirements, albeit within the context of strict qualifying criteria. This is a wise concession by regulators, recognizing that while there is no established measurement methodology agreed within the industry today, there is a huge interest in developing a risk sensitive and transparent quantitative approach to support more focussed risk management.

Many industry participants are excited by the opportunity to create a meaningful operational risk measurement approach, which is in alignment with our own risk management frameworks. Nevertheless, the creation of such an approach is not a trivial exercise, with new challenges emerging as quickly as new ideas blossom. The Industry Technical Working Group (ITWG) was founded in 2000 by a small group of operational risk practitioners from financial institutions around the world, who were interested in developing and sharing practical new ideas for the quantification of operational risk. Over the past three years these ideas have evolved to become an agreed core approach for risk measurement based on the actuarial modeling of operational losses. During this period we have identified and resolved many challenges with this core model. However, research and development continues in many areas, including the testing of the core model with real data, as well as the extension of the model to incorporate scenario analysis inputs and more “forward-looking” factors.

The purpose of this paper is to share our learnings to date. We choose to do so in a manner that focuses on emerging best practices, issues we have encountered and criteria that must be addressed, rather than prescribing a single approach for the resolution of these issues. Nevertheless, we also share ideas about solutions, where these have been developed.

The paper has been structured around the four basic elements of an Advanced Measurement Approach (AMA) which have been described by regulators in the narrowly circulated draft “Operational Risk - Rules Language” paper released last Fall (paragraph 21 d):

- Internal data
- External data
- Scenario Analysis
- Factors reflecting the business environment and internal control systems

As we began work for this paper, these elements seem to naturally split into two groups:

- internal and external data capture and modeling, and
- implementation and integration of scenario analysis and factors reflecting the business environment and internal control systems.

For each of these areas we have tried to articulate:

- Emerging "best" practices for the capture and integration of each element – these are practices which have been implemented in many banks, and where there is a high degree of agreement as to what is required and what works;
- Key challenges still facing practitioners, with some criteria for resolution of these issues, as well as potential solutions

Underlying Assumptions

The Industry Technical Working Group shares a common view that loss data should really be the *foundation* of an LDA-based AMA approach. We believe that loss data is the most objective risk indicator currently available, which is also reflective of the unique risk profile of each financial institution. It is for this reason that ITWG banks are all undertaking operational risk loss data collection. It is not just to meet regulatory requirements, but also to develop one of the most important sources of operational risk management information.

We acknowledge that internal loss data also has some inherent weaknesses as a foundation for risk exposure measurement, including:

- Loss data is a "backwards-looking" measure, which means it will not immediately capture changes to the risk and control environment
- Loss data is not always available in sufficient quantities in any one financial institution to permit a reasonable assessment of exposure, particularly in terms of assessing the risk of unexpected losses.

These weaknesses can be addressed in a variety of ways, including the use of statistical modeling techniques, as well as the integration of the other AMA elements, including external data, scenario analysis and factors reflective of the external risk and internal control environments, all of which are discussed in the next chapters.

However, the reader should be aware that this paper has been prepared on the assumption that loss data will be used as the primary input to creation of a loss distribution, which will, in turn, be the foundation for the AMA approach. All the signatories also agree that this assumption still permits a wide degree of variation in the emphasis placed on each AMA element, and also that the weighting of the various elements may change over time as we achieve more robust data collection and gain more experience with operational risk measurement.

Attribution of Ideas

This paper has been prepared by a group of operational risk practitioners from the financial institutions listed on the title page. Much of the information presented herein represents concepts that are in actual "production" within some of our organizations, and we have tried to identify these wherever possible. However, much of the content of this paper is composed of the thoughts and ideas of the individual participants, which may not or may not reflect the position of their financial institutions.

II. INTERNAL AND EXTERNAL DATA

As noted in the Introduction, the ITWG shares a common view that loss data is the most sound basis of an AMA approach, as it is the most objective risk indicator currently available. Ideally, given a relatively stable environment, we would rely strictly on internal loss data, which is reflective of the unique risk profile of each financial institution. However, we recognize that even with perfect data collection processes, there will be some areas of the business that may never generate sufficient internal data to permit a comprehensive understanding of the risk profile. It is in these cases that we believe external loss data can shed light.

Further, we believe that the best way to assess internal and external loss data for purposes of estimating the amount of operational risk capital required, is to use standard statistical and actuarial modeling techniques. This approach, called the Loss Distribution Approach, is in varying stages of implementation at most of the ITWG financial institutions.

Since the ITWG first articulated this approach in the summer of 2001, we have all continued to conduct research and development using our own data as it has begun to accumulate. Collectively we have gained a significant amount of knowledge regarding emerging best practices, and the key challenges and issues to be resolved. The purpose of this section is to provide a broad outline of some of these new practices, including the Loss Distribution Approach, internal data collection, and also current options for sharing and obtaining external loss data. We will also share some thoughts on the "burning issues" in the use of internal and external data to develop a meaningful loss distribution, including how to determine an appropriate loss capture threshold, how to determine when there is sufficient loss data to build a credible distribution, and how to incorporate external loss data.

Overview of the Loss Distribution Approach (LDA)

The fundamental premise underlying LDA is that each firm's operational losses are a reflection of its underlying operational risk exposure. A well-managed firm will have a well established "risk culture", supported by strong operational risk management processes and tools, as well as good controls. These mechanisms serve to minimize the organization's operational losses – both expected and unexpected. Conversely, a firm that is not focussed on strong operational risk management may tend to experience higher losses.

The LDA uses standard actuarial techniques to model the behaviour of a firm's operational losses through frequency and severity estimation to produce an objective estimate of both expected and unexpected losses. This estimation is considered to be a reasonable 'baseline' indicator of the riskiness of the firm.

For purposes of this estimation, we define an operational loss to be *the amount charged to the Profit & Loss statement net of recoveries in accordance with GAAP*[1]*, in the resolution of an operational risk event*. In other words we are only capturing 'out of pocket' costs, revenue reversals, and asset write-downs as formal operational losses. This narrow definition makes it easier to get a consistent and objective measurement of losses. Near misses, opportunity costs and contingent liabilities are excluded unless recognized by GAAP. Similarly, time spent by staff during regular hours identifying and rectifying a problem is not considered a loss, although many financial institutions may choose to collect this information separately, for management purposes.

The first step to generating meaningful loss distributions, is to organize loss data into categories of losses and business activities, which share the same basic risk profile or behaviour patterns. We expect that fraud losses in the Cards business will share a unique loss distribution, which may be quite different from employee claims in the Investment Banking business. If all losses are lumped together it may be difficult to discern a pattern, whereas if they are separated it becomes

[1] GAAP - Generally Accepted Accounting Principles

easier to describe a unique risk profile and be confident that it is a good picture of potential exposure. The Basel Committee has specified a standard matrix of risk types and business lines to facilitate data collection and validation across the various AMA approaches. The ITWG expects that firms using LDA will be required to map their loss data to the standard matrix, and prove that they have accounted for losses in all aspects of their operations, without being further restricted as to how they actually model the data. In other words, any given firm may choose to collapse or expand the cells in the matrix for purposes of building a specific loss distribution.

Once loss data has been collected, sorted and scrubbed, LDA involves modeling the loss severity and loss frequency distributions separately, and then combining these distributions via Monte Carlo simulation or other statistical techniques to form a total loss distribution for each loss type/business activity combination, for a given time horizon, as illustrated below. In the case of a Monte Carlo simulation, the first step is to choose a business unit/loss type combination, and draw a random sample from its loss frequency distribution. For example, one such selection might be a frequency of five events for a given loss type and time bucket. This value is then used to determine the number of events to be randomly drawn from the corresponding severity distribution. For example, we might simulate five events of size 100, 12503, 8938, 1030692 and 2211. These severity samples are then added together to generate a point on the total loss distribution.

The next step is to 'fit' the distribution of 'observed' total loss points to a curve, which best describes the underlying pattern of total loss occurrences. It is this curve that will allow extrapolation from the data points to determine the likely amount of total maximum losses or minimum capital required at any given percentile. To establish this curve, a standard statistical distribution type is selected from a well-established set of such distributions. To assess the relationship between the observed events and the estimated distribution i.e. to ensure that the selected distribution is a reasonable fit, standard statistical techniques are used, such as Parson's Chi-Square test, and the Kolmogorov-Smirnov test.

In practice, the biggest challenge is selecting the distribution that best fits the *tail* of the observed data, due to the inherent scarcity of low frequency, high impact operational loss events. One of the key objectives in the modeling process is to ensure that the tail estimate is *relatively* robust and stable, because this is where the estimate for required capital is derived. A firm should not be comfortable with the result, if this estimate shows a high degree of variance from one simulation to the next.

The loss distribution generated in the above fashion represents the range of possible total operational losses associated with that particular loss type/business activity, for a predetermined period. The distribution is then used to determine the level of capital required at a desired percentile. The choice of percentile is an expression of the firm's risk appetite – the higher it is, the more capital the firm is holding to protect itself from unexpected losses. Thus, if the distribution reflects a time horizon of one year, and the 99.9 percentile, a firm will be able to determine the amount of capital required to protect itself from a total annual loss so high that it would only be expected to occur once in 1000 years.

Figure 1

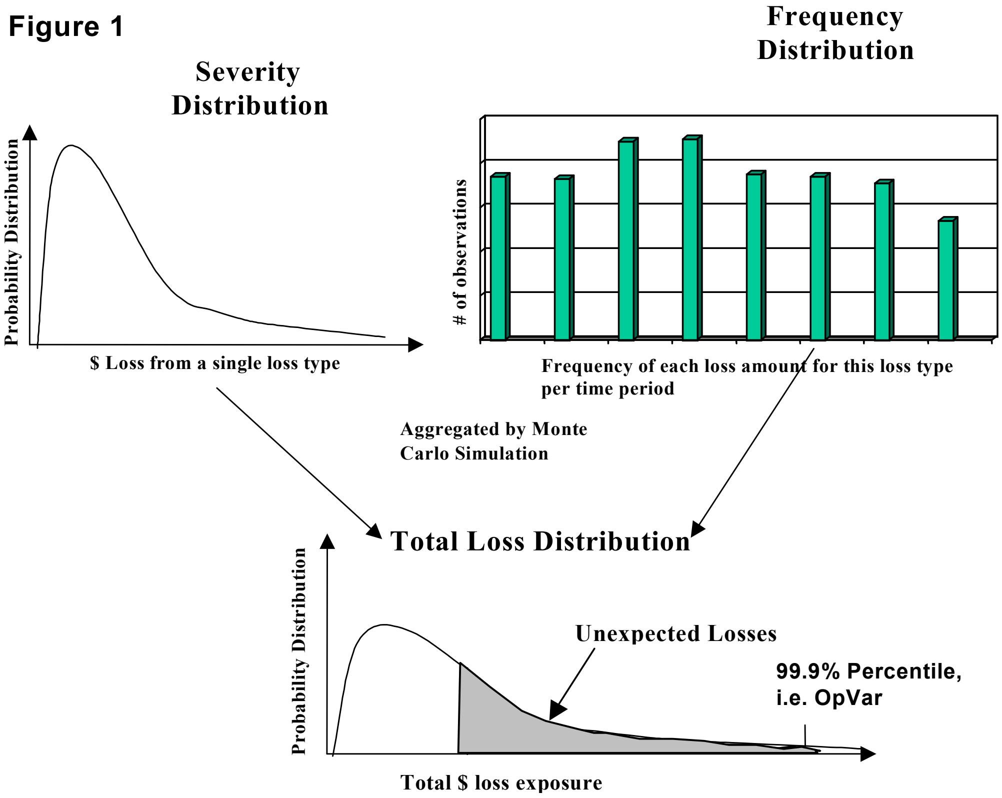

Regulatory capital will be determined using the above methodology, but with the time horizon and percentile likely prescribed by regulators. The Basel CP3 document (Consultative Paper) issued in April this year suggests supervisors will set a one year time horizon and a percentile equivalent to that used for credit and market risk assessment. Capital is required for "unexpected" losses only, *as long as* the bank can demonstrate that it has adequately captured "expected" losses in its internal business practices. Assuming this is the case, a firm would simply read the total loss amount on the curve at the appropriate percentile, and then subtract the mean or expected losses to determine the capital required.

The overall capital charge for the firm is calculated by aggregating the capital assessments generated in the above fashion, ideally in a way that recognizes the risk-reducing impact of less than full correlation between the risks in each of the business line/risk type combinations. In other words, the total capital required should be less than the sum of the parts because of the benefits of risk diversification across various businesses.

To summarize, the Loss Distribution Approach requires the firm to use standard statistical techniques to estimate a function that shows the probability of the full range of possible total losses that could be experienced in any given year, for each business line and loss event type. From this, the firm simply selects the amount of losses for which it wishes to protect itself, and holds a commensurate amount of capital. Clearly the fundamental building block for this approach is internal loss data, and in fact the Basel Committee has stipulated that we must begin with at least three years of internal loss data, building to a minimum level of five years of data once firms are past the initial years of implementation. In the next section we will describe what we have learned about the kind of information that needs to be collected, some of the standard data collection processes in place at many ITWG financial institutions, as well as key challenges that need to be resolved in data collection.

Internal Loss Data Collection

A clear data collection policy is an essential element for robust collection of internal loss data. This policy should clearly articulate what data is to be collected, as well as standards, roles and responsibilities for its collection. The following information is broadly representative of emerging best practices in data collection policy and practice at ITWG financial institutions.

What to collect

The scope of the data collection effort is directly related to the bank's definition of operational risk, and both should be clearly stated in the firm's operational risk policy. Assuming the definition mirrors the Basel Committee definition, the following data elements should be considered for collection, as they may be helpful, or indeed essential, in modeling loss distributions and improving operational risk management:

- Date(s) of event occurrence
- Date(s) of event discovery
- Date(s) of event write-off
- Country(ies) of event occurrence
- Organizational entity(ies) responsible for the event
- Legal vehicle(s) and country(ies) in which the loss is booked
- Regulatory and internal lines of business (level 2) which bear the loss
- Event category (level 2)
- Amount(s)of the loss(es) (local currency)
- Exchange rate(s) and/or exchange rate date(s) (if not the reporting currency)
- Recovery amount(s) and recovery date(s) and exchange rate and/or exchange rate date(s)
- Type of recovery (e.g. insurance, customer, counterparty, other)
- Indication as to whether the loss is associated with a credit or market risk loss (needed to avoid double counting), and if so, the amount attributable to the operational risk event
- Indication as to whether the loss related to a larger event (e.g. is this one loss in a larger disaster?)[2]
- Description of the event describing root cause(s) and failed/missing controls
- General ledger account number to which the loss was booked
- Person and organization to contact for follow-up

In addition to determining what information is required for each loss, it is also important to set a data collection threshold, below which no losses will be collected, or at least much less information will be collected. This particular challenge is discussed in more detail later in this paper, however, the ITWG generally advocates that the threshold should be as low as possible given the constraint of data collection costs.

[2] There may be events that unfold over time, and also events with multiple effects. An example of the former is an event where an employee steals several times over a period of months or years. In this case, individual losses may be discovered and written off at different times as the investigation continues. An example of the latter: an earthquake destroys buildings affecting several bank businesses. In designing an operational loss data collection process, a provision must be make to link these related losses together.

There are a number of common “scope” issues that arise as a financial institution begins to implement a data collection program, and to the extent possible, these should be anticipated and incorporated into the policy. Some of the common issues are discussed below:

Near Misses: While collecting near misses may be an attractive proposition from the standpoint of improving operational risk management, it does present significant challenges from a data collection perspective, particularly if the objective is to capture a complete dataset (important from a modeling standpoint). For example, if a well-defined transaction processing system ends with a manual quality check of the output (e.g. the supervisor does a quick scan for errors), and an error is caught, is that a near miss? On the one hand it could be argued that the manual check is part of the process. On the other hand, the need for someone to check the output at the end of the process could suggest an error-prone process, and there may be some concern that a quick scan by a supervisor is not the most effective control. One possible way to try and capture some of this information for management purposes, given that it is unlikely we can capture it all, is to simply flag these losses as “near misses” and make them optional rather than mandatory.

Boundary Losses: “Boundary” losses are losses that would traditionally be accounted for as credit or market risk losses but which involve an operational risk event. Examples are:

- a borrower defaults on its loan and the documentation required to perfect the bank’s security interest in the collateral was faulty, resulting in no recovery on the collateral;
- an incorrect index is selected to mark to market a trading position and when the correct index is used a loss is reflected in the bank’s books.

The most recent CP3 consultative paper released by the Basel Committee requires collection of credit-related operational losses in the loss event database for those financial institutions interested in qualifying for Advanced Measurement Approaches, as this is viewed to be important management information. However, these losses are to be excluded from capital modeling as it is assumed they are captured under Basel’s credit risk provisions. It is less clear from the most recent document as to whether it is necessary to capture market-related operational losses. Regardless of regulatory policy, clear guidance must be given by the firm to its businesses as to how to record these “operational risk” losses.

Other definitional issues: Businesses may also seek guidance as to whether refunds, accounting adjustments, or customer accommodations[3] are to be treated as operational losses for data collection purposes. Generally a good rule of thumb to follow is that a loss should only be booked when the business believes an error has been made, or is found by the courts to have done something wrong. Thus, a decision to reverse a fee simply to accommodate a valued customer who is unhappy, should not be booked as a loss. Nevertheless, there will always be some judgement required in these types of cases.

Under the Threshold Losses: Should these be collected, and if so what information should be collected (event category, line of business, aggregated amount)? If the decision is not to collect these, then the firm needs to determine how it will address regulators questions regarding provisions for expected losses. Some possible solutions for estimating expected losses are provided later in this paper.

Timing: When should an event be entered - when the reserve is established or at the time of actual write-off? When is a loss followed by a recovery a loss or a near miss?

[3] There is a dispute, the bank maintains there was no operational event on its part, but chooses to “make the customer whole”)

It is important to supplement the Operational Risk policy and its definitions with additional clarifications which emerge over time as the result of questions.

How to Collect

In an ideal world, financial institutions' accounting systems would collect all operational loss information required. Since this is a very unlikely scenario, especially in large internationally active firms, many banks have created or purchased web-based data collection systems which lend themselves to a decentralized data collection effort. Some banks, perhaps most, combine centralized data collection for some losses (e.g. insurance losses, employee health/safety, legal settlements/fines/fees) and a decentralized collection process for others (e.g. losses arising in operations.)

Centralizing input where the processing is centralized promotes completeness, and accuracy (the function inputting is typically the in-house "expert" and is familiar with the details of the loss and recoveries). Organizations may also choose to centralize input based on geographical clusters of countries. For example, all countries in Latin America may forward their data to a central person who is responsible for reviewing and inputting it. This allows a review of the quality of the data and a check to ensure all countries have reported prior to preparing regional and corporate operational loss reports. This is especially useful in the early stages of data collection but it may slow down requests for additional information as the central person will likely not be familiar with the details of each event and an additional step in the data collection process may cause errors of its own.

Where operational losses are not processed centrally, decentralized data collection may result in better information about the loss and more timely reporting. Some businesses have integrated loss data collection into their accounting process by using the web-based system to generate the write-off form required to pass the entry to the financial records. This ensures that the loss database and the financial records are "in sync" except for end-of-month timing differences (e.g. a write-off form may have been generated by the loss data collection system, but all approvals haven't been secured so it can't be entered into the financial books.) This approach eliminates the need to reconcile the database to the financial statements, which can be extremely difficult and time-consuming. When coupled with centralized functional input and widespread training as to the definition of an operational loss, this approach should result in a materially complete and robust database.

Other Issues

There remain a number of other issues around loss data collection that need to be considered, and may be handled in different ways depending on the capability and culture of the firm:

Security: Data must be secure from internal and external hacking which may require the loss data to be encrypted. The bank may also want to limit internal access to the data by establishing user specific read/write privileges to data. The bases used to grant entitlements must be sufficiently granular to reflect the desired level of protection but not so granular that granting entitlements is overly burdensome. The entitlement scheme must accommodate users' needs to view the data be geography, by business, by function, and by legal vehicle at a minimum.

Discoverability: Many banks are concerned about the potential for this confidential loss data to be "discovered" through the legal process, becoming available to individuals who might use this information against the firm. Some banks may want to establish guidelines as to what information should *not* be input into the system. For example, it may be desirable to ensure that there is no reference made to client names in the loss event database.

Impact of taxes on the magnitude of the reported loss: Ideally, a loss that is not tax deductible, should be "grossed up" to a before-tax value to be comparable to other operational losses. This raises questions as to what tax rate to use (statutory or effective) and practical concerns regarding required tax knowledge of staff reporting losses to the database.

Completeness: The Basel Committee has indicated that regulators should seek evidence to assure themselves that AMA firm's data is complete, and this will likely be a challenge for most financial institutions. Most firms have created an operational loss data collection system to supplement loss data extracted directly from accounting systems. Reconciliation between the bank's financial systems and the operational loss data collection system is extraordinarily difficult if not impossible due in large part to the fact that operational losses are booked in many different accounts, including contra revenue and salary accounts, and are frequently not the only items booked in those accounts. Potential solutions for this challenge include more structured general ledger accounts, and segregation of those losses that cannot be reconciled from those that can, so that the problem can at least be minimized.

Consistent classification by event categories and line of business: Consistent classification of losses into the regulatory business and event "hierarchies" is probably one of the biggest challenges facing firms with respect to future usability of the data. Training in this regard is obviously important. Some banks have also developed aids to promote consistent loss classification where data collection is decentralized. These include decision trees and web-based "wizards" which classify losses based on a user's answers to a series of questions.

This list of issues and unresolved questions around loss data collection is likely to grow over the near term as banks gain experience with operational loss data collection. Nevertheless, significant progress has been made and industry standards will undoubtedly evolve over time.

As financial institutions grapple with the challenges of collecting internal loss data, it becomes increasingly clear that there will always be data insufficiencies in certain business lines and risk event types. As a result, many ITWG member firms have sought sources of *external* loss data both to supplement their AMA model, as well as to support better operational risk management through review of significant loss events at similar businesses. The next section discusses the options currently available for sourcing external data as well as some things to consider in each case.

Sourcing External Data

There are a number of sources of external data, including data collected by the firm itself through its own direct efforts, as well as commercial vendors and industry data pools. We will discuss each of these in turn, together with user requirements, advantages and disadvantages.

Internal Collection Efforts

Many firms begin this search by simply devoting some internal resources to scanning the newswires and other public sources of event information. Similarly, there are often opportunities to leverage industry forums such as financial crime prevention committees, to capture external loss events. Obviously, the main advantage of collecting external data in this fashion is that it is inexpensive.

The main drawback is that it is very time consuming. Regardless of the source of the data, to be useful, the firm must attempt to collect as much information about each loss as it would for an internal loss event, which means reading through the article, often making educated guesses as to what is really going on in each story, and entering the data into the firm's database. Clearly, there will usually be information gaps for the event that will be impossible to fill. Also, data collected this way is subject to many reporting biases – the media tends to report frauds, lawsuits, and physical disasters, and the cost estimates provided cannot be relied upon for accuracy. Data collected via industry groups might be more accurate, but will be limited in terms of business and risk types. Finally, there can never be any sense of having a *complete* set of loss data from these sources. These information gaps and biases will make it more difficult to use this data effectively in an AMA model, although it may be quite helpful from an operational risk management standpoint.

Commercial Vendors

There are a limited number of commercial vendors of operational risk data. Vendors include OpRiskAnalytics and OpVantage, a subsidiary of Fitch Risk. These vendors capture data from a variety of sources, including various media sources, court records and other information sources. These commercial vendors appear to have achieved a degree of maturity and stability in their databases.

The main advantages of acquiring external data from a vendor is that it is less time-consuming, and users will probably end up with more data than they would if they source it themselves directly. Nevertheless, external vendors suffer from most of the same disadvantages already discussed in terms of biases in the data, albeit the data will probably be *relatively* more robust.

Industry Data Pools

Recently industry data pools have received more attention. One of the earliest data pools operates under the name of GOLD (Global Operational risk Loss Data) and is co-ordinated by the British Bankers Association. More recently has been the development of the ORX (Operational Risk data eXchange). A number of industry associations are considering the pooling of data at a national level.

The objective of these data pools is to provide a complete profile of loss data from each participant at a relatively low common threshold. This data is then anonymised and re-distributed to the contributing firms. Possible uses of pool data will be dictated to some degree by the rigor with which data is collected, especially in terms of completeness and accuracy. We discuss later in this section the approaches for integrating external data into an AMA. In many cases it is necessary to make some assumptions about the distribution and parameters associated with the data in any given business line and event type, which would clearly place some high demands on the quality and standards of a data pooling operation. Further, most ITWG members believe that external data should ideally be "scaled" to better fit their own unique risk profile. This demands the collection of "exposure indicators" for each participating organization for each data point (this could be done quarterly or monthly), whether they be size indicators (as will initially be the case) or risk environment indicators (the ideal for the future). Again this makes the data sharing effort more complex and challenging.

Thus, the main advantages of a data pool lie in the degree of rigour associated with the collection effort, and the completeness, accuracy and coverage of the data, and this will be of greater interest for firms who wish to qualify for AMA. Disadvantages lie in the high standards imposed on participants, who must ensure their own collection efforts make the grade, and who must also collect the exposure indicators in addition to the loss information. All this rigour, plus the additional concerns participants have regarding data security, makes data pools a potentially expensive option, and one that may take some time to materialize into a source of significant data volumes.

Having now completed a description of a high level Loss Distribution Approach together with a discussion regarding internal and external data collection, it is appropriate to focus on some of the key challenges associated with developing this portion of an AMA. We have highlighted three key challenges and make some suggestions as to how they can be handled.

Summary of key challenges in developing a Loss Distribution:

Challenge #1: Completeness Of Data

Given that the LDA-based AMA approach places emphasis on loss data as a basis for an estimation of operational risk, the ITWG has considered how banks can satisfy themselves that the data used to generate the estimates of capital requirements is sufficient. It has considered two related dimensions to this issue: the size and number of loss data points available and the thresholds used in data collection. For both of these issues the ITWG believes it has developed credible assessments and solutions.

It is worth noting that the concerns over the completeness and accuracy of loss data are of relevance to banks considering any type of AMA approach. This includes the scenario and scorecard based approaches, as whether data is used to build distributions directly, or to validate other AMA inputs, banks (and supervisors) will need to be assured that the data is credible.

The ITWG has focussed on two cases where the testing of the sufficiency of data, in terms having confidence in resulting capital estimates, is necessary. The first is where the bank has a large number of losses for a particular event type, but few large losses (e.g. credit card fraud). The second is where the bank a small number of large losses. In this second case, there are two variants: where the losses are of a predictable (limited) size (e.g. physical damage to assets is limited to the value of the asset/cost of replacement) and where they are of unpredictable size (e.g. legal liability, where a punitive award of damages is unpredictable). In each of these two basic cases, there is a danger that the LDA-based AMA calculation will be in breach of the requirements set out in CP3. Specifically, there is doubt that severe tail loss events are captured and/or that the soundness standard of 99.9% and a 1-year holding period is met.[4] Nevertheless, in each of the two situations described above (including both variants of case 2), the ITWG believes that proven statistical techniques may be employed to test the sufficiency of the data and hence the level of confidence in the resulting capital estimations.

In the case of a large number of small losses, the ITWG proposes that the properties of the Maximum Likelihood Estimator (MLE) may be used to give assurance that the resulting capital calculations are credible. The MLE technique is suited to situations where the number of data points is large. In the situation of few large losses, the MLE becomes less reliable as a basis for

[4] Consultative Document: The New Basel Capital Accord: para. 627, *Basel Committee on Banking Supervision, April 2003*

assessing the credibility of capital estimates, particularly where the size of losses is unpredictable, and bootstrapping sampling techniques may be employed. Bootstrapping is well suited to situations where the number of data points is small and the underlying distribution assumed to be non-normal.

The choice between MLE and bootstrapping leaves open the question of when it is appropriate to apply each technique as a basis for assessing credible estimates of capital. This issue is particularly relevant for the case where there are relatively few large losses, but of predictable size. At present, the ITWG has not determined criteria for assessing where the transition exists between the use of MLE and bootstrapping, although it is continuing to work in this area and is seeking to develop guidance in this regard.

Challenge #2: Setting the Data Threshold

The second key issue with regard to data collection and sufficiency that the ITWG has considered is the setting of the threshold for data collection. CP3 sets out requirements with regard to the threshold for data collection, and notes that the choice of threshold should be appropriate and that excluded activities, both individually and in combination, have no material impact on overall risk estimates.[5] The ITWG believes it is possible to demonstrate the impact of the choice of threshold on the confidence that may be placed on capital estimates.

In theory, the choice of the threshold does not impact the ultimate capital assessment, providing there is sufficient data available to allow the bank to understand the nature of the distribution and hence the level of expected loss. Nevertheless, there are concerns that the choice of threshold can influence the sufficiency of data and hence the credibility and stability of capital estimates. In order to test the impact of the threshold on capital estimation, the bank should ensure that the fitted distribution resulting from its data sample passes a goodness of fit test. Where a loss data collection threshold is set above zero (as will invariably be the case), there are two options open to the bank to adjust its measurement methodology.

The first is to estimate the loss distribution below the threshold. This has the advantage of ensuring that an estimate of losses below the threshold is provided, which assists the bank in complying with the CP3 requirement to demonstrate 'that it has measured and accounted for its expected loss exposure'.[6] The second is simply to truncate the distribution of individual losses at the chosen threshold.

Whilst the first of the two issues described in this section (data sufficiency) is almost entirely a concern of the statisticians in a central operational risk management function, it is important to note that the choice of a loss data collection threshold is not a purely statistical issue: it will impact all businesses and units within a bank that generate loss data. The choice of threshold is therefore tied to a cost-benefit analysis of collection and, more specifically, to the level of data that local management finds useful. Therefore, it is not possible to endorse a single data collection threshold (i.e. that proposed in CP3 of €10,000) and within banks it is highly likely that different thresholds will be adopted. This is entirely reasonable, providing that the choice of threshold is not significantly impacting the credibility of capital estimates. The ITWG believes that there are techniques available to banks to satisfy themselves (and supervisors) that the

[5] Consultative Document: The New Basel Capital Accord: para. 633, *Basel Committee on Banking Supervision, April 2003*

[6] Consultative Document: The New Basel Capital Accord: para. 629 (b), *Basel Committee on Banking Supervision, April 2003*

threshold choice is reasonable and does not adversely materially impact the credibility of capital estimations.

Challenge #3: Incorporating External Data

There are many ways to incorporate external data into the calculation of operational risk capital. Here are some of the ways in which external data can be used. External data can be used to complete an incomplete internal loss data set, to modify parameters such as expected and unexpected losses derived form the internal loss data set, and to improve the quality and credibility of scenarios. External data can also be used to validate the results obtained from internal data or form benchmarking.

There are two main issues with incorporating external data. These are a) relevancy of the data and b) scalability of the data.

Relevancy:

The business and control environment, in which a particular bank operates, is of one the key drivers of the bank's exposure to operational risk, and therefore the loss experiences of other banks that operate in a significantly different environment may not be relevant to that bank. For example if a particular bank operates primarily in a geographic location which due to labor laws or the reliability of its technology infrastructure, is subject to frequent business disruptions, then supplementing its sparse internal data with external data which reflects more stable environments, would result in a significant underestimation of the bank's actual exposure. Likewise if an institution carries out very little activity in investment banking, whereas as the external data comes form banks that have significant investment banking activities, then using external data in this case would result in a significant overstatement of that bank's risk. The same distortions result by using external data that reflects significant differences in the control environment. For example, the extensive reviews conducted by regulators of the large and spectacular unauthorized trading reveal pervasive and serious weakness in internal controls. Including such loss data in arriving at the capital for a bank with very strong internal controls would result in a substantial overestimation of the required capital. All these examples show that the relevancy of each component of external data to a particular bank is very important. Some of the ways that banks have tried to incorporate only relevant external data are to segment external data in peer groups and only use data form the corresponding peer group. Another alternatives is to use expert judgment on the individual external data point to determine if that point, from the perspective of a particular bank, is an outlier, and remove those outliers. There are a variety of different approaches, each with advantages and disadvantages, with no clear superior approach. However, it remains important to recognize that each bank should develop an approach to determine the relevancy of external data and that approach should be well documented and well reasoned, and subject to periodic review.

Scalability:

Scalability refers to the fact that operational risk is dependent on the size of the bank, i.e. the scale of operations. A bigger bank is exposed to more opportunity for operational failures and therefore to a higher level of operational risk. The actual relationship between the size of the institution and the frequency and severity of losses is dependent on the measure of size and may be stronger or weaker depending on the particular operational risk category. Nevertheless as with relevancy, not taking into account the scale of operations or size of the particular bank relative to other banks included in the external data set will significantly distort the calculated operational risk for that bank. One way to deal with scalability, similar to the way relevancy is dealt with, is to define a similar size peer group and use only data from that peer group. Another way is to use regression analysis to determine the relationship between size and frequency and another relationship between severity of losses and size.

Once again there are a variety of different approaches, each with advantages and disadvantages, with no clear superior approach. However, it remains important to recognize that each bank should develop an approach for scaling external data and that the approach should be well documented and well reasoned, and subject to periodic review.

III. SCENARIO ANALYSIS

Definition

For the purposes of this paper scenario analysis is defined as the forecast of operational losses and events that cause them, based on the knowledge of business experts.

Introduction

Several ITWG members have experience in the use of scenario analysis. Their experience has found that the process can yield invaluable insight into risks and mitigation, and provide quantitative input into capital calculations that can be used with a certain degree of confidence in concert with other inputs. The primary set of issues related to scenario analysis is the inherently subjective nature of the process. However, some techniques have been identified to attempt to reduce the subjectivity and increase the repeatability. Overall, the quality of results that have been generated confirms scenario analysis as a key tool for operational risk management.

Scenario analysis has both quantitative and qualitative aspects to it. In this paper we limit the discussion to the use of scenario analysis solely for risk measurement purposes. At a macro level we define three types of scenario analysis:

1. Supplementing insufficient loss data
2. Providing a forward-looking element in the capital assessment
3. Stress testing the capital assessment

Each of these types is discussed in the following sections. The information is presented as a series of "use cases" that have been practised by ITWG member institutions. We further discuss the processes that have been employed for the collection and validation of scenario information.

Supplementing Insufficient Loss Data

The most common use to date of scenario analysis has been to provide a means of supplementing insufficient loss data in order to generate complete frequency and severity distributions that can then be modeled to generate a capital requirement.

Use Case 1

The objective here is to incorporate scenario-based losses into the internal loss database with the view that scenarios would be used to supplement the historical loss distribution, particularly at the tail, for the purpose of LDA.

A potential loss event could arise under three types of scenarios: 1. *Expected Loss* (optimistic scenario), 2. *Unexpected Serious Case Loss* (pessimistic scenario), and 3. *Unexpected Worst Case Loss* (catastrophic scenario). We record the loss frequency and severity of the potential event under each scenario type. The gathered information on loss frequency and severity would be based on

- expert opinion provided by business managers and risk managers during face-to-face interviews,
- internal sources such as internal loss databases, risk assessments, key risk

indicators, management accounts, and

- external sources such as external loss databases obtained from industry consortia or from the public domain.

In generating the potential loss events we consider the 20 (level-2) event types defined by the Basel Committee. This is one way to ensure that scenario analysis captures an exhaustive list of operational risks. One or more potential events could fall under a particular regulatory event type. For example a business manager may expect two potential events to occur under the "systems" event type, one related to an electricity blackout and one related to a system software failure.

One would expect the loss severity values for a given potential loss event to increase and the frequency values to decrease from scenario type 1 to scenario type 2 to scenario type 3. Recording the severity values is straightforward. Frequency values could be recorded in two ways the choice of which depends on the number of individual loss events of a given scenario type per year, or holding period of a different time-length. The alternatives are as follows:

- ♦ If one or more potential loss events are expected to occur within a year for a given scenario type, then frequency should be recorded in terms of the number of loss events per year.

- ♦ If less than one potential loss event is expected to occur within a year, then frequency is recorded as the number of years between successive occurrences. In this case, the frequency values are pre-defined: a potential loss event would be expected to occur every 2 to 3 years, 4 to 5 years, 6 to 10 years or more than 10 years.

Where the recorded severity and/or frequency values have been influenced by historical events (either internal and/or external), the business should also indicate the occurrence date, amount of final (or estimated) loss, as well as the name of institution involved and other relevant information.

The information described above could be summarised in the following example of a scenario output template:

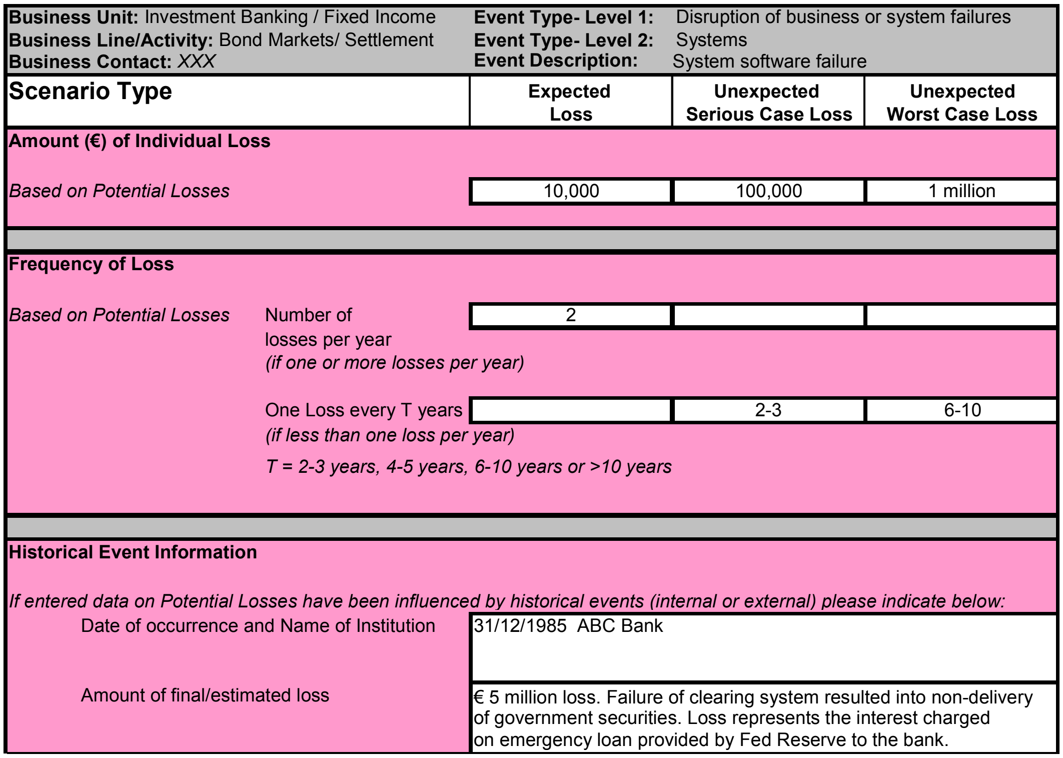

Business Unit: Investment Banking / Fixed Income	**Event Type- Level 1:** Disruption of business or system failures		
Business Line/Activity: Bond Markets/ Settlement	**Event Type- Level 2:** Systems		
Business Contact: *XXX*	**Event Description:** System software failure		
Scenario Type	**Expected Loss**	**Unexpected Serious Case Loss**	**Unexpected Worst Case Loss**
Amount (€) of Individual Loss			
Based on Potential Losses	10,000	100,000	1 million
Frequency of Loss			
Based on Potential Losses Number of losses per year *(if one or more losses per year)*	2		
One Loss every T years *(if less than one loss per year)* *T = 2-3 years, 4-5 years, 6-10 years or >10 years*		2-3	6-10
Historical Event Information			
If entered data on Potential Losses have been influenced by historical events (internal or external) please indicate below:			
Date of occurrence and Name of Institution	31/12/1985 ABC Bank		
Amount of final/estimated loss	€ 5 million loss. Failure of clearing system resulted into non-delivery of government securities. Loss represents the interest charged on emergency loan provided by Fed Reserve to the bank.		

*The fact that we are dealing with potential events means that their loss value is uncertain (hence the need to express them by three scenarios where each scenario assigns a different severity and probability to the potential event). This poses the question how we incorporate this information to an internal database of historical losses. The answer would be that we shouldn't add three data points to the internal loss database otherwise a potential event would be triple counted. We should incorporate the potential event into the database as a single data point. This means we need to calculate an average loss value of the potential event. In order to calculate the average loss value of the potential event one should use the statistical concept of expectation. This means if a potential event has loss severity X1, X2, X3 under scenario types 1,2,3 respectively and the corresponding probabilities for each scenario are p1, p2, p3 then the expected loss value of this potential event would be equal to (X1*p1) + (X2*p2) + (X3*p3). This could be seen as taking the weighted average loss of the potential event under the three scenario types, where the weights are given by the probability of each scenario (estimated by the corresponding frequency values).*

For example, consider a potential system software failure. Suppose the expected loss impact is €10,000 (under scenario type 1), the serious case unexpected loss is €100,000 (under scenario type 2) and the worst case unexpected loss is €1 million (under scenario type 3). Also suppose the frequency values attached to each scenario type result to the following probabilities: 50% for scenario type 1, 35% for scenario type 2 and 15% for scenario type 3. The expected loss value of a potential system software failure would then be equal to (10,000*0.5) +(100,000*0.35)+(10^6 *0.15) = €190,000. So the data point to be added in the loss database would be €190,000.

Having supplemented the database, one would then perform an LDA analysis assuming the loss distribution is now complete. So the next step would be to model the severity and frequency of individual losses. The simulated loss severity distribution would be combined with the simulated loss frequency distribution using a convolution technique such as Monte Carlo simulation. The operational risk capital would then be extracted from the aggregate loss distribution.

Use Case 2

The objective is to use scenario-based losses alone to estimate capital. The result may be combined subsequently with available loss data to generate a capital requirement.

The framework is the same to what has been described above. So we use the scenario analysis output as described above. Each potential event is described by the three types of scenarios, and attached to each scenario is a loss severity and frequency (or probability) value. These three scenarios can be considered as the three points of a discrete loss distribution of a particular potential event. If we have generated 30, say, potential events it means we have 30 discrete loss distributions.

We are interested in the distribution of all potential events, i.e. the aggregate loss distribution. Under the assumption that the generated potential events are independent we use a convolution technique to simulate the distribution of aggregate losses. For example, if we have 30 potential loss events, this means we have 30 distributions each made of three points, which would yield an aggregate loss distribution that consists of 3^{30} grid points. The operational risk capital would then be derived from the aggregate loss distribution.

We note here that although we use the same scenario analysis output as in Use 1 there is a marked difference between the two uses. In Use 1 we compute the expected loss for each potential event which is then incorporated into the internal loss database as an additional data point. In Use 3 we don't compute the expected loss of potential events but use instead their discrete distributions in an aggregation/convolution process for simulating the aggregate loss distribution.

Use Case 3

The objective in this case is to derive distribution parameters from the scenarios that can be combined with similar parameters from historic data to generate a capital requirement. A by-product of this approach is the calculation of capital directly from the scenarios on a stand-alone basis that can be used for validation. This is similar to the objective in Use Case 2 described above.

In this case business experts in each business line are asked to forecast a future loss profile by assessing the frequency of events that could occur within different loss ranges ("buckets"). In addition the experts are asked to estimate the maximum possible stress losses by value, and the circumstances in which they could occur. The template used to capture the results is shown below.

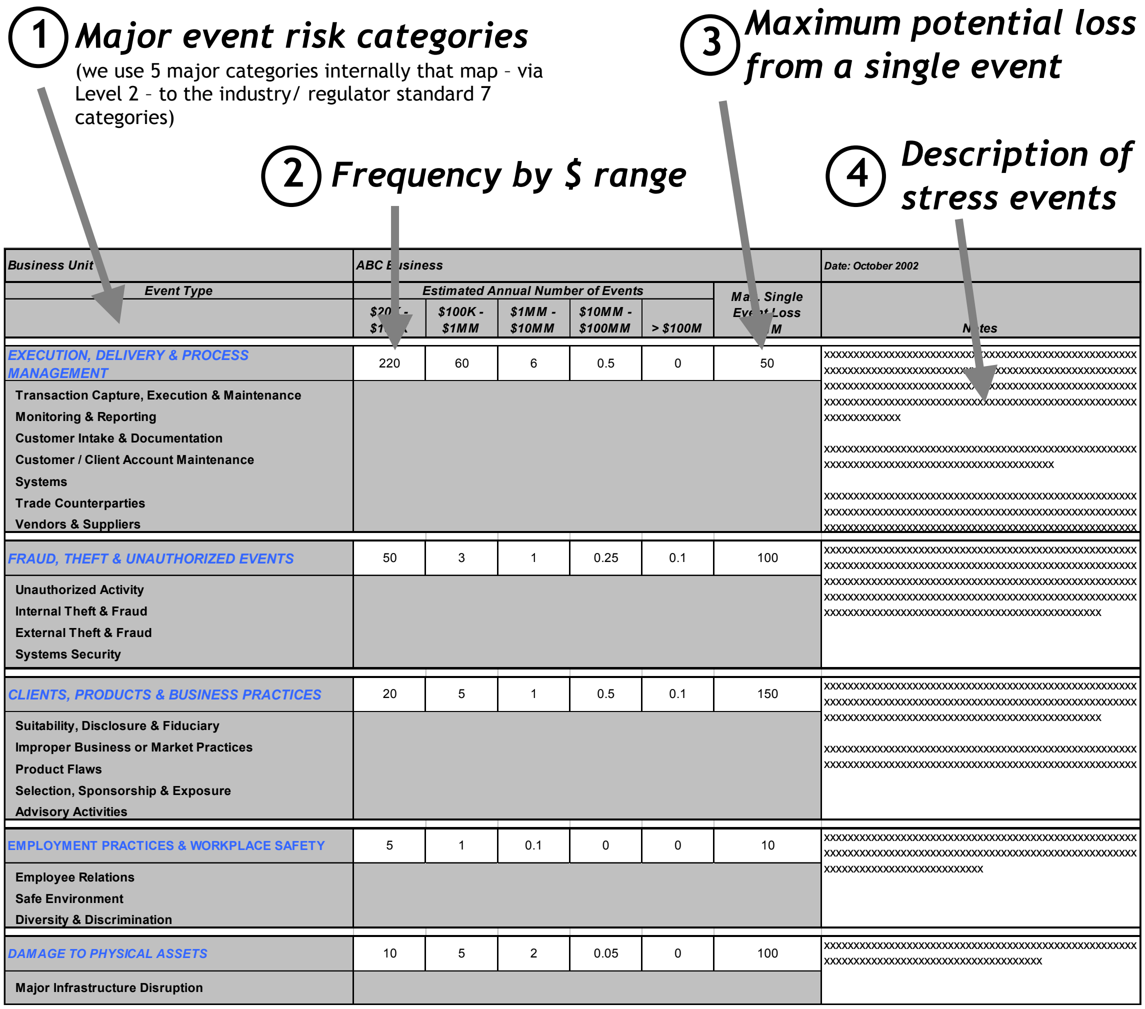
1
Major event risk categories
(we use 5 major categories internally that map - via Level 2 - to the industry/ regulator standard 7 categories)
3
Maximum potential loss from a single event
2
Frequency by $ range
4
Description of stress events

Business Unit	ABC Business						Date: October 2002
Event Type	Estimated Annual Number of Events					Max. Single Event Loss $MM	Notes
	$20K - $100K	$100K - $1MM	$1MM - $10MM	$10MM - $100MM	> $100M		
EXECUTION, DELIVERY & PROCESS MANAGEMENT	220	60	6	0.5	0	50	xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
Transaction Capture, Execution & Maintenance Monitoring & Reporting Customer Intake & Documentation Customer / Client Account Maintenance Systems Trade Counterparties Vendors & Suppliers							
FRAUD, THEFT & UNAUTHORIZED EVENTS	50	3	1	0.25	0.1	100	xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
Unauthorized Activity Internal Theft & Fraud External Theft & Fraud Systems Security							
CLIENTS, PRODUCTS & BUSINESS PRACTICES	20	5	1	0.5	0.1	150	xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
Suitability, Disclosure & Fiduciary Improper Business or Market Practices Product Flaws Selection, Sponsorship & Exposure Advisory Activities							
EMPLOYMENT PRACTICES & WORKPLACE SAFETY	5	1	0.1	0	0	10	xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxx
Employee Relations Safe Environment Diversity & Discrimination							
DAMAGE TO PHYSICAL ASSETS	10	5	2	0.05	0	100	xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
Major Infrastructure Disruption							

Based off this data, frequency and severity distributions are derived and modeled. The results in this case are combined with loss data distributions using credibility theory to generate the required capital.

Providing a forward-looking element in the capital assessment

The objective of scenario analysis in this case is to use the opinion of experts to turn a purely loss-based calculation of capital based on historic data into a more forward looking capital assessment. Although we separate this objective from the supplementation of loss data, the use cases described above can all be considered to provide this forward-looking element. In addition, we can also consider using the results to modify the parameters of a loss-based distribution.

For example, we may simply modify the frequency of event occurrence by taking the weighted average of historic losses in a given business or risk category, by estimated frequencies captured in the scenario analysis process. Similarly, the mean loss parameter from the loss data may be modified by the average loss information captured in Use Case 1. The weighting given to the parameters from the data and the scenarios reflects the degree of confidence attached to each of them.

A more technically sophisticated approach uses a Bayesian method to combine information from the scenarios to adjust the loss data parameters.

An example of combining distributions from loss data and scenarios to generate a forward-looking estimate of capital is illustrated below. The distributions are weighted according to the degree of confidence in each.

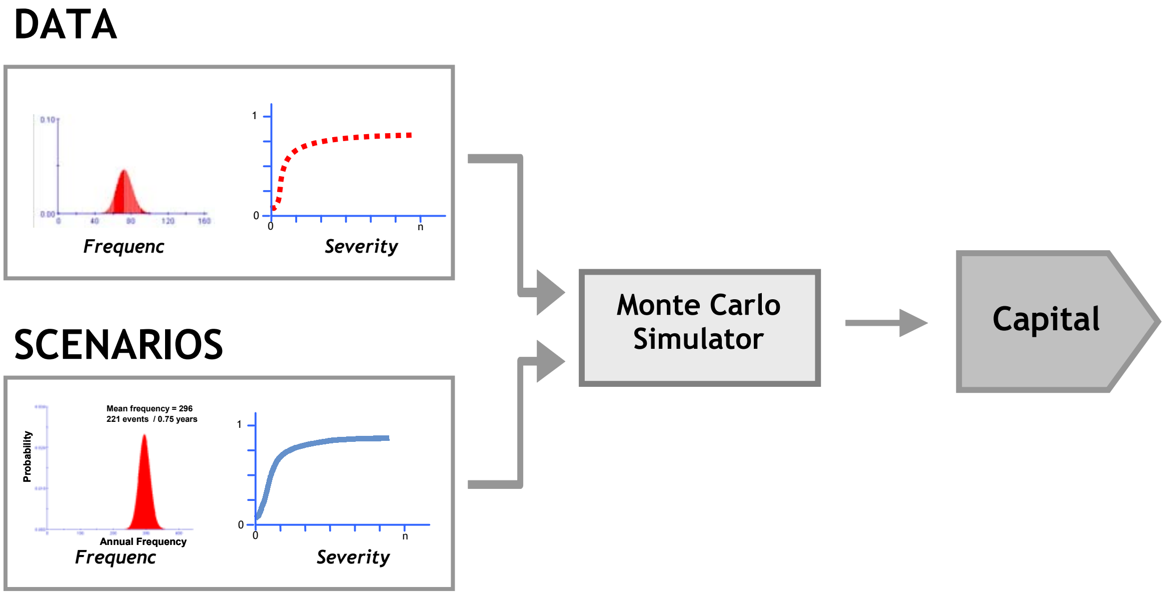

Stress testing the capital assessment

Stress testing is an instance of scenario analysis which is particularly appropriate for considering high impact/low frequency loss. Its scope is just to capture consequences of extreme potential events. There are basically two ways of performing stress testing, which are relative to two quite different tasks.

One option is to directly ask for an assessment of so-called stress scenarios. Along this line, the analysis is not different from the one for standard scenario. So the answer to the question is used as an input to the op risk capital quantification model, since the information is used to calibrate parameters of the model.

As a second option, we may use stress testing as a technique which measure robustness of already obtained results. We can label this issue as "sensibility analysis" since its aim is to quantify how model result (i.e. op risk capital) changes as the determinants vary.

The process is as follows. After we generate a model estimate we can adjust the values of its parameters and its inputs by setting them at some other level, different from the one currently taken, and see how the model output changes. For example, in the context of causal modelling we can come out with a model of the kind of

$$Y_n = \hat{\beta}_0 + \hat{\beta}_1 X_{1n} + \ldots \hat{\beta}_p X_{p\,n}$$

where Y represents annual loss, the X's are p operational risk drivers (factors) and the β's are the estimated parameters; then we can generate a stress scenario by setting one of the coefficient or one of factors equal to, let's say, the 99% quantile of its distribution. The resulting value for Y is an indication of the output level in presence of extreme conditions.

This is the most technical approach to stress testing, and note that it doesn't necessary involve expert opinion, but it can take account of it. However it is best designed for an institution that has already collect risk drivers, since the more drivers are available, the more scenarios can be constructed and the resulting risk description is fuller and more comprehensive.

Collecting Scenario Information

Assembling scenario analysis information can be considered separately as a top-down and as a bottom-up exercise. Both approaches capture business judgment. On a top down basis, business managers are asked to identify operational exposures ranging from day-to-day losses up to stress events ("what keeps you awake at night?"). The bottom-up process can start with a detailed process analysis or risk assessment, and assign probability and severity of losses due to individual failures or events. *The remainder of this section is limited to top-down assessment, based on the experience of members of the Scenario Analysis working group.*

Who is involved?

The continued use of the term "business judgment" in the context of scenario analysis implies that expert knowledge is required. A broad range of input and views will produce a more balanced result. The following roles/functions can be considered as sources of input:

- Business heads and/or senior business managers on the revenue side of the business
- Heads or senior managers of support functions
    - Operations
    - Technology

- Legal
- Insurance
- Risk management
- Internal audit
- etc.

How is data collected?

Information may be collected by a variety of means:

Interviews: One-on-one or small group interviews, typically conducted by a member of operational risk team. Questions should be open-ended and notes should be taken. In practice it is useful to have more than one member of the operational risk team present – one to ask questions and one or more to take notes.

Workshops: A series of facilitated workshops that assemble the individuals defined above for a particular business area. The same, or similar, set of questions can be used and again it is helpful to have more than one member of the operational risk tem present at the session. The workshops can incorporate an element of brainstorming to describe stress events.

Questionnaires: Data may also be conducted by questionnaire. This approach does not lend itself to open-ended questions; instead the questions must be narrowly defined and the information sought needs to be very specific, though judgment may still be involved in the answers. E.g. questions may involve risk-ranking businesses or functions. Despite their limitations, questionnaires can be useful in several instances, e.g.:

- Getting input from geographically distant individuals
- Getting a broader set of views to validate interview/workshop results

What are the inputs?

Preparation prior to interviews and workshops is necessary to generate a more consistent set of results. Preparation for workshops may include some or all of the following:

- ***Background:*** A definition of the objectives of the scenario analysis process, including a brief background, who will be attending, how the workshop will be conducted, what results are expected and how they will be used.
- ***Internal Loss Information:***
  - Internal loss statistics: whatever data is available to describe the profile of past operational losses. This may be a combination of complete data from recent data collection efforts, or anecdotal data.
  - Description of large internal losses
- ***External Loss Information:***
  - External loss statistics, from a pre-defined peer group
  - Description of large external losses sustained by members of the peer group

Preparation for interviews and questionnaires will normally be a subset of the above.

What are the outputs?

Outputs from scenario analysis may capture descriptive information in the form of defined stress scenarios and/or specific quantitative information, e.g.:

- Estimated expected losses
- Estimated “serious case” losses
- Estimated “extreme case” losses, or maximum estimated loss from a single event
- Profile of losses by frequency by size (bucketed)
- Ranking of businesses, or process, or risks etc., in order of risk or exposure

Quantitative information from the scenarios can be used as described elsewhere in this paper to include into the Loss Distribution modeling. Qualitative data can be used both to help validate modeling results, and directly linked to risk assessments.

How can repeatability of the process be ensured?

Strong efforts must be made to ensure repeatability of the process by consistent preparation, and by consistent application of the quantitative and qualitative results of the scenario process. Repeatability may be defined as the ability to repeat the scenario process for a given business at a point in the future, but with a different set of individuals performing the same functions, and generate a similar result.

Generating outputs for a given business may be the result of an iterative process of several rounds of interviews/workshops/questionnaires and validation, until a reasonable and defensible set of results is achieved.

Further considerations in the construction of loss scenarios

A common feature of this segment is the use of risk events, individual or multiple, to drive the scenarios. This segment can be further sub-divided:

- Micro focusing upon an event, an issue, a control, etc that affects an activity, or a business in a location
- Macro an event, an issue, a control, etc that affects multiple businesses in one location or a business in multiple locations
  e.g. Earthquake in Tokyo and interaction with Credit & Market Risks

<u>Input Sources for Micro</u>

Internal Risk Assessments	Business Level
	Location
	Activity
	Change e.g. New Product, Systems, Organisational

External Events experienced by competitors, firms with related exposures
Specific Control Issues e.g. Lack of BCP
Involves experts / risk specialists / business specialists / location specialists

The input sources for the micro event can be formal, for example the results of internal control self-assessments, or responses to the informal question "“What keeps you awake at night?" However, as with other aspects of the OR framework the use of the results may influence the willingness of managers or experts to participate.

The results of such micro event scenarios can be related to the activity and consideration of the risk mitigating actions, for example review the insurance coverage, or consider re-engineering a process, or outsourcing the activity, or ultimately accept the risk.

Micro Risk Event Scenarios could be extremely time consuming in terms of creation and follow-up discussion. It should be noted that the vast majority of such scenarios are unlikely to produce a risk estimate that would appear in the tail of the loss distribution.

Input Sources for Macro

Regional / Global events imagined
External events experienced by competitors or firms with related exposures
e.g. Governmental confiscation of assets, Trade Sanctions etc
Developed Regionally / Centrally using input from experts / specialists
Boundary between plausible extreme events vs. catastrophic events

The creation of Macro Risk Event Scenarios can be extremely time consuming to create (months not days) and only an order of magnitude or "quantum" of OR implications obtained. The range of these scenarios could be extensive and it is not clear how much insight they will actually provide. In particular, there is the boundary between plausible extreme events and catastrophic events to consider.

The benefit from constructing such scenarios may appear long before a risk number is produced. For example, the scenarios used to consider Business Continuity requirements do not often rely upon a loss quantum to be generated before risk mitigating actions are taken.

The Risk Event Scenarios may interact with other parts of the OR framework. For example, the result of a risk assessment may generate a scenario that is meaningful for a business or activity. In order to monitor the change in risk from period to period the firm could devise a risk indicator. This risk indicator could be a component reflecting the Internal Business Environment and Internal Control Factors. Under these circumstances the firm needs to take care that it is not double counting.

Validating Scenario Information

Validation of scenario analysis results is best accomplished by looking at several different types of validating information, both internal and external.

Internal

- *Validation against internal loss data:* Over time this represents the best method of validation of results of the scenario analysis process, at least for the expected losses and for the middle part of the distribution. The level of validation will obviously improve over time as a greater amount of time series data is collected.
- *Reasonability checks:* The scenario outputs can be viewed for reasonability across business lines, in terms of the frequency and severity of losses that have been defined. The results of modeling the scenario outputs can also be validated in like fashion.

- *Multiple inputs:* If more than one source of information has been collected for the same business e.g. workshops and questionnaires, the results can be cross-compared for reasonableness

External

- *Validation against commercially available industry loss data:* The quantitative results can be compared to peer groups from the industry. Though commercially available databases do not lend themselves to statistical comparison due to incompleteness of data sets, and scaling issues, they represent an important source of information for reasonability checks.

- *Validation against industry consortium loss data***:** Validation of scenario outputs, and more generally benchmarking of operational losses, will be increasingly practical in the near future with the establishment of industry loss data sharing consortia. The level of validation – statistical vs. reasonability – will depend on both the completeness and more importantly the quality of the consortia data. In the case where consortium data is both complete and high quality, this will allow the best external source of data for validation.