Fundamental Issues in OpRisk Management

Misconceptions about certain key concepts are causing confusion throughout the industry. By Ali Samad-Khan, Armin Rheinbay and Stephane Le Blevec


Ali Samad-Khan

ORA OpRisk Advisory

SOUND operational risk management (ORM) begins with a comprehensive understanding of certain fundamental concepts, some of which are badly understood and many of which are completely misunderstood. The purpose of this article is to explain some of these misconceptions and to shed light on the issues causing the most confusion throughout the industry.

The meaning of risk

Understanding the meaning of the term risk is the most fundamental prerequisite to developing an ORM programme. Many people don't realise that the term risk – as it is used in the risk management profession – is very different from the term used in informal conversation, as shown below:

Informal: I am exposed to fraud risk
Formal: I am exposed to the risk of loss from fraud

In casual conversation, risk is simply a type of incident – for example, a fire, fraud, reputational damage, a lawsuit, or something that could cause an adverse outcome, such as not having enough resources to complete a task or insufficient training. In formal expression, risk is a metric used to describe the uncertainty surrounding an event such as a fraud.

The best way to explain the meaning of risk is through a generic 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

> **"RISK IS A METRIC USED TO DESCRIBE THE UNCERTAINTY SURROUNDING AN EVENT SUCH AS A FRAUD"**

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 recognise 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 surrounding an adverse consequence – not the adverse consequence itself – and the adverse consequence need not be an actual 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 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 recognise 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 or uncertainty. Only where there is certainty is there no risk[2].

Expected and unexpected loss

Two very important metrics in ORM are 'expected loss' and 'unexpected loss.' While these terms are ubiquitous in the risk management industry, there is still much confusion about what they really mean. Some still believe the expected losses are the 'small' losses and the unexpected losses are the 'large' losses. Obviously these definitions do not yield any metrics and, in fact, are potentially misleading.

In formal expression, expected and unexpected loss are, respectively, the amount of money a firm loses on average in a year and the amount above the average that a firm could lose in a very bad year (at a pre-specified probability level). To restate this in more technical terms, the expected loss[3] is the arithmetic mean of an aggregate loss distribution, with respect to a certain time horizon – say, one year. The unexpected loss is the value-at-risk (VAR), which is described in conjunction with a confidence level (for example, 99%). VAR at the 99% level represents the amount of money one could lose where there is only a 1% probability of a larger loss (that is, where one is 99% confident that the aggregate loss in any given year will not exceed this amount of money). If the 99% level VAR were calculated correctly, one would expect to see an aggregate loss over that value only once every 100 years or, more reasonably, 10 times every 1,000 years. VAR is generally calculated in excess of the mean. See figure 1 for an illustration.

The terms expected loss and unexpected loss have important practical applications. Since the expected loss is the amount of money a business loses on average in one year, it is also the amount a business should budget to cover its annual cost of operational failure. The unexpected loss or the VAR is the amount a business could lose in a near worst-case situation and is the amount the business ought to reserve as capital. The expected loss is used to calculate profitability; both variables are used to calculate risk-adjusted return.

While most people realise it is hard to calculate VAR using internal data alone (because of the small sample size), many are unaware that because op risk is characterised by fat-tailed distributions, even the expected loss cannot be estimated using just internal loss data. This is because in fat-tailed distributions, the mean is affected by outliers and therefore one needs many years of data to arrive at a stable estimate. Consider a simple example: how many tsunami drownings take place in a year on average? Suppose a large tsunami occurs exactly once every 100 years and causes 200,000 deaths, then this would impact the mean or expected loss by 2,000. Therefore, the view that 'expected losses' are the small losses is not only wrong, it is potentially misleading, because the most efficient way to reduce the expected loss is to prevent the large losses, not focus on the small losses.

1. The total loss distribution illustrates the concepts of expected loss and unexpected loss

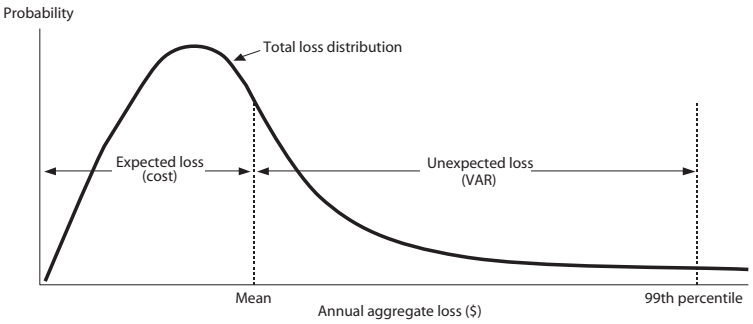

> **"THE VIEW THAT 'EXPECTED LOSSES' ARE THE SMALL LOSSES IS NOT ONLY WRONG, IT IS POTENTIALLY MISLEADING, BECAUSE THE MOST EFFICIENT WAY TO REDUCE THE EXPECTED LOSS IS TO PREVENT THE LARGE LOSSES, NOT FOCUS ON THE SMALL LOSSES"**

Risk assessment

There are many standards for risk assessment in ORM. One such standard, the traditional Coso[4] framework, is widely used in the US. The Coso ERM framework endorses a view that risk be assessed based on likelihood and impact, whereby risk is calculated as the product of these two factors. For example: a 10% likelihood and a $10,000 hypothesised impact would give you $1,000 worth of risk. However, this traditional method of calculating 'risk' does not actually give you the level of risk. Instead it gives you the probability weighted (expected) damage from a single hypothetical incident. This alone demonstrates why many traditional ORM methods cannot be used in modern ORM, because traditional ORM uses as a foundational element an entirely flawed conception of risk[5].

The difference between the traditional ORM

[1] In ORM we are only concerned with the risk of loss.
[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.
[3] The term expected loss has its origins in the field of probability and statistics, where the term expected value is used to describe the arithmetic mean of a distribution.

2. Traditional ORM produces false positives and false negatives

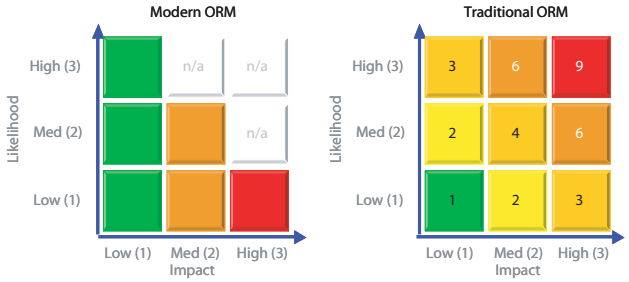

view of risk and the modern ORM conception is shown in figure 2. As can be seen, the traditional ORM approach characterises high risk as high-likelihood and high-impact, not low-likelihood and high-impact. This is a problem, because a business environment characterised by such catastrophic turbulence could never exist. (So under traditional ORM, no business could ever be described as being high risk.)

Traditional ORM risk assessment programmes put managers in an awkward position. To understand why, consider an example. It is widely known that unauthorised trading is a very significant 'risk'. But, because unauthorised 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 figure 2, 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 unauthorised trading to be classified as a high risk, respondents must answer incorrectly and falsely classify unauthorised 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 is clearly not conducive to promoting a good ORM culture.

How could such a flawed method have survived as the industry standard for so long? To find the answer one has to understand the roots of traditional ORM. 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[6]. 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.

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 hypothesised 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-

[4] Coso is an acronym for the Committee for Sponsoring Organizations of the Treadway Commission. For more information on Coso, please visit its website at www.erm.coso.org.
[5] For a full discussion of this topic please refer to *Why Coso is Flawed?*, by Ali Samad-Khan, *Operational Risk* magazine, January 2005 (the Coso article).
[6] Op risk is not the risk of a failure to meet one's business objectives. It is the risk of operational loss.
[7] Classification is still an evolving science and much work remains to be done in this area.
[8] Process analysis within each business line is an important aspect of control assessment.
[9] While internal data represents the character of the organisation, it is not sufficient for comprehensive risk assessment. Only aggregated industry (external) data, which provides a large sample size, can reveal the risk profiles of the different businesses.

ORA OpRisk Advisory

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 make public disclosure of such information in order for market participants to have full knowledge of their risk management practices and capabilities. At a minimum, banks should be asked 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 could be validated during on-site, regulatory examinations.

Operational vs operations

Many banks have already begun implementing ORM programmes as part of the Basel II requirements. In their rush to meet regulatory deadlines a large number of organisations failed to recognise the difference between the words 'operational' and 'operations'.

In fact, there are huge differences between ORM and operations management.

First of all, operations management is primarily a back-office management task involving the processing and systems functions. ORM has a much broader scope than just operations management. Op risk manifests itself in all the activities of an organisation, including the head office, corporate functions, the legal department and the activities of the board of directors. Second, operations management is primarily about managing operations or process efficiency. ORM is fundamentally about managing risk, specifically preventing operational losses, particularly the large ones. Third – though this varies from region to region – in the US banking industry, the major operational risks are primarily driven by events such as fraud, sales practices violations and unauthorised activities, which may not be high-priority issues in operations management. Lastly, the level of op risk in the operations area of a bank is significantly lower than that in the front offices.

The evolution of modern ORM

Traditional ORM was based on the assumption that intelligent, educated people could, through their own intuition, identify their organisation's significant risks, corresponding controls and associated metrics. Modern ORM is based on the view that intuition is not sufficient and that this process must be based on historical loss data, and rigorous, scientific analysis.

Therefore, the path towards modern ORM began with an entirely new question: what framework will make the legitimate use of aggregated historical loss data feasible and practical? The answer was a matrix, specifically a two-dimensional matrix, structured along the lines of 'generic' organisational unit and 'risk' class. The invention and initial use of the data matrix caused a paradigm shift in ORM.

Establishing a bold new theoretical framework is one thing; making it a practical reality is something altogether different. To make modern ORM workable, the industry had to find a meaningful way of finding structure in the 'risk' universe. This meant partitioning the risks into a set of unique classes, which were useful for management purposes, easily understood (to ensure consistent classification) and which also represented homogenous characteristics.

This posed a daunting task, because the disparate set of operational 'risks', which include fraud, fire, sales practice, key man, legal and reputation, as well as improper training and lack of supervision, appeared to defy structure. Necessity is the mother of invention, and the initial problem was eventually solved. But it took several years to even appreciate the complexity of this problem. This is because the 'risk' universe consists of three independent dimensions: causes, events and consequences, meaning that every loss consists of at least one element that is a part of each of these dimensions. An early conception of the risk universe (circa 2001) is shown in figure 3 (above). Since a matrix must

3. The universe of operational risks has three dimensions: causes, events and consequences

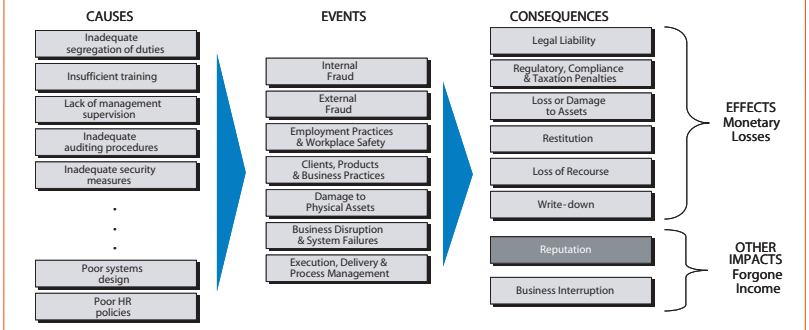

"THE IMPROVEMENTS IN MANAGEMENT PRACTICES, NOT JUST THE DERIVATION OF A CAPITAL FIGURE, ARE THE TRUE LEGACY OF RISK MEASUREMENT"


Stephane Le Blevec

ORA OpRisk Advisory

> **FUNDAMENTALLY, MODELLING IS ABOUT ANALYSING DATA SETS NOT MANIPULATING DATA POINTS. MODELS THAT ARE BASED ON SPURIOUS DATA MANIPULATION TECHNIQUES ARE NEITHER ART NOR SCIENCE**

consist of mutually exclusive and exhaustive classes it became necessary to pick one dimension. After much deliberation, the event dimension was found to be the optimal choice. However, finding a meaningful way of describing the different hierarchical elements in an event-based framework took yet another year[7].

For the second dimension of the matrix, the organisational structure dimension, the two most obvious choices were the process and the business line. For a variety of other reasons, the business line structure was deliberately chosen over the process structure[8].

In summary, modern ORM is based on a two-dimensional matrix approach, in which the unit of analysis is a cell within the matrix. Modern ORM requires both internal and external (industry) data[9]. The data within each cell represent a distribution of losses representing a class of events within a business line. By studying the causes of loss for each event class, one can identify common elements within and across classes and their relevant corresponding controls. By using a common matrix for risks and controls, one can use the modern ORM structure to identify and continuously track legitimate risk and control metrics side-by-side. This information is critical for effective risk management.

Traditional ORM is instead based on a one-dimensional process approach in which the unit of analysis is the audit issue within the universe of business processes. The methodology is based on a process of identifying issues based on control weaknesses and estimating the damage that could result if these issues are not resolved. While this approach is inappropriate for risk assessment it could be the starting point for control assessment, but traditional practices would have to evolve significantly for such a process to yield legitimate metrics representing the quality of the internal control environment.

Issues with traditional ORM

Many of the problems with the traditional ORM approach have been documented in the Coso article, but there are a few others that bear mentioning.

First of all, as described above, under the traditional likelihood-impact approach one assesses potential damage resulting from a specific issue, not the risk associated with a class of events. In order to be able to legitimately assess risk at the process and issue level, one requires a matrix of industry data mapped to the process and issue structure. No such data is available today, and it is unlikely that any such data will be available in the foreseeable future. Without such data it is very hard to identify relevant 'risks', let alone assess them. In addition, the very low-likelihood events that drive risk, because they are not well known, generally do not make it into process analysis. This leads to the problem of over-controlling the known issues (typically the low risks) and completely ignoring the unknown issues (generally the high risks).

Second, since traditional ORM does not require as part of its 'risk' taxonomy a disciplined, mutually exclusive set of risk classes, these undisciplined risk assessments can lead to double and triple-counting. For example, sales practices, customer and legal risk could be identified as separate risks, yet they often mean the same thing. (In fact, it is theoretically possible to identify an infinite set of risks.)

To a large extent, the way organisations structure their approach to ORM – in other words, how they state the problem – determines whether they will succeed or fail. Any effort to incorporate legitimate modern ORM methods directly into an issue/incident-orientated (process-based) approach or vice-versa will result in confusion (as many have discovered), because loss data is meaningful only when it is aggregated into classes of events. It is simply not possible to objectively use operational loss data to assess likelihood or impact at the process/incident level.

This approach is far too granular to be supported by the type of operational loss data that exists today, or that is likely to exist in the foreseeable future. Firms that attempt to do so are unknowingly trying to solve an unsolvable problem.

Modelling: art, science or nonsense

One senior US regulator recently observed that even though there were vast differences in methodologies and data being used by banks to quantify op risk, most banks were arriving at similar VAR figures. This is easily explained. Given the preponderance of highly subjective, even arbitrary, assumptions being used in op risk modelling today and the sensitivity of the results to these assumptions, it is not difficult to back into virtually any desired number. Banks generally want to 'pick' a 'result' that doesn't stand out. However, forcing a politically expedient result – one that is close to the regulators' expectations – proves nothing, and reveals very little about the robustness of the bank's underlying methodology.

Consequently, it is not difficult to see why some organisations have politely concluded that modelling op risk is more of an art than a science. While there is an element of art and science in all modelling, many of the op risk models in use today are based on such arbitrary assumptions and unscientific methods that this pseudo-science is giving operational risk modelling a bad name. For example, some organisations 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 1,000 or more.

Armin Rheinbay

ORA OpRisk Advisory

External loss data is essential for op risk modelling, but incorporating external data into the modelling process requires an objective, scientific approach. Directly combining internal and external data violates one of the fundamental precepts of op risk modelling 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. Extracting a loss data point from its original data set, causes it to lose all informational value. Fundamentally, modelling is about analysing data sets not manipulating data points. Models that are based on spurious data manipulation techniques are neither art nor science, they are just plain nonsense, and they erode the credibility of the honest and diligent people who work in this field.

Banks that want to apply for the advanced measurement approach under Basel II must establish higher internal standards for their quantification models. To encourage sounder thinking in this area, as part of the Pillar III requirements under Basel II, regulators could require that banks disclose not just their expected and unexpected loss estimates, but also the surrounding confidence intervals (which the regulators could validate through stress testing). These confidence intervals should represent the minimum and maximum values that could be calculated by varying any weights and assumptions based on 'expert' opinion.

Measurement vs management

Some people contend that modern ORM is about measurement and traditional ORM is about management. Let us examine this assertion.

Well managed organisations have discovered that effective ORM goes beyond simply building 'awareness' in the hope that sound risk management practices will emerge spontaneously. Pragmatists know that effectively managing op risk involves creating the right culture or, more specifically, a culture and framework designed to turn awareness into appropriate action.

Getting managers to act optimally requires the right set of incentives, because people do what they have an incentive to do and generally do not do what they don't have an incentive to do. But in order for incentives to work properly, they must be based on the right metrics.

An effective ORM programme requires a sound framework, one that must be able to provide accurate, reliable metrics that identify within each business the most significant risks as well as the quality of their corresponding internal controls. This information must be made transparent and provided to managers on a periodic basis, so that they are able to – and have an incentive to – make educated decisions when developing risk management, risk mitigation and risk transfer strategies. Hence, managing op risk requires a process for accurately monitoring (measuring) each business' changing risk and control profile.

A modern ORM programme, if implemented correctly, can achieve all these objectives. However, a traditional ORM programme cannot. In fact, traditional ORM is more likely to lead to op risk mismanagement because the downplaying of major risks (which as we have seen is an inevitable consequence of traditional ORM) can leave organisations unknowingly exposed to catastrophic operational failure.

Separately, arguing that measurement is only about calculating a capital figure misses an important point. Measurement raised the standard, and it was measurement that turned ORM into a science. It brought comprehensiveness, structure and discipline to the process. It led to the development of a much more efficient management framework for ORM. It forced the industry to probe the op risk definition and classification issue, which in turn brought greater clarity to the analysis. And, perhaps of greatest importance, it revealed that traditional ORM was based on an incorrect conception of risk, and that furthermore, the entire framework had serious issues. These improvements in management practices (not just the derivation of a capital figure) are the true legacy of modern risk measurement.

> **"ORGANISATIONS THAT HAVE TRIED TO BUILD MODERN ORM PROGRAMMES ON TOP OF THEIR EXISTING TRADITIONAL FRAMEWORKS HAVE FOUND ORM TO BE A VERY CHALLENGING TASK"**

Summary and conclusions

Traditional ORM was developed at a time before loss data existed, which precluded it from rising to the level of a science. Loss data and advanced risk measurement techniques turned ORM into a science. While exploring data and measurement issues it became clear that there were many flaws in traditional ORM. The need for a robust method of addressing these problems is what led to the development of modern ORM.

Modern ORM is very different from traditional ORM. In many ways modern ORM is incompatible with traditional ORM. Organisations that have tried to build modern ORM programmes on top of their existing traditional frameworks – leveraging existing terminology, processes and procedures, without probing the core issues – have found ORM to be a very challenging task. This remains the underlying source of much of the confusion in the industry. OR&C

Ali Samad-Khan is president, and Armin Rheinbay and Stephane Le Blevec are principals at OpRisk Advisory, a consulting firm, specialising in operational risk management. They can be reached at their respective offices in the USA, Switzerland and France. For contact details please visit www.opriskadvisory.com