Why COSO is flawed

COSO not only fails to help a firm assess its risks, it actually obfuscates the risk assessment process. By **Ali Samad-Khan**

Ali Samad-Khan

Operational risk is one of the most significant risks that businesses face in today's complex global economy. For most of the world's leading institutions it has become more than apparent that implementing an effective operational risk management programme can help reduce losses, lower costs associated with fixing problems and increase customer and employee satisfaction, thereby improving financial performance and enhancing shareholder value.

Basel II may have forced banks to review their approach to managing operational risk, but for most leading institutions the question was never whether to establish such a programme, it was how. But many institutions are still unsure of the benefits. Some are still struggling to decide whether to comply with the BIS basic indicator, standardised or advanced measurement approach. Nevertheless, compliance issues aside, most banks have come to the conclusion that if they are going to have to establish an operational risk management programme then they want it to be based on a sound framework. What is perhaps surprising though is that while we are many years into this process, there is still no industry consensus on what shape or form this framework ought to take. And while there has been much heated debate on this issue, much of it has been based on personal opinion and not fact. This is because, even today, a number of fundamental misconceptions exist about the true meaning of operational risk management in its modern conception. The purpose of this paper is to shed light on one of the main issues that is driving this confusion.

Many people believe that managing operational risk can be accomplished by following the Committee for Sponsoring Organizations of the Treadway Commission (COSO) approach.[1] The recently released COSO framework sets the standards for enterprise-wide risk management (ERM). COSO views ERM as a process aimed at helping organisations identify potentially adverse events and subsequently manage the associated risks in furtherance of the entity's business objectives. When applied to operational risk management this is often translated to mean: begin with a comprehensive survey of the organisation to identify, define and assess the full spectrum of 'risks' in each business' underlying processes. Then define a series of responses or controls to mitigate the risks that threaten to prevent the entity from meeting its objective. This is often accomplished by establishing a list of issues and follow-up action plans to ensure compliance with this programme can be verified over time through the audit process.

At a macro level, this approach appears both comprehensive and sound, but the devil is in the details and the specious logic underlying COSO becomes evident during implementation. While COSO may help organisations identify and resolve some of their more obvious control weaknesses, in our view, it is completely inappropriate for use in operational risk management. Fundamentally, COSO is inappropriate for use in operational risk management because the definition of risk used under this approach is wholly inconsistent with the definition of risk used in the risk management industry and by the BIS (see next section for a full explanation of this point). In addition, the method COSO prescribes for an organisation to assess its risks is highly subjective, overly simplistic and conceptually flawed.

COSO not only fails to help a firm assess its risks, it actually obfuscates the risk assessment process. Because risk assessment is a foundational element in the risk management process, and because COSO yields an entirely counterfeit set of risks, the spurious and misleading results of the flawed risk assessment stage contaminate every subsequent stage of the process. As a result, the recommended risk mitigation strategy – the set of controls and action plans designed to mitigate the identified risks – is likely to be non-optimal at best. In the worst case, it may lead organisations to expand and intensify control structures in areas where they are already over-controlled, while completely ignoring areas of major control weakness, leaving the organisations both oblivious and vulnerable to huge operational losses that could hit them like a bolt from the blue.

In our view, COSO – as it is currently applied – is a wholly inappropriate approach for managing operational risk; it is a huge waste of resources and is very likely to do more harm than good.

One obvious issue with COSO is that it is hugely resource-intensive. This is because COSO requires that all processes be assessed, irrespective of their individual contributions to the organisation's total risk (because one cannot know the level of contribution to total risk without first conducting a risk-assessment). Identifying and documenting the risks in each and every process could take many person-years. One mid-sized bank recently estimated that it would require 192 person-years to complete such an assessment across the entire organisation. Clearly, the cost of such massive resource commitment is not something an organisation can easily absorb, particularly if the exercise needs to be repeated on an annual basis, which is necessary because for operational risk management to be effective it must be implemented through a dynamic process with continuous monitoring. Still, this is not a problem as long as the cost can be justified.

A serious problem with COSO has to do with the way the 'risk' information is collected. The starting point under COSO in a typical implementation is the identification, definition and assessment of risks in a business process. In general, the persons interviewed are business managers. While these persons may be well qualified to run their own businesses, they do not necessarily know anything about risk. Yes, they can probably come up with a long list of potential risk scenarios, but that's only half the story. To know which risks are real risks the manager would also have to know the relative probability of each 'risk event' that could affect his or her business. After all, a tsunami and a wire-transfer error are both risks, but without knowing whether a 99% level tidal wave or 99% level fat-finger error could do more damage – in the context of their existing control environment – they cannot know which risk poses a greater threat. And as it turns out there is often a major discrepancy between perception and reality.

The only way one can identify one's real risks is by studying historical loss data. A risk manager, whose job it is to know about historical losses, is much more likely to be aware of the full range of potential risks affecting a business (and their relative probabilities) than is a business manager. Having a qualified risk manager ask a business manager where his or her major risks are is similar to having a doctor ask his or her patient: to which major diseases do you think you are most exposed? Some patients will know the answer, but most will not, which is why they went to see their doctor in the first place. In a well-managed organisation, the risk professional should serve as the doctor and the business manager as the patient.

For those who still believe the right approach is to ask business managers to self-assess the risks within their organisation's underlying processes, we ask what would they have recommended to the Governments of India, Indonesia or Sri Lanka, prior to the recent tragedy, considering that tsunami risk probably was not a recognised risk in any of their processes?

Another major problem with COSO is that a typical risk-assessment implementation generally produces a huge catalogue of risks – often in the thousands. Thus, when it comes to actually managing these risks across an organisation, ie, determining which risk mitigation strategy is optimal, it is very difficult to prioritise actions because without a 'normalised' rank ordering of risks one cannot know which controls should be given precedence in implementation.

To address this problem, COSO developed the likelihood-impact method of risk assessment. Under this approach, businesses calculate the magnitude of their risks based on a mathematical formula, where risk is equal to the likelihood that a given event will occur multiplied by its effect (impact), should it occur, such that, Likelihood x Impact = Risk.

FOR those who understand the concept of risk, as it is used in the risk management industry, it is clear that there is something fundamentally wrong with this approach. Using the COSO formula the worst-case outcome is characterised by high likelihood and high impact; however, under the risk management approach, the worst-case outcome is characterised by a low probability (low frequency) – high impact (high severity) event, such as a $1 billion dollar unauthorised trading loss. In fact, there is no such thing as a high likelihood (high frequency) – high impact (high severity) event. This would characterise a risk (type of loss) that occurs hundreds of times a year and each time causes

[1] COSO's objective was to help standardise procedures for enterprise risk management by developing a conceptually sound framework providing integrated principles, common terminology and practical implementation guidance. For more information on COSO, visit www.erm.coso.org

billion-dollar losses. This is clearly a phantom risk. What's even worse is that COSO also completely understates the one area of real risk. In summary, the COSO approach to risk assessment will tell you your risk is very high in areas where you have no risk, and will also tell you that you have moderate risk in the very area your risk is of the highest order. Simply stated, COSO produces both false positives and false negatives.

The contrast is illustrated in figure 1 (right).

Some advocates of COSO have suggested that this problem only exists when the analysis is qualitative or high level. They argue that likelihood and impact analysis works well when the inputs are expressed in more quantitative terms, such as percent probability and dollar magnitude. To examine this argument, let us express it in the context of a simple business problem. Suppose you want to know the risk associated with your having a car accident during the coming year. If you know that you have a 10% chance of having an accident and you expect that accident will cost $10,000, then you would calculate your risk as follows:

Likelihood x impact = risk
Risk 1: 10% x $10,000 = $1,000

But as you further consider this matter you realise that the problem is more complex than originally perceived. After all, the 10% likelihood relates only to a $10,000 event; and there is also a possibility, let's say a 1% probability, that you could have a very bad accident, which could result in the total destruction of your $50,000 car. Therefore you have two possible ways of estimating your risk, as shown below:

Likelihood x impact = risk
Risk 1: 10% x $10,000 = $1,000
Risk 2: 1% x $50,000 = $500

What becomes immediately apparent is that two completely valid assessments can yield different risk results. In fact, upon further consideration, it becomes evident that the problem is still more complex because there are, in fact, multiple 'solutions', because there are potentially an infinite number of likelihood and impact combinations, as shown below:

Likelihood x impact = risk
Risk 1: 10% x $10,000 = $1,000
Risk 2: 1% x $50,000 = $500
. . . . . .
Risk 999: 5% x $25,000 = $1,250
Risk 1,000: 20% x $6,000 = $1,200

From the outcome, one can clearly see that all the 'risk-results' are banded together (from $500–$1,250) with little differentiation. This is because the higher the impact the lower the likelihood (an incremental gain in likelihood offsets any corresponding reduction in impact). The major differences in the risk-results are due to the fact that the product of two figures near their respective means is greater than the product of two figures at opposite extremes. But this is an idiosyncrasy of the arithmetic process and is not reflective of any legitimate difference in the level of risk; in fact, the opposite is true. For example, if we were to think of a 'risky event' as a 99% level (1% likelihood) event, then from the table shown above one can see that this would correspond to a $50,000 loss. But when one defines risk as the product of likelihood and impact one can see the $50,000 x 1% ($500) event would imply less risk than the $25,000 x 5% ($1,250) event. This wholly absurd rank-ordering clearly demonstrates that, far from improving operational risk management, COSO obfuscates the process of determining an organisation's true risk profile.

1. COSO vs reality: COSO produces false positives and false negatives

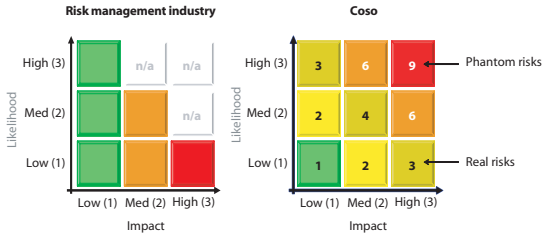

So what should we do now? Suppose now we were to take a weighted average of all the risk-results drawn from the table above? What would that answer represent? The mean of all the risk results would equate to probability weighted severity (which seems to equate to mean severity). But if this is true then we have a problem, because mean severity is somewhat similar to expected loss (mean aggregate loss), whereas the risk management industry and BIS definition of operational risk equates to the unexpected loss. Without knowing anything else about COSO it is clear that the meaning of risk under COSO is altogether inconsistent with the true meaning of operational risk. By following the COSO definition of risk, one is shooting at the wrong target, one that is not even a close approximation!

What is fundamentally wrong with the COSO-based risk assessment approach is that the question is flawed. Instead of looking for the product of likelihood and impact we should be taking as the results of this process the full set of likelihood and impact combinations. And if we were to plot them on a graph we would get something similar to what we see in figure 2 (right).

When you connect the dots, the full set of combinations would represent a set of points on a continuum. This is known in actuarial science as a severity distribution.

And what does one do with this severity distribution? Will the severity distribution give us the answer we are looking for? If we look at the 1% probability event on this distribution, will that not tell us our level of risk? No, not quite yet. As it turns out the severity distribution is just one piece of this puzzle. Returning to our example, the severity distribution is a distribution of single-event losses, showing the full set of losses and corresponding probabilities associated with a single car accident. But this is not what we want. We want to know our operational risk in terms of the total amount of money we could lose from all the car accidents we could have in the next year. For this we also need

to know how many accidents we could have in one year – or more precisely a probability distribution for the number of accidents we could have in a given year. This is known in actuarial science as a frequency distribution.

Under the risk management industry and BIS definition, operational risk – as shown above – is described in the context of an aggregate or total loss distribution, which is a convolution (a mathematical combination) of both a frequency and a severity distribution, where the relevant points are the expected loss and, more importantly, the unexpected loss. The expected loss is the total amount of money one expects to lose in a year, on average, and the unexpected loss is the total amount of money one could lose in a very bad year (at a specified confidence level) in excess of the average. (For a technical explanation of the terms expected and unexpected loss, please refer to the BIS guidelines).

From the above discussion it should be evident that the risk result under the likelihood-impact approach equates to mean severity, which is completely unrelated to the term risk as it is defined by the risk management industry and the BIS. In fact, mean severity multiplied by mean frequency gives you the mean aggregate loss – the expected loss. Whereas the real measure of risk is the unexpected aggregate loss.

The likelihood-impact analysis approach to risk assessment can be summarised as follows: it is based on a process whereby one asks the wrong people to answer the wrong question, which in any case is a flawed question, because it has an infinite number of different, but theoretically valid answers. And even if you were to ignore the answers and take only the potentially useful information from this process – the full set of input pairs – you would still only have one part of the solution. No matter how you sum it up, four wrongs don't make a right.

In summary, implementing COSO requires a gargantuan effort, and, in the context of operational risk management, it produces spurious and misleading results. Acting on this information may divert managers' attention from their real risks and instead focuses their attention on phantom risks, while at the same time providing them with a false sense of security. Furthermore, any risk mitigation strategy based on this flawed risk information is likely to focus attention and resources on the wrong controls. It is highly conceivable that this approach could lead to an intensification of controls where a business is already over-controlled, while completely ignoring areas of control weakness. As we all know, the consequences could be disastrous.

Operational risk management in its modern conception

An effective operational risk management programme requires a sound framework. The goal of this framework should be to provide reliable information to key decision-makers so that they are aware of their most significant risks as well as the quality of their corresponding internal controls, information that will allow them to make educated decisions when developing risk management, risk mitigation and risk transfer strategies. Managing operational risk fundamentally revolves around the process of optimising the risk-control relationship in the context of cost-benefit analysis. This, in turn, requires a process for accurately monitoring (measuring) each business' changing risk and control profile.

> "SOME MAY STILL ARGUE THAT COSO IS USEFUL BECAUSE IT IMPROVES BUSINESS PROCESS MANAGEMENT; EVEN IF THIS IS TRUE, IT SHOULD NEVERTHELESS BE CLEAR THAT BUSINESS PROCESS MANAGEMENT IS NOT OPERATIONAL RISK MANAGEMENT"

To accomplish this goal four things must be done correctly. First, the risk management department must be able to provide managers with objective information to help them better understand where their risks really are, not ask them to guess where their risks might be. Fundamentally, one cannot manage one's

2. Severity distribution

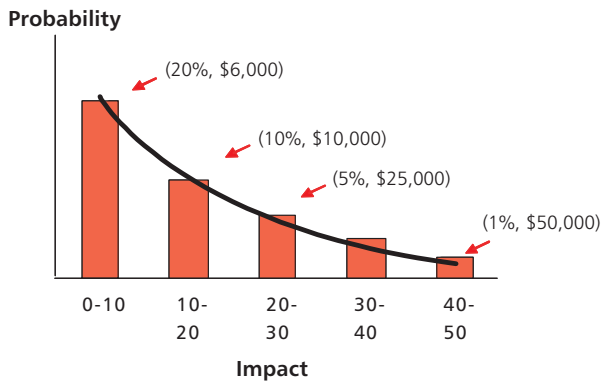

operational risks without measuring one's operational risks. It is very difficult to be able to differentiate between major risks and minor risks and real risks and phantom risks without being able to accurately measure these risks in the first place. It is also impossible to develop an effective risk management programme without knowing which risks must be dealt with as a top priority.

Second, one must help managers understand how well their real risks are being managed through their existing set of controls, so they can know where they are over-controlled and where they are under-controlled in the context of their overall operational risk strategy and risk (loss) tolerance. One cannot have a zero-tolerance policy towards operational risk, just as one cannot institute perfect controls. An organisation has to be realistic in establishing a level of risk and loss tolerance.

> "MANAGING OPERATIONAL RISK FUNDAMENTALLY REVOLVES AROUND THE PROCESS OF OPTIMISING THE RISK-CONTROL RELATIONSHIP IN THE CONTEXT OF COST-BENEFIT ANALYSIS"

Third, one needs to determine what level of controls is appropriate after having conducted a circumspect analysis of the associated costs and benefits of each risk mitigation and transfer strategy.

Fourth, one needs to institute a comprehensive and fully transparent monitoring and reporting process with built-in incentives to encourage desired behavioural change.

It is difficult to think of ways one could even begin to manage operational risk without having these foundational elements in place. Best-practice calls for an integrated operational risk measurement-management programme, whereby objective, transformed (normalised) measures are used to identify levels of risk and internal control quality. But for these measures to be meaningful they must be based on reliable information, specifically: internal and external loss data, theoretically valid risk measurement and assessment, objective control self-assessment, validated risk indicators, appropriate follow-up action results, disciplined scenario analysis and well founded VAR calculation.

Can this really be done and is it practical? The answer is yes to both questions, but only if the underlying framework is based on sound reasoning, which must in turn be based on a comprehensive understanding of the issues. And these issues must be addressed logically and objectively, one issue at a time.

Conclusions

COSO was initially conceived in the early 1990s, and for a long time represented best practices in enterprise risk management. Then banks began collecting historical loss data, and we entered the dawn of a new age. As the process of collecting loss data became more widespread, thanks to the bold

3. The actuarial approach

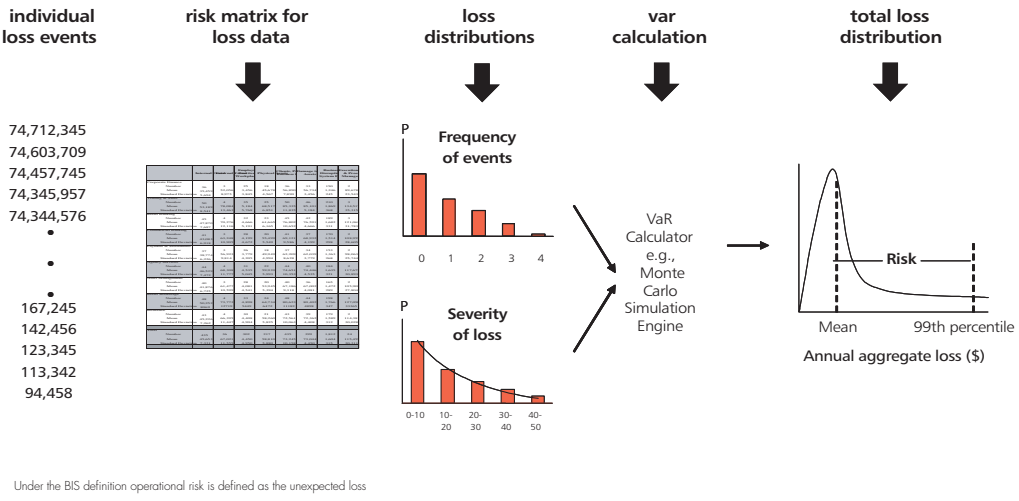

Under the BIS definition operational risk is defined as the unexpected loss

insistence of the BIS, loss data began fuelling an entirely new and more scientific way of thinking about what came to be known as operational risk management. It was the analysis of this data and the issues subsequently raised that eventually led to the development of modern operational risk management as an objective discipline.

Some may still argue that COSO is useful because it improves business process management; even if this is true, it should nevertheless be clear that business process management is not operational risk management.

There are also those who speak of operational risk management as independent from operational risk measurement. In our view measurement is an integral part of the management process. After all, what is risk management other than the mitigation of major risk in the most cost-effective way. It's difficult to see how one can accomplish this in a large organisation without reliable measures.

Basel II was introduced to encourage banks to improve their operational risk management. But following COSO does not improve operational risk management; instead it promotes phantom risk management and does more harm than good. Furthermore, any organisation that applies COSO-based risk-assessment to this end will clearly be demonstrating to its regulators, to its investors and to the rating agencies that it has not yet grasped even the most basic understanding of operational risk management – ie, operational risk management is about managing risk. In our view, far from meeting the standards of the advanced measurement approach or even the standardised approach, a COSO-based operational risk management framework may only just barely meet the minimum standards of the basic indicator approach – which has no standards at all!.

One of the biggest problems we face in the operational risk management area is that many of those professing to be experts in this field actually know very little about operational risk management in its modern conception. By continuing to espouse their outmoded and impractical views on the subject these individuals are unknowingly doing more harm than good, as their flawed guidance is steering the industry in the wrong direction. Based on the advice of these individuals many organisations have invested millions of dollars implementing frameworks and software that they will soon discover have neither improved their management of operational risk nor achieved any level of BIS compliance. While it is easy to see how many banks could have fallen into this conceptual black-hole, if immediate steps are not taken to lead them out, they are likely to fall into an even deeper abyss.

There are no shortcuts to developing a comprehensive framework for managing operational risk. And one cannot get on the right track without confronting the difficult issues head on. If an organisation's operational risk management framework is not founded on fundamentally sound reasoning the entire programme will eventually unravel at the seams. An ill-conceived operational risk management programme is also likely to leave an organisation vulnerable to major operational losses. The damage from even one major loss could be far greater than the cost of establishing a state-of-the-art, integrated operational risk measurement-management programme. Just think how little a very simple global-early warning system would have cost to build and maintain relative to the lives lost and property damage that resulted from the recent Asian tsunami.

> "ONE OF THE BIGGEST PROBLEMS WE FACE IN THE OPERATIONAL RISK MANAGEMENT AREA IS THAT MANY OF THOSE PROFESSING TO BE EXPERTS IN THIS FIELD ACTUALLY KNOW VERY LITTLE ABOUT OPERATIONAL RISK MANAGEMENT IN ITS MODERN CONCEPTION"

The operational risk management industry has been plagued by disinformation and methodology. The industry would be much better served if instead of expressing the personal opinions of the 'experts' it made an effort to understand the issues. Without doing so, it's hard to see how anyone could become qualified to address this challenging problem. We certainly don't pretend to have all the answers, but we do think we have hit upon many of the right questions. It is important to recognise that one can never arrive at the right answers without probing the most important issues. Only by analysing and re-analysing these issues can one begin to shed light on what may be the right questions. Finding the answers is the easy part. Discovering what are the right questions is the major challenge. OpRisk

Ali Samad-Khan is president of OpRisk Advisory. He has eight years' experience in operational risk management, having previously worked at Bankers Trust, OpRisk Analytics (which was acquired by SAS) and PricewaterhouseCoopers, where for over three years he headed the operational risk group within the New York FRM practice. He can be reached at ali.samad-khan@opriskadvisory.com