

Measurable Value Creation Through an Advanced Approach to ERM

Greg Monahan

Copyright 2009 by the Society of Actuaries.

All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

Abstract

This paper presents an advanced approach to enterprise risk management (ERM) that significantly improves upon current approaches, largely due to two fundamental elements, namely: a focus on the management of strategic plans and, secondly, a heavy reliance upon probability theory. This paper is structured as follows. Section 1 reviews essential probability theory. Section 2 introduces the “Strategic Objectives at Risk” (SOAR) Methodology. Section 3 describes the SOAR process that is the risk management process at the heart of the SOAR Methodology. Section 4 discusses the direct measurement of stakeholder value added by the SOAR Methodology.

1. Review of Probability Theory


In this section we review some basic probability theory and statistical measures.

A probability distribution *function* describes the relationship between an outcome and its probability. Probability distributions are very often presented graphically and are usually most easily understood when presented in this way. Some common probability distributions are described here.

The Uniform Distribution

A uniform distribution is one in which every outcome has the same probability. A well-known example of the uniform distribution describes the possible outcomes from a roll of a fair, six-faced die, as shown here:

Figure 1
Uniform Distribution



The Normal Distribution


Arguably the most recognizable probability distribution is the normal distribution. The much less recognized normal distribution *function* is:

Equation 1.
Normal Distribution Function

$$P(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2}$$

The normal distribution is used to describe countless variables, such as height and weight of people and test scores. When values that are normally distributed are plotted, the graph looks something like this:


Figure 2
Normal Distribution



The Empirical Distribution

It is sometimes impossible to define a function to relate outcomes and probabilities, even though we can observe historical outcomes and, from those, estimate probabilities. In such cases, it is possible to work with the empirical distribution, which is just a histogram of observed outcomes. An example of a distribution that you might judge too hard to attempt to define a function for is shown here:

Figure 3.
Empirical Distribution



Whilst it may not be possible to define the function behind the empirical distribution shown in Figure 3, it is very easy to use the historical data to predict future outcomes, provided you apply the assumption that the distribution of future outcomes will follow the distribution of past outcomes. On the basis of this assumption, the possible future outcomes have the probabilities stated in

TABLE 1.

At this point, I deviate slightly to make a general comment about models. Albert Einstein warned us not to make things simpler than they ought to be. I am afraid that, in finance, we too often ignore this advice. Indeed it can be very easy to build a model to predict future outcomes based on historical observations, but the limitations of the model must be acknowledged. Making the model simple (for example via the application of some assumption like “the distribution of future values will match the distribution of past values”) does not change the reality. A model made simple via assumption is, in my view, a very dangerous tool, even when the assumptions and the model’s limitations are known. Have you ever heard the terms “estimated probability” and “theoretical probability”? I suspect that not many people hear these terms regularly, but most of the probabilities we deal with in finance are *estimated* or *theoretical*. Consider, for example, these two statements:

“There is a one in six chance of rolling a six.”

“The one-year probability of default for a AAA-rated bank is 0.03 percent.”

That the probability of rolling a six (from a fair, six-faced die) is one in six is a *fact*. That the one-year probability of default for a AAA-rated bank is 0.03 percent is an *opinion*. Model outputs are very often stated and accepted as if they were facts, and this supports my belief that models are commonly misunderstood. Take, for example, estimates of movements in equity prices. Movements in equity prices are quite often considered to be normally distributed, and it is common to generate a probability distribution of the future movement in a stock price using Monte Carlo simulation based on the mean and standard deviation of historical movements. The generated probability distribution represents an opinion, but it is almost always treated as a fact. From the generated probability distribution, people will very likely make some statement like “the probability that the stock will fall by more than 10 percent is less than 5 percent.” This statement suggests that it is a *fact* that a fall in the stock price of more than 10 percent has less than a 5 percent probability, but it is not a fact; it is an *opinion* based on a model. A more appropriate expression would be “the *estimated* probability that the stock will fall by more than 10 percent is less than 5 percent.” By including the word estimated in the statement, the speaker is reminding the listener that the statement is based on the outcomes of a model, which has been designed to provide *estimates*. I will close this aside by stating that I believe the word “estimated” is underutilized.

TABLE 1
Probability Distribution

Outcome	Probability
-3	5%
-2.5	20%
-2	1%
-1.5	10%
-1	30%
-0.5	1%
0	15%
0.5	2%
1	1%
1.5	1%
2	2%
2.5	1%
3	11%

Summary Statistics

The management of risk requires a good understanding of a few basic statistics used to describe or summarize a set of numbers (such as the set of numbers behind a probability distribution), and these statistics are discussed here. A “good” understanding is that level of understanding which allows the numbers to be correctly interpreted.

Mode

The mode of a set of numbers is the number that occurs most frequently. In Figure 3 above, the mode is -1. In Figure 2, the mode is 0. In Figure 1, there is no mode; no outcome appears more frequently than any other. A good way to remember the mode is to think of it as the outcome that has the highest probability of occurring.

Mean

The mean, or average, of a set of numbers is calculated by dividing the sum of the numbers by the count of numbers, as expressed here:

$$\mu(x) = \frac{\sum_{i=1}^n x_i}{n}$$


The mean is quite often referred to as the “expected’ value; however, risk managers must be careful that the term “expected value” is not confused with the term “most expected value,” which means most likely, and is the mode of the distribution. Consider the case where the outcomes in

TABLE 1 above represent the number of telephone calls you expect to receive tomorrow in relation to the number you received today. In this case, -3 means you expect to receive three fewer telephone calls. If asked to guess how many telephone calls you think you will receive tomorrow, what would you say? The answer that gives you the highest chance of being correct is “one less than today” as -1 is the mode of the distribution. A common mistake is to provide the mean as the answer to this question. From the data in Table 1, it is easy to calculate the mean of -0.7. Note that the mean is not necessarily the most likely value and that it is very often not even a possible outcome! I try to remember this; the average outcome is highly unlikely.

Standard Deviation

Standard deviation is a measure of how widely spread values are from the mean. In terms of risk management, the greater the standard deviation, the higher the risk.


$$\sigma(x) = \sqrt{\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n - 1}}$$



Skew

Skew is a measure of the symmetry of a distribution. A positive skew indicates that values greater than the mean are more widely spread than the values below the mean.


$$skew(x) = \frac{n}{(n-1)(n-2)} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{\sigma(x)} \right)^3$$



Kurtosis

Kurtosis is a measure of the heaviness of the tails of the distribution, and larger values represent heavier tails—values that are a long way from the mean.

$$kurt(x) = \frac{n(n+1)}{(n-1)(n-2)(n-3)} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{\sigma(x)} \right)^4 - \frac{3(n-1)^2}{(n-2)(n-3)}$$



“at Risk” Measure (or Percentile)

The “at Risk” measure, or percentile, is the value of the outcome that has a certain probability. The statement of the “at Risk” measure usually takes one of the following forms, and these statements are equivalent:

- “The 99th percentile outcome is -3.”
- “The worst outcome, expressed at 99% confidence is -3.”
- “The outcome at risk, expressed at 99% confidence is -3.”

One way to calculate the value at a certain percentile is to rank the values from lowest to highest and then select the n^{th} value, where “n” is determined by the percentile sought and the number of values. If you know the probability distribution function, for example as per Equation 1, the value at a certain percentile can be returned from the function (or its inverse).

2. SOAR Methodology


A review of papers presented at past symposiums quickly reveals that a commonly accepted definition of ERM describes a framework that aims to ensure that the individual risks faced by an organization are adequately managed, ideally in a coherent fashion. The apparent acceptance of this definition may be due to the fact that many past papers discuss ERM in the context of financial institutions, within which distinct risk classes have been separately managed for many years; it is only relatively recently that these institutions have sought to measure these risks in an integrated fashion, often referring to the models they put in place as ERM. Whilst very likely a valuable tool, a framework for the coherent measurement of individual risks does not meet my own definition of ERM and I prefer to refer to such frameworks as examples of enterprise wide risk management, or integrated risk management. I propose that ERM be thought of as the process of managing risks that impact the value of the organization and are not addressed as part of the risk management function of individual business units. Under the definitions proposed here, a bank that seeks to measure credit and market risk via the execution of a single Monte Carlo simulation process is conducting integrated risk management, but not conducting ERM.

The remainder of this part is dedicated to introducing the SOAR Methodology, which has been designed specifically for ERM. Though designed to be applied to the management of risks associated with strategic objectives, the SOAR Methodology boasts no inherent limitation in its application. For example, the calculation of VaR on a bank trading portfolio can be considered a very specific application of the SOAR Methodology. One can use the SOAR Methodology when preparing for a job interview, or deciding where to go on holiday.

The core principle of the SOAR Methodology is to use data to forecast the possible outcomes that may eventuate as a result of the pursuit of an organization's strategic objectives. Through disciplined, data-based management, the risk manager can favorably influence the probability distribution of possible outcomes. A probability distribution that is more favorable to the organization is one that is taller and thinner and located closer to the desired outcome than the original distribution. By way of example, consider the case where an organization sets as one of its strategic objectives to achieve revenue of \$100 million over the next 12 months. Following an analysis of the variability in annual sales over the last five years, you show that the possible values for revenue over the next 12 months are \$70 million, \$80 million, \$90 million, \$100 million and \$110 million and that these outcomes have probabilities of 10 percent, 20 percent, 30 percent, 25 percent and 15 percent respectively. The probability distribution would look like the one presented in

Figure 4.

Figure 4.
Probability Distribution of Revenue




The facts associated with this forecast include the following:

1. The probability of achieving revenue of \$100 million is only 25 percent.
2. The probability of achieving revenue less than \$100 million is 60 percent.
3. The most likely outcome is revenue of \$90 million.
4. The expected outcome is revenue of \$91M (to get this result you have to do a bit of math, rather than read the graph).

Now imagine that by strict adherence to the SOAR Methodology, you are able to create a set of circumstances that alter the possible outcomes and their probabilities, such that the probability distribution now appears like the one in Figure 4.

Figure 4
New Probability Distribution of Revenue



The facts associated with the new forecast include the following:

1. The *estimated* probability of achieving revenue of \$100 million is 40 percent.
2. The *estimated* probability of achieving revenue less than \$100 million is 35 percent (and is equal to the *estimated* probability of achieving revenue of more than \$100 million).
3. The most likely *estimated* outcome is revenue of \$100 million.
4. The *estimated* expected outcome is revenue of \$100 million.

3. SOAR Process

The SOAR process is the driving force of the SOAR Methodology. The SOAR process is a management process that imposes discipline and data-based decision making. The SOAR process comprises the following steps:

Set


The “set” step involves the determination of metrics associated with strategic objectives and target metric values. When correctly defined, achievement of the target metric value is equivalent to achievement of the strategic objective. In some cases, such as when the strategic objective is stated in terms of some convenient unit of measure, the determination of the metric and the associated target value can be quite straightforward, whilst in other cases some imagination may be required. Consider the following examples. If the strategic objective is “to increase annual revenue to \$100 million,” then annual revenue *could* be used as the metric and \$100 million *could* be the target value. Note the deliberate use of the word “could” in the previous sentence, recognizing that other metrics and other target values could be defined. If the strategic objective is “to be the world’s premier alternative investment platform,”¹ the metric and, hence, its target value, are less obvious. The word “premier” needs to be converted to something numerically quantifiable. In order to define the appropriate metric, the meaning of “premier” needs to be determined. Does it mean “the fund that produces the highest returns”? Does it mean “the fund that produces the most stable returns”? Does it mean “the fund that has the largest number of investors”? Or does it mean something else? Once the meaning of “premier” is determined, by consulting with the person responsible for achieving the strategic objective, the enterprise risk manager will be able to determine an appropriate metric and its target value.

Observe

The “observe” step involves the regular observation of metric values. As with setting metrics, observing metric values will sometimes be mundane and will sometimes require a little more effort. Consider again the first of the two examples used above. In that example, annual revenue could have been chosen as the metric and \$100 million chosen as the target value. Although annual revenue can only be observed annually, monthly revenue can be used as a predictor (or indicator) of annual revenue. In some cases, weekly or even daily revenue might be suitable indicators. Year-to-date income could also be suitable. Imagine you found yourself in the following situation six months after setting the objective:

¹ Citigroup Annual Report, 2005, p 21.

Figure 5
Graph of Year-to-Date Sales



In this case, one can imagine feeling a little pessimistic about achieving total sales of \$100 million by month 12.

Analyze

The “analyze” step involves the regular analysis of observed metric values to derive an understanding of their behavior in order to more accurately predict their future possible values. To some extent, the level of sophistication of the analysis will be a function of the volume and nature of the data. By simply looking at Figure 5 above, one can make justifiable estimates of the value of “sales_YTD” in future periods. An estimate of 100 in month 7 would, based on the historical data alone, seem hard to justify.

An important task within this step is the validation of the data. Before you attempt to make a prediction, you should be sure you have correctly understood the past. Data should not be accepted as fact; rather it should be questioned as if it is an opinion—this approach, although pessimistic, will prove much more valuable in identifying errors in the data. Consider the case where you ask your IT department to provide you data on staff turnover over the past two years and they provide you the following:

Year	Month	Number Who Left	Number Who Joined
2009	Feb	3	4
2009	Jan	2	1
2008	Dec	1	5
2008	Nov	100	3
...
2007	Apr	3	3
2007	Mar	2	4

In this case it would be quite reasonable to question the figure for departures in November 2008. Less obvious would be the following:


Year	Month	Number Who Left	Number Who Joined
2009	Feb	3	4
2009	Jan	2	1
2008	Dec	1	5
2008	Nov	0	3
...
2007	Apr	3	3
2007	Mar	2	4

The figures above should be considered with the same skepticism as the data in the previous table. Generally speaking, values of 0 should attract your attention because 0 is quite often a misrepresentation of the fact, which is that the data was not recorded. It is not just the presence of a 0 that justifies the time taken to validate the data—it is the importance of the accuracy of the data that justifies the effort of validation. Accurate historical data is fundamental to your analysis, and your analysis is fundamental to the achievement of the strategic objective, so validation of the data is vital and must be considered mandatory.

React

The “react” step relates to the action taken in response to what is revealed in the “analyze” step. Two people should react to the analysis—namely, the enterprise risk manager and the owner of the strategic objective. Quite often, the initial reactions of these two people will be quite different. The role of the enterprise risk manager is to help the strategic objective owner fully comprehend the implications of the analysis. Consider the case presented by the following graph:

Figure 6
Observed and Forecast Metric Values




The immediate response from the enterprise risk manager should include concern over, firstly, the volatility of the (observed) metric value and, secondly, the discrepancies between the observed and forecast values. The likely initial reaction from the strategic objective owner will be delight at the fact that the observed value exceeds the forecast! So the enterprise risk manager must communicate to the strategic objective owner the risk of not achieving the objective implied by the historical data. The following facts can be extracted from the observed metric values:

1. The change in metric value has been negative in two of six occasions, which could be interpreted as a probability of 33 percent that the value will fall in any month.
2. The falls in metric value have been greater in magnitude than the rises.

It is the enterprise risk manager's responsibility to help the strategic objective owner understand the implications of the volatility of the metric value. In this example, one very significant implication is that it is hard to determine a trend, or derive a forecast, from the historical observations. Depending on the method you use, Microsoft Excel can produce a range of forecast values, as shown in Figure 7.

Figure 7
Forecast Values Based on Trend in Observed Values




In this case, Microsoft Excel has produced forecast values of the metric at month 12 ranging from around -2 to around 23!

4. Measurement of Stakeholder Value Added by Application of the SOAR Methodology

The value added by application of the SOAR Methodology can be directly measured as the difference between the value corresponding to the same percentile from two probability distributions—one representing the distribution of the dollar value of outcomes before the SOAR Methodology was applied and a second representing the distribution of the dollar value of outcomes after applying the SOAR Methodology. Consider the probability distributions presented in Figure 99.

Figure 9
Measurement of Value Added



The distribution labelled “without_SOAR” shows the probabilities of various outcomes (measured in dollar terms) before the SOAR Methodology was applied. This can be referred to as the “inherent risk” distribution. The other distribution, labelled “with_SOAR,” shows the probabilities of various outcomes following application of the SOAR Methodology. This distribution can be referred to as the “residual risk” distribution. Summary statistics of the two distributions are presented in TABLE 2.

TABLE 2
Inherent Versus Residual Risk

Statistic	Inherent Risk	Residual Risk	Difference ²
Mode	5	6	
Mean	4.7	6.1	1.4
Standard Deviation	0.65	0.66	
Skew	-0.45	0.29	
Minimum	3	5	2
Maximum	6	7.5	1.5

² Difference = Residual Risk – Inherent Risk.

The value added by the SOAR Methodology is shown in the “difference” column and has been calculated at several percentiles, namely the 50th (i.e., the mean), the 99th (taking the minimum as a proxy) and the 1st (taking the maximum as a proxy). The choice of the percentile at which to measure the value added is somewhat subjective. In this example, the value added by the SOAR Methodology is estimated at between \$1.4 million and \$2 million. Whilst it might be tempting to consider measuring the value added as the difference in the *modes* of the two distributions, I do not recommend this approach as the modes of the two distributions might have very different probabilities. The main goal of the SOAR Methodology is to create a residual probability distribution that is taller and thinner and located closer to the desired outcome than the inherent probability distribution. Taller and thinner are summarized by standard deviation, while proximity to the desired outcome is represented by the mean of the distribution, so to measure the value added as the difference in the means and including an adjustment for the change in standard deviation is appropriate. This is expressed in Equation 2.

Equation 2
Value Added by SOAR Methodology

$$V_{SOAR} = (\mu_R - \mu_I) \frac{\sigma_I}{\sigma_R}$$

where

V_{SOAR} is the value added by the SOAR Methodology

μ_R is the mean of the residual distribution

μ_I is the mean of the inherent distribution

σ_R is the standard deviation of the residual distribution

σ_I is the standard deviation of the inherent distribution

Conclusion

This paper has presented an advanced approach to ERM that significantly improves upon current approaches, largely due to two fundamental elements—namely, a focus on the management of strategic plans and, secondly, a heavy reliance upon probability theory.

The SOAR Methodology offers the following significant benefits over existing ERM methods:

1. By design, the method focuses on the management of risks associated with strategic plans and strategic objectives. This demands that the ERM officers and strategic objective stakeholders apply the same focus and a degree of discipline that they would otherwise be unlikely to apply.
2. The SOAR Methodology recognizes the difference between ERM and similar concepts like enterprise wide risk management and integrated risk management, and a less similar but often confused concept, namely performance management, and has been designed specifically for ERM.
3. The methodology includes an iterative process that involves data-based decision making and this process is unique to the SOAR Methodology
4. The value added by application of the methodology can be directly estimated, thus allowing the organization to determine whether or not its application is economically justifiable.

Greg Monahan is the founder of SOAR Advisory (www.soar-advisory.com) based in Sydney. He can be reached at gmonahan@soar-advisory.com.