A practical guide to understanding operational risk and how to manage it

MASTERING OPERATIONAL RISK

second edition

- Provides an invaluable framework for the management of operational risk
- Helps you identify and manage risk appetite
- Provides a practical approach to applying stress testing to operational risk
- Gives you a business approach to modelling operational risk
- Shows you how to change your culture effectively

TONY BLUNDEN
JOHN THIRLWELL

FT PUBLISHING
FINANCIAL TIMES

ALWAYS LEARNING

PEARSON

Mastering Operational Risk

(twice during the working week, once every three working months) than to percentages (40%, 1.67% respectively). It is often easier for a firm to articulate its likelihoods as time periods and for the operational risk manager to convert them to the percentages, which will be necessary for modelling the risk and its controls.

Ranges and single figures

Many organisations find it difficult to assess the risk likelihood and impact (and the control design and performance) as a single figure. The problem stems from the character of operational risk, which is naturally imprecise and variable. It is therefore problematic to attribute a single value to a dimension of operational risk assessment. A common way around this is to consider a range of values (£5m–£15m, rather than £10m). Single figure assessment can then be introduced after the firm has gained experience of the assessment process.

How to set the range

Once the decision to choose a range and time horizon has been made, the question arises as to the base used for the ranges. Many firms prefer to use gross revenues on which to base the range of values. These are useful because they can be directly affected by the business (the first line of defence) and therefore encourage embedding of the process. If net profitability is used, it must be borne in mind that it is more difficult for business heads to influence the costs allocated to them and they are therefore likely to be less willing to accept the ranges. However, particularly in industries with a very low profit margin, gross revenues may be inappropriate and profitability may be more relevant.

The beginning point of the highest range is often set at three or four months of gross revenues or profitability, whichever is appropriate. The top end of the lower ranges can then easily be set at one month and one week if four ranges are being used, the full set of four being: above three months; three months to one month; one month to one week; below one week. If five ranges are used, it is common to have an additional small range of two days, making the bottom two ranges one week to two days and below two days. You will notice that the top of each range is a multiple of around three to four of the one below. This is a useful rule of thumb when setting ranges.

One pair (impact/likelihood), two pairs (average and worst case) or three

Whilst most organisations start their risk and control assessments with assessing one pair of scores for a risk (such as impact/likelihood), some choose two pairs (adding average and worst case for each of impact and likelihood). A third level can be added, by taking the assessment of impact and likelihood to an 'extreme worst case'. In mathematical terms, average represents a 50% confidence level, 'business as usual worst case' takes the assessment to a 95% or 1 in 20 confidence level, whilst the third level reaches up to 99.5% or even 99.9% confidence (1 in 200 or 1 in 1000).

Two pairs help the firm to think about the expected and unexpected parts of a risk's distribution without making the mathematical part explicit. Three pairs enable a firm to begin to specify the distribution which constitutes the risk. A very occasionally used alternative is to specify a mean and a standard deviation, although this requires an understanding of mathematics which is generally beyond most managers.

Use of losses to back-test impacts and likelihood

Losses provide a real-life view to guide and challenge the subjective assessments which are made when considering the likelihood and impacts of risks. However, it is important to ensure that the causal analysis of the losses correctly identifies the risks to which the losses relate. The mapping of the risks underlying the losses to the risks which have been identified during the risk and control assessment is difficult but worthwhile as it enables a methodology to be established which tests the assessors' judgement against actual losses experienced by the firm.

This 'back-testing' technique should take account of any changes to the firm. An example might be an increase or decrease in staff numbers, which might affect the relevance of both the frequency and severity of losses to the firm as it currently exists. So when using losses, make sure that the risk profile when the losses occurred is the same as the risk profile at the time of the assessment. If not, adjustments will have to be made.

Heat maps

Heat maps give readily accessible and visual representation of the risk profile of a firm. They are often the first risk report seen by the board and, as such, must be positioned as the start of risk reporting and *not* the final risk report. They are helpful in allowing management to focus on the most significant risks to the firm, in the absence of any further data.

Figure 5.4

Heat map of residual risk assessment

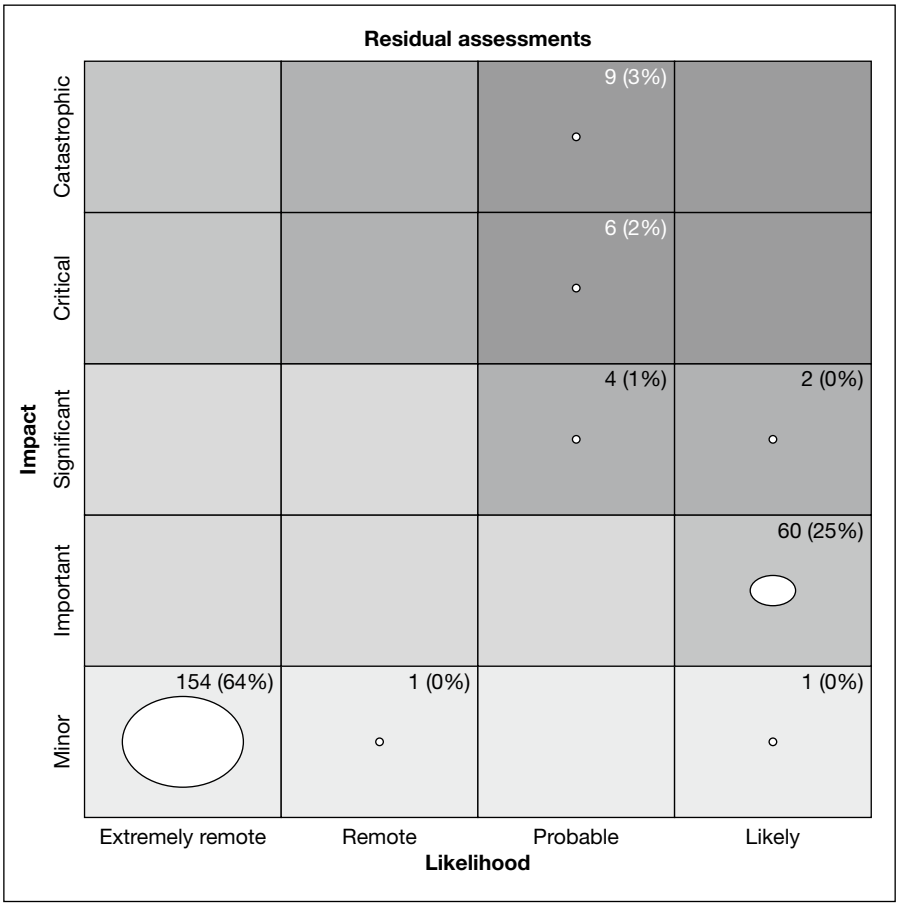

Source: Courtesy of Chase Cooper Limited

Heat maps often start as net or residual risk heat maps (see Figure 5.4) and then expand to cover both gross and net risk as the firm develops its expertise in risk assessment, as shown in Figure 5.5 later. Illustrating, by means of a heat map, the reduction to a risk due to the mitigating effect of the controls is helpful in visualising which controls are fundamental to reducing the risk profile.

OWNERS

Different levels

There can only be one ultimate owner of a risk and that person must be at board level. However, the board director responsible for the risk may delegate the management of the risk to another person who in turn may

delegate further. This can lead to confusion over who owns the risk. It is likely that those to whom the risk is delegated own only a part of the risk for which the board member is ultimately responsible. However, the risk and control assessment process will decompose the risk down to each level – strategic, process, activity. As a result, the person responsible for a particular process or activity which contributes to the strategic goal of the firm for which the board member is responsible will be clearly seen as the person who owns the risks inherent in the process or activity.

It is often also the case that the CEO is nominated as the owner of most of the risks when a board first considers its strategic risk profile. This must be challenged. Board members should take responsibility for their own risks, for example the sales and marketing director must take responsibility for risks relating to sales and marketing, such as mis-selling.

Risk owners

Risk owners may exist at several levels, although there is only one ultimate risk owner. The owner of the risk is responsible for measuring, monitoring and mitigating the risk, at the relevant level, within the risk appetite set by the board. The actual tasks of measuring, monitoring and mitigation are generally given to another member of staff. This does not reduce or remove the risk owner's responsibility for managing the risk, which is carried out through receiving and actioning reports from the staff to whom the tasks have been delegated.

Control owners

These are the people responsible for managing the mitigation of the risk through the operation of internal controls. Control owners are vital both in designing appropriate controls to mitigate the risk and in ensuring adequate performance of the control in line with the board's risk appetite. They are responsible for identifying any action plans necessary to increase the effectiveness of the control and are also responsible for implementing the action plans.

Liaison between risk owners and control owners

It is essential that there is good communication between the risk owner and the control owner. In an initial risk and control assessment it is often the case that the risk owner scores the control as less adequate than the control owner. This can be due to the risk owner not fully understanding the control, or to the control owner being too optimistic about the control's effectiveness due to a misunderstanding of the risk's likelihood and full impact. Both of

these are probably due to a lack of communication between the risk owner and the control owner. With good communication it is possible to design and perform controls to a level which matches the firm's risk appetite, i.e. at the most efficient level for the firm. If necessary, the head of operational risk will challenge and resolve the different scores.

IDENTIFYING CONTROLS

Suitable level

Just as identifying a suitable level of risk can be a challenge, so too can identifying the appropriate level of control. However, as controls are typically identified after risks, it is often easier to set control identification to the appropriate level. If the risk identification has been set, for example, at a business objectives level, the controls which are identified should be at the same level.

It is very easy to identify controls at a departmental or activity level and relate these to the business objectives of a firm. However, this should be avoided as there will be a mismatch between the level of the risks and the level of the controls. Additionally, it is important to identify and then score the strategic controls which are in place to mitigate the risks to the business objectives. If this is not undertaken a firm can be lulled into a false sense of security, believing that its business risks are well controlled by a considerable number of activity or departmental controls.

Independent controls

When identifying controls, we are seeking to identify the independent controls which mitigate a risk. Although there is some point in identifying linked controls, far more business benefit will be achieved through identifying and scoring controls which are independent of each other. Controls which are linked to each other, perhaps in a sequence, are only as good as the preceding control. This means that if the first control in the sequence fails, none of the other controls gives any benefit in mitigating the relevant risk(s). It is therefore vital that controls are checked to ensure that they are independent, otherwise they become another source of false security.

An example of three typical independent controls are those which might be considered to mitigate the risk of 'Failure to attract, recruit and retain key staff', as shown in Figure 5.6 later: 'Salary surveys', 'Training and mentoring schemes' and 'Retention packages for key staff'. Linked controls within this example may be 'Salary increases' and 'Title changes', both of which are linked with 'Salary surveys'.

Mitigating more than one risk

It is often said that a single control mitigates more than one risk. In principle this may well be true. In practice it is unlikely that the application of the control is exactly the same. Often the control is the same, but applied differently by different departments. For example, a staff appraisal is a very common control which mitigates the risk of 'Failure to attract, recruit and retain key staff'. However, the control is likely to be applied differently in different departments and the effectiveness of the control will vary considerably around the firm. The head of operational risk should therefore challenge, whenever it is suggested that a control mitigates more than one risk, in order to avoid two similar controls being mistaken for the *same* control. The risk and control assessment shown in Figure 5.6 provides examples of similar but different controls, such as staff training as a control to improve competencies, and training and mentoring as a control to mitigate key staff turnover. It is important to define not just the control but its purpose.

Controls are only one form of mitigation

Controls are the most common continuing method of mitigating risks. They are completely within the management's sphere of influence and in a firm practising good operational risk management they will be increased or decreased to reflect the sensitivity of the firm to a particular risk. In practice this rarely happens, in part because of inertia. Firms should be alive to change and accept it as part of everyday life, all part of a culture of continuous improvement.

Another method of mitigation for the firm is to transfer the risk to another party entirely, for example through insurance (see Chapter 12, Insurance). This enables a clear cost to be attached to the mitigation through the premium charged by the insurance company. It also explicitly limits the exposure of the firm to the excess, or deductible, applied to the insurance policy. However, there may also be a limit to the loss which the insurance company is willing to suffer and the firm will once again be exposed above this limit.

Another method of mitigation is to remove the risk altogether from the firm, for example by ceasing business in the particular product to which the risk is attached. This is, of course, an extreme move to take but may be justifiable in the circumstances. For example, it was reported that Goldman Sachs withdrew from some markets before its peers during the 2007–9 financial crisis and by doing so considerably reduced its losses.

Types of controls

Controls can be divided into four types: directive, preventative, detective and corrective.