THE
INSTITUTE OF
OPERATIONAL RISK

Institute of Operational Risk
Operational Risk Sound Practice Guidance

Risk Appetite

December 2009

Table of Contents

Title: Risk Appetite	Date issued: 01 Dec 2009
	Version: 1
File name: Risk Appetite Sound Practice Guidance V1	Update date: 01 Dec 2009

1. Introduction

In common with a number of aspects of operational risk management, risk appetite is an area that attracts differing views among practitioners. One of the reasons for this may be the relative immaturity of the discipline. Another may be the wide variety of contexts e.g. size and structure of organisations, complexity of product/service offerings, regulatory jurisdictions etc.

For these reasons the following summary makes no attempt to suggest a one-size-fits-all solution to any of the practical challenges an organisation faces. Rather, it aims to outline a variety of good practices from which may be drawn a collection of appropriate, relevant and proportional ideas.

Additionally, whilst the focus of this paper is the management of operational risk, it must be recognised that for an organisation to manage its risks holistically there would need to be an understanding of the inter-relationships between different risk types. For instance, it may well be the case that the emergence of an operational risk could precipitate a consequence involving another risk type and the combination could in turn lead to a reputational impact.

An aspect of operational risk appetite that is sometimes lost in the debate about bottom-up and top-down approaches is that operational risk is important to all organisations. It is therefore essential that whatever approach is taken, the Board and senior management are properly engaged in agreeing and monitoring the appetite for operational risk and setting acceptable, and unacceptable, boundaries for business activities and behaviours.

1.1. Definitions

Definitions can vary according to the context: industry sector (e.g. financial services, building, health); stakeholder perspective (e.g. external investors versus internal executive and management); risk type (e.g. operational risk versus credit or market risk).

Suggestions have been published by a wide spectrum of commentators including regulators, industry bodies, consultancies and academic establishments. One of the more generic definitions of risk appetite was published in BS31100: “the amount and type of risk that an organisation is prepared to seek, accept or tolerate.”

In theory, accepting any type of risk may carry some aspect of reward but it is important for operational risk practitioners to be aware of the view that this risk type differs from, for example, credit risk and market risk. Operational risks (e.g. IT systems failures or external fraud) may be inherent in operational activities but are rarely intentionally sought and have no material upside in terms of return/income generation. There are, however, cost/benefit decisions involved in defining an appropriate balance between accepting potential losses on the one hand and incurring costs of mitigation on the other (including associated operational inefficiencies that introducing a new control could involve).

Other distinctions are that:

- Whereas taking credit or market risks is likely to be encouraged *up to* the stated appetite level, operational risk is more likely to be mitigated *downwards* as long as the cost of mitigation does not exceed the expected loss (or the associated benefits). Risks may also be accepted for strategic reasons, for example operational risks associated with launching a new product, because the revenue stream should exceed any associated losses.

- Credit or market risks can be capped, including by systems, but operational risk cannot be constrained in the same way. Having entered a business, a firm is committed to managing the associated operational risk unless it withdraws from the marketplace altogether.

On this basis an organisation's definition of Operational Risk Appetite (ORA), which could be documented explicitly in a statement or within policy or strategy, or implicitly within operational standards, might more accurately be described as the operational risk it is prepared to tolerate.

1.2. Regulatory context

Whilst the management of risk should be positioned first and foremost as a means of supporting the achievement of company objectives, those are likely to include compliance with regulatory requirements. There has been much published by way of "guidance" (e.g. within the FSA Handbook) which provides a clear indication of the regulatory expectation that firms establish an ORA and that the ORA is likely to provide an important mechanism for demonstrating compliance with "Senior Management Arrangements, Systems and Controls" requirements and the 'use test'. Regulators take a particular interest in risk appetite because of its importance to governance.

1.3. Appetite and tolerance

In simple terms, expressing ORA is a question of defining what is acceptable to an organisation and what is not. This could be achieved by deciding, for each type of risk, what is acceptable, what is unacceptable, and the parameters of the area between those two i.e. what is tolerable.

It is common practice when monitoring performance against ORA to assign a "RAG" status (Red, Amber, Green). When doing so, the definitions are generally accepted as:

Status	Meaning	Required action
Green	Acceptable	No action required but continue monitoring
Amber	Tolerable, but action required to avoid a Red status	Investigate (to verify and understand underlying causes) and consider ways to mitigate / avoid within a specified time period
Red	Unacceptable. Urgent attention is required.	Investigate and take steps to mitigate or avoid within a specified short term

This approach can be applied across the range of operational risk framework components (including risk and control self assessment, internal loss event reporting and scenario analysis) and provides a clear indication of proportional response to the perceived materiality of the associated risk. Specifying a timeframe for resolution can emphasise the perceived urgency and significance of the underlying issue. This makes good business sense and promotes a consistent understanding across the organisation.

1.4. Objectives and purpose

From a business perspective there are a number of benefits to be accrued by defining ORA:

- Enabling the Board to exercise appropriate oversight and corporate governance by defining the nature and level of risks it considers acceptable (and unacceptable) and thus setting boundaries for business activities and behaviours;

- Providing a means of expressing senior management's attitude to risk which can then be communicated throughout the organisation as part of promoting a risk aware culture (for example clarifying the relationship between risk and profitable business);
- Establishing a framework for business/risk decision making (which risks can be accepted/retained, which risks should be mitigated and by how much) which ensures an appropriate balance between being risk seeking and risk averse. ORA can represent a powerful tool for managing the business, for instance not only where a breach occurs but also where a potential breach can be predicted and averted;
- Improving the allocation of risk management resources by bringing focus to higher priority issues (i.e. areas where appetite thresholds are under threat);
- Ensuring an enhanced view of risk expenditure so that the cost of risk does not exceed the benefits;
- Aligning strategic goals and operational activities through optimising the balance between business development/growth/returns and the related risks inherent in pursuing those goals. This will enable the strategy to be put into effect;
- Encouraging more conscious and effective risk management practices e.g. prioritising risk related issues for escalation and for response.

All of which can help to enhance performance and thus enhance value to stakeholders.

2. Process for setting risk appetite

There are a number of considerations involved in the approach to setting ORA, which can be expressed in a variety of ways e.g. through key risk indicators, risk and control self assessment, losses as well as broad qualitative statements. Good practice might well involve some combination of all the following alternatives.

2.1. Top-down and bottom-up

In the context of sound corporate governance it is clear that ORA must be owned by the Board and established with their full engagement. The remainder of this section deals with the ways in which a comprehensive ORA can be constructed.

Conflicting views have been expressed as to whether operational risk appetite should be set using a top-down or bottom-up approach. A number of surveys of operational risk practitioners (e.g. Marsh and AIRMIC 2009) have reported wide support for a hybrid approach.

In any case it would seem sensible to start with a top-down cascade from the Board (which has an enterprise-wide perspective) in order to set the cultural context for the organisation, to provide a basis for oversight and governance, and to facilitate alignment to strategy. This will often be expressed in qualitative terms but may also include quantitative measures e.g. relationship between expected/unexpected losses and Profit Before Tax.

A number of factors favouring a complementary bottom-up approach, where limits are defined at lower levels in the organisation in line with operational activities, are:

- Ensuring coverage of locally relevant risks and factors;
- Facilitating the setting of proportional tolerance thresholds;

- Promoting buy-in by involving management at all levels;
- The potential to forge links to personal performance objectives and related rewards (in itself a recognised way of embedding operational risk management).

2.2. Qualitative and quantitative

Qualitative expressions of ORA (i.e. without any reference to quantification) can emphasise the relationship between risk and business management. This is often regarded as the best way to describe the attitudes and behaviours of the organisation as a whole – in other words, its "risk culture". This would be achieved through a series of statements, for example:

- Recognition that some risks, however unwelcome, are unavoidable (e.g. terrorism, natural disasters, consequences of economic downturn). It is therefore accepted that some level of such risks has to be tolerated to avoid stifling or curtailing business operations.
- It is sensible to accept risks where the cost of mitigation/avoidance exceeds the expected loss provided the resultant risk is not too high.
- Risks will be accepted when the estimated losses are within prescribed tolerance levels.
- Unacceptable behaviours might include: knowingly breaking the law; knowingly breaching regulatory requirements; knowingly breaching company policy; damaging the environment; disrupting service to customers etc.
- Unacceptable risks could include operating within specific countries, selling particular products etc.
- As difficult as it is to define damage to reputation it can be useful to use qualitative measures to describe events that may lead to unacceptable damage to reputation or loss of trust with stakeholders.

Sources of reference for qualitative ORA statements can include communications from the CEO and Board (aimed at internal and external audiences), business strategy and policy papers.

Quantitative expressions of ORA, on the other hand, involve hard data, usually having roots in business management information which could be any combination of KPIs (key performance indicators), KRIs (key risk indicators) or KCIs (key control indicators).

Such measures are usually accompanied by thresholds so that it is immediately apparent when a breach has occurred or is imminent. The concept of setting zero thresholds may seem impractical but they can have a cultural purpose in reinforcing the message that it is not appropriate to accept avoidable losses without question.

Examples of quantitative measures include:

- The amount of economic or regulatory capital allocated to operational risk;
- Delegated limits of authority beyond which subordinates have to escalate for approval;
- Performance levels e.g. no more than xx% chance any business critical system is unavailable for more than one day in any one year;
- Each component of the operational risk framework:
  - Losses – based on budgeting, aggregate annual amount by business area/loss type and/or sensitivity i.e. an adverse trend of 5% may be acceptable, 10% tolerable, but 15% unacceptable. Note that minimum reporting thresholds imply single loss limits,

but there may be also be aggregate limits as in the case of an annual budget. The aim is to cover both high volume/low value and low volume/high value types of events. Reporting and escalation thresholds imply clear expressions of appetite i.e. below the threshold = acceptable, above that level = tolerable or unacceptable.

- Risk/control assessment: by establishing boundaries on a matrix of likelihood and impact to distinguish acceptable/tolerable/unacceptable levels of residual risk;
- Key Risk Indicators – where thresholds would be set in units relevant to the KRI metrics e.g. number, financial value, percentage, variance etc linked to measurements.

By embracing all aspects of the framework (including forward looking as well as historical perspectives), an organisation can establish longer term as well as current/short term ORA settings.

The RAG status described earlier can also be applied to all the foregoing to achieve an aggregate view of past/present/future performance against risk appetite.

2.3. Absolute and relative

The distinction here is that absolute measures are fixed and relative measures are variable, moving in proportion to some other benchmark. For example:

- Qualitative statements might include a zero tolerance for breaking the law (absolute) and a tolerance for financial crime that is expressed not in value but in comparison with a peer group (relative). The implication of the latter is that the organisation does not on the one hand want to suffer greater losses than its competitors (possibly attracting negative reputational impact in the process) but on the other hand does not want to encourage operational inefficiency (possibly impacting customer service) by being over-controlled.
- Quantitative measures could include monetary values, volumes, times etc (absolute), or losses as a proportion of PBT (profit before tax)/capital; the ratio of complaints to active customers; project over-run as a percentage of plan (relative).

Either/both approaches can be considered more appropriate for different metrics. In many cases a relative measure can be useful to impart a sense of context at the same time as identifying an adverse trend.

2.4. Setting thresholds

A key element of the setting up process is to establish agreed thresholds. These provide specific definitions, for each expression of appetite, of what constitutes “acceptable” becoming “tolerable” or “unacceptable”. For qualitative expressions of appetite this could simply be a matter of stating explicitly what is acceptable and what is not.

There are various ways of depicting a distribution curve for operational risk. The following example is for illustrative purposes. Thresholds can be set at any point along the curve of risks from high likelihood low impact to low likelihood high impact, that is:

- The value of expected losses that can be sustained within normal trading (i.e. within existing budget or P&L capacity) and will be tolerated, being within stated risk appetite;
- The value of expected (and to some extent unexpected) losses where the risk is insurable and the amount not covered (either before cover is provided or above the maximum amount covered) is tolerated within risk appetite;

- The amount of capital allocated as a buffer against severe unexpected losses, unsustainable within normal P&L resources.

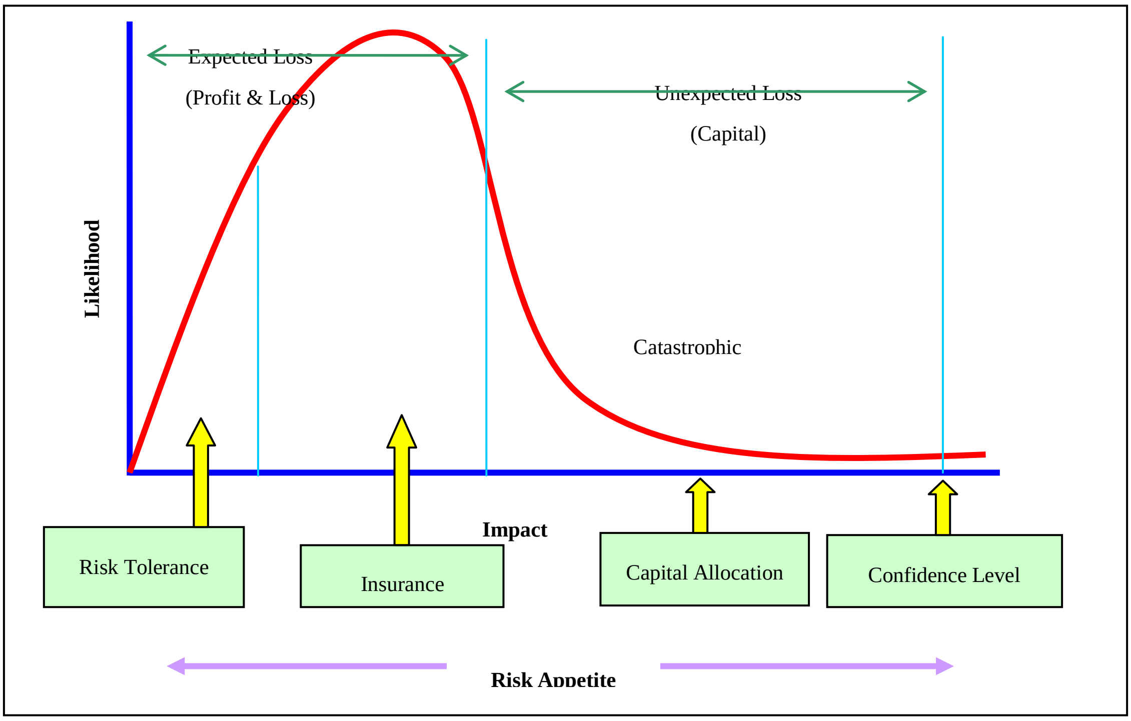

Note: Setting confidence levels is also an expression of appetite.

Where quantitative data is involved it may be appropriate to express tolerance within a range of values. This may include both positive and negative variance. For example, if a business is monitoring the number of employees, a significant variation above or below the target optimum may be an indication of different kinds of adverse consequence. Too many employees could signal inefficiency, wastage and unnecessarily high costs, whereas too few might lead to failures in procedures and controls or a decline in customer service standards – see illustration below. Being alerted to either possibility is helpful from business and risk perspectives.

Target no. employees	Amber threshold	Red threshold	Actual						
			Mth 1	Mth 2	Mth 3	Mth 4	Mth 5	Mth 6	Mth 7
987	+/- 5%	+/-10%	**989(G)**	**935(A)**	**964(G)**	**996(G)**	**1041(A)**	**1012(G)**	**etc**

Identifying the “right” thresholds is ultimately a matter for the respective business line to determine, drawing on practical experience of the context and expected future developments. But the decision can be informed by reference to any relevant data, be it historical or predictive, internal or external (i.e. benchmarking against comparable organisations or industry standards). Even so, if the procedure for review/approval of thresholds is sufficiently flexible, initial attempts can be fine tuned in the light of use in practice. Thresholds should be sufficiently sensitive to provide “early warning” of potential appetite breaches, but not so hypersensitive that alarm bells ring needlessly.

As an example of how historical data could be used as a start point, a review of the previous 12 months' recordings (to account for seasonal fluctuations) could be translated to tolerance thresholds as follows:

- The mean of recorded data could be adopted as the threshold for moving from Green to Amber on the basis it indicates above "normal" levels and is worthy of investigation.
- The worst recorded position could represent the threshold for moving from Amber to Red, indicating there is no appetite for the position to be even worse in future than previously experienced.

If thresholds are set on a bottom-up basis it would be prudent to ensure an appropriate level of governance by means of review and approval by some higher authority (e.g. a Risk Committee). This can achieve a number of benefits:

- Informing senior management of the levels of risk that subordinates are proposing to accept and providing an opportunity for challenge of such positions;
- Enabling a wider perspective of the individual business/department tolerances (for consistency where appropriate) and aggregate tolerances being adopted across business lines;
- Ensuring that bottom-up tolerances align to or are within (but do not exceed) top down parameters.

3. Implementation and practical application

Once ORA has been defined, documented and communicated to decision makers at all level in the organisation, the focus can turn to the practical application of related procedures.

3.1. Monitoring

There would be little practical purpose in defining ORA if the operational reality is not then checked against the defined tolerances. A key aspect of monitoring is the objective to provide early warning of emerging issues and, to be effective as a management tool, it is more than a mechanical procedure, it also requires moderation e.g. arbitration through interpretation.

There are two distinct steps involved in monitoring procedures:

- The first is arranging for the required data to be reported by the appropriate party at an agreed frequency. From the outset it is important to take all reasonable steps to ensure the integrity of the data i.e. in respect of completeness, accuracy and timeliness. It would be helpful to validate ORA reports against Key Risk Indicator and internal loss event information, where these are reported separately, to ensure consistency.
- The second is the crucial stage of converting data to information by adding context and interpretation (e.g. how the data compares with business performance metrics, whether the data is suggesting the emergence of increased or reduced risk i.e. whether movement is relatively positive or negative). This entails the identification and investigation of adverse variances and trends and in particular analysing the underlying causes. Some key considerations include:
    - Whether recurring “ambers” are reflecting a static or worsening position;
    - Whether a cluster of “ambers” represents an overall “red” in aggregate;
    - Whether recurring “greens” may suggest thresholds are not sufficiently sensitive and should be reviewed

Clearly, monitoring , performance against qualitative statements of ORA is more challenging, but should be attempted to achieve the benefits of early warnings. Often monitoring aspects of the ORA may just provide the focus to ensure that the right conversation takes place. The value of this should not be underestimated.

3.2. Aggregation and reporting

One of the challenges in aggregating upward-flowing information is the potential for distortion or misrepresentation. A “Red” status in a small business unit may be of little or no significance to the Group Board and there is a danger that the meaning and value of the “unacceptable” flag will become diluted. On the other hand it would be completely inappropriate for the business unit to simply adopt group level tolerances – because then everything would be perpetually “Green” and apparently require no attention whatsoever.

One solution is to construct a conversion table (based perhaps on business scale criteria) so that a business unit Red will always be so locally, but in reporting through successive layers of senior management will become “Amber” or even “Green” to provide more accurately a sense of proportion in the changing context in which it is viewed. For example:

Risk:	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Business RAG	**R**	**A**	**A**	**R**	**A**	**R**	**R**	**A**
	Business A = 80% of Division X		Business B = 20% of Division X		Business C = 20% of Division Y		Business D = 80% of Division Y	
Division RAG	**R**	**A**	**G**	**A**	**G**	**A**	**R**	**A**
	Division X = 80% of Group				Division Y = 20% of Group			
Group RAG	**R**	**A**	**G**	**A**	**G**	**G**	**A**	**G**

From a Group perspective, the most significant risk is Risk 1, in the largest business of the largest division. But at business unit level all high risks (Red) – the most significant locally - would receive an appropriate level of attention.

A second point about reporting appetite-related information is the need to be clear about the objective, which could be one or more of the following:

- For information only, as a source of reference in case of future need.
- For governance purposes, to enable and demonstrate oversight.
- To secure approval to a proposed/recommended position.
- To obtain a decision or request action.

It is important to ensure that reporting of ORA information is not perceived as a vehicle for presenting too optimistic an interpretation of positions and trends. The real value is the provision of early warnings which can encourage timely management intervention and action to avert emerging issues.

3.3. Management and decision making

Last and by no means least is where ORA processes and procedures reach a logical and meaningful conclusion – the point at which business and risk management is exercised.

"Ambers" and "reds" have to drive action of some kind and the decision to be reached is a choice between:

- Accepting the breach. After weighing all the evidence, it may be the case that a particular breach could involve a truly one-off exception. In other cases it may be appropriate to review and re-calibrate previous tolerance levels if they are believed to be too sensitive. It is recommended that such acceptances should be recorded and revisited regularly.
- Taking steps to mitigate/avoid and prevent a recurrence. This is likely to be the most appropriate response to a breach of ORA and will require approval to implement some additional or alternative control measures.
- Some intermediate management action – for example, conducting extended or more intense monitoring, undertaking additional root cause analysis or investigating the cost/benefit of mitigation options.

In each case the business, risk and regulatory expectations will be met. The organisation’s senior management will be aware, be informed and be involved in the decision making process. It should then be a straightforward matter to assemble evidence of such activity to demonstrate adherence to the “Use Test”.