Actuaries Summit

Take the Lead

17-19 May 2015
Grand Hyatt Melbourne

Actuaries Institute

Dark risks versus ERM upside:

A collection of risk management thoughts

Prepared by Greg Martin and Simone Leas

Presented to the Actuaries Institute
Actuaries Summit
17 – 19 May 2015
Melbourne

Keywords: Risk Management, Risk Classification, Risk Measurement, Risk Reporting.

Institute of Actuaries of Australia
ABN 69 000 423 656
Level 2, 50 Carrington Street, Sydney NSW Australia 2000
t +61 (0) 2 9233 3466 **f** +61 (0) 2 9233 3446
e actuaries@actuaries.asn.au **w** www.actuaries.asn.au

Synopsis

Risk management has come a long way in recent decades to modern enterprise risks management (ERM) concepts focused on adding value through risk management optimisation, including "capturing upside risks". However, the authors see significant "clunkyness" in a number of areas of common ERM practice. Three of these areas they address in this paper. Firstly, the authors propose a different perspective on considering the universe of risks faced by a business distinguished by the risks' particular characteristics (including value adding versus non-value adding risks, systemic versus non-systemic risks) which they have found useful, inter alia, in discussing and defining Board risk appetites, understanding issues around risk behaviours and modelability, and what "capturing upside" infers from this perspective. Secondly, the authors outline a bottom-up approach to operational risk measurement that they suggest is practical for day-to-day risk management activities, and they provide a reconciliation between this and a top-down capital reserving approach (including the "dark risks" gap) that others may find useful in considering the difference between risk management, risk measurement and capital reserving in this complex area. Finally, succinctly communicating an organisation's risk profile can be challenging and something with which many practitioners seem to struggle. The authors share one of their reporting tools – a high level risk map "on a page" - that they have found useful in communicating overall risk profiles to Boards and senior management and which they hope others might find of interest.

Table of Contents

1 Introduction

1.1 A brief history of risk management (context)

Risk Management has implicitly been around for thousands of years but it was only fairly recently in the 1960s that the subject was "formally named, principles developed and guidelines established"[1].

Initially hazard risks (e.g. fire) were the primary risks managed by organisations. Risks were identified, assessed and managed, and ultimately mitigated, usually by taking out insurance. Often insurance brokers were the first risk managers. Financial risks grew in importance in the 1970s as the impact of volatility in asset prices, interest rates and exchange rates caused organisations to focus on these risks. Tools for managing financial risk were developed and experts in these tools (often investment specialists) became the risk managers of these risks. Within insurance companies, managing insurance risk has always been a core discipline traditionally carried out by actuaries and underwriters.

By the 1990s risk management had grown into vital parts of some organisation's operations and strategy, with regulators emphasising and progressively focusing on these risk areas (due in part to high-profile cases of organisational failures such as Barings bank and Enron).

Up until this point, while risk management was growing in application, its implementation was typically somewhat siloed and disjointed.

1.2 Rise of ERM (our aim)

Progressively, organisations, experts and regulators have recognised the importance of more holistic risk management. This recognition gave rise to the modern concept of Enterprise Risk Management (ERM). A key idea of ERM is the coordination of risk management which aims for risks to be consistently managed across an organisation, to a consistent standard, with senior management oversight across the whole business.

A fundamental aim of ERM is to add value through optimising overall risk management across an enterprise relative to its business objectives and risk appetites.

1.3 Where we are now? (some issues)

Nonetheless, it seems to the authors of this paper ("we", "us") that much current ERM practice remains somewhat "clunky" relative to ERM's aim.

On the one hand, we see a tendency for risk management to still be based on, or rooted in, much of its (bottom-up) siloed or piece meal evolutionary history. Many different risk types and categories of risks are identified, with a range of frameworks applied, with much discussion, analysis and modelling of individual risks, and shoehorning of risks into existing risk categories. While considering risks by cause aids practical discussion and management:

- There often seems to be less analysis or discussion of what differentiates risks or unites them in terms of their behaviour, the way organisations treat them or what risks may most appropriately be grouped together for consideration.
- It does not help top down consideration of how or if the universe of risks are captured, nor help Boards express their risk appetite in a logical, top down reasoned fashion.

On the other hand, we also see (top-down) approaches aimed at managing all risks "equally", or where risks are differentiated (in term of appetite and/or management) a lack of a coherent rationale underpinning the differentiation – which at the least limits generalisation of a Board's approved risk appetite statement and management approach.

On two other fronts, it seems to us that:

- There is often some confusion in some risk areas between "risk management", "risk measurement" and "capital reserving". We see this in particular in the area of what is generally referred to as operational risk management.
- Communicating risk status and themes can be a challenge. A large, complex financial institution inevitably must manage a large number of risks. It is not uncommon to hear of "200 page" risk reports. We have heard more than the occasional mention of dissatisfaction from directors and senior management trying to see the wood for the trees.

1.4 Outline of this paper (some thoughts and suggestions)

This paper comprises three areas of thought on the above in three sections.

Firstly, in Section 2 we outline a perspective on risk examination and classification that applies an alternative top down approach that we have found helpful in our work with:

- Discussing and distilling Board risk appetites;
- Constructing overall risk management frameworks;
- Considering the effective mix of processes, culture, controls, "safety nets" and mitigants in managing the different categories of risk;
- Understanding and explaining how and why different approaches to risk measurement and modelling are effective and relevant to different risks;
- Where and what the "upside risks" are to capture.

Secondly, in Section 3, we set out a bottom-up approach to measuring operational risks that we suggest may help in considering (assessment, measuring and managing) operational risks in a day-to-day practical sense.

We also outline how this bottom-up assessment approach can be aggregated and adjusted to a total operational risk capital reserving approach. This includes allowance for unknown unknowns (or "dark risks") which a bottom up approach can't foresee, and how this may be reconciled to a top-down (industry) data based capital reserving model approach.

Finally, in Section 4, we illustrate a simple risk status map "on a page" that we have developed and find useful in communicating the overall risk status and key risk themes, trends and connections for an institution. We would be pleased to receive comments and discussion on this.

1.5 Notes and Appreciation

We do not claim originality or invention of all of the content of this paper. Much is based on the thoughts and research of others that we have seen over our time as risk managers. Indeed, as practitioners, it is difficult to distinguish one's own thoughts from those collected from others, text books, seminars, committee discussions, papers reviewed, etc over the years. Nonetheless, some of it reflects our invention and we hope it provides some different or at least novel views or perspectives that others may find useful.

It is noted that in the interests of ease of explanation and discussion, this paper has mainly been written from a shareholder risk perspective. It is acknowledged that other stakeholder perspectives are also relevant, especially depositors, policyholders, unit-holders and members with respect to financial services entities. The propositions in this paper can be applied equally to all stakeholder perspectives and indeed the differing perspectives of different stakeholders (e.g. one stakeholder "value adding risk" may be another's "non-valuing adding risk") can be instructive and relevant, for example for risk appetite construction.

The authors would like to thank Rob Curtis for the invaluable comments he made on our draft paper which helped us considerably. Nonetheless, the views expressed in this paper are those of the authors alone and in respect of the specific subjects of discussion. The views may not represent those of our employer or the Actuaries Institute.

2 Risk Identification & Classification

2.1 Introduction

There has been much work over the past circa half century to identify and name risks, and this has typically been based on identifying risks by their particular source. For example:

- The risks associated with the market (changes in asset values, currencies, interest rates, and exchange rates) have been identified and are commonly referred to as market risk.
- The risks arising from failure to collect funds from creditors and counterparties were identified and called credit risk.
- The risks arising from losses due to actual experience being different than that assumed when an insurance product was designed and priced were identified and called insurance risk.
- The risks arising from losses resulting from inadequate or failed internal administration and manufacturing processes, people and systems are typically collected under the title operational risk.
- The risks arising from failing to meet the organisation's strategic objectives was called strategic risk.

Identifying and classifying risks this way, by their source, is natural and has advantages, including ease of understanding, but there are some limitations to using this underlying approach to ERM.

The current common approach is fundamentally a bottom up approach to identifying and classifying risk and can result in ambiguity. Robert Chapman in his book ***Simple tools and techniques for Enterprise Risk Management*** commented on this limitation (page 131) *"Due to the nature of risk, the boundaries between classes are sometimes not clear and each business must decide for itself where sources of risk will reside in their bespoke taxonomy."*

More recently, the Actuaries Institute Life Insurance Risk Appetite Working Party's paper titled ***Developing the Risk Appetite Framework of a Life Insurance Business*** also observed (page 15) that *"In some cases, certain risks may be classified under several categories. For example, aspects of Market Risk related to counterparty default may be considered as part of Credit Risk instead. In relation to Operational Risk, it is important to recognise that the above definition is very broad and as a result may overlap with other risks unless defined more clearly. For example, any failures which may occur in the claims processes could be defined as part of Insurance Risk or Operational Risk."*

Another example of the consequence of this approach is the Canadian Actuarial Society's Operational Risk paper which takes 15 pages (page 15 to 30) to define its interpretation of operational risk[4]. It is not clear to us how this ultimately helps in risk identification or management for a specific situation, but does serves as an example of how complex the industry has made risk classification and how much time is spent on this matter which might arguably be used otherwise.

Another issue with the historic approach is that as new risks arise or existing ones rise in importance, for example cyber risks and terrorism risks, there is a tenancy to return to amend classifications to shoehorn the new risk into one existing classification or another. Or for others wanting a new hobby horse or activity opportunity to create a new "risk classification" resulting in an endless array of "frameworks" and countless "treatments" with no uniform common language. We are not suggesting that new material risks should not be considered or receive focus and attention, far from it, but ask if the current approach and some of the activity is optimal.

In doing all this it is not clear the universe of risks is even covered.

Furthermore, this approach does not seem to directly help with issues such as setting and articulating risk appetites; what the connections and

commonalities of the different risk "classes" or "types"; or help explain the "why" or "what" are the connections, linkages, different appetites, approaches to measurement or management.

This complexity begs the question: Is the industry looking at all this the right way? Is the industry asking the right questions? Would a different perspective help?

This section of our paper sets out an approach to considering risks that we have found helpful for financial services businesses when discussing risk appetites with Boards and considering approaches to risk assessment, measurement and management. The approach also helps generalise ("extrapolate") risk appetite statements and preferences to apply across a broader range of situations.

2.2 *Another perspective on risk categorisation*

The universe of risks of an organisation can range across:

- Different characteristics, e.g. different risk-return aspects;
- Different importance to different stakeholders (shareholders vs. customers);
- Different considerations in terms of risk appetite and tolerance; and
- Management at different levels within the organisation.

An alternative view of risks could be to consider the universe of risks categorised by their different particular traits.

2.3 *Value Adding Versus Non Value Adding Risks*

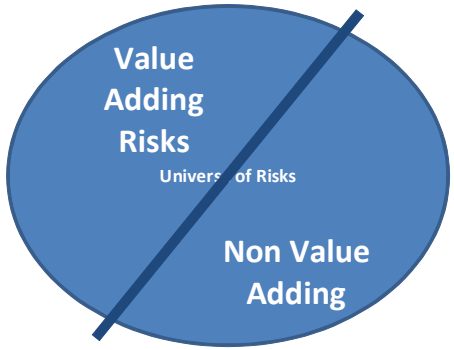

One risk trait that we find particularly interesting and relevant is to consider risks categorised by those that provide an economic return to an organisation for bearing them, versus those that do not earning an economic return for bearing. These two risk traits and the management approach to them would seem to us should be fundamentally different:

- Those that generate an economic return are typically related to "the business" of the organisation. For example an insurer taking on insurance risk or a bank taking on credit risk. While these risks need to be managed, generally speaking more of these "wanted" risks comes with more business and so "more is good".
- Those risks that generate no economic reward are typically related to those "unwanted" risks that unavoidably come with being in any business. Management of these, in our view, should be focused on reducing them subject to cost-benefit constraints. Irrespective of

"appetite", if the value gain from reducing a non-value adding risk is more than its cost of reduction, it should logically be reduced. For these "less is good".

Aspects of these risks and their management vary as a consequence of the above, often having almost diametrically opposite characteristics of each other. One should expect risk appetite statements to, explicitly or implicitly, recognise this risk dimension in a difference in approach to these two fundamental types of risk. Indeed, one may suggest the identification of the value adding risks would seem to be pertinent to the articulation of the objectives and business strategy of an organisation with respect to which its ERM framework should operate.

There would also seem to be some other interesting observations about value-adding versus non-value-adding risks.

2.3.1 Value Adding Risk Features

It seems to us that the risks that can be harnessed to provide economic return are those that can be reasonably priced. For an insurer, what is typically categorised as "insurance risk" would fall into this category.

- These risks tend to have a substantial degree of predictability and repeatability, with much of the risks reasonably identified, controlled, diversified (pooled), modelled and priced. One may call them, generally speaking, tractable.
- It is of course arguable if there is not a "chicken and egg" aspect here. The risks one can build a sustainable business or industry on would generally need to be tractable. Equally, if one has a business or industry based on certain risks, one would allocate considerable resources over time to control and manage them, and develop complex models for them, so they are tractable. Indeed, one would seek to see these risks not change overtime (and especially not "reduce").
- Insurance is based on people dying, becoming disabled, having car accidents, etc on an ongoing basis so insurers can collect, pool and price the risks effectively for the community. Insurers' reaction to a car insurance claim is not to call it an "incident", investigate its "root cause" nor seek to "stop it happening again". The behaviour is the exact opposite.
- The underlying risk process is generally not materially related to the entity incurring the risk. Any insurer insuring someone against death with the same sum insured would have essentially the identical risk. The administration system used by the insurer doesn't change the underlying risk.

So for these risks, insurers do not want the underlying risk process to change. Insurers typically have good data, good models and good measurement tools. They typically have good risk management/hedging tools which allows them to "dial up" or "dial down" the risk, and they seeks to optimise their risk exposure. Industry data is typically fairly directly relevant and scalable to the individual institution.

2.3.2 Non-Value Adding Risk Features

Non-value adding risks would seem to have virtually the opposite features on the dimensions considered above. What is generally categorised as "operational risk" would fall into this category.

- Organisations constantly seek to remove these risks on a cost-benefit basis. Any significant risk incident invariably leads to some change, so tomorrow is different to yesterday. Historical data and outcomes cannot easily and reliably be applied to the future.
- The risks often change from business to business, and scale with the size of a business. A fraud can only be as big as the money in the bank account. The bigger the business, the bigger the bank account, the bigger the potential fraud.

- If value adding risks seem to display a self-fulfilling "chicken and egg" outcome, non-value adding risks seem to display more of the "travelling back in time paradox". The better one identifies a risk, the more it is changed/controlled/mitigated and less clear the residual risk becomes.

Non-value adding risks are typically not easily modelled, measured or priced. They can be ad hoc, unpredictable and non-repeating at an organisation level. Indeed, we would suggest that these risks by their nature are not so much "statistical" in their behaviour but are more mathematically "chaotic". Attempting to apply statistical modelling concepts based on historic observations to a chaotic process is arguably problematic.

This is not to say that examining non-value adding risk outcomes and data at a macro industry or economy level, and/or building simple models of potential outcomes, is without merit. Rather, we question:

- The reality of scaling industry data/models to individual entities using anything more than the most basic approach. For example if the range of historic frauds in the insurance industry at the 99.5% confidence is x% of premium of an affected entity, then a crude individual insurer risk measure may be x% of their premium.
- The reality of scaling up entity derived risk measurements in this area based on entity only data on historic risk events.

2.3.3 Risks in the Twilight Zone

One may ask if for an insurer most of what is typically labelled "insurance risk" would fall into the value-adding category, and most of what is generally labelled "operational risk" is non-value adding, where does things such as "asset risk" or "asset-liability mismatch" sit? Is there a third category here?

We do not believe that this is the case. The twilight zone or horizon seems relatively sharp to us. However, the answer as to what they are may be different for different entities. Furthermore, there may be some interaction (correlation) in risk outcomes between value adding and non-value adding risks (e.g. a valuation fraud that crystallises with a large market movement).

Businesses that believe they have skills to add shareholder value via asset management may see "asset risk" as a source of value add. This should be reflected in their business model, business objectives and risk appetite. Other organisations may see asset risk as something that simply comes with investing their reserves and may not seek to add significant shareholder value (or they have shareholders that seek that value add elsewhere). This results in a different business approach and risk appetite.

The key issue here is to consider exactly what risks add value (and therefore all the rest do not). The business plan, business objectives, risk appetite and risk management framework should then respond accordingly.

2.4 Systemic Risk and Entity Specific Risk

Another risk trait that we find interesting is to consider risks by those that impact an organisation in isolation and those that are likely to impact a number of entities concurrently and/or similarly.

Entity specific risks are those that are likely to impact an organisation in isolation, independently of other organisations. They can be risks that come

from within the organisation or risk outcomes that come from outside but which impact the organisation alone. They can relate to value adding or non-value adding risks:

- Value adding: The part of claims experience that is simply statistical (random) noise. The large, single retained claim.
- Non-value adding: an isolate fraud event, a unit pricing error.

Systemic risks are those which are likely to impact a number of entities, the broader industry or economy overall in a similar way. They are, essentially by definition, risks that originate from outside of the organisation. Again, they can be value adding or non-value adding:

- Value adding: systemic claims variation (e.g. disability claims in an economic down turn), pandemic, regional earthquake, or bad car insurance claims in a very wet weather year.
- Non-value adding: damaging legislative or regulatory change, economy wide market or credit event (for those entities that do not see asset management as value adding).

A board's risk appetite for systemic versus entity specific (non-systemic) risks can vary significantly. This may not be so much a willingness to take on systemic risks, but rather a lack of comfort with material non-systemic risk.

While sometimes this is portrayed as the "regulator test":

Q: "What is your appetite for being called down to the see the regulator", A:"Depends; are we alone or with the rest of the industry?"

In reality we see that there is more to this risk trait.

- Consider the impact of two risk events: A severe pandemic on a life insurer: all life insurers would be affected. A likely regulator and government response will help the industry through (as it should). While there would be short term losses, the individual company is not isolated, and the overall industry will likely recover together. It would be economically costly for governments to ask the whole industry to hold capital for 100 years to deal with the most extreme event.
- A very large fraud that involves material financial loss: while the financial loss maybe less than the initial pandemic loss in the example above, the impact on the entity in terms of reputation, credit rating and regulator scrutiny, may well mean the eventual shareholder loss (in absolute or relative terms) could be seen as greater. The ability to recover in an industry otherwise unaffected by this event may be challenging.

The nature of the impact of a risk and relative risk appetite means this risk dimension should, directly or indirectly, be reflected in a business's risk appetite and risk management framework.

2.4.1 Risks in the Grey Zone

Unlike the value adding versus non-value adding risk, there is arguably a material grey zone between wide spread systemic risks and narrow entity specific risks. It is also not always obvious which trait a risk has. An insurer providing a unique product in the market may be exposed to the effects of wide spread economic conditions. However, if the insurer is the only one offering that product it could be adversely impacted in isolation. Should it regard that risk as systemic or entity specific? We would be inclined to suggest the later, but this view would not be without argument.

Notwithstanding the material grey zone, the authors think that considering the risk profile of an organisation in terms of this dimension provides considerable insight in terms of risk appetite, actual risk nature and risk management approach.

2.5 Convolutions of these 4 traits

In considering the four risk traits discussed above (value-adding vs. non-value-adding, systemic vs. entity specific), there are some perhaps interesting observations on the four risk trait combinations.

2.5.1 Value Adding, Systemic Risks

For insurance enterprises (and indeed many financial services entities) these risks substantially comprise the systemic insurance (liability), asset, liquidity and asset-liability mismatch risks taken on. These risks generally come with a market price for the risk (or "beta" in the Capital Asset Pricing Model).

One might call value adding, systemic risks **Beta Risks**.

- The industry has considerable modelling, measurement, management and pricing theories, frameworks and skills in this space. Economic Scenario Generation (ESG) models are common and well established, especially for asset (market) related risks. Although, it is less clear if insurance models in terms of underlying systemic risk (economic cycle risk etc) are as well developed or how well they directly address this risk.
- The industry has a range of established tools to manage risks in this area, from traded tools such as options and futures, to proprietary tools such as reinsurance, OTC instruments, securitisation and various other risk sharing devices.

Nonetheless, this is an area where the bets are often large and there is a long history of both mainstream and peripheral businesses (individually and collectively) that have mis-modeled these risks and the extent of the outcomes that can emerge can be severe. Exactly what the underlying systemic risk or risk driver actually is can sometimes be opaque and misunderstood. For large organisations, how a systemic risk might impact across the whole organisation can be difficult to assess.

2.5.2 Value Adding, Entity Specific Risks

There would seem to be two distinguishable subcategories of this combination:

- The entity specific statistical or random experience attaching to the Beta Risks discussed above. For example, the individual claim or asset performance or liquidity risks within the insurance and asset risks. One might call these **Alpha Risks**.
- The entity specific, value adding risks that relate to the businesses uniqueness, differentiation and strategy in the market, which one might call **Strategic Risks**.

For **Alpha Risks**, like Beta Risks, the industry has considerable modelling, measurement and management frameworks and skills in this space. The risks are typically well managed via scale, diversification and via tools such as reinsurance and other risk sharing devices.

In contrast, **Strategic Risk** is an area where there is little ability to hedge or off load risk. The industry has few measurement tools, and risk frameworks tend to focus on "risk monitoring", constructing "risk logs" and establishing "exit strategies". This is a risk management area that is often significantly bespoke. Yet this area, Strategic Risk, can accompany the principle source of value creation or destruction for a business. For shareholders of a business, Strategic Risk probably typically represents the largest bet.

It is interesting to note:

- In terms of consequence for policyholders, unitholders and depositors, the largest area of bet is probably typically Beta Risk.
- It is arguable that some aspects of Alpha Risk, while relating to value adding activities (e.g. insurance underwriting) might more correctly be categorised as non-value adding. An insurer is not explicitly rewarded for holding a higher per claim retention. However, a higher retention captures more profit margin. On balance we think Alpha risks are appropriately counted under the value adding category.

2.5.3 Non-Value Adding, Entity Specific Risks

As with the value adding, entity specific risks, there would seem to be two distinguishable subcategories of this combination:

- The risks that substantially comprise the internal business risks arising from the entity's execution of its business. The risks that arise from its systems, processes, project execution, change processes, day-to-day business decisions, resourcing levels and stresses, compliance standards and culture (to name a few). Errors, omissions, internal fraud, failing of systems. One may call these **Operational Risks**.
- Risks that affect a business in isolation (or as part of a small group) that come from outside agents, but affect a business because of its profile. External cyber-attack/breach, local terrorism or utility failure. One may call these **Resilience Risks**.

Many existing definitions of Operational Risks would include both of these subcategories. Indeed, they have much in common:

- As discussed in section 2.3 above, organisation's appetite and approach to these risks is fundamentally different to Beta, Alpha and Strategic Risk.
- As discussed above, we see the industry's risk modelling and measurement ability in this area as weak and indeed, because of the changing world and behaviours, modelling and measurement based on historic analysis and data is problematic in fact.
- Other than insurance for some business risks, there are few hedging, sharing or risk transferring tools available.
- Organisation's management approach often tends to be one of risk mitigation (restrictions, "locks", review processes, separation of duties, sign-off, "redundancy" backup in its various forms). However, in many respects the industry is constantly playing a guessing and catching up game in this space.

However, an organisation's management constructs for these two risk types are usually different (or at least we argue they should be).

More often than not, while mitigations as above are necessary and appropriate, risk management of **Operational Risks,** at its heart, comes down to overall business culture and the behaviour, capabilities and vigilance of individuals. Balancing controls (which encourage people to not think, but rely on the safety net) and accountability (which encourage vigilance without a safety net) is complex and without a comfortable answer. Operational risks are a problem for everybody in the business, every day.

Resilience Risks on the other hand, seem to us, to be all about mitigation. Strong locks, well maintained and tested fences, effective backups and fail-overs. The achievement of this risk mitigation state for the whole business is usually the accountability of specific individuals.

2.5.4 Non-Value Adding, Systemic Risks

These risks are largely those that result in compliance change costs and risks for an entity, and similar items that arise from external action, most often by governments and regulators.

One might call these risks **Imposition Risks**.

We note that in our view this category does not include competitor activities or market disruptors. Risks from these drivers are part of Strategic Risks as they go to an entities ability to create value.

This is the area of least ability to model, measure, manage or price. There are generally no insurance or risk transfer mitigations typically readily available. Any realistic solutions will be bespoke.

Mitigations primarily involve activities such as monitoring and lobbying governments and regulators either individually or via industry bodies.

2.6 Alternative Classification Map

Following the above discussion, the resulting universe of risks would be split as follows:

A reconciliation between a number of common risks a business faces and the top down construct outlined above is as follows:

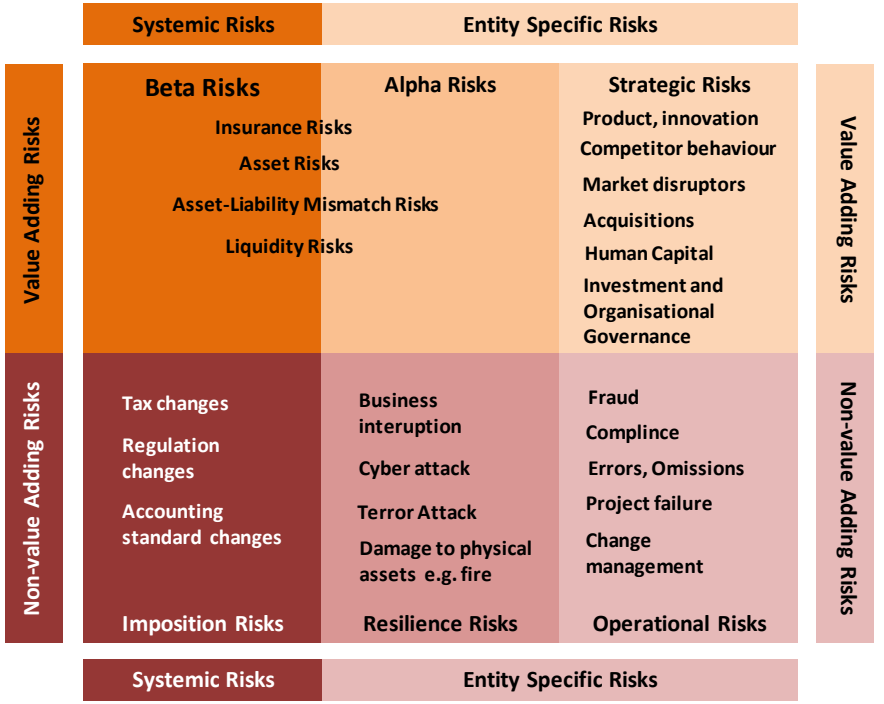

2.7 *So What? You ask*

The purpose of this discussion has not been to suggest a replacement risk definitional structure for risk management professionals. The current list of risk categories many of us refer to (per the graphics above) has a practical use in discussing individual risk exposures and managing them. Furthermore, in practice the risks organisations face are sufficiently varied and nuanced that there is probably no simple "universal definitional construct" that could be developed.

Nonetheless, we have found the perspective on various risks set out in this section helpful in our work as risk manages on a number of fronts, with:

- Discussing and distilling Board risk appetites. We have found it practical to discuss and set risk appetites having regard to the traits discussed in

this section, which provides a generalised construct that can easily be applied/extrapolated across the overall universe of risks;

- Constructing overall risk management frameworks. Where and how the universe of risks is managed within a business, and that it is in principle covered;
- Considering the effective mix of processes, culture, controls, "safety nets" and mitigants in managing different types of risks, given the traits that apply to the risks;
- Understanding and explaining to others how and why different approaches to risk measurement and modelling are effective and relevant for some risks and not for others, and indeed help ourselves understand what we should be trying to achieve with any such modelling or measurement within our risk management model.

2.7.1 Capturing Risk Upside

Another interesting observation we would make is around "capturing risk upside". ERM is said to be concerned as much about upside risk capture as downside risk avoidance.

We would contend, considering the discussion in this section, that upside risk exposure is realistically about optimising the value-adding risk exposures – that is, the Alpha, Beta and Strategic risks – and by "optimise", perhaps that means "maximise within the risk appetite".

There would seem to be little opportunity to capture upside within the non-value adding risks. For example, there would seem to be little upside in saving money by under investing in cyber risk mitigation.

3 Risk Rating (Measurement) of Op Risks

3.1 Introduction

Risk assessment and measurement is a vital step in the risk management process. If an organisation cannot measure a risk, then it cannot manage it.

As discussed in Section 2, risks in the "Alpha" and "Beta" categories tend to be reasonably tractable in terms of applying traditional and/or more modern statistical and modelling approaches to their measurement. There is often good historic data for these risks which is relevant to the way the risks behave and the industry generally has models that can apply these behaviours to the actual risk profile of a specific entity, e.g. an asset model that uses the actual asset profile held, including hedging or gearing instruments such as derivatives, combined with an appropriate economic scenario generator to produce an entity specific asset model. The industry has reasonably well established and generally accepted models for these risks that can allow for the specific treatments an individual entity may apply to these risks. Notwithstanding the cautionary comments in 2.5.1, we do not intend to discuss the assessment and measurement of these risks in this section.

The focus in this section is on the risks in the "Operational" and "Resilience" categories where, as discussed in Section 2, the application of historic data to measure future entity specific risk exposure is more problematic. We briefly look at two common approaches for measuring these risks (stochastic models and qualitative models) and describe a particular method we propose that builds on a bottom up approach that have found useful.

3.2 Two Typical Measurement Approaches

There would seem to be two common approaches used to measuring "Operational" and "Resilience" risks (which for ease of reference we will refer to as Op Risks for the balance of this Section), "Stochastic Models" and "Qualitative Models". Although we note at the outset, in our view their respective uses should typically be for different purposes; aggregate risk reserving (stochastic) versus individual risk assessment/treatment and management (qualitative).

3.2.1 Stochastic Models

Technical risk managers interested in quantifying an entity's overall Op Risk exposure (e.g. for capital assessment purposes or simply aggregate risk measurement) often seek to model Op Risks by fitting distributions to Op Risk event frequencies (e.g. using a Poisson or negative binomial distribution) and events severity (e.g. using a Weibull or Lognormal distribution). These are then combined into a total loss distribution. Depending on the available data and/or approach, this can be done by major risk categories or on some overall aggregate Op Risk basis.

Such approaches can be useful in attempting to quantify measures such as potential "1 in 200 year event" dollar losses. They can be useful in assessing losses and values in the tail of the distribution for an entity in total.

However, as well discussed by others (and briefly mentioned in Section 2 of this paper):

- While many large organisations will have some Op Risk loss data, this invariably relates to small loss amounts. Few organisations, still in business, have loss data for the "tails" where stochastic assessments are typically interested. Furthermore, the data available will typically be for a few individual risks (e.g. frauds, WHS injuries) but not all.
- While there may be other industry wide data, it is typically incomplete and how it can be scaled down to the individual entity is often unclear, although approximate approaches have been developed.
- As discussed in Section 2, the very nature of Op Risks and the industry's behaviour towards them fundamentally makes historic data less

applicable to the future than for other risk types (such as Beta and Alpha Risks). The advancement and progress of the world and industries means new Op Risks are constantly arising, while at the same time the industry identify and treat old or known risks to reduce their recurrence.

- While these models may deal with dollar quantified loss events, how they aggregate risk exposures that may not, at least in the first instance, be measured in terms of dollars is unclear, e.g. a "reputational" damage outcome or the non-dollar "cost" of the death of a staff member.

Putting these difficulties aside and assuming a risk manager does their best to develop the best model practical, in practice stochastic models are useful for estimating the potential risk profile of an entity as a whole. However, these models are by their nature inherently top-down models; they are not easily applied to:

- Consideration of individual risks that a risk manager in their day to day work may examine and consider for "treatment"; and
- In particular, attempting to assess the impact on the stochastic model of the various (bottom-up) entity specific treatments applied to its individual risk profile. Any such adjustment for this is typically largely qualitative and high-level judgment based. This is also before one starts to consider the impact of individual entities relative "risk culture".

In the end one has a very complex model, with qualitative overlays, that is opaque to all but the most qualified user. It is useful for addressing certain questions (e.g. capital reserving), but is not very helpful to many non-technical business managers and risk managers "at the coal face" assessing individual Op Risk exposures and their management.

3.2.2 Qualitative Approach

It would seem to us this is why simpler qualitative approaches, typically based on assessing Op Risks using simplified frequency and impact ratings remain in common use.

These simple "likelihood" and "impact" grid based models have a significant history of use in internal audit and compliance assessment frameworks. Many business users are familiar with them and they are a practical way to attempt to collect entity specific risk profile data to understand current Op Risk exposures and exposure trends. They can also readily deal with non-dollar risk outcomes and exposures.

Nonetheless, based on our experience:

- When rating risks most non-technical managers struggle to think about likelihoods and the corresponding impacts. They seem to think of these concepts independently. Often this results in managers, at least initially, thinking of the worst impact and linking it to how many times the risk could occur creating a vastly overstated risk rating.
- Having then overcome this issue, managers typically have insufficient knowledge or experience to envision how big a big loss event could really be. Their experience will be mostly based on the high frequency low impact events that they have seen which does not assist a complete view of the risk. As discussed in Section 2, any historic large impact event will almost certainly have been mitigate post event – so will now be "off the table" in the mind of most managers.
- In particular, evaluating very low probability "1 in 200 year" events which capital assessments and Board risk appetite limits may be concerned about, is very challenging.
- Most of the common risk measurement approaches are basically "linear" and do not take account of the tail of the risk which is typically more exponential.
- While these approaches may be helpful to indentify individual risk exposures of concern and for treatment (e.g. via a Red-Amber-Green, or RAG, status indicators), it is somewhat problematic to aggregate a series

of red, amber and green boxes into an aggregate entity risk measure, and indeed into a dollar based measure for capital.

In summary:

- The strength of stochastic models is that they can provide meaningful measurements of tail risks, on a "scientific" basis, and they implicitly allow for future "unknown unknowns", in addressing the big question of an entity's aggregate Op Risk exposure. Their weakness is that they struggle to help with mundane risk management activity involving assessing individual risks or the potential impact of a given treatment – these models don't scale down well, they are too complex and non-technical managers struggle with them, and risk modification impacts are typically only allowed via vague qualitative adjustments.
- The strength of the common qualitative models is their ease of understanding and use by non-technical managers, and their direct application to mundane risk management tasks. Their weakness is that they don't directly address the tail risks or measurements, are not good on "unknown unknowns" and often do not make best "scientific" use of the loss data that may be available (although in our view, greater use can be made of available data than often is).

3.3 Alternative qualitative approach

We note that the limitations of one method above are roughly the strengths of the other and we have sought to implement a somewhat different approach to the common qualitative approach in an attempt to retain its strengths but reduce some of its weaknesses.

The approach suggested uses data points provided by non-technical managers and staff to fit a distribution for the identified Op Risks. The advantage of this method is that it uses the knowledge of the management and staff at the coal face to assess current, actual risk profiles. Having fitted a distribution to each risk, an entity level distribution can then be generated for an aggregate exposure assessment.

As mentioned in the Introduction, much of this is not claimed to be new or unique, or "invented" by us.

3.3.1 Finding data points

We illustrate a two-step approach.

Instead of asking managers and staff for both the impact and the likelihood, instead select a range of pre-selected impacts and ask the managers and staff to consider how often each impact is expected to occur. For example, asking relevant managers and staff to think of an internal fraud with the following impacts (the scale would need to be relevant to the particular situation and entity):

Risk Assessment						
Impact	$0-$-.1m	$0.1m - $0.5m	$0.5m - $3m	$3m - $8m	$8m - $15m	>$15m
How often?						

The yellow boxes would be completed as far as is practical or the manager thinks is realistic and plausible. The amount of data points found from this method can be variable. While the more the better, provided at least three data points are obtained, some frequency/loss distribution can usually be approximately fitted. Obviously, if there is any loss data available for the risk, it can be directly considered and/or incorporated into this process.

As one is often concerned with lower frequency, higher loss events, and this is likely to be where the managers' estimations are least reliable (and/or data scarce), as a second step the manager is asked to specifically think about a large event (e.g. a large internal fraud). Once they come up with the event themselves you could question the impact and likelihood. With this as a base, you could ask for a 'worse' event and again question the impact and size of this event. Keep probing until you get the 'worst of the worse'. The

derived loss distribution above can be modified based on this input to give a final estimated loss distribution for the particular risk.

A classic RAG status can be defined based on loss distribution parameters relative to defined risk appetite limits in a similar way to the traditional Impact x Frequency score based RAG indicators.

As would be typical, Red and Amber risk RAG statuses would suggest specific risk exposures worthy of further consideration for management and treatment/mitigation.

One of the features of this approach is that it does not need to only relate to dollar impact outcomes. For example, the above loss points could be accompanied by, or exchanged for, other loss dimensions such as:

Risk Assessment						
Financial Los	$0-$-.1m	$0.1m - $0.5m	$0.5m - $3m	$3m - $8m	$8m - $15m	>$15m
Regulation	No breach	Minor breach	Regulator Reprimand	Private Directions	Fine, Public Orders	Loss of Licence
People	Increased staff turnover	High staff turnover	Significant loss of talent	Substantial loss of talent	Significant Exec staff loss	All Senior Exec's leave
How often?						

Other non-financial dimensions can be added (e.g. Work Health & Safety, Reputation damage outcomes, operational impacts, etc).

In applying this approach across different risks, we have tended to maintain the same loss scale so that results can more easily be benchmarked from one risk to the next by non-technical risk managers.

3.3.2 Aggregation to the Entity Total

Having fitted statistical distributions to each indentified Op Risk by business unit etc, it is then possible to aggregate these to total business unit levels, subsidiary levels, and the total entity. Depending on the statistical distributions fitted, this could be via direct mathematical aggregation or relatively simple stochastic simulation.

In doing this, the issue of correlation between the individual risks needs to be considered and allowed for. Unless there is established quality data for some risks, this will inevitable involve substantial professional judgment. In the authors' view, a general "0.2" correlation allowance is probably a good judgment place to start before considering if risks may have greater correlation. One should be very wary of assuming independence.

In adopting this approach some points are noteworthy:

- The individual risk assessments will no doubt involve significant estimation error. In terms of the aggregate result, the law of large numbers, by adding multiple estimates together, can help reduce aggregation mis-estimation.
- To the extent all risk assessments may tend to under estimate the "unknown unknowns", then so will the aggregate assessment.

Notwithstanding these issues and mis-estimations, adopting this approach does provide a way to:

- More explicitly track overall risk exposure trends within the entity;
- Directly benchmark risk exposures and trends by underlying risks and business units;
- Identify the risk hot spots for focus and attention; and
- Directly link bottom up risk assessments to aggregate entity risk profile.

3.4 Top Down – Bottom Up: Fitting it Together

In considering top-down stochastic models based on available entity and other "scaled" data (e.g. from industry) and other speculations, it seems to us that such approaches are subject to:

- Normal estimation error in terms of parameter uncertainty;
- In particular, while using past data applied to a current situation where some past events may well have been mitigated but are still reflected in

the data, this implies that there will be some allowance for future "unknown unknowns" in the results. However, it would be far from clear whether such an implicit allowance is reasonable and adequate or not;

- The need for adjustment for the extent to which the specific entity varies in its risk profile from the "average" situation implied by data used and the end model derived, including risk management framework quality and relative risk culture dimensions.

In the case of the bottom-up qualitative approach:

- This will also include estimation uncertainty and quite possibly a net under-estimation bias;
- Mis-estimation (underestimation) will especially reflect the extent to which the quality of the risk identification process and estimations are honest and ethical, which also reflects the quality of the risk management framework (RMF) and risk culture. We suggest the mis-estimation "gap" will likely be greater the weaker the quality of the RMF and risk culture.
- The result will directly relate to the "known unknowns" of the actual current risk profile, but will likely largely ignore the "unknown unknowns", that is the unforeseen and unidentified risks and outcomes - the so called "black swans", or the term we prefer, the "dark risks". The risks and outcomes that are there but that you can't see them, similar to the concept of "dark matter" and "dark energy" in physics.

Putting these elements together suggests a reconciliation between the two approaches (in terms of aggregate risk measurement) such as the below graphic.

This also suggests an approach to determining an entity's overall Op Risk capital charge that can explicitly take into account estimated entity effects such a risk management strategies adopted, underlying risk profile changes and trends, and potentially allow for and/or imply measures related to RMF quality and risk culture.

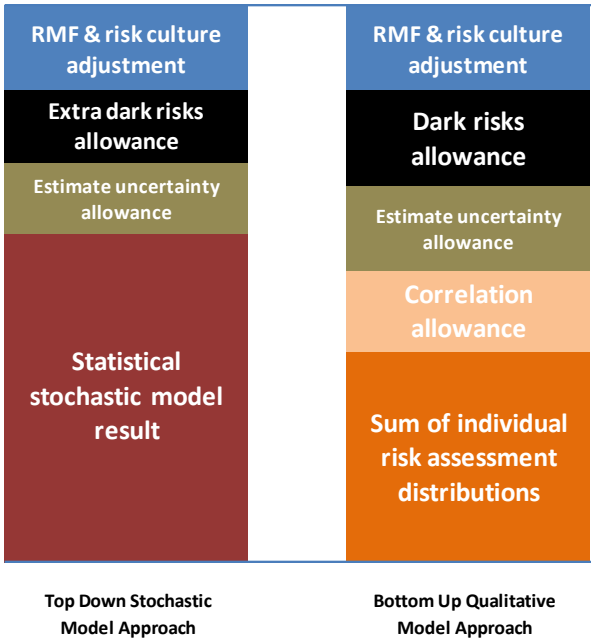

We make the following notes on this graphic:

- While the graph above implies all additive adjustments, some may instead be subtractive. For example, for the risk culture adjustment to the top down stochastic model, this could be negative if it is believed (or assessed) that the entity's risk culture is (on average) better than that reflected in the source of the data and information upon which the model is based. If the top-down model is based on some industry average Op Risk reserving data or benchmarks, then the implied capital reserving result could arguably be reduced for an entity that believes it has a better than average RMF quality and risk culture.
- Whereas, the RMF and risk culture adjustment for the bottom up qualitative model can probably only be positive. This may involve a

smaller adjustment for an excellent RMF and risk culture (inferring quality, risk aware, ethical and honest exposure estimates), and a larger adjustment for a less excellent RMF and culture.

- The relative size of the adjustments illustrated in the graphic are not intended to be indicative of any particular situation.

3.5 Imposition & Strategic Risks

We have not directly discussed "Imposition" or "Strategic" risk assessment or measurement in this section. However, it seems to us that the qualitative approach outlined in this section could also be applied to these risks, although we have not attempted this to date.

4 Risk Communication

4.1 Introduction

A large, complex financial institution inevitably must manage a large number of risks. As discussed in the Introduction to this paper, it is not uncommon to hear of "200 page" risk reports that step through the business risk by risk, providing commentary, assessments, overall measurement quantification and analysis of movement; stress testing analysis and scenario testing analysis; and consideration of emerging and "key risk indicators" and measures. We have heard more than the occasional mention of dissatisfaction from directors and senior management with these large reports.

In observing this, we fully acknowledge we have seen some good reports with RAG status indicators and trend arrows and the like that significantly help with communications. We are also not suggesting that risk officers simply dispense with their substantive reports!

However, we note that:

- In practice, a large part of such report contains largely static information. Last month's key scenario risk exposure is usually pretty much the same this month;
- In amongst this forest of detail and mostly "sameness", what is changing, any common themes across the business, the risk connections or inter-connectiveness of the risks, can be hard to distil and identify.
- While the risk officer preparing the report to senior management or board will typical seek to provide an executive summary or overview to try to draw out the key messages, this is only as good as the reliability of the risk officer's own perceptions and identification of patterns.

It is not our intention to attempt to solve the overall dimensions of this problem and its issues in this brief paper. However, we would like to share one simple tool that we have found of assistance.

4.2 A Risk Map on a Page

We have developed a one page "risk map" which we have used to communicate overall risk status, outlook and trends. Consideration of the map itself has helped us to stand back and see some patterns ourselves, which has also helped ensure (as a sort of a visual check list) that the overall risk profile is being considered and not some areas overlook due to habit or familiarity.

A number of directors have commented favourably on this map as a high level snap shot and a guide to the interpretation and review of a larger, detailed report.

The approach adopted is to depict (albeit stylised) the overall business structure on a page, left to right along the value chain and operational chain of the business, from governance at the top to external drivers at the bottom. RAG status and callout comments highlight the risk "heat" areas and changes and trends.

4.3 Illustration of the Risk Map

The following page provides a simplified example of the authors risk heat map for a fictitious insurance company.

It sets out on one page a diagram of the key areas of the organisation and rates each risk using an abridged RAG status which is set with reference to the Risk Appetite Statement. Brief descriptions of the adverse risk statuses and trends are provided in the right hand boxes. Any changes to the risk status since the last report can be highlighted by putting the previous colour in a box behind.

The authors would be very interested in the view of other actuaries and risk professionals, and alternative approaches and experiences in this area of high level risk status communication.

5 References

1. Stephen P. D'Arcy, 2001, *Enterprise Risk Management*
2. Robert Chapman, 2006, *Simple Tools and Techniques for Enterprise Risk Management*
3. Institute of Actuaries of Australia Life Insurance Risk Appetite Working Party, 2015, *Developing the Risk Appetite Framework of a Life Insurance Business*
4. Canadian Institute of Actuaries (KPMG), 2014, *Research Paper on Operational Risk*

Simplified example of a Risk Status Map

Risk Heat Map for ABC Ins Co

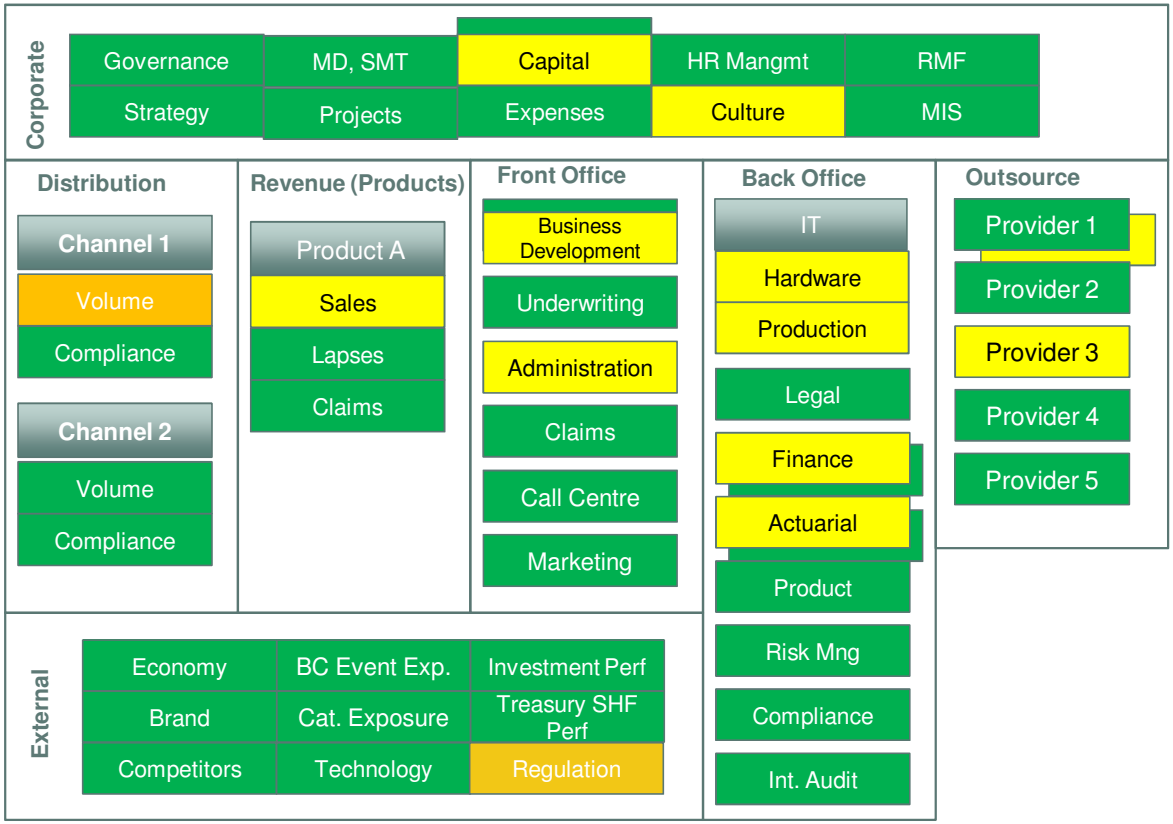

Amber Risks

1. Sales volumes significantly below target for distribution channel 1
2. Large volumes of law reforms increasing risks of regulatory breaches

Yellow Risks

1. Product A sales volumes below budget
2. Stretched staff in Finance and Actuarial due to financial year end work
3. Stretched resources in Admin due to recent high rate of staff turnover
4. In April, Provider 3 did not meet its SLA. Requirements. Output is being monitored more frequently.

	Description
	Risk is substantially above tolerance* and requires immediate focused action#
	Risk is significantly above tolerance* or above tolerance with material adverse trend requiring action# in the short term
	Risk is at or moderately above tolerance* or inside tolerance with an adverse trend and likely to go outside tolerance if trend continues.
	Risk is inside tolerance* with a steady trend

•As defined by the Board in the Risk Appetite Statement

Actions may include change in controls, mitigation, focused management oversight/monitoring or contingency/response plan development.