

# Broadleaf

Creating value from uncertainty

Broadleaf Capital International Pty Ltd

ABN 24 054 021 117

[www.Broadleaf.com.au](http://www.Broadleaf.com.au)

## Discussion paper: Unknown unknowns

In the ten years since the United States Secretary of Defense at that time, Donald Rumsfeld, brought the term “unknown unknowns” to prominence, its use has spread throughout the risk management community. The term itself and the confusion it creates will never go away but that very lack of definition presents an opportunity to think about what we leave out of our risk assessments. There may be things we can do to reduce these gaps and improve the quality of our assessments and the final results.

This material was initially published, in two parts, in Risk Management Today issues 32 and 33, August and September 2013.

Version 2, 2014

© Broadleaf

1 of 14

## Contents

<b>1</b>	<b>Unknown unknowns</b>	<b>3</b>
<b>2</b>	<b>How do we interpret the words?</b>	<b>3</b>
<b>3</b>	<b>What do we know?</b>	<b>4</b>
<b>4</b>	<b>Generalising</b>	<b>5</b>
<b>5</b>	<b>Extending</b>	<b>7</b>
<b>6</b>	<b>Searching</b>	<b>8</b>
<b>7</b>	<b>Early detection</b>	<b>8</b>
<b>8</b>	<b>Conclusions</b>	<b>13</b>
<b>9</b>	<b>Acknowledgements</b>	<b>14</b>
<b>10</b>	<b>References</b>	<b>14</b>

## Figures

Figure 1: Known and unknown	5
Figure 2; Cynefin framework	10

## 1 Unknown unknowns

In the ten years since the United States Secretary of Defense at that time, Donald Rumsfeld, brought the term “unknown unknowns” to prominence, its use has spread throughout the risk management community. Specialists and novices drop it into conversation and puzzle over what it means for them. It has absorbed large amounts of time in casual and professional discussions where opinions about unknown unknowns and interpretations of the term differ widely.

The term itself and the confusion it creates will never go away but that very lack of definition presents an opportunity to think about what we leave out of our risk assessments. There may be things we can do to reduce these gaps and improve the quality of our assessments and the final results. Some of these are discussed here.

## 2 How do we interpret the words?

No attempt will be made here to provide a sound definition of unknown unknowns as this has absorbed a lot of energy for many years with no clear resolution in sight. Rather than try to lay that challenge to rest, the effect it has on the way we think about risks will be examined to see if we can learn anything from the ideas it stimulates.

One interpretation of unknown unknowns is that they are things we haven't thought of that could affect us, that is there are limitations to our knowledge or awareness. It is more or less assumed that, if we haven't given these unidentified factors any thought and they could affect us, it will be bad for us if they do occur although we can't be certain because we are not sure what they are.

With this mindset, unknown unknowns become a container that we can load up with all our anxiety about having overlooked something important. The absence of a really solid definition makes it difficult to think clearly about the reasons for this anxiety. Anxiety is uncomfortable and it is not a very good foundation for clear thinking or sound decision making.

The futility of trying to include in plans things of which you are not aware has led some to conclude that the concept of unknown unknowns is little more than an academic distraction. Even trying to talk about the subject is difficult.

It can be argued that, unless we are aware of gaps in our knowledge, we cannot act on them so worrying about what to do with unknown unknowns is a waste of time.

Others use unknown unknowns in quite a different and much less acceptable way. They use it to mean the net effect of all the things that we effectively choose not to analyse in detail even though, in principle, they could affect us. Given sufficient time and imagination, we can contemplate a large number of uncertainties that we would not generally take the time to analyse because they seem remote, very unlikely to affect us or too far beyond our control. This interpretation seems to be closer to an excuse for limiting analytical effort. There may be good reasons to limit that effort but throwing in an impressive sounding term that is poorly defined will often close down the discussion.

Whether it is as a repository for anxiety about how comprehensive our risk identification and analysis has been or as an excuse for leaving aside risks we effectively choose not to include, the term unknown unknowns is a useful catch all description. Having no really solid definition, its use is difficult to challenge and attempts to do so will often descend into fruitless and rambling discussions that absorb a lot of time without achieving very much.

## 3 What do we know?

Common sense tells us that we do not know all there is to know about everything we do. No one can argue credibly that “absolutely nothing has been left to chance” as is sometimes claimed. They might say that “absolutely nothing we have thought about has been left to chance” although even that is generally an exaggeration.

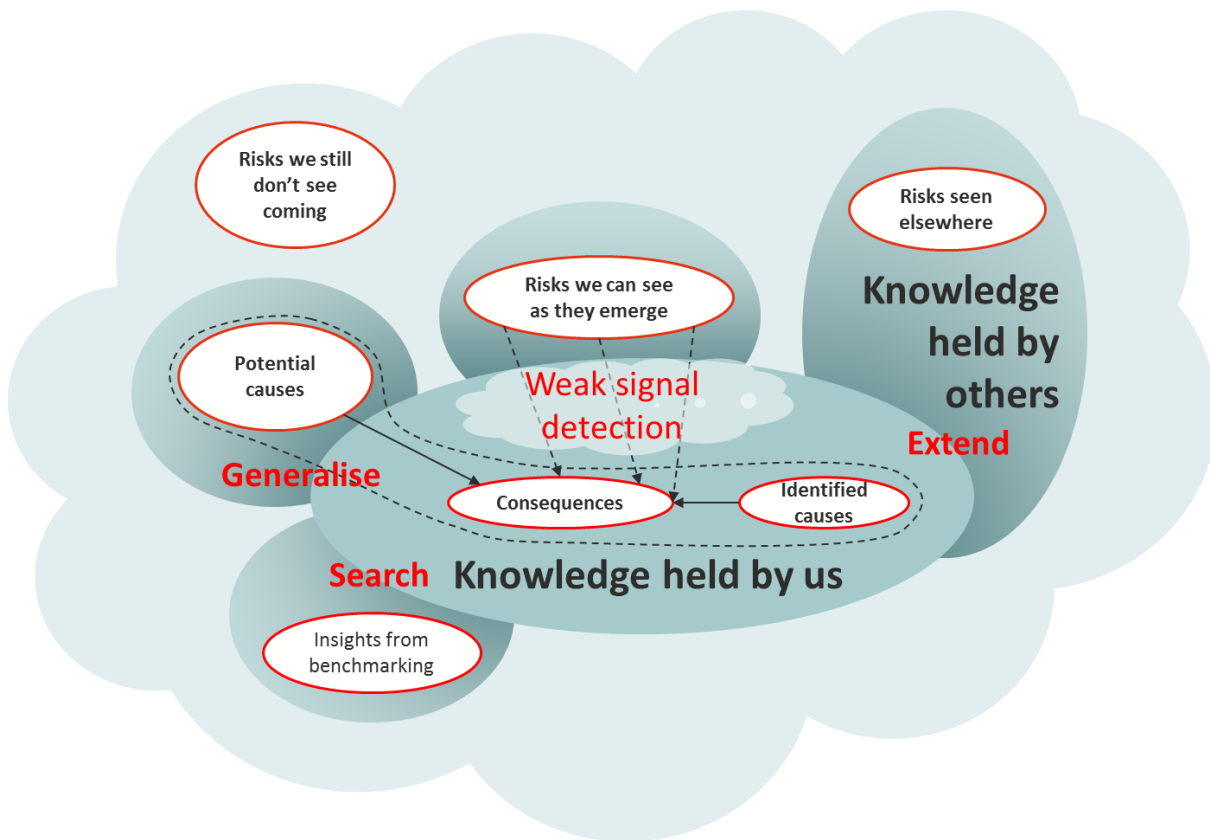
It is not uncommon for people to declare that, because limits on our knowledge are inevitable we should give up trying to overcome them. However, there may be ways to make some progress in this area and a few strategies for doing so, that have been stimulated by considering the concept of unknown unknowns, are outlined here:

- Generalising known risks to encompass multiple possible causes;
- Extending the range of the knowledge we draw upon;
- Searching for fresh insights into unanticipated developments; and,
- Early detection of the emergence of unforeseen circumstances.

Each of these is illustrated in the diagram in Figure 1.

# Broadleaf

Figure 1: Known and unknown



## 4 Generalising

This suggestion is not extraordinary, yet nor is it widely exploited. It can be illustrated with the following example dating from times when formal risk management was less widespread than it is now.

A project to develop a new aircraft had been delayed several months because, the first time a propeller was mounted on the prototype and run up in static tests, the propeller disintegrated and tore apart the fuselage, one wing and a lot of measuring instruments. The team's management had not anticipated the delay and demanded a risk assessment of the work to restart and complete the development. The project team saw this demand as a sign that management lacked confidence in them.

The team was very keen to explain that the propeller had been well designed and thoroughly checked. As far as they were concerned, no one could have foreseen the propeller failing or the impact this would have on the schedule. The risk assessment facilitator asked them if they had experience of other

### Unknown unknowns

© Broadleaf Capital International

airframe developments and was told about many projects the team had worked on in the past. He then asked them if it was unusual for something dramatic to happen that caused a major delay and they happily told him about several catastrophic and exciting disasters they had witnessed, all totally unpredictable.

The facilitator then suggested that, if it is routine for a major event to occur on a challenging airframe development, where the designers are pushing to improve on the performance of existing planes, perhaps it would be reasonable to consider the risk of such an event in the risk assessment of these projects even if the precise cause might not be known in advance. It is not necessary to know exactly what might blow up, catch fire, disintegrate or fail in some other way. The nature of the work is such that something often does fail causing a delay and some unplanned costs.

By assuming that the detailed causes of a risk have to be spelled out in order to include it in a project's risk management arrangements, the team prevented themselves from thinking about a major risk that was otherwise fairly predictable. The length of the delay a project might suffer from such an event may vary depending on the precise cause and no one can tell if a catastrophic failure will occur or not on any particular project but this is the case for any risk. The likelihood of it happening is neither zero nor one and the magnitude of the consequences is uncertain.

After being caught out by a major unforeseen event, it may be comforting to be able to label it an unknown unknown. Presenting it, retrospectively, as something that was beyond our grasp takes away some of the responsibility for having failed to foresee it. However, as the example above shows, by generalising the description of a risk from a specific detailed cause and effect to a broader statement about a type of disruption that is foreseeable, we may be able to prepare for the risk even without being able to pin down the precise cause in advance. This is not to say that it is easy to see these events coming but they might not be as mysterious as the label unknown unknowns suggests.

There are parallels here with scenario based planning. It is not necessary to define in detail how a scenario might arise to be able to understand that it is plausible and think about how you would respond if it did arise. One approach to preparing for some of the risks we have not identified in detail may be to see if we can generalise our analysis in some areas and work on plausible high level expressions of uncertainty and its consequences while recognising that we might not have spelled out every possible root cause.

This is sometimes seen in the way cost and schedule implications of safety issues are assessed for large infrastructure projects. A project will generally have done all it reasonably can to ensure that people working on it are safe but experience shows that, from time to time, usually with a relatively low likelihood, safety incidents can still affect the progress of a project even if it is just a near miss that causes a delay while work practices are reviewed. A project can provide for this in its risk management plan without going into anywhere near the detail found in a safety risk management plan where many individual triggers would be considered.

## 5 Extending

Most risk identification processes tap into several sources of information. Many people might be interviewed or brought together in workshops, for instance, to identify and analyse the risks affecting a system. In this way, some of the gaps in one person's experience may be covered by the knowledge of others.

The need for diversity in the identification of risks is nothing new but, if there is serious concern about failing to spot important sources of uncertainty, perhaps it deserves more than the casual attention it often receives. Risk identification exercises may be deferred if people whose knowledge is known to be crucial are not available but it is uncommon to see a conscious search for diverse inputs and opinions that might really push the boundaries of accepted knowledge. Even when diverse inputs are available, poor processes can lead to useful insights being overlooked as participants censor or fail to stretch themselves.

Linstone and Turoff [1] draw attention to the fact that the Delphi method was not originally intended to derive a consensus but rather to pursue and understand diverse views among a group of knowledgeable people. Disagreements were seen as beneficial as they opened up the subject and challenged initial assessments. More recently developed methods, such as headstand brainstorming, which involves thinking about how to make something fail and then using those ideas to see what is required to achieve success, seek to help individuals or groups open up the boundaries of their own thought processes.

If we can bring this to the fore and remind ourselves and others that our knowledge is inevitably limited, perhaps a little of the precious time allocated

to identifying risks may be used to push against these limitations. There is a vast difference between the staid routine risk analysis exercises commonly used by many organisations and some of the more creative methods that might enhance them. Time is always a constraint but conservatism, and a sense that some methods are not formal enough to be taken seriously, also limit the chance of drawing on a wider range of inputs or stimulating participants to extend themselves.

## 6 Searching

Searching the unknown for uncertainties you don't know exist may seem ridiculous and, without something to give it direction, it would be little more than day dreaming. However, exploring how things work with people undertaking comparable activities can open up new ideas and provide insights that we might not gain by simply looking harder and longer at our own situation.

Benchmarking is sometimes viewed as a purely quantitative exercise focused on comparing ratios and other metrics from one case to another. A more subtle approach is to explore not just numerical characteristics of related systems but also the important cause-effect relationships that are at work and the reasons why people in one organisation adopt a different approach to people in another organisation when faced with essentially the same challenges. This approach is at the core of a small number of international benchmarking networks, most notably in the work on very large projects led by IPA (Independent Project Analysis Inc.).

In one sense this may be regarded as another way to extend our knowledge base but active comparisons carried out through discussion between peers do more than graft the existing knowledge of others onto our own. They offer a setting within which fresh insights can be developed that were not previously available to any of the parties engaged in the exchange.

## 7 Early detection

No matter how far we stretch known risks to see what new forms they might take, extend the boundaries of the knowledge available to us or seek insights from comparisons between related systems, we will never be able to make sure absolutely everything that could affect our future has been considered in full. Even where we have identified what we are concerned about and analysed it as



thoroughly as we can, unexpected situations can emerge in complex systems [2]. This is consistent with findings described by Gardner and Tetlock [3] suggesting that the success rate of forecasting is a lot lower than most people think, a finding that holds across a wide range of subject areas with both qualitative and quantitative forecasts.

“Philip Tetlock assembled a group of some 280 anonymous volunteers ... the experts made some 28,000 predictions ... the veracity of the predictions was determined ... to be only slightly more accurate than random guessing”

The same essay paints an interesting picture of two modes of individual behaviour that were seen to affect forecasting efforts. The inclusion of diverse sources of information and an acceptance of complexity and uncertainty appeared to improve the reliability of forecasts although, even then, many forecasts still failed. This is consistent with the behaviour of complex systems described in the Cynefin framework [2] which is outlined below – some things that matter to us cannot be forecast reliably and the best we can do is catch them as they begin to emerge by detecting what are referred to as weak signals or early indications of an impending development.

A comprehensive discussion of the Cynefin framework, developed by David Snowden and his colleagues [2] [4], is not feasible here but it provides assistance with thinking clearly about what we do or do not know, in fact what we can or cannot know, about a system’s behaviour. One of the framework’s strengths is the explicit consideration of complexity from which unexpected and indeed unpredictable situations can emerge, perhaps the real unknown unknowns.

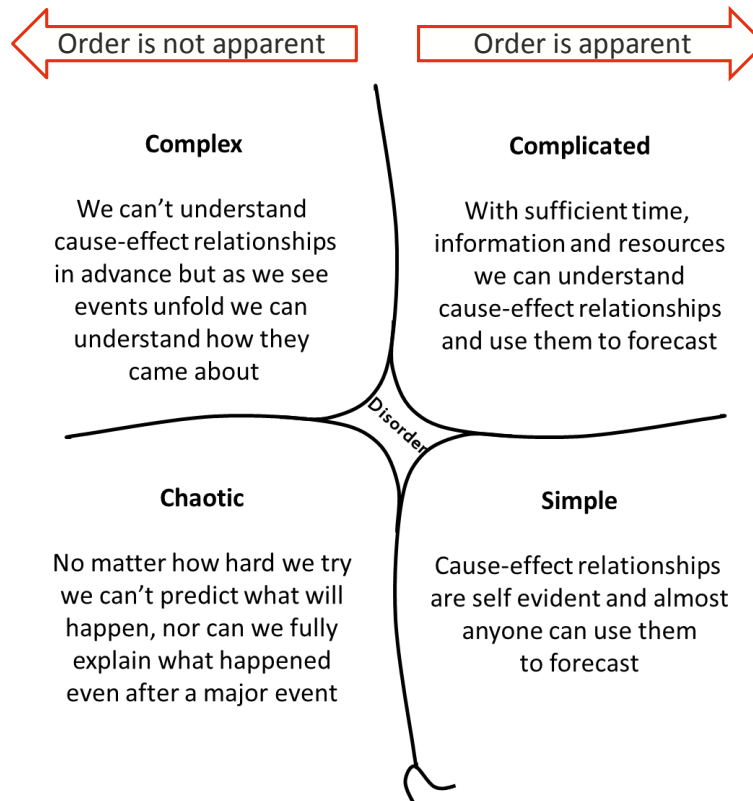
Among the many insights we can take from the Cynefin framework and associated ideas are that:

- Not everything that is important can be exposed and analysed in advance; but,
- There are strategies we can use to help us work with what we cannot predict.

The first point takes some of the pressure off. Without abdicating responsibility for whatever we can influence, we might as well just get used to the fact that some things that we care about will not be controlled in advance by tighter procedures or more intense analysis. The second point means that we should not give up in the face of this unpredictable behaviour. We can take steps to spot emerging situations before they overwhelm us.

The framework is illustrated in Figure 2.

**Figure 2; Cynefin framework**



This is not a general exposition of the Cynefin framework but, for the purposes of this discussion, it is introduced to help think about what we can and cannot know about aspects of systems we are trying to manage. The Cynefin framework divides the ways we understand a system, such as the subject of a risk analysis, into four characteristic domains with one overarching condition, as illustrated in Figure 2.

The overarching condition in the centre, labelled Disorder, represents not being conscious of the fact that different ways of understanding a system exist and have important implications for the way we seek to manage it. The framework divides these ways of understanding into situations in which we can expect to:

1. Readily assess all that the future might throw at us, the Simple region, where we will generally rely on standard practices and established procedures to guide us;

2. Explore future possibilities as comprehensively as we can afford to by investing specialist effort in analysis, the Complicated region, where expertise, studies and investigations will help us;
3. Be able to detect events as they unfold, even though we could not predict them, the Complex region, where we need to experiment and sense emerging patterns as they develop while recognising that we will never be able to understand all the factors at work no matter how hard we try;
4. Have to accept that events will surprise us and we might not be able to learn anything useful to reduce the chance of being surprised again, the Chaotic region.

There appears to be sufficient overlap between the characteristics of the complex domain and some challenging aspects of risk management to make it worth exploring the connection between the two. The Cynefin framework is relatively new and its relationship to risk management, as conceived in ISO31000 and its predecessors, does not appear to have been developed in any depth to date but it is hoped that this discussion might stimulate further interest in doing so.

No matter what the reason, we must accept that surprises cannot be prevented completely nor can we ever be certain that we have even reduced them to negligible levels. This means that there will always be merit in being able to spot unexpected developments as they emerge and while there is still time to respond effectively. This applies as much to undesirable developments that we want to dampen down as it does to beneficial developments that we want to encourage and support.

This might sound like a call for the use of leading indicators. However, leading indicators are inevitably framed by what we already know we need to be concerned about. We only look for what we expect might get out of control. In addition, true leading indicators are the exception rather than the rule. Quite often, we watch the train wreck as it reaches its crescendo rather than spotting the warning signs in time to prevent it.

To obtain early warning of an emerging situation in time to avoid being surprised, we need to tap into a lot of real time information that can be captured, aggregated and interpreted swiftly without undue expense. It is also important that, as far as possible, the sources of this information be free of the bias and blind spots that generally accompany preconceived frames of reference. One promising approach that has been used, although not as far as

the author knows in conjunction an ISO31000 risk management framework, is the SenseMaker® method [5].

It is not intended to go into this method in detail here but, applied to risk management, it could involve some of the personnel of a business or other organisation taking a few minutes once a week to describe something, briefly, that they have observed happening, a narrative fragment, and indicating its significance using a predefined set of characteristics, a signifier framework. The narrative fragments might be prompted by a question or other stimulus and could, for instance, be as simple as a single sentence describing something that seemed interesting, irritating, out of place or surprising. The signifier framework might allow the contributor to describe, again as a loose example, how this observation relates to the organisations' policies and procedures, management behaviour, workforce attitudes, activity within the organisation, customer requirements, the personal wellbeing and satisfaction of the contributor and other factors.

When summaries of the results from all or selected subsets of such inputs are examined by an analyst, experience elsewhere suggests that it will be possible to identify:

- Patterns that indicate interesting relationships between the factors at work in the organisation, possibly a particular procedural problem always being associated with a certain group of customers;
- Anomalies in the data where some of the inputs show different relationships to the rest of the contributions, such as everyone in the organisation seeing good alignment between personal and organisational goals except for one team; and,
- Of particular interest for risk management, changes from one round of inputs to the next as something new emerges, possibly something we have not anticipated but wish we had.

Key features of this method are its diverse and information rich inputs, the fact that the contributors interpret their input themselves without the constraint of a predefined analytical framework, and the short time and few processing stages the information has to traverse from its source to the decision makers who use it. The inputs are not lost in the morass of statistical analysis and reconciliation with preconceived structures that result in conventional surveys having turn around times denominated in months and processes that tend to obscure any insights that are at odds with the survey design.

While this specific method does not appear to have been applied to mainstream organisational and project risk management yet, there are echoes of such a process in the engineering practice of starting formal meetings with a “safety moment”. Someone will recount a recent experience or observation that made them think about safety and link it to practical lessons for the people in the meeting even though the narrative might have nothing to do with the work environment. It may be as simple as describing how collecting a teenage child from a late night party made the presenter aware of the dangers of driving while tired and using this to reinforce the need for fatigue management on construction projects where heavy machinery, very large trucks, ordinary passenger vehicles and pedestrians interact with one another.

Weak signal detection may offer a new means of spotting emerging risks, one that is qualitatively different from existing methods. Risk management systems will always need to maintain a watch on the factors we know we need to control and the relationships we know may be important. However, a creative means of tapping into large amounts of diverse information may be the only way to tackle risks in truly complex systems where even the most diligent risk identification process carried out before work starts, which is the usual pattern, will be unable to forecast everything that can emerge after work is underway.

## 8 Conclusions

The ill defined concept of unknown unknowns can absorb a lot of time for little real gain but thinking around the subject of what we might not be aware of can help improve existing risk management practices. None of the ideas proposed here is completely novel but it may be worth considering whether fresh impetus should be given to:

- Generalising identified risks to incorporate additional and possibly as yet unknown causes, focusing a little higher up the cause effect chain;
- Extending the knowledge we draw upon by deliberately incorporating people with diverse points of view into our processes and valuing diversity in the information we generate rather than trying to force consensus; and,
- Deliberately seeking fresh ideas by benchmarking and contrasting our work with that of others who are close enough to permit comparisons while being different enough to reveal interesting insights.

In addition, no matter how diligent our efforts to expose the uncertainties that might affect us, we have to accept that unanticipated situations can emerge. However, we can take steps to obtain early warning as these emerge. One

method for doing this has been proven in other settings and appears to offer benefits for risk management in enterprises, organisations and major projects. In each of these, we face truly complex behaviour, as described in the Cynefin framework, and significant numbers of people are available to provide observations that can throw light on emerging situations. With an early warning system in place, we should be able to enhance our ability to manage some of the risks we will always be unable to identify in advance.

## 9 Acknowledgements

The thoughts set out here have been refined in conversation with the author's colleagues in Broadleaf Capital International ([www.Broadleaf.com.au](http://www.Broadleaf.com.au)) as part of the team's continual endeavour to advance the science of risk management.

Descriptions of Cognitive Edge materials and the Cynefin framework are based on the author's as yet limited involvement with this relatively new body of work and anyone interested in exploring them is strongly recommended to make their own enquiries via the Cognitive Edge web site at <http://cognitive-edge.com/network>.

## 10 References

1. H.A. Linstone, M. Turoff, Delphi: A brief look backward and forward, Technol. Forecast. Soc. Change (2010), doi:10.1016/j.techfore.2010.09.011
2. David J. Snowden, Mary E. Boone, Harvard Business Review (Nov 2007)
3. Dan Gardner and Philip Tetlock, Overcoming our aversion to acknowledging our ignorance, CATO Unbound (Jul 2011), <http://www.cato-unbound.org/2011/07/11/dan-gardner-and-philip-tetlock/overcoming-our-aversion-to-acknowledging-our-ignorance/>
4. <http://cognitive-edge.com/network>
5. <http://www.sensemaker-suite.com/smsite/index.gsp>