The true measure of risk management: control effectiveness or residual risk level?
Ross Liston (Associate Director, Advisory, KPMG - Auckland, New Zealand).

This paper puts forward an approach to reposition risk management practices to incorporate greater attention on measuring control effectiveness. The approach draws on the premise of critical control management, demonstrating that focusing risk management on those controls that are most critical will provide the best means to gauge how well the risk is being managed – perhaps better than what current risk registers and risk heat maps can provide.

(Insights presented are drawn from the author's first-hand knowledge of global implementation of operational risk management and the associated critical control monitoring and management in high hazard industries. Such operational risk practices are primarily focused on the first line of defence, but lend themselves to scaling up through the horizontal layers of an organisation and providing reliance information for assurance purposes.)

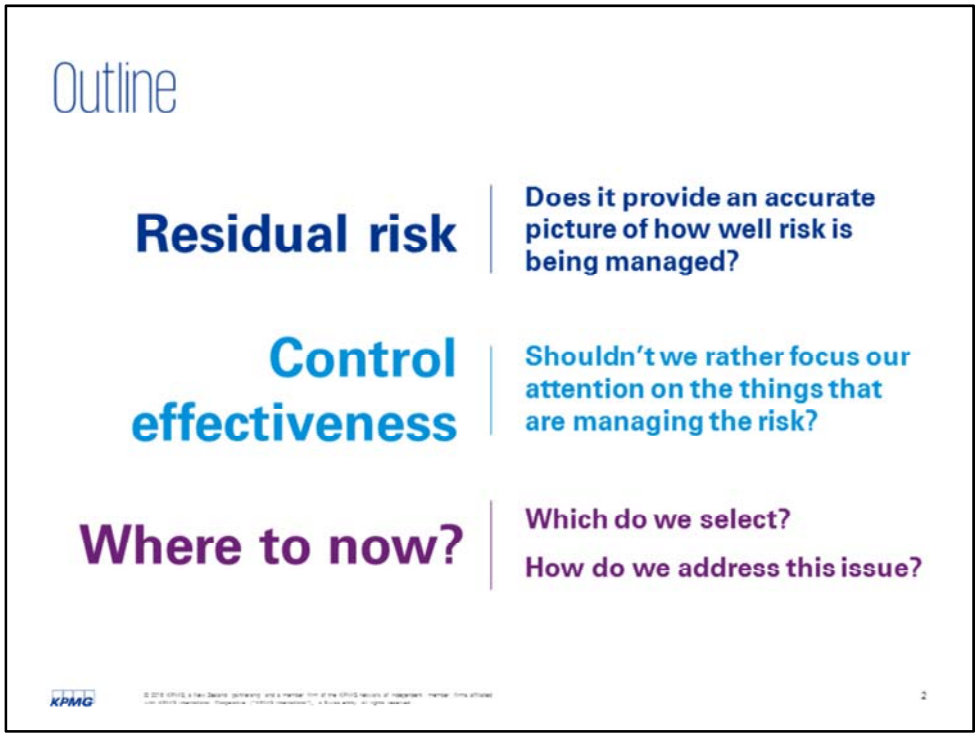

The presentation is based around three main sections:

1. How do we currently use 'residual risk' in risk management? *Does it provide an accurate picture of how well risk is being managed?*
2. How might we shift our focus to consider 'control effectiveness'? *Shouldn't we rather focus our attention on the things that are managing the risk?*
3. How do we move forward? *Which do we select? How do we address this issue?*

When organisations use risk registers and their associated risk management matrices/heat maps there is an overwhelming tendency to strive to move risks into the 'green'.

But that is not necessarily correct or of value.

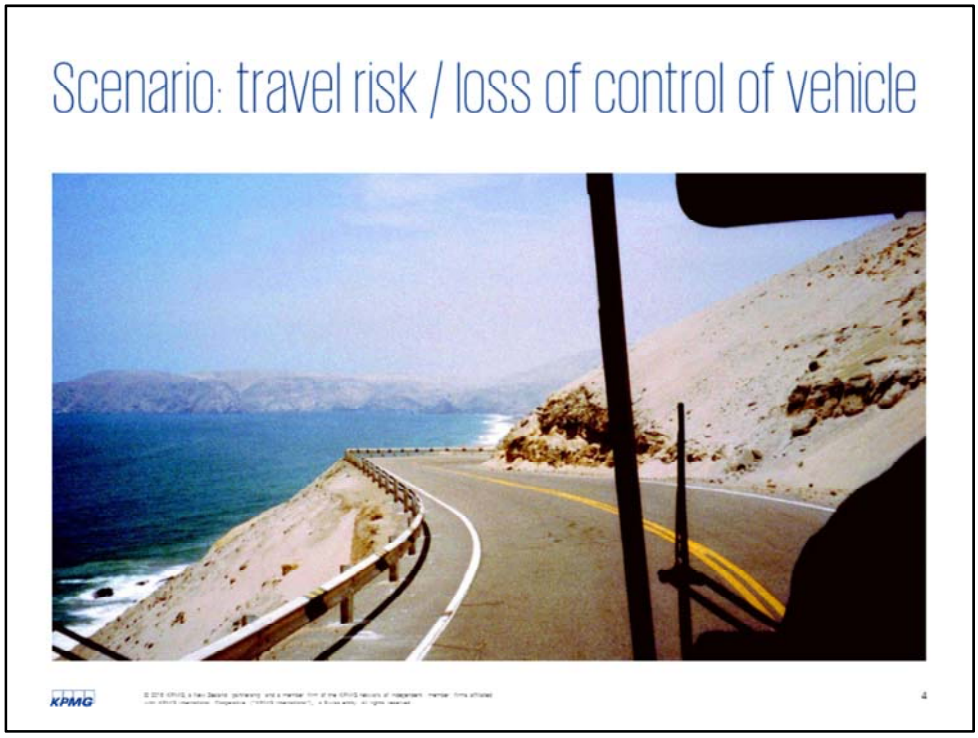

Let's assume you're looking at your organisation's risk register. More often than not a transport or travel-related risk will be included. It may either relate to a safety consequence where the vehicle occupant dies or is injured; or an environmental consequence if we have a spillage; or a business interruption consequence if we can no longer use the vehicle for transport or travel purposes. Regardless of the consequence the source of risk relates to 'loss of control of the vehicle'.

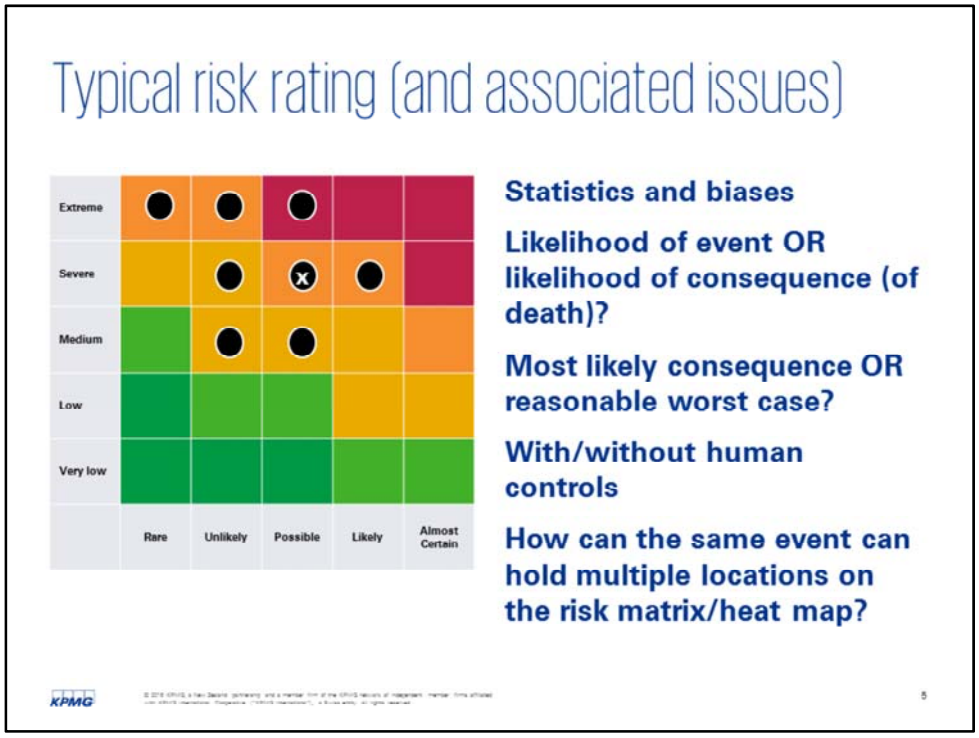

Continuing with the same scenario, we would wish to determine the rating for this risk – i.e. its likelihood and consequence. It can be assumed that the average response of group of people would be a risk rating that has a 'severe' consequence and 'possible' likelihood (see black dot with the 'x' above) – particularly if they are thinking about a safety consequence. However if we reviewed the range of responses across the same group (for the same scenario regarding a safety consequence), we would probably see a much wider range of options – as illustrated by the position of other black dots. What causes such variations? The reasons for this are:

- We are not good at understanding statistical concepts very well, and this is often exacerbated by our cognitive biases. (Refer to any behavioural economics text to explore this further).
- Do we assess the likelihood of the event (i.e. loss of control of vehicle) or likelihood of consequence of death? The latter will be of a much lower probability. What would happen to the risk rating if we were interested in other consequence types likes cost, the environment, reputation, etc.?
- Are we assessing the 'most likely consequence' or the 'reasonable worst case'? In the case of safety the former may be a collision with some sort of injury, whereas the latter may be single or multiple fatalities.
- Are we assessing the risk with or without human controls – since we know human factors in control application will have an impact on the risk.

<u>Nevertheless, how can the same event can hold multiple locations on the risk matrix/heat map?</u>

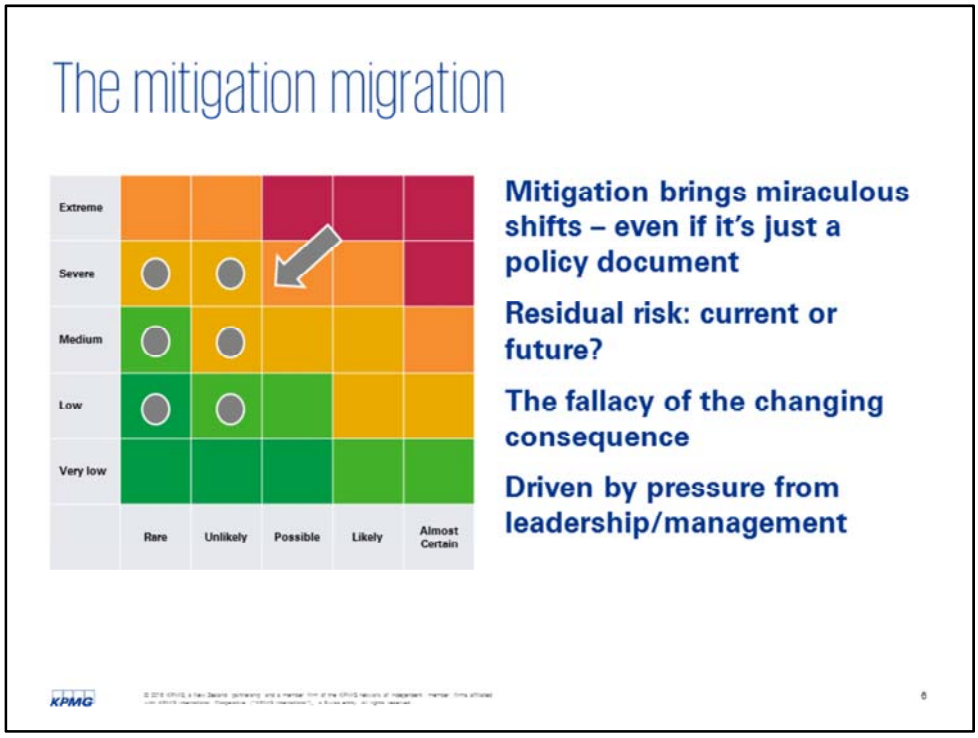

If we assume the previous slide was the inherent risk rating (i.e. before controls are in place), a separate set of issues arise when we determine the residual risk rating (i.e. with controls in place), namely:

- To begin with, the application of controls sees the risk rating miraculously shift down (in severity) and left (in likelihood) – even if the control listed is merely a policy document (dependent on people for implementation).
- Is the residual risk shown the current value, or where we anticipate it will be in the future (after all controls have been implemented)? Is this future next week, next month, next year, etc.?
- How can a consequence change so dramatically – if at all? (If you consider a vehicle accident in the 1970s there would a chance that someone might die; fast forward to today, event with a range of controls in place (like airbags and crumple zones) people still die in vehicle accidents. It's a fallacy to change all consequence levels willy-nilly).
- The drive to move risk into the green (or yellow) often comes from leadership and management pressure to do so. This can lead to the wrong behaviours in an organisation.

Base on the preceding argument we should stop using the risk register process to demonstrate management of risks, and rather use it for the identification and prioritisation of risks.

To address the shortcomings of the risk register we should be turning our attention to measuring control effectiveness. We should opt for the equivalent of ‘kicking the tyres’.

The risks we experience are seldom static. In our vehicle example: i) the condition of the vehicle can deteriorate; ii) we can expect impacts from other road users; iii) there will be driver issues (like fatigue), and iv) exposure to ever changing road and driving conditions. These changes cannot be reflected in the risk register. Not only that; if we choose to accept or harness any of these risks (i.e. the upside of the risk) how do we glean this from the risk register or position on a heat map.

Would it not be better to know the objective of a control and then assess whether it is being achieved?

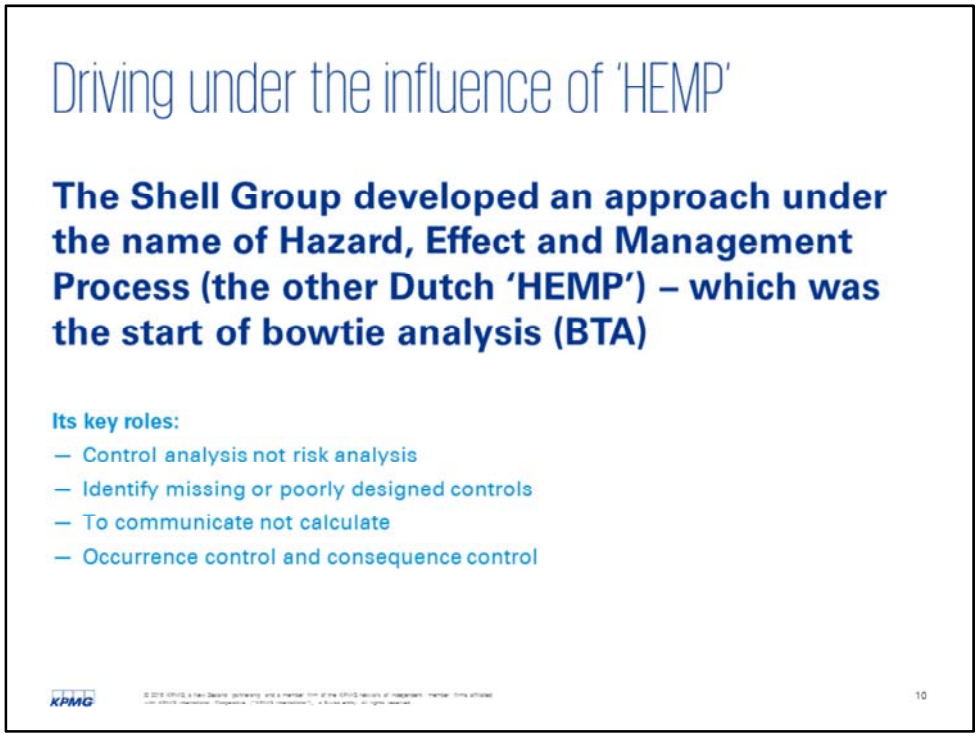

So how then do we determine how well the controls are working? There are a number of ways to address this, but the Hazard, Effect and Management Process (*the other Dutch 'HEMP'*) or bowtie analysis (BTA), as its more commonly known has proved to be effective. BTA was developed for use in process safety of particularly high-hazard industries, but has proved to be a useful tool for a wide range of applications.

Amongst other things, BTA's key roles (particularly when you first adopt it) include:

- Using it primarily for control analysis as opposed to risk analysis. Although the latter is also feasible, as your adoption of the technique becomes more detailed and technical. (For our purposes it is important to first analyse the risk before developing a BTA; otherwise it can be a clumsy and frustrating technique to follow).
- In so doing it is very helpful in identifying missing and/or poorly designed controls.
- Due to its visual representation, BTA is an extremely useful tool to communicate risk management - more so than to calculate risk. (Although this too is possible when your adoption becomes more detailed and technical).
- Understanding how to control the occurrence of the event and the consequence of the event

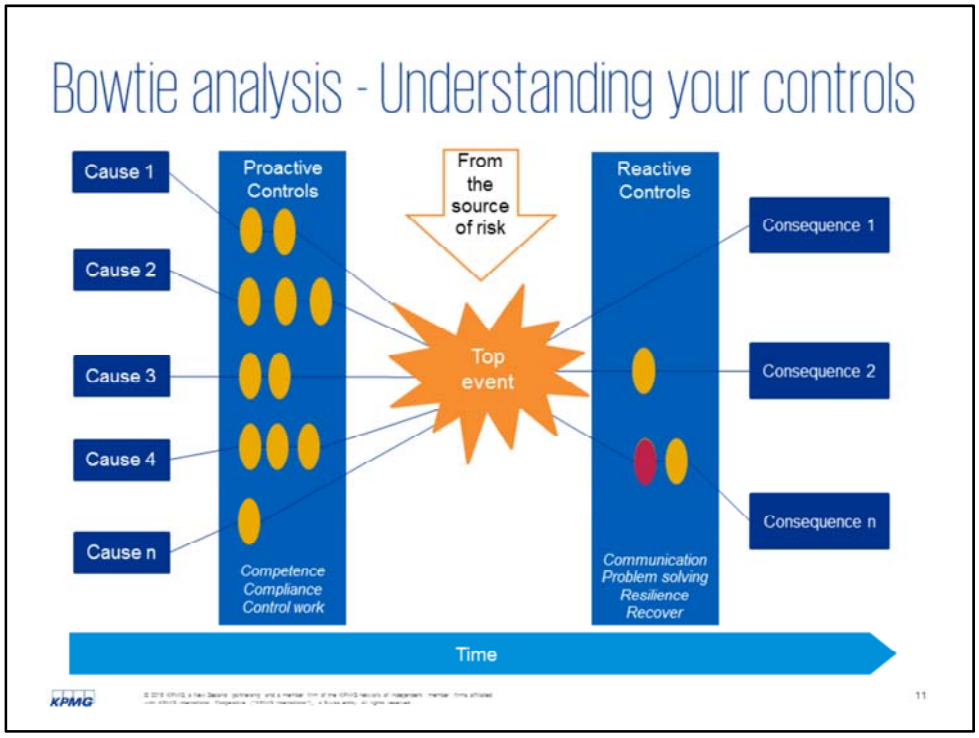

The starting point for a BTA is always the *top event* (also referred to as the *initiating event*) which is derived from the source of the risk (which in our scenario is the moving vehicle). The top event is the transition point (over time) from control the source of risk to the source of risk being out of control. We define the various causes of the event (on the left); which result in a range of consequences (on the right). We identify the controls (viz. the yellow ellipses) used to *proactively prevent an unwanted event* (which typically take the forms of competence development, compliance to standards, and controlling the work through careful planning and execution); and those that *reactively mitigate and recover from the consequences of the unwanted event* if it occurs (which typically include communication that notifies systems and/or people how to respond, in the moment problem solving, personal and organisational resilience and recovery to get back to operational state). Some of these controls may be considered more important than others (viz. red ellipse). Although the critical control is shown on the right in the above example they can be located on either side of the bowtie, depending on the context.

(NOTE: In its more technical and detailed application a BTA may comprise a 'fault tree' on the left, and an 'event tree' on the right).

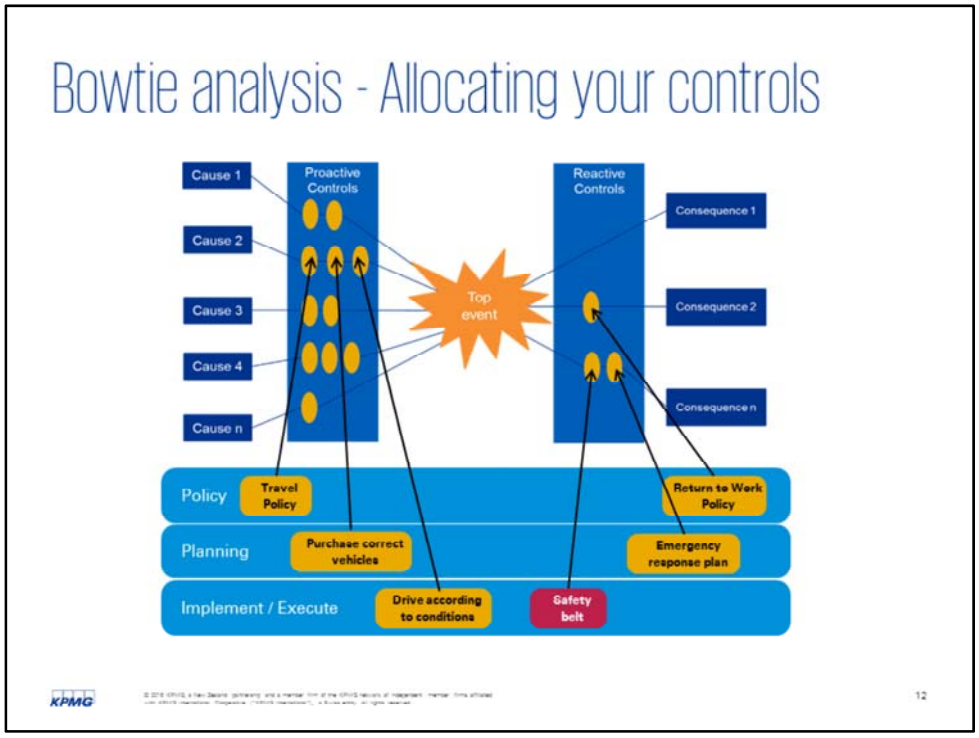

The controls included in a bowtie are located in the sequence in which they would be applied and are typically spread across all levels of the organisation: from those responsible for developing policies and standards (which may often be second line of defence), to those responsible for planning and scheduling work (which are the more senior levels within the first line of defence) and finally down to actions implemented and/or executed in relation to the task. The people who are responsible for the latter are predominantly employees in the first line, and are most often the front-line operators. Once we know the full suite of controls, we identify those which we consider to be 'critical' from the perspective of the operator through to management. They serve as a reliable proxy to demonstrate that the risk is being effectively managed. The example over the page will show how the critical control scales up from the operator to include a wider range of supporting mechanisms.

'Critical' control identification is not simple, however one may start with the following rules of thumb. The critical control is typically a control:

- *which one would not want to remove from a suite of controls, as its exclusion would have the greatest impact on the risk rating; OR*
- *would be the preferred control to adopt if there were no controls in place.*

When we think of a control (refer to red line and boxes), we need to consider what outputs and outcomes we expect from a control, and consequently the inputs that would be required to achieve this. If we determined the seatbelt to be the critical control we would assume the following:

- Output: hold the vehicle occupant in place during the loss of control of the vehicle/collision;
- Outcome: vehicle occupant is not injured during a loss of control of the vehicle/collision;
- Inputs: Quality seatbelt installed and maintained in vehicle (as defined in policy and specified as part of vehicle procurement), vehicle occupants trained to use seatbelt, etc.

We are then in a position to monitor these requirements as we 'move up' from operator level (where we are only monitoring control requirements, i.e. seatbelt characteristics that the operator is aware of and able to assess), through manager level (where we monitor the tasks associated with the use of the control), up to the risk owner who conducts performance assurance of the control and its inputs and outputs (i.e. extending beyond the seatbelt itself).

NOTE: This critical control (i.e. the seatbelt) is primarily used to manage a safety consequence, but we could have just as easily selected *adherence to vehicle maintenance requirements* as our critical control. In this case the critical control will have been on the left of the bowtie and by default we would be managing all potential consequences (cost, interruption, reputation, etc.).

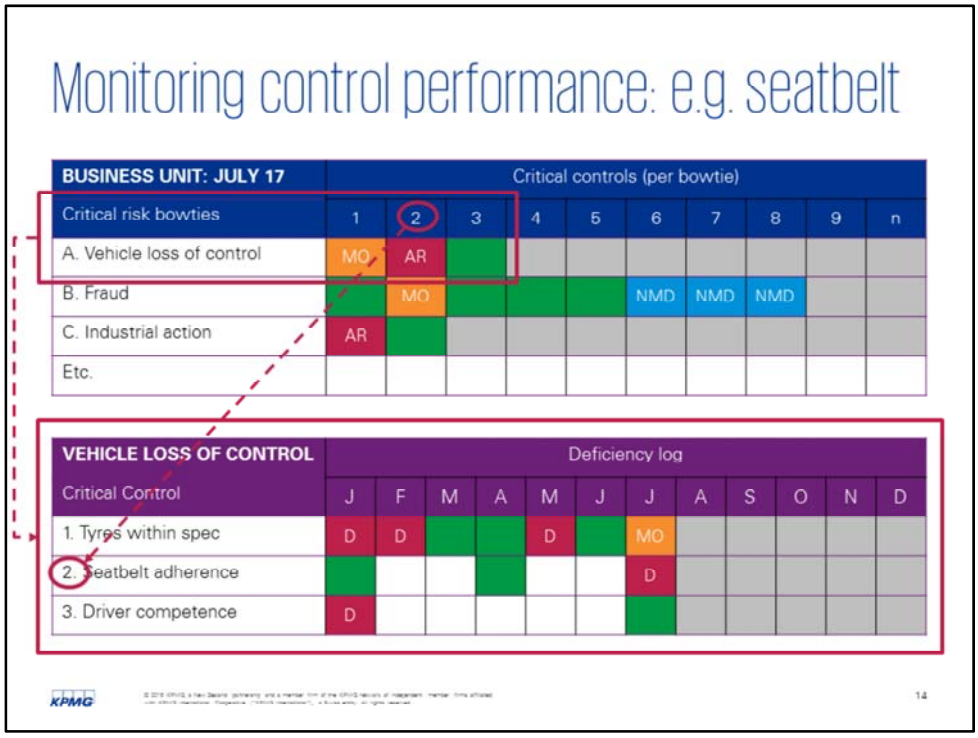

The technique can be used to keep track of:

- All the critical risks across a specific location for a specific time period (which for our example is Business Unit critical risk of 'A. Vehicle loss of control' for July 2017. This critical risk has 3 critical controls, whereas the 'B. Fraud critical' risk has 8 critical controls and 'C. Industrial action' has 2 critical controls); and
- The performance of the critical controls over the year depending on the frequency of the monitoring range selected (which for our example comprises a monthly assessment of '1. Tyres within specification', a quarterly assessment of '2. Seatbelt adherence', and a half yearly assessment of '3. Driver competence – up until July 2017).

(NOTE: The data provided not only shows the current state of the controls but also allows for further interrogation and trend analysis (such as sessional patterns, manager and/or asset performance in a specific area of the business, etc.)

Key

- *Green block – critical control functioning within specification*
- *Red block: Top table = (AR): action required – the critical control is not functioning according to specification. Bottom table = (D): deficiency – the critical control did not functioning according to specification.*
- *Orange block (MO): monitoring overdue - We don't know how the critical control is currently performing as monitoring has not been conducted.*
- *Blue block (NMD): no monitoring define – we've selected a critical control in a bowtie but have not defined how we will monitor it.*

For any item that is not within specification (i.e. not green), we would expect to see additional explanation and information of actions to address things.

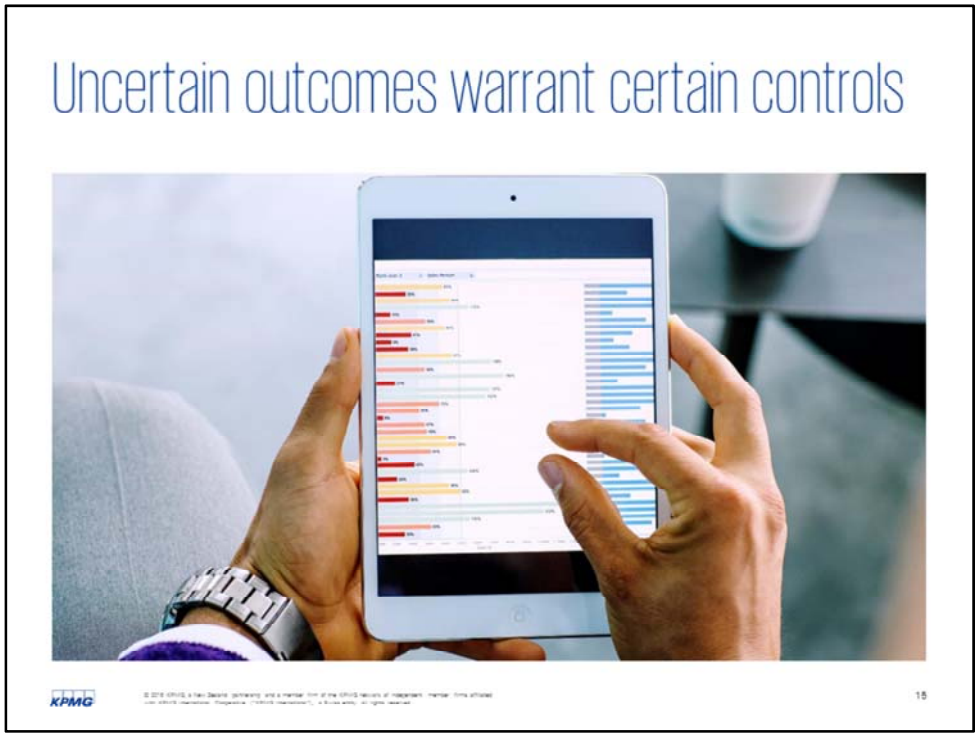

If we've determined that risks are dynamic and can lead to uncertain outcomes (whether good or bad), than we would be in a better position to assess this if we had some certainty in our controls.

KPMG

Where to now?

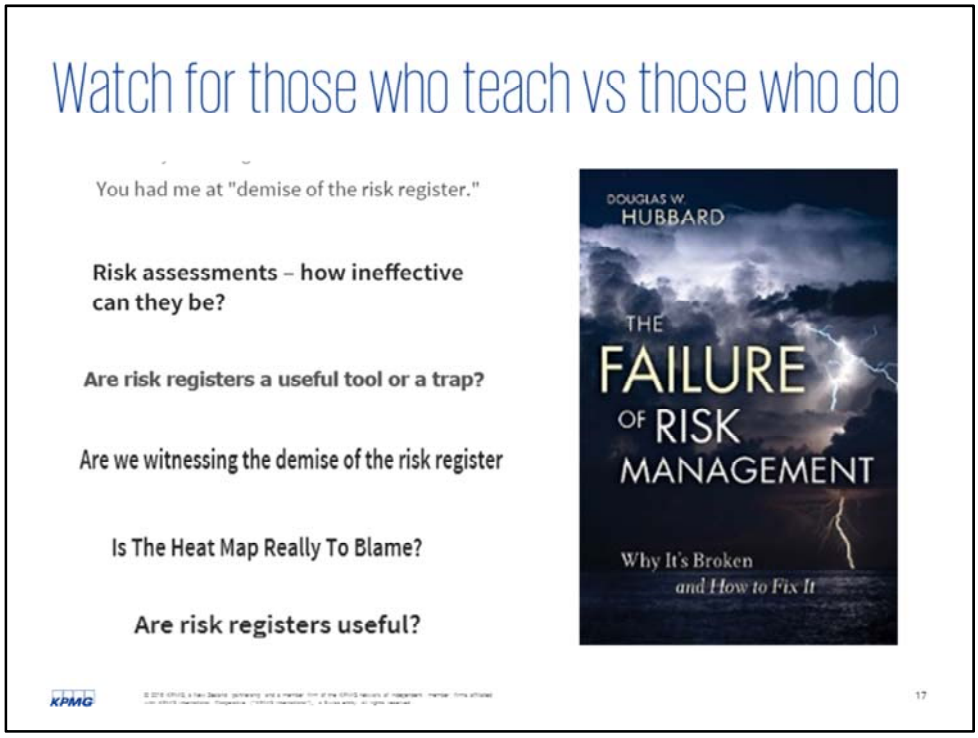

So where to from here? Well firstly, I'm not trying to add to 'noise' around any supposed faults with risk management not being an objective science, nor am I riding on the coat tails of those specifically challenging risk registers and heat maps (as the LinkedIn excerpts above illustrate). Instead, since I have been fortunate enough to straddle teaching risk management and performing risk management and I can assure that I've seen the critical control management process work. But, admittedly I have also seen the that risk registers are the de facto default of risk management practices across a range of businesses – and most often this has been poorly done.

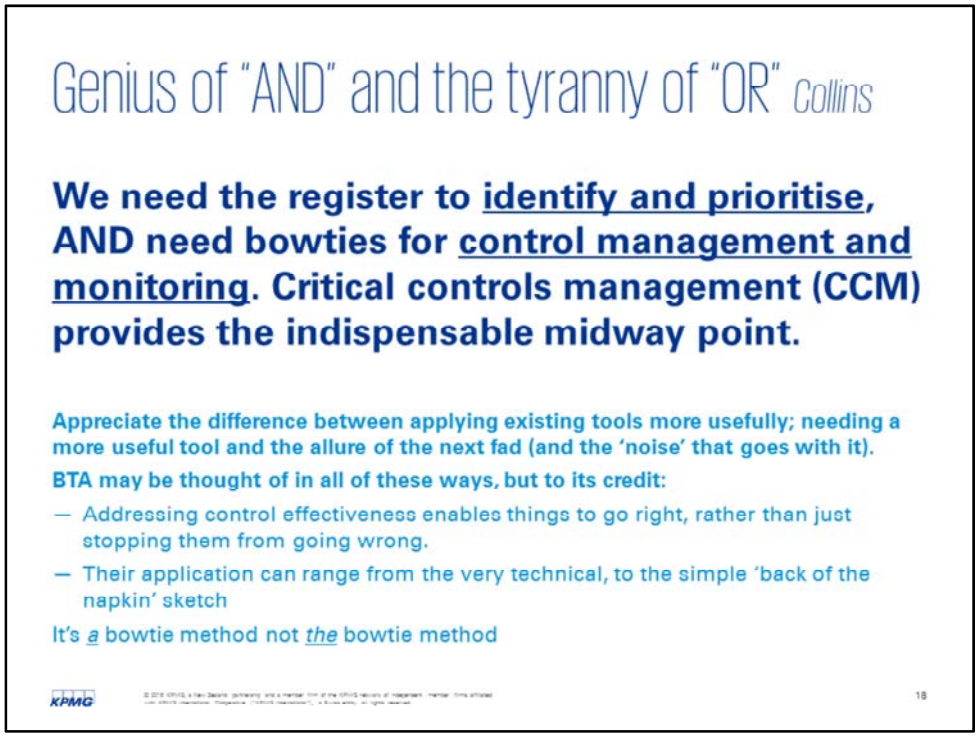

Consequently I am not advocating the one tool over the other but rather suggesting the use of both the register and critical control management.
To do this effectively, we need to appreciate the difference between applying existing tools more usefully; needing a more useful tool, and being cautious of the allure of the next fad (as may be punted by those who 'teach' risk management).
Although BTA may be thought of in all of these ways, it does have a number of benefits, such as:

- Being a technique that can be used when needing to define control effectiveness, to enable things to go right, rather than just stopping them from going wrong.
- The tool is scalable and can be applied across a range of complexities, from the very technical (with a fault tree on the left and event tree on the right), to the simple 'back of the napkin' sketch (that helps us understand the trajectory of a source of risk from in control to out of control).

Remember it's a bowtie method not the bowtie method – so find what suits your needs and adapt/mature it as you go.

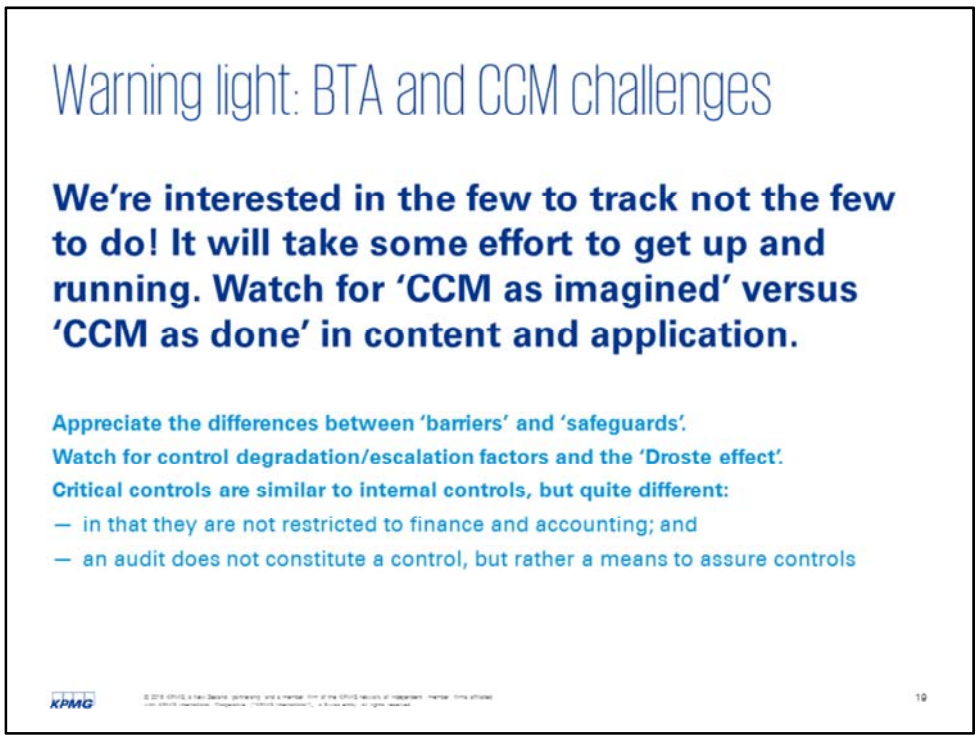

However, like most good things, there's fine print. In this case one should be aware of the challenges of adopting the BTA for critical control management (CCM) – it will take some time to get set up. We recommend that you have a deep narrow focus on a few risks to begin, rather than a shallow wide approach that demand much effort before yielding value.

Once you've selected the risk/s you wish to focus on it is critically importantly that everyone involved understands that CCM does not mean you only have to implement a few controls; instead it is the few controls we select to keep an eye on. (NOTE: All your controls should be considered necessary – if not, you're wasting resources).
When getting going with CCM guard against being swept away by an exquisitely rigorous solution, and rather adapt it to the type of information that is readily available for tracking (i.e. the content) and how the process of tracking this information will actually be performed (i.e. its application). In other words narrow the gap between 'CCM as imagined' versus 'CCM as done'.
If you're going to get into BTA and CCM watch for some of the following hurdles (as the author can attest from first hand learnings):

- There is a difference between a barrier (which 'halts' the progress of a risk) and a safeguard (which supports a barrier). All too often the latter is incorrectly given elevated status, and consequently we infer a greater sense of control than warranted.
- Identify and appreciate that controls can be undermined by degradation/escalation factors (e.g. the vehicle occupants behaviour in using or not using a seatbelt will determine its effectiveness). Know which is which and how to insert them into the BTA.
- The presence of the 'Droste effect' (https://en.wikipedia.org/wiki/Droste_effect), for example if you deem employee capability to be an important component of management of a risk, the capability of the person/s providing the education to enhance

employee capability is equally important to that outcome. As a result know to include in the critical control monitoring requirements.

Finally, you'd be correct in seeing similarities with the management of internal controls as described in the COSO framework; but there are a number of differences, such as:

- This version of operational risk management is far broader than financial and accounting controls and can be applied to almost anything in the business;
- An audit or audit process does not necessarily constitute a control itself but is rather a means to assure (verify and/or validate) the controls.

When will it click? Focus on controls not colours

Thank you

Ross Liston
Risk, Safety and Sustainability

kpmg.com/nz