Don't worry, we've got that under control…

"Don't worry, that's not a risk, we've got that under control!" I've been told this earnestly and confidently so many times, but so often without any serious thought or critical reflection.

I am an unapologetic 'broken record' on a particular hypothesis – the majority of major risk realisations throughout history have been caused by overconfidence in ineffective controls[1]. Honest analysis shows that if we even pay rudimentary lip-service to thinking about risk, we will rarely be blindsided by unforeseen or truly unforeseeable risks. Rather, the correct response would be "We knew that was a risk, but we thought we had it under control".

The RMS Titanic sinking, the Space Shuttle Challenger disaster, the deaths during the Commonwealth Home Insulation Program, the events of the evening of the 2016 Census, the 9/11 terror attacks, the Global Financial Crisis and the Japanese Fukushima nuclear disaster are all examples of foreseen or foreseeable risks that were thought to be adequately controlled.

By foreseen risks, I am not suggesting we can predict every risk realisation at the specific level of detail to which it ultimately unrolled. The disintegration of the Space Shuttle Challenger on a particular lift-off in January 1986 began after a joint in a solid rocket booster failed. O-ring seals used in the joint were not designed for the unusually cold conditions at this launch, allowing burning gas from the rocket motor to reach critical structures and the external fuel tank. Now, I accept this is a very specific 'risk realisation' however it represents a type or class of risk event that was well understood. The special commission appointed to investigate the accident found NASA's organisational culture and decision-making processes had been key contributing factors to the accident, with the design of the boosters containing a potentially catastrophic flaw in the O-rings, but they had failed to address the problem properly. NASA managers also rejected warnings from engineers about the dangers posed by the low temperatures of the launch morning and failed to adequately report these concerns.[2] Overconfidence in and breakdowns of trusted systemic (engineering management systems to ensure design integrity)

[1] Controls are those things we believe we have in place to help us manage a risk. They can be systems, resources, procedures and processes, people capability or infrastructure.

[2] Adapted from Wikipedia article on the Challenger disaster.

and specific (pre-launch warnings) controls contributed to this disaster and were not unique to the specific failure of those O-rings on that morning.

Building on this hypothesis, I am also a firm believer that inherently (pre-control) severe and currently well controlled risks are the most dangerous risks we manage. These risks are often referred to as 'control-critical' risks highlighting their critical reliance on current controls. Because of the effectiveness (perceived or otherwise) of these controls, these risks may have a moderate or low residual severity[3]. However, over-confidence in these controls, or changes in context, can quickly move these risks from control-critical to out-of-control. They become 'watermelons' – reported as green, but red in the middle. History demonstrates that these risks pose more danger than the risks which are explicitly acknowledged as being currently severe.

Douglas Hubbard in his book *The Failure of Risk Management* refers to 'catastrophic' overconfidence. "Perhaps one of the most pervasive, exhaustively researched, and thoroughly confirmed phenomena discovered by [judgement and decision making] psychologists is that almost everyone is naturally overconfident in their predictions." "They will underestimate real risk systematically."

So, if for the sake of the argument you accept the premise that overconfidence in controls is a major cause of risk realisations how do we engineer our risk frameworks and culture to address it?

I suggest you can do four things:

1. require the explicit assessment of control effectiveness during Risk Analysis,
2. implement mechanisms to report and manage control-critical risks,
3. assess controls against a holistic maturity model, not just the "are they working today?" test, and
4. focus attention on your most critical or pervasive controls.

Assessing control effectiveness

I love watching workshop groups endlessly debate the likelihood of a risk event on a poorly described five-point scale. Is it "probable" or "likely"?!?! Firstly, humans are terrible at estimating the likelihood of anything outside their day-to-day experience and secondly, risk events are often broadly or poorly defined, making likelihood estimates hit or miss at best. There is one conversation though that doesn't reach the point of diminishing returns – challenging the effectiveness of controls. The critical questioning of control effectiveness is an essential element of a good risk culture.

To encourage this, I recommend you include the explicit assessment of control effectiveness in your risk analysis step. By this, I mean a non-negotiable requirement to rate the effectiveness of the controls for each risk on some scale. This can be "1-5", "good-bad" or "strong to weak", it doesn't really matter how you implement it. What is important is that it requires people to have

[3] I'm using the term "residual risk" to refer to today's risk considering the effect of current controls. This is sometimes also referred to as current or real risk.

the conversation, rather than simply writing a list of aspirational or assumed controls and then moving on with misplaced confidence.

You can assess control effectiveness against each individual control or collectively as a set. Assessing each control individually is perhaps more robust, but it does add process. Also, some controls aren't really meaningful in isolation and considering the set as a whole is my best-compromise recommendation.

The higher the inherent risk severity, the more rigorously control effectiveness needs to be considered. Nothing in life is ever completely "under control" nor is anything uncontrollable. We can exercise a degree of control over any risk. The control may not be strong, but saying a risk can't be controlled and therefore can't be managed is a fatalistic cop-out. Somethings we can't prevent, but we can seek to mitigate the most undesirable consequences.

Don't neglect control critical risks

I've seen risk frameworks that only capture or report severe (high or very high for example) residual risks. This, by rule, removes control-critical risks from the formal risk management and assurance process. Control critical risks may be currently low or medium residual severity today, but they remain critically dependent upon the continued effectiveness of the controls in place.

If these risks are removed from the risk framework, what management process will ensure this continues? The risk framework should capture risks which are both residually high, and risks which have been assessed as inherently severe and control-critical, so may be residually low today. We need to move away from the flawed premise that low residual risks are safe or good – they're not necessarily either!

Holistic control maturity assessment

Often the assessment of the effectiveness of controls can be as narrow as "does it work now?", or "well, the risk hasn't happened yet!". The O-rings in the Space Shuttle had been reassuringly effective controls for many launches. We need to think beyond this narrow lens to determine if our controls are truly effective.

I suggest controls should have six attributes - they should be sustainable, timely, holistic, well defined, assured and resilient. Some of these attributes are relevant to individual controls or only to sets of controls.

They are summarised in the diagram below and some key questions you might want to ask to test them are provided.

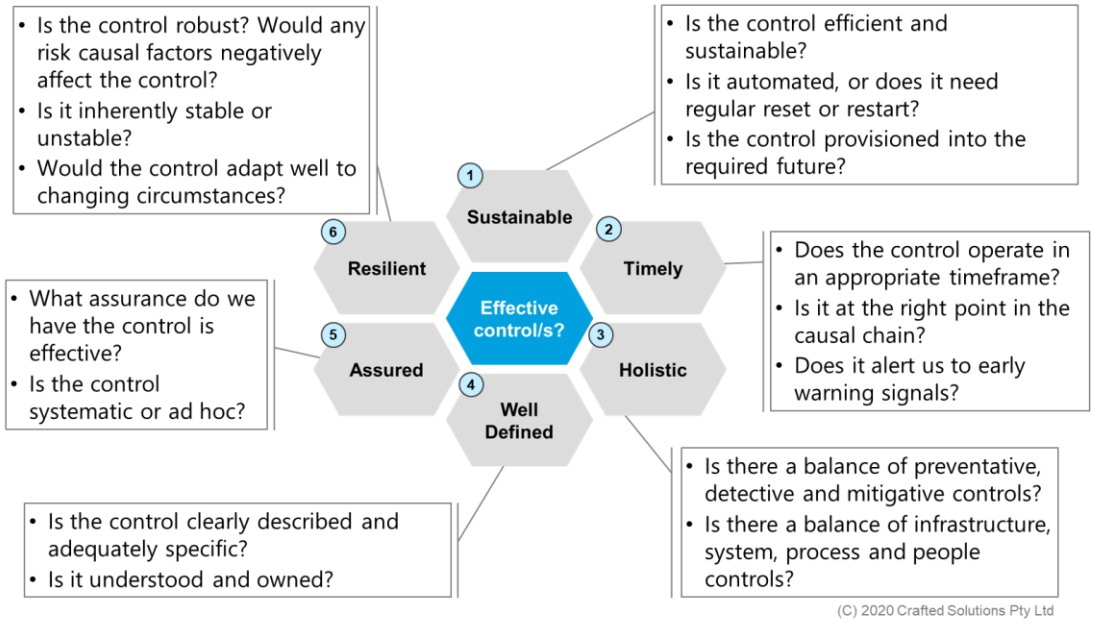

1. Sustainable

- Is the control efficient and sustainable? We live in a real world where controls have to be paid for and if that cost is deemed excessive or unsustainable, the control may be removed or weakened, sometimes without warning or notification. Do we understand our 'pervasive controls' – those controls that act against many risks? Funding these well can be disproportionately rewarding. Is the control provisioned into the required future or at the very least the time-scope of your risk assessment?

- Is the control automated, or does it need regular reset or manual restart? I once had a house with a rough-built basement storage area. In heavy rain, this would flood and sometimes quite quickly, so I put a sump-pump in there to pump out the water, especially if the rain began when I wasn't at home. Unfortunately, the circuit breaker on the pump would randomly trip, rendering the pump useless until I thought to check and reset it. A good control in theory, but I couldn't rely on it because it would turn itself off unpredictably until I intervened. Is your control built into a work program, or do you need to constantly kick it off again every few monthly management meetings?

2. Timely

- Does the control operate quickly enough? Are you trying to mitigate short or long-term impacts? Audit is often described as a risk control, and it does have important role in that regard. However, it can be untimely. Who wants to be told in an annual audit report that they've had 12 months of systemic fraud in their program?

- Is the control at the right point in the causal chain? (A causal chain is the series of events leading to event leading to further events etc.) The debate about causes vs risks vs consequences is a long one and I'll write on that in a future article. In summary, an event can be a cause, a risk or a consequence depending on your frame of reference. However, a reason for trying to pick the right point in the causal chain to define as your risk event is to maximise your ability to exercise control. If you are too far downstream in the causal chain

it is often difficult to implement effective control. The "brace position" is not my preferred aviation safety control.

- Does it alert us to early warning signals? Good controls provide a detective function (refer below), in that they let us know when they are being tested or are in jeopardy. Employed appropriately, IT cyber-security systems can be quite good at this, acting not only to prevent successful cyber-attacks, but also alerting us to the scale and nature of the evolving threat. Does the fire-exit door which is designed to allow safe egress but prevent access into the building from outside (a security control) let the security control centre know it's been chocked open by staff going for a smoke or a coffee?

3. Holistic

- Is there a sensible balance of preventative, detective and mitigative controls? Controls can be thought of in three main forms. **Preventative controls** reduce the likelihood of the risk event occurring. (For example, a security fence). **Detective controls** help understand and predict the risk event and how well other controls are working. (For example, security cameras or audit). **Mitigative controls** act to reduce the consequences should the risk event occur. (For example, insurance or a business continuity plan). It is a human bias to prefer preventative controls, but do we have an appropriate balance in case our preventative controls are overcome?

- Is there a balance of infrastructure, system, process and people controls? Policies are inherently weak controls unless backed up with other attributes. Good people controls are essential – people have to know how and want to do the right thing, even when they think no one is watching.

4. Well-defined

- Is the control clearly described and adequately specific? I'm not a smart man, but if I read a risk register and cannot understand what the described control actually is, I doubt the people implementing it do either. There is a balance here between specificity and clear language. Phrases such as "good communication" and "staff training" by themselves are not controls, they are generalities that are not specific, measurable or able to be assured.

- Is the control understood and owned? Many controls will be implemented by others on your behalf. It sounds obvious, but do these control implementers know and accept they have that responsibility and the effect on the risk should they stop doing it? What mechanisms do you have to ensure this in place and doesn't drop away next time the control implementer has a big restructure, cost-reduction, or crisis of their own?

5. Assured

- What level of trust or confidence do we have that the control is effective and operating as intended? Do we just trust the person implementing the control or do we have some means of testing that? The more critical the control for the bigger the risk, the more assurance we need.

- Is the control systematic or ad hoc? Often a control will be described as "Mary always does that". Mary does something to control a risk (whether she understands that link or not), because she chooses to do it or has just always done it. It's not documented, it's not

trained, it's not assured, it's not measured. If Mary is sick, moves on or just decides the control is too much effort, what happens to the control? Is the control part of a system or just something that "Mary always does"?

6. Resilient

- Is the strength of the control proportionate to the inherent risk? If your main defence against a tsunami is a sea wall, and modelling or experience suggests that a tsunami could be X metres high (with the same likelihood you estimated for the risk), then don't be calling your 'half-X' metre tall sea wall an effective control.
- Is the control resilient and robust? Controls can be vulnerable to the same causal factors as the risks they manage. If your back-up IT server (a control for the risk of a system outage) is located in the same basement as the primary server it may be an effective control against motherboard failure, however it provides no resilience against the risk of monsoonal flooding.
- Is the control inherently stable or unstable? A well-functioning modern swimming pool gate is an inherently stable control, if it is left open it closes and latches under its own weight. Of course, this can be circumvented, but it provides a higher level of control than an ordinary gate which remains in whatever position it was last left. Old school pre-ABS car brakes under the foot of an unskilled driver were also unstable – applying the control too hard could result in a skid which would actually significantly reduce the effectiveness of the control.
- Would the control adapt well to changing circumstances? Adaption to change is a key attribute of good risk management. Is the control highly reliant on the circumstances that exist today with little ability or a plan to adapt? Will the O-ring still work if we launch on a cold morning?

Focusing attention on your most critical controls

Sometimes, documented controls can be long lists of people's day jobs or of marginal relevance to the risk. They are padded with theoretically relevant, but mostly ho-hum stuff. However, in amongst this are sometimes a few really important or critical controls which if they fail would have little backup or directly trigger the risk. Or, as mentioned earlier, they could be pervasive controls which work against several of your risks simultaneously.

Like all lists of risk 'things', controls should be prioritised. The most critical controls should be first. They should also be the most carefully designed and described, and the most deeply challenged, tested and assured. As you are designing, documenting or reviewing the controls for a risk, think like an experienced hiker packing their controls into their backpack for a long difficult hike. Pack carefully considered, contextually relevant, good quality and reliable gear as your life may depend on it, but remember you'll have to carry it.[4]

[4] Refer to my very brief risk parable "Pack up your risk controls in your old kit bag…" for the full version of the hiker analogy.

Summary

In summary, I believe the critical and systematic challenge of control effectiveness is perhaps the most powerful attribute you can build into your risk framework and culture. The effectiveness of controls should be challenged continuously and against a holistic set of attributes, not just "is it working now?".

First published 9 January 2020, updated as v1.2 on 28 September 2022

Sal Sidoti changes how people and organisations think about risk. He has over 30 years' experience working with public and private sector organisations across Australia and further afield. Sal works with his clients to design and implement tailored risk arrangements that support decision making in practice. Approaches that go beyond cut-and-paste compliance and death by spreadsheet. Sal is the Director and Principal Consultant of Crafted Solutions Pty Ltd. He is home based near Canberra, Australia.

https://au.linkedin.com/in/sjsidoti
enquiries@craftedsolutions.com.au
www.craftedsolutions.com.au

Title image courtesy of Rallis Kourmpetis via Unsplash.