A cape doesn't always make a superhero:

Understanding our unrelenting overconfidence in risk controls

I recently spoke on how you can 'systematise' or embed a culture of control effectiveness into an organisation or program. My talk was in-part based on my article *"Don't worry, we've got that under control"*[1], and my long-held view that the majority of major risk realisations throughout history have been caused by overconfidence in ineffective controls[2]. We are consistently and sometimes catastrophically bullish in our ability to manage the risks we know we face.

But, why?

The talk prompted me to further explore the causes of this overconfidence, and there are many of them. Human and organisational biases, and process or system limitations, all contribute. It can often be more than simple optimism bias. In my previously mentioned article, I talk about things you can do overcome this overconfidence, but to enable that, it is helpful to explore why it occurs. Knowing the possible root causes helps you recognise them and tailor fit-for-purpose solutions to overcome them in your context. My favourite top 12 causes of overconfidence in risk controls are illustrated and explored below, including some suggestions on how you can address them.

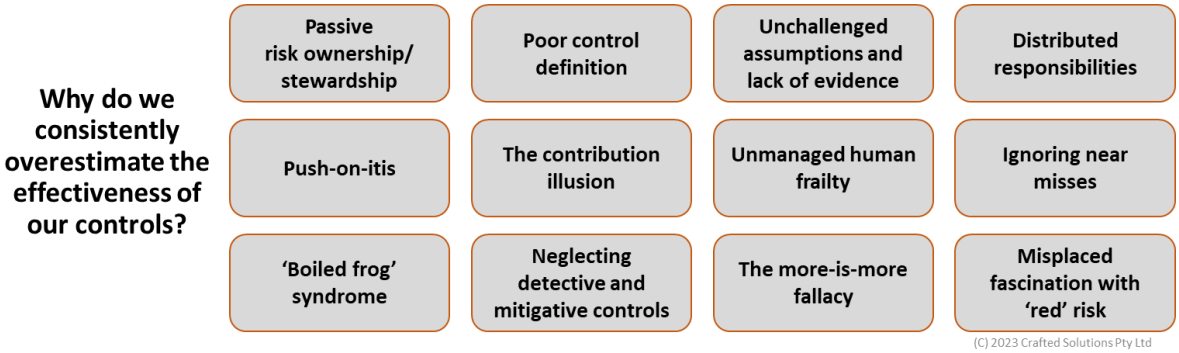

[1] https://craftedsolutions.com.au/wp-content/uploads/2022/09/Under-Control-v1.2-Crafted-Solutions-FINAL-2022_09_27.pdf

[2] Controls are 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.

Passive risk ownership or stewardship

I once reviewed a risk framework which allocated the treatment action "active management" to a subset of risks. My response was "ok then, what is *passive* risk management? Just sitting back and waiting for it to happen?"

All risks require active management, although the nature of that management will vary. The champion of this is the risk owner or steward - the person assigned responsibility to oversee and manage each risk.

Their primary role is to **actively** monitor all aspects of the risk and raise the alert if any of these move out of agreed parameters or if something new or concerning emerges. The key word here is active, and this includes actively ensuring that agreed controls remain effective. Failing to do this is perhaps the most common reason why controls become neglected, ineffective or irrelevant to the risk they were managing.

Ensure risk owners or stewards are appointed for every risk and
they understand one of their most important tasks is the active monitoring and
challenge of the controls in place.

Poor control definition

Controls that aren't adequately defined can't be effectively implemented or assured. Too often, controls are described as two or three word nick-names or summaries, and these are so inexact, that it is hard not to be sceptical as to what is actually being put into action.

For example, I often see "staff training" as a critical control for a significant risk. In isolation, these two words are largely meaningless. What staff, what training, when do they get it, who tracks completion, who ensures the training remains relevant, etc, etc? While we can't write "War and Peace" for every control, they should be sufficiently defined that they can be tested and assured. If the control is as meagrely defined as "staff training" then all you need is for some staff to have done some kind of training and you can consider yourself good-to-go.

'Row-per-risk' risk register templates can encourage this excessive brevity. Risk tools and 'control libraries' can also contribute to this problem. Drop-down selections of common controls might appear efficient, but can encourage people to toss in a bunch of generic controls that aren't tailored or even relevant to the risk.

Finally, each control should clearly identify who is actually responsible for doing it in practice. And (it is unfortunate to have to write this) the implementor should be aware of this allocation, understand the requirement, and have acknowledged the responsibility. Unowned controls are almost guaranteed to be works of fiction.

Ensure control definitions and descriptions are adequately
comprehensive and include clear responsibilities for
control implementation and oversight.

Unchallenged assumptions and lack of evidence

It is a persistent human frailty that we are systematically overconfident in any list of controls we ourselves author. I've seen people defend an obviously rubbish list of controls like it was their first-born child. There is sometimes the belief that if I've written this list of controls, it is *your* job to prove to me that they *aren't effective*.

That is completely backward. The starting assumption should be that controls are ineffective until proven otherwise, and evidence is available to justify that belief. To support this, clear objectives and performance measures should be agreed for each important control. These will also enable the monitoring of ongoing effectiveness over time.

Controls should also be routinely pressure tested and assumptions challenged. Give that swimming pool gate a good hard shake once in a while to make sure the latch really is still locking solidly (even in winter when it is easy to neglect). Have a staff member submit a dodgy expense claim and see what really happens. For particularly critical controls, a level of independent assurance should be available.

Make a control effectiveness assessment an explicit element of every risk analysis. Require risk owners or control authors to explain why they think their controls are effective and against what objectives and performance measures they made that judgement.

Increasing distribution of responsibilities

The world is becoming less monolithic. The age of large organisations (public and private) who do everything in-house is over. Outsourcing, partnering, off-shoring, shared services, contracting and matrix management are the way things are done. One of the consequences of this is that controls are often not implemented by the people who 'own' the risk. This is magnified for highly shared risks – those large complex risks which have no natural single owner.

In a monolithic enterprise, everyone arguably answered to the same ultimate boss and hence in theory would be working to the same priorities and goals. A common task united them, and controls could be relied upon because they were implemented 'by our own people'. Ok, it was never completely quite as rosy as that, but today's highly distributed operating models make the task of ensuring control effectiveness even more difficult. More sophisticated governance models are needed.

Every control should have an assigned owner or implementor.
If that is outside your team, organisation or project, the local link/liaison/contact for the control should also be identified.

Push-on-itis

Most commonly talked about in aviation safety, push-on-itis refers to our tendency to want to continue to the destination, regardless of problems we might be experiencing. A pilot who is almost half-way through a flight might want to continue to the destination, even if a potential problem should cause them to return to their origin or to an alternate diversion point.

The same phenomenon occurs in projects and programs nearing completion. Many of their risk controls can become irrelevant or ineffective. However, rather than pausing to review, adjust or reinforce the risk controls, the program manager pushes on, hoping to get across the line before any of the associated risks can be realised.

Controls suitable at project initiation may not be appropriate later. Similarly, there can be a focus on sunk-cost – 'we've invested so much already, let's just keep going!'

Be particularly aware of the potential for decayed or irrelevant controls as projects near completion. Have specific transition reviews to discuss the right controls as projects move through their life-cycle phases or gateways.

The contribution illusion

When a risk hasn't been realised for a period of time, it can convey the impression that the controls or a particular control are effective. The illusion of a positive contribution is created.

Risks are by definition uncertain, and hence they can be realised or not. It can be difficult sometimes to understand the extent to which a given control is influencing that. For example, if I wear a traveller's lucky charm, and have never been in a plane accident, what is the contribution of the charm as a control? If I remove the control, will I be immediately now doomed? Unless I understand how the risk works as a system (i.e. an interconnected and interrelated set of parts), I won't be able to know. It is seductive to assume that the control must be effective simply because the risk hasn't happened yet.

The contribution illusion can also be used to justify unnecessary or irrelevant controls. The control is in place, the risk hasn't occurred yet, hence the control must be essential. Countless security and safety controls are justified on this basis every day.

To overcome this, you need to understand how each control works, how it interacts with other controls, and its effect on causes or consequences. There are a host of analytical methods to assist with this. Even a simple bow-tie analysis can rapidly and visually (remember, humans are visual creatures) illustrate control gaps or potential redundancies. The more expensive or inconvenient the control, the more effort should be invested in understanding the impact of the control being in or out of the system.

Justify your investment in controls by analysing the role and contribution of each to the risk 'system'.

Unmanaged human frailty

Almost all controls rely on people. People do things, people check things, people put something in place, or people designed and built the control in the first place. Yet the average person can't push a shopping trolley down the aisle of a supermarket without crashing it into something. I was once told by an eminent accident investigator that every time he heard the phrase "human error" his response was "humans will *always* make mistakes - who allowed that human to make this mistake and for it to have the consequences it did?"

If people are a control or part of a control, we need to put a wrapper around them to help manage their inherent frailty. They need to be part of a system that supports them. Even the most highly competent pilots in the world are part of a system of aids, colleagues, enablers, checks and assurances.

Whenever you see a control reliant on a human being ask the question – how do we ensure the person can and will do what we need and want them to – even on a really bad day? If your key fraud control is a weekly reconciliation by the CFO, who or what ensures that person does that reconciliation to an appropriate standard? Even when they've got a sick child and its budget week, or their Internet gambling debts are piling up?

Controls reliant on a person should also explain
how that person's actions are supported and assured.

Ignoring near-misses

Risks are being nearly realised all around us, every day. Often this is the result of one control failing and others coming into function. For those familiar with the 'Swiss cheese' model, it may have been one last 'slice of cheese' that blocked the risk from being realised. Each of these near misses should be an opportunity to better understand our controls and their weaknesses or potential vulnerabilities. If we are oblivious to them, it is easy to continue on assuming our controls are all effective when they aren't.

Cyber-security is a great example. Modern enterprises are being subjected to hundreds or even thousands of attempted intrusions every day. The vast majority are blocked, and each is an opportunity to better understand the risk if they are studied and learnt from. Good cyber defences are continually evolving, but only because we acknowledge that every day 'it could have happened'.

When an undesirable event occurs or nearly occurs, ask the questions - had we identified that as a risk, had we assessed it correctly, and were the controls we thought we had in place effective? If the answer to any of these is no, it is an opportunity to improve your management and control of that type of risk.

Ensure you have a formal process to review incidents and
near misses through a risk lens.

Boiled frog syndrome

There is an urban myth (now debunked I understand) that if you place a frog into a pot of hot water it will jump out, but if you put one into a pot of cold water that you then heat, the frog will not jump out and eventually perish. The analogy being that if the situation around us worsens slowly enough, we won't make a hard decision to do something about it.

Control effectiveness can be like that. The relevance and effectiveness of controls can decay slowly over time and it's easy to put off doing anything about it. The tread on your car tyres is a control for skidding on a wet road, but when do you replace those expensive things? This month they're only a tiny bit balder than they were last month, so no problem, I can go to the tyre shop next month…

Hard metrics or Key Risk Indicators (KRIs) can provide a trigger point at which a decision is forced. Modern tyres have tread wear indicators that provide a visible signal that the tread depth has reached a point that the tyres are unroadworthy and need replacement.

Critical controls should have agreed objectives and performance measures attached that provide a visible health indicator. Monitor control performance against these and have hard triggers to force corrective action when needed.

Neglecting detective and mitigative controls

Controls can have a combination of three main functions:

- Preventative controls reduce the likelihood of the risk event occurring. (For example, a security fence or password access to a sensitive system)
- Mitigative controls act to reduce the consequences should the risk event occur. (For example, insurance or a business continuity plan)
- Detective controls help understand and predict the risk event and how well other controls are working. (For example, security cameras, or audit and review).

It is a strong human bias that we focus on preventative controls (and are overconfident in them) at the expense of detective and mitigative ones. The result is a fragile set of controls reliant on prevention.

Each risk has its own context. Some must be prevented, some can only be mitigated (their consequences reduced), but each control function should be considered as part of a holistic self-reinforcing set. Again, bow-tie analysis or similar visualisations can be very useful for seeing how controls sit around a risk and act on causes or consequences.

Ensure a sensible balance of preventative, mitigative and detective controls are in place for every risk.

The more-is-more fallacy

Less is almost always more in everything to do with risk management. Yet, some lists of controls seem to be judged on their length rather than their quality. Impressively long lists of largely irrelevant padding. This is bad for two reasons:

- to a disinterested, inexpert or time-poor reader the long list might convey the impression the risk is well controlled, and
- the padding distracts from the possibly few really important controls hidden in the list, making them less likely to be implemented or challenged as a priority.

But, what makes a control more important? Some examples include:

- the criticality of the control – the level of reliance upon it or lack of alternatives,
- the fragility or vulnerability of the control, or
- the pervasiveness of the control – does it control many risks?

Control lists should be concise and prioritised.
The few most critical controls should be first and clearly flagged.

Our misplaced fascination with red risk

'Green' risks are not necessarily safe or good. More organisations are burnt by 'green' risks they thought they had under control than 'red' risks they knew were uncomfortably severe.[3] Yet, we spend so much time focused on the red risks, sometimes removing the green risks (by rule) from our risk reporting and discussion. Why, because the green ones are under control and hence, they can look after themselves! Without love, the control effectiveness these risks are reliant upon will inevitably decay unnoticed and they will become watermelons – calmingly green on the outside, but dangerously red inside.

Whether it is an explicit part of your framework or not, understanding the inherent (or pre-control) position of each risk, even in broad terms, in addition to its residual or current (post control) severity can be helpful in identifying control-critical risks. These then need to be subject to constant review and challenge, not parked for a disinterested six-monthly gloss-over.

Even low risks require regular challenge and review,
especially those that are control-critical.

[3] I am referring here to the common convention used in typical likelihood/consequence 'heatmaps' that severe risk is red and mild risk is green. Yes, these heatmaps are rubbish, but a lot of people still use them.

Summary

Nothing in life is ever "completely under control" or "completely uncontrollable". But, overconfidence in ineffective controls is the #1 cause of risk realisations.

The critical and systematic challenge of control effectiveness is perhaps the most powerful attribute you can build into your risk framework and culture. It helps ensure our risks are neither over or under controlled.

To achieve this, it is important to understand why overconfidence in controls is so common. It isn't always the simple optimism bias written about so often - there are other potential causes.

First published 2nd February 2023.

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 Jeremy Perkins via Unsplash.

