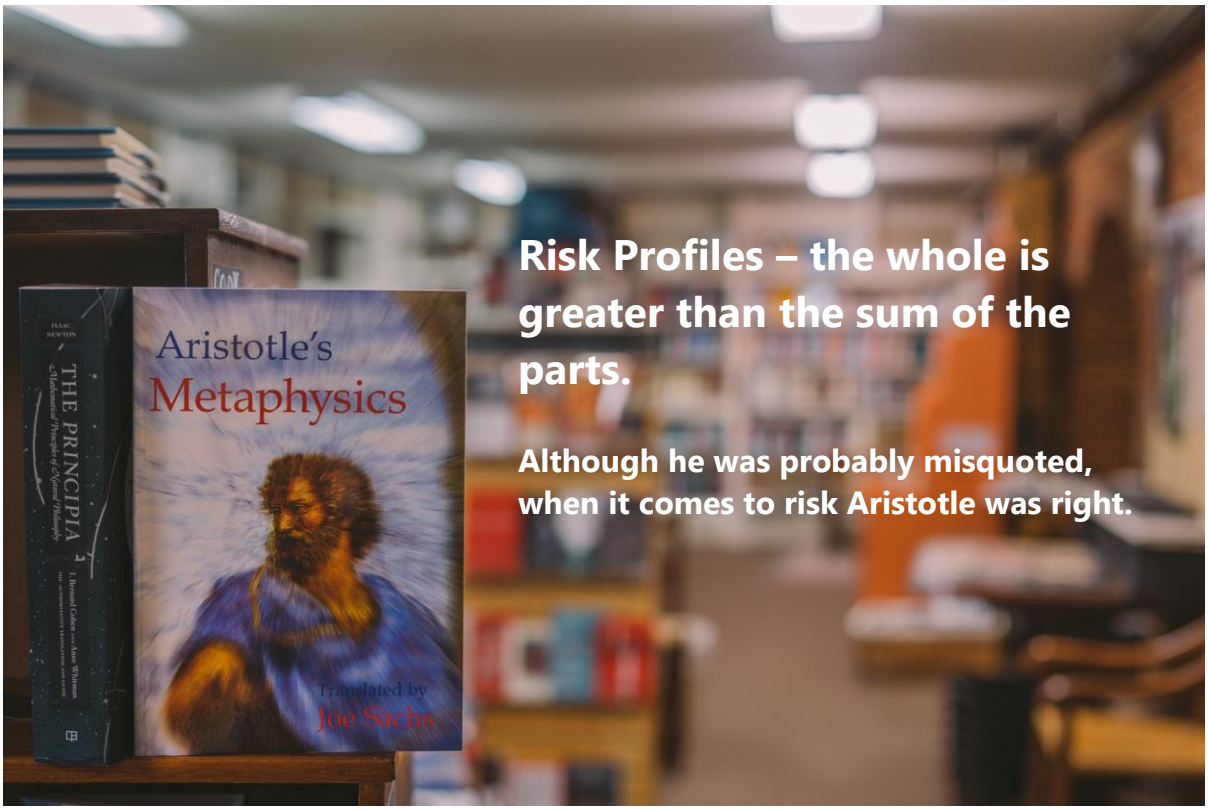

We need to think about our risks as profiles - something more than lists of individual, unrelated uncertainties. Failing to systematically consider aggregations and interdependencies between our risks is a major vulnerability and weakness.

I've used the term 'risk profile' for a long time, and people sometimes ask why. Why not just call it a risk register or risk list? There are several reasons, but the most important is that a profile conveys the impression of a whole, not just a set of independent parts. Regardless of how you derive them or what approach you use, any set of risks for an organisation or activity needs to be considered as a set, as well as individually. In my experience, this remains a major failing in many organisational approaches to risk management.

I once conducted an independent review of a major project. Based upon the project's own risk register I concluded it was in danger and likely to fail. This finding was not initially welcomed, and I was reminded that none of the project-level risks were of "high" or "extreme" severity. They were all "medium" or "low". How could the project be in jeopardy?

Now, I hadn't redone their original risk assessments, nor am I ever a fan of simplistic likelihood/consequence/severity analysis, but that is what I was given to work with in this case. The reason I made this judgment was because of my concern over aggregation and interdependence across and between their risks. While in isolation (given their own framework and business rules), each risk was acceptable, I felt the profile wasn't. The project's risk problem was systemic.

The challenge compounds when we consider that business units, organisations or projects are rarely islands divorced from everything around them. Each operates in an environment or ecosystem of partners, competitors or neighbours. Our lives and work increasingly rely on networked systems with complex interactions.

You may need to consider the risk exposures and strategies of your ecosystem partners as they may compromise yours. Projects operating within programs are breeding grounds for sometimes unidentified risk aggregations and interdependencies.

To understand our true total exposure, we need to understand these intra and inter relationships as well as the individual risks themselves.

These relationships have two main forms – aggregation and interdependence. Aggregation refers to accumulations of risk and interdependence to relationships between them. The key things I look for are summarised in the table below.

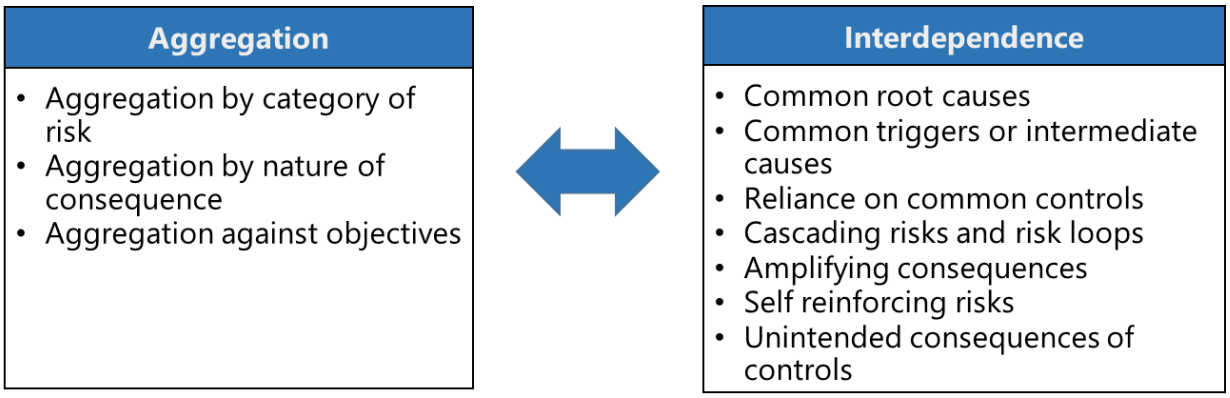

Failing to consider risk aggregation and interdependence is a significant structural failing in many risk programs.

Aggregation

Aggregation asks what the true consequence would be should several risks occur at once or in quick succession. It helps understand the actual risk exposure for all or a part of an organisation or project and can be more complex than just adding together your individual risks. There is the chance that if two or more risks are realised simultaneously or in a short period of time, that the consequences might be multiplied (i.e. more than the simple sum) or be unacceptably damaging.

You might argue that the chances of two or more risks being realised simultaneously is low enough to ignore. However, there are several reasons why this is more common than you might think.

- Real-world risks are rarely either "realised" or "not realised", there is big grey area between a pure risk and a pure issue[1] They are more often realised on a spectrum of

[1] This reflects one of the failings of point-value likelihood/consequence analysis models – it encourages people to accept that there is a single X % chance of the risk being realised with Y consequences. That's rarely how the real-world works. Refer to my article "The Risk Management Drinking Game" for further explanation of this and several other risk myths.

consequence at any given point of time. Some risks might always be part-realised with some consequences being felt.

- Risks are rarely truly independent of each other and for many reasons often 'arrive together'. I expand on this in detail below.

There are three main forms of aggregation I suggest you consider – aggregation by category, by consequence and by objective.

Aggregation by Category. Many organisations group their risks in categories – financial, safety, people, etc. Just say a major project had six people category risks in its 'top ten'. This suggests that risk to people and people capability is a major concern. Collections of risk against a single category suggest a systemic vulnerability. Unsurprisingly, such vulnerabilities should be addressed systemically. What is the organisation doing to manage its people capability into the future, is this project merely a reflection of broader vulnerabilities or concerns?

Aggregation by Consequence. Each of our risks will likely have several consequences should they be realised. Many organisations use this anticipated level of consequence as a major factor in determining the severity of each risk.[2] But what happens if several risks have the same type of consequence, and these could be realised simultaneously? These consequences would then sum or multiply and may significantly exceed our ability to tolerate.

The simplest example is with financial or monetary consequences. We might invest in the stock market with five different portfolios, which we might optimistically consider five independent risks. Each portfolio, should it badly tank, could cost us $10,000, an amount we could cover. But what if three or four portfolios simultaneously lost $10,000? This would exceed our ability to cover and could be financially devastating. We rely on the fact that these portfolios are independent and will not collapse simultaneously – something which is often not the case. (GFC anyone?)

Aggregations by Objective. AS ISO 31000 defines risk as the 'effect of uncertainty on objectives' and despite the efforts of some to twist and misquote this to invalidate it, I believe it is a fine definition.[3]

Several of your risks could simultaneously jeopardise the same objective. For example, one of an organisation's strategic priorities might be made uncertain by several of the organisation's enterprise risks. While individually the uncertainty they create to each objective might be acceptable, collectively, they might not.

Often it is this 'objective-aligned' level of uncertainty that senior decision makers are most interested in. "You can talk about your individual risks all you want, what I really need to know is how certain I am that I can achieve my most importance value-preservation and value-creation objectives." Objective-centric approaches to risk management make this linkage very explicit.

[2] Unfortunately, in some frameworks, alongside a point estimate of likelihood, a point estimate of consequence is one of only two narrow dimensions of risk severity.

[3] Having lost sight of this definition is one of the main reasons I believe risk management in some sectors has lost its way, but that is a story for another day…

Interdependence

Interdependence considers the relationships between risks which might lead to their aggregation. Life works in causal chains – something happens that causes something to happen that causes something to happen. We might have been quite happy to have the first something happen, but are devastated by the last something/s. These somethings can be risks. We need to understand, communicate and if necessary, manage these relationships and not just the individual risks themselves.

Common sources and root causes. Many environmental and contextual factors lead to the existence of our risks. For example, increasing global temperatures and weather instability is a significant and shifting causal factor for many real-world risks such as crop failure or flooding damage. Understanding the root causes for our biggest exposures helps us identify and influence them systematically rather than trying to blunt their effects risk by risk.

For example, we might identify that higher-than-average staff churn rates have reduced tenure and experience in our organisation. We also note that low levels of competence or currency is a contributing factor for many of our big risks. In addition to introducing enhanced training job-by-job, we could also seek to understand and address why our churn rates are so high. If we are successful, this could have an enterprise-wide impact across our entire risk profile. Post event reviews of significant risk realisations (e.g. major accidents or disasters) almost always reveal underlying root causal issues that had not been identified or addressed.

Common triggers. Risks usually don't just happen - they are almost always triggered by a preceding condition or event. Triggers emerge from root causes and if unmanaged lead to risk events occurring. Poor quality control during aircraft manufacturing could lead to micro-cracks in holes drilled in airframe structures. These then are the triggers of intermediate causes such as major cracking failures and ultimately the loss of the aircraft. Again, the world works in causal chains – stuff happens that make other stuff happen. Should several of your risks share a common trigger, then it is highly possible they may all occur at once. Good risk frameworks will identify common triggers and put additional controls in place to manage them or block their ability to initiate the risks.

Reliance on common controls. Controls are things that manage our risks. They can be preventative, detective or mitigative. I am a broken record on the importance of controls and believe that the biggest cause of risk realisations in practice is overconfidence in ineffective controls[4]. Sometimes however, a single control might be important to many risks and the failure of this control could jeopardise all of them. Should our confidence in this control be overstated (a massively common human bias) we now have simultaneous and multiple risk occurrences with the further potential for aggregation and compounding consequences. We need to identify these 'pervasive controls' and ensure they are particularly robust and well monitored.

Cascading risks and risk loops. A particular form of interdependence is cascades between risks, where one risk in your profile could trigger another (or many others). In risk loops, this process closes back on itself, with a cycle of reinforcement and repetition. A cursory review of any risk profile you might have access to, will likely quickly reveal such cascades and loops – but where and how are they analysed and communicated in a structured way?

[4] Refer to my article "Don't worry, we've got that under control."

Self-reinforcing risks. Although not strictly a risk interdependence, self-reinforcing risks should be mentioned here. These are risks that as they begin to be realised can cause themselves to get worse as their consequences can circle back as causal factors. Staff exoduses and stock market runs are good examples - once they start, they build a momentum and become self-reinforcing. This ability needs to be reflected in the assessment of the risk and a full range of preventative, detective and mitigative controls deployed to break the feedback loop.

Unintended consequences of controls. Sometimes, treating one risk can make others worse. Many medications have this effect - they reduce the risk of one condition but increase the risk of others. Explicitly identifying these 'unintended consequences' in your risk analysis can serve as a warning that particular controls should be used sparingly or with caution.

Considering aggregations and interdependencies systematically

Being aware of the potential for aggregation and interdependence between risks or within networks is the first step in understanding and managing them. Review your risks (and those of your neighbours) and consider each of the pathways described above. You then need a means to record and communicate these relationships, something that may not fit naturally within many traditional risk frameworks. Visual representations of risks, rather than tables of rows, can be useful.

In more complex environments, there are a range of technical approaches that can be used to model or quantify these relationships. Fault tree analysis and social-network analysis are a few examples that can be explored for applicability. Each has different techniques which can help us understand and communicate the strength and criticality of the relationships between risks and risk profiles.

Summary

Any set of risks is more than the sum of the parts. We need to understand and be able to communicate the interdependencies and potential aggregations between them. Regardless of how you derive your risks and how complex your environment, three steps are key:

1. Understand the potential for interdependencies between your risks and risk profiles, how this can be manifested, and how important it can be.
2. Workshop, model, analyse or otherwise understand these relationships for your risks and risk profiles.
3. Identify fit-for-purpose mechanisms to record, communicate and manage them.

v1.0, originally published September 2022.

Sal Sidoti changes the way people and organisations think about risk. He has over 30 years' experience working within and advising public and private sector organisations across Australia and further afield. He works with his clients providing tailored risk management advisory services that support decision making in practice. Outcomes not templates, approaches that go beyond cut-and-paste compliance and death by spreadsheet. Sal is the Director and Principal Consultant of Crafted Solutions Pty Ltd and is home based near Canberra, Australia.

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

Title illustration courtesy of Tbel Abuseridze via Unsplash

