Cybersecurity - a complex and ill-structured domain

Cybersecurity is an ill-structured, complex domain, rife with uncertainty. Although it's got some unique aspects, by and large these problems are not totally unique to what Human Factors and Safety Science have been working on:

- Risk and safety in complex systems
- Accident causation
- Coping with complexity
- Automation surprises
- Coordination of man-machine systems

Large cybersecurity incidents are organisational incidents, meaning that it's the working of the organisation, the dynamic of the socio-technical system that couldn't control their production processes and the hazards well.

Looking at the past events, we need to learn and go beyond "distancing through differencing", as at the core, looking at the complex socio-technical systems, cybersecurity problems are not that dissimilar from the patient safety problem in healthcare or aviation safety, to name but a few. Both domains undergo rapid changes, both are hypercomplex and interconnected in their own ways.
Consider the similarities:
In 1987 a state of the art ferry called Herald of Free Enterprise (HoFE) departed from the Belgian port of Zeebrugge. Due to a number of small "failures" and conditions combining together the ferry departed with their bow door open. As the ferry went past the port boundary, the bow wave increased to 4 metres high, flooding the main deck. The main deck had no bulkheads by design, to increase the efficiency of loading and unloading the ferry. The ferry took in water and capsized, killing 193 people, - the worst maritime disaster since World War II.
The ferry was seaworthy, the company had the licence to operate, the captain and other crew had been properly licenced. This accident led to the creation of the UK Maritime Accident Investigation Board (MAIB).

In 2017 Equifax effectively left their internet "front door" opened, - the organisation didn't notice that one of the Apache application servers remained unpatched despite Equifax InfoSec department receiving a US CERT notice, disseminating it internally to get it patched and doing several scans..
Effectively, as the ferry in the example above, **Equifax went sailing across the sea of the internet not noticing that their "bow doors" were open**. As the Apache Struts 2 vulnerability was actively exploited, no wonder that someone found their unpatched server and exploited it.
The investigation revealed a lot, but most staggering is that Equifax, as their CEO testified:

- Spent 250 million dollars in 3 years on cyber security
- Had 225 cybersecurity professionals employed
- Was ISO27001 certified

Although Equifax's breach was investigated by a Congress committee, no cybersecurity incident investigation board has been created.
Looking at the numbers provided by the Equifax CEO and having looked at the patterns of the organisational breakdown (e.g. fragmented distributed problem solving and anomaly left in limbo, weak signals missed), it's clear to me that more of the same approach is not going to be of any help.
It's not that we need more money, more "cybersecurity" people and more normative specifications (standards), - we need a **qualitatively different approach**.

Both of the events, - HoFE capsizing and Equifax are complex organisational disasters, they cannot be understood by just looking at the technology and physical assets. In both cases, the investigation shows how complex systems - although weak signals were abound, the organisations were drifting into failure. People, organisation and technology need to be looked at as a joint unit of analysis with their relations and interactions understood.

So, what to do? Here are some broad suggestions.

Move away from linear simplifications to understanding complex systems.

It amazes me that although everyone says that **cybersecurity is a complex problem**, **but we typically use linear models**! Defence-in-depth, kill chain, MITRE ATT&CK, - we have plenty. Most of the time, there isn't even a debate as to what models of failure we use in cybersecurity, linear or complex? In order to understand cybersecurity incidents in complex organisations we need to use complex models that reflect what's going on in the system better. AcciMap (https://en.wikipedia.org/wiki/AcciMap_approach) developed by Jens Rasmussen and STAMP developed at MIT (http://psas.scripts.mit.edu/home/) can be good candidates.
Acknowledging that cybersecurity is a complex problem means doing away with the term "root cause". There is no **"the root cause"** in a complex tangled layered network, it's not found as there is no **objective** stopping rule for the investigation, it's assigned based on various constraints, including politics and power. Falls are caused by gravity, drill rig blow-out are caused by the hydrocarbon pressure, and cybersecurity breaches are caused by cybercriminals, but that does not help much as the root cause either.
Rasmussen in his 1997 paper "Risk management in a dynamic society" noted "Savage and Appleton (1988) talk of 'second generation management applied to fifth generation technology' in manufacturing."
The generation of complex technology that we use for information processing is far beyond the generation of cybersecurity risk management using linear models.

Cybersecurity is a multi-disciplinary problem

Some of the largest advances in understanding of accidents and safety were done not by technologists, but by sociologist and other non-technical researchers, consider the theories that stood the test of time:

- Normal Accident Theory by Charles Perrow (sociologist)
- High Reliability Organisation (a multi-disciplinary group)
- Normalisation of deviance by Diane Vaughan (sociologist)
- Practical drift by Scott Snook (organisational theorist)

Focusing on **STEM is not enough**, cybersecurity and safety are complex and multi-disciplinary. A **little bit more technology and normative prescriptions (standards) are not going to work** (they evidently haven't moved the needle much in the last 20 years), we need multi-disciplinary programs aimed at cybersecurity. Cybersecurity needs to actively involve other disciplines, i.e., human factors, sociology, organisational theory and safety science.

Understanding and learning from failure

One of the reasons for the commercial aviation to have achieved the ultra-safe status is their relentless learning from failure and disseminating it across the industry.
Cybersecurity is a very young domain and is getting a lot of attention lately due to high-profile breaches, this means that we might not have had the time to go through the learning cycles. But other industries have! Aviation, oil and gas, maritime, petro-chemical, healthcare, - have all had the attention of society and regulators for their accidents. Various investigation boards have been established around the world.
To me it means this:

- IT and cybersecurity may need **an investigation board** and learn from incidents and near misses
- We need to learn how other industries have been managing complex risks

Academics working in Australia provide plenty of insight into the work and failure of complex systems, - Andrew Hopkins from ANU and Sidney Dekker from Griffith University, to name a few.

Tale of two stories. Escaping hindsight bias

When a cybersecurity incident happens, like in Equifax's case, exposing over 140 million personal records, the **first story** that emerges is a story of human error. In Equifax's case, the SVP that didn't forward the email was blamed and fired by the CEO. The CISO was all of a sudden bullied in the media for having a degree in music composition. Such judgements fuelled by **hindsight bias** and are unhelpful in learning
To make progress on cybersecurity in complex systems, we need to look for the **second story**, as Richard Cook and David Woods present in their seminal paper on patient safety "A Tale of

Two Stories: Contrasting Views of Patient Safety" (1998) (https://www.researchgate.net/publication/245102691_A_Tale_of_Two_Stories_Contrasting_Views_of_Patient_Safety). As in healthcare (a hypercomplex domain!), to understand incidents in complex systems, we in cybersecurity need to look for the second stories and try to escape the hindsight bias.

> The detailed investigations are second stories revealing the multiple subtle vulnerabilities of the larger system which contribute to failures, detecting the adaptations human practitioners develop to try to cope with or guard against these vulnerabilities, and capturing the ways in which success and failure are closely related.The second stories examine how changes in technology, procedures,and organizations, combine with economic pressures to create new vulnerabilities and forms of failure at the same time that they create newforms of economic and therapeutic success. (Cook and Woods, 1998)

The second story of the Equifax breach revealed several failure patterns in the organisation, trade-offs and ill-formed adaptations directed at coping with complexity of the IT operations and the business of the company. It is important to go beyond the level of technology that often dominates in cybersecurity investigations, e.g. lack of segmentation, passwords saved in plain text, expired certificate, etc **and move to understand the patterns of work, trade-offs, goal conflicts and constraints** in the organisation that lead to such disastrous outcomes.

I wish the public and the industry knew the second stories of the recent high-profile breaches in Australia.

Cybersecurity is a complex, socio-technical problem. Managing Complex Risks

One of the founding fathers of the contemporary safety science - Jens Rasmussen (https://en.wikipedia.org/wiki/Jens_Rasmussen_(human_factors_expert)) in 1997 offered a way to model complex organisational risks: Risk management in a dynamic society, a modelling problem (https://backend.orbit.dtu.dk/ws/portalfiles/portal/158016663/SAFESCI.PDF)
In his paper he discussed the entire socio-technical system as a unit of analysis for risk management, going from the government and regulations all the way to the company managing a hazardous process (information processing in the case of cybersecurity). At the end of the day, cybersecurity is about managing a hazardous industrial process, but unlike, say, the petro-chemical industry, it does not involve flow and processing of hydrocarbons, it involves flows and processing of data.

Rasmussen suggested that all complex organisations in the process of their optimisations drift into failure. Economic forces together with the workload create a pressure for the operations of the company to move towards danger, in other words, a drift.

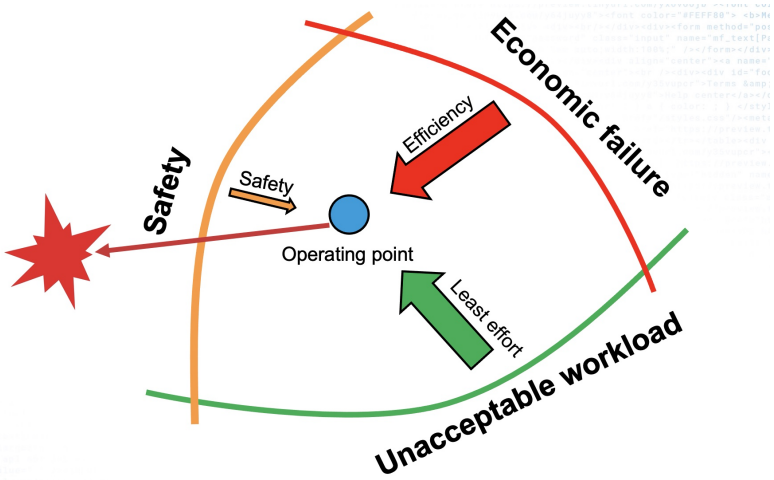

Rasmussen's suggestion was instead of fighting normative deviations (standard violations), make boundaries explicit, provide ways to cope with the breakdown at the boundary and provide ways to rebound.

In Equifax's breach example, the organisation was drifting towards the safety boundary under the forces of economic performance and workload pressure.

Rasmussen argued that "a top-down, system oriented approach based on control theoretic concepts is required":

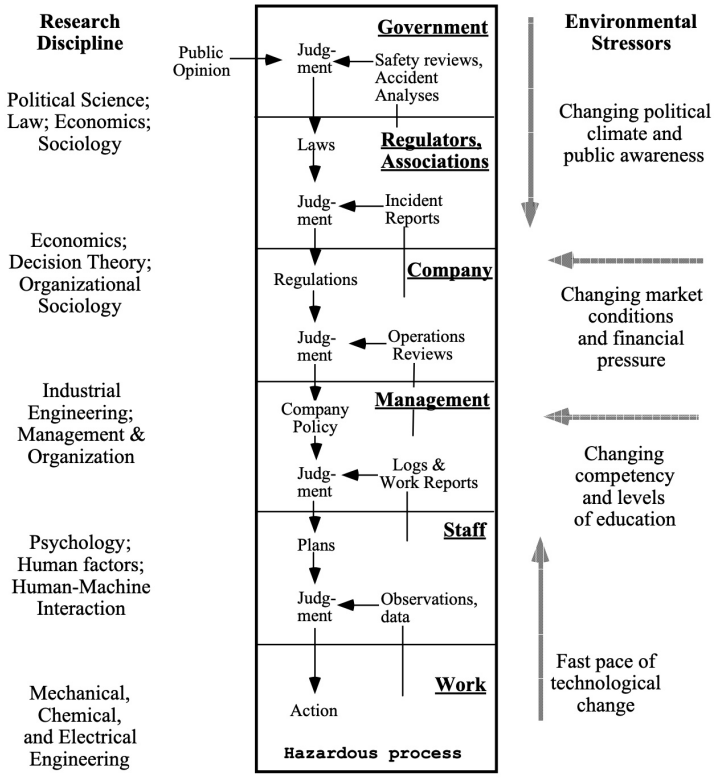

Resilience Engineering

In the beginning of 2000 there was a realisation in the Safety Science that a new paradigm was needed, as the accidents that the industries were experiencing were manifesting in new ways. Future incidents are not going to be like the ones we're experiencing today.
Resilience Engineering (https://en.wikipedia.org/wiki/Resilience_engineering) as a subfield of safety science was born to understand how complex adaptive systems cope with surprise. In other words, the focus is on what we need to build-in today that will play out in the long run as sources of resilience.
David Woods, one of the founding fathers of Resilience Engineering points out that there are three common patterns to the failure of complex adaptive systems:

decompensation: exhaustion of capacity when encountering a disturbance.
working at cross purposes: when individual agents in a system behave in a way that achieves local goals but goes against global goals.
getting stuck in outdated behaviours: relying on strategies that were previously adaptive but are no longer so due to changes in the environment. This means that relying on **"best practice" can be dangerous**. Constant update and recalibration is required.

All of the above points are relevant to cybersecurity:

- Organisations exhaust the capacity to deal with a hazard (e.g. ransomware).
- Different groups at the organisation work at cross purposes - (faster, better, cheaper vs safer).
- When the world around us is changing at a rapid pace, getting stuck in the models that were previously adaptive, but no longer help is a sure way of experiencing a **fundamental surprise** (ransomware and perhaps AI will be one of new threats).

Engineering resilience means building adaptive capacities into the systems that play out in the long run. In the short run, however, resilience looks like inefficiency, under the pressure of economic performance, maintaining resilience is hard!

Public policy: Cybersecurity is a 'wicked problem'

Cybersecurity is a wicked problem (https://en.wikipedia.org/wiki/Wicked_problem), it resists resolution, there is no single solution to the problem, there is no determinable stopping point, an effort to solve one aspect of the problem reveals or creates other problems. Although inherently technology-heavy, cybersecurity is a social problem, as vividly demonstrated by the recent breaches in Australia, affecting nearly half the population.

Australian Public Service Academy (https://www.apsacademy.gov.au/sites/default/files/2022-04/22-016%20-%20Toolkit-TheoryBites-31Mar22_04_muddling.pdf) echoes Charles Lindblom and his classic 1959 paper “The Science of Muddling Through” (https://www.jstor.org/stable/973677) that suggest:

> Making policy is at best a very rough process. Neither social scientists, nor politicians, nor public administrators yet know enough about the social world to avoid repeated error in predicting the consequences of policy moves. A wise policy-maker consequently expects that his policies will achieve only part of what he hopes and at the same time will produce unanticipated consequences he would have preferred to avoid. If he proceeds through a *succession* of incremental changes he avoids serious lasting mistakes in several ways. (Lindblom, 1959)

Making a public policy for cybersecurity is “the science of muddling through”, as Lindblom indicated. There is a great risk that the dynamism of the domain will outpace and outmanoeuvre the goals of the policy without significant feedback and calibration.

Josephine Wolff in her “You’ll see this message when it is too late. The legal and economic aftermath of cybersecurity breaches” outlines the interplay of various ex-ante prescriptions, ex-post liability, the cybersecurity reporting obligations, and cyber insurance that creates a Gordian knot as complex as the domain itself.

Wolff suggests that one way to achieve progress is to target the end goals of the perpetrators that go beyond the context of compromising a computer system. The conflict of interests of different stakeholders has shaped the present state of the art in the cyber domain, created certain recurring patterns of behaviour. “It is these recurring patterns that are most helpful in trying to identify the bottleneck stages of incident lifecycles. They indicate which elements of cybersecurity incidents remain static across time because perpetrators are so dependent on them and have so few alternatives. … The variety of different intermediaries and competing interests involved in cybersecurity incidents has largely served to undermine cybersecurity efforts and progress. ” (Wolff, 2018)