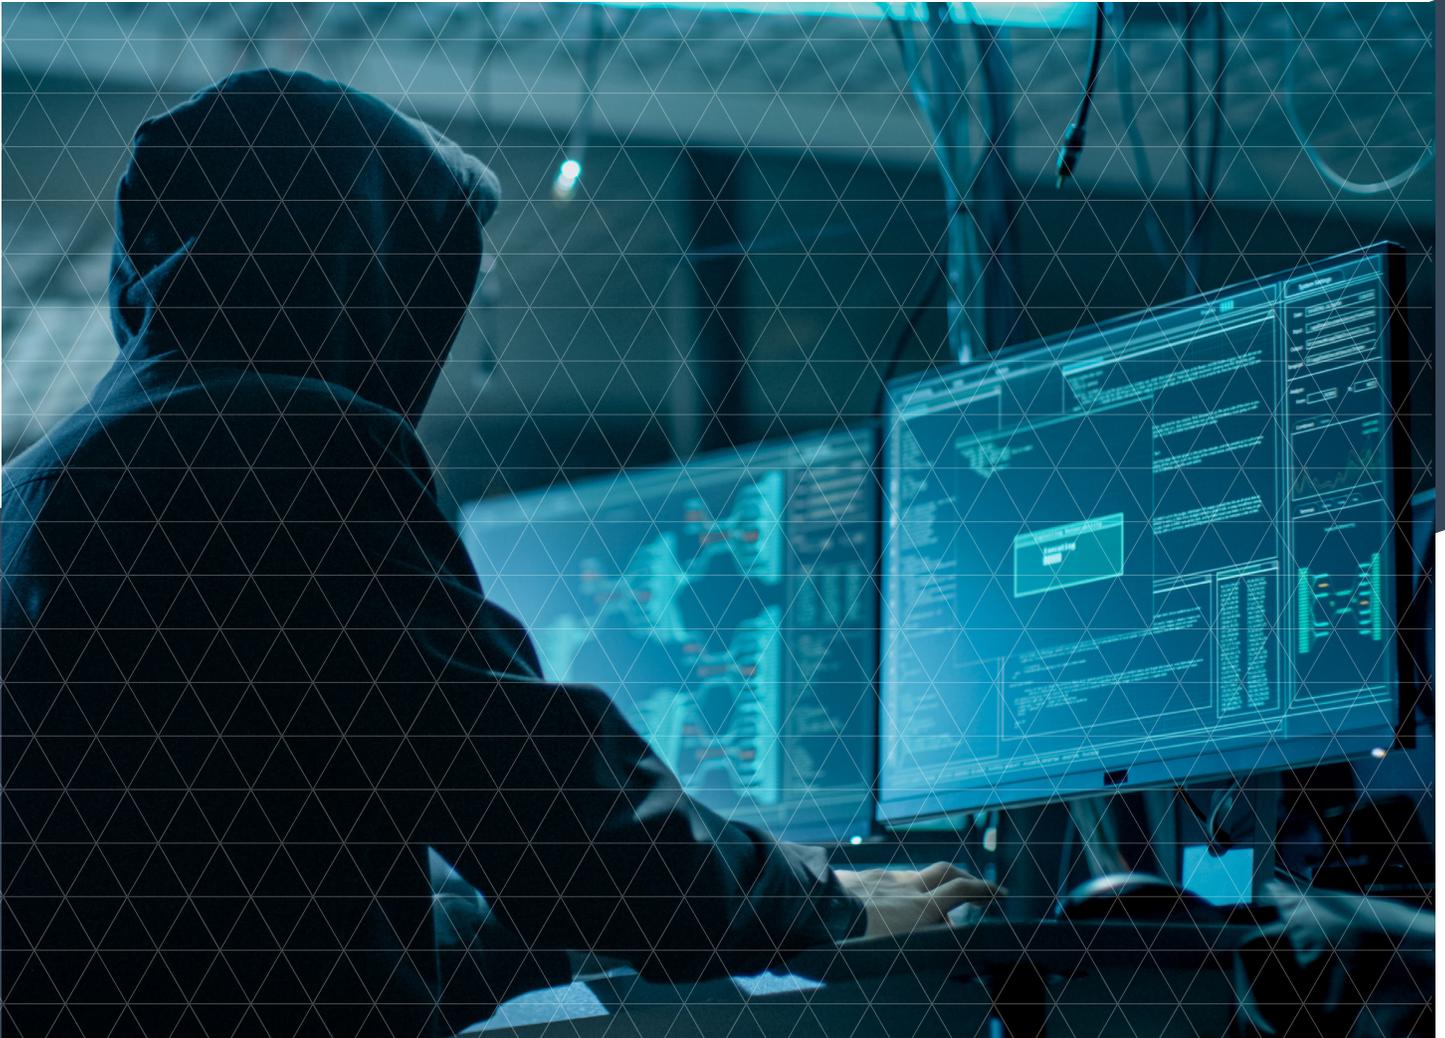


WITHSTANDING CYBER-ATTACKS: Cyber-resilience practices in the financial sector



Authors

Benoît Dupont
Université de Montréal

Clifford Shearing
University of Cape Town & Griffith University

Marilyne Bernier
Université de Montréal

ABOUT THIS REPORT:

According to PricewaterhouseCoopers (2016), 61% of Canadian CEOs believe that cyber security is the biggest potential business threat to their organization's growth prospects. A recent study by Scalar (2017) shows the negative impact of cyber attacks on productivity: in 2016, 53% of Canadian companies reported an incident that resulted in the loss of sensitive information, with an average of 44 events per year and a marked increase in the sophistication and severity of attacks. On average, organizations spent approximately \$7.2 million to re-mediate cyber security compromises, which includes clean up or remediation costs, lost user productivity, disruption to normal operations, damage or theft of IT assets and infrastructure, and damage to reputation and marketplace image.

Although many security vendors and consultants claim that their technologies and methodologies offer high levels of protection against cyber-risks, they tend to underestimate their manufactured nature, meaning that attackers constantly innovate to identify new exploitable vulnerabilities and that the asymmetry of the cyber security adversarial landscape plays in their favour. Consider the case of Kaspersky, a prominent cyber security company, whose systems appear to have been compromised by Russian intelligence operatives and used to steal confidential NSA documents that were illegally stored on the home computer of one of the agency's contractors running Kaspersky's antivirus application. The case came to light when Israeli intelligence operatives hacked into Kaspersky's systems and watched their Russian counterparts launch their attack in real time (Perloth and Shane 2017). The leaks of other secret tools belonging to the CIA and the NSA by Wikileaks and a group called the Shadowbrokers illustrate how even the most security conscious and best resourced organizations will at one point or another be compromised by determined and persistent adversaries.

The empirical approach proposed in this research project will therefore rest on the study of how the concept cyber-resilience is understood and applied by risk managers in financial institutions in four major industry hubs, and what lessons can be learned from the experience of those who have been faced with shocks, no matter the final outcome.

The main objective of this research is therefore to contribute to our understanding of cyber resilience in the financial sector by studying the processes, decisions and inter-dependencies that foster a state of resilience. Through case studies that approach cyber-resilience as a dynamic and fluid process instead of as a final state of equilibrium, the aim is to identify key principles, norms, cultural features, technologies and practices that have demonstrated their effectiveness in making financial institutions better prepared to manage and adapt to their growing and complexifying cyber risk portfolio.

Withstanding cyber-attacks: Cyber-resilience practices in the financial sector

Benoît Dupont¹, Clifford Shearing^{2, 3, 1} and Marilyne Bernier¹

Plans are worthless, but planning is everything. There is a very great distinction because when you are planning for an emergency you must start with this one thing: the very definition of “emergency” is that it is unexpected, therefore it is not going to happen the way you are planning.

Dwight D. Eisenhower (1958: 818)

Introduction

Over the past 25 years cyber-risks have morphed from mere annoyances into potentially catastrophic events that threaten the survival of technology-dependent organizations. The General Board of the European Systemic Risk Board, which oversees the EU financial system, has identified cyber risk as “a source of systemic risk to the financial system” (ESRB 2020), based on fear that a cyber-security incident could escalate, creating a liquidity crisis that would erode the confidence of financial actors and destabilize the whole system. Such catastrophic scenarios may sound extreme but a 2018 report from the International Monetary Fund, using data provided by the Operational Riskdata Exchange Association, estimates that aggregate losses generated by cyberattacks on 7,947 banks worldwide amounted to \$97 billion yearly (9% of net income), with value-at-risk (VaR) oscillating between \$147 and \$201 billion (14% to 19% of net income). The most pessimistic models show annual losses of up to 51% under the most adverse circumstances (Bouveret, 2018: 20-21).

To understand how cyber-risks can generate such disruptive shocks to the critical infrastructures on which we depend, consider the well-publicized case of the Danish shipping giant APM-Maersk. What is of interest here is not identification of the attackers or what this attack says about the current state of cyber conflict (Arquilla and Ronfeldt, 2007; Rid, 2012; Rid and Buchanan, 2015) but how disruptive such incidents have become for even the largest and most mature organizations, and how victims are upgrading their conventional cybersecurity risk management model to embrace a cyber-resilience paradigm.

APM-Maersk, which operates more than 800 vessels and manages 76 port terminals around the world, is responsible for 20% of global container traffic, making it a key hub in the world trading system. On June 27, 2017, the company’s computer systems were infected by a malicious piece of software, later titled NotPetya by the cybersecurity industry. NotPetya propagated at unprecedented speed and was able to bypass the security features of even recently updated

¹ Université de Montréal.

² University of Cape Town.

³ Griffith University.

systems and encrypt their contents, making it one of the most aggressive digital threats ever observed (Greenberg, 2019: 183). Attributed to Russian state-sponsored hackers, NotPetya was initially used against critical infrastructures in Ukraine before it spread to a much broader set of corporate victims, including some in Russia (NCSC, 2018). Although Maersk's infection and contagion took less than seven minutes, it took almost two hours to disconnect the company's entire global network in a vain attempt to ward off the attack (Greenberg, 2019: 152).

Maersk's management provided a candid and useful account of how it dealt with the attack's aftermath. Adam Banks, the company's Chief Technology and Information Officer, painted a bleak picture of the damage: 49,000 computers and printing devices as well as 56% of the 6,200 servers and 83% of the 1,200 applications needed to run the business were instantly inoperable. The fixed phone lines needed to coordinate the response were unavailable because they relied on the computer network, and mobile phones were crippled because all contacts that used Microsoft Outlook to synchronize had been wiped (Ritchie 2019). Backup copies of the encrypted data could not be restored as all the domain controller servers needed to rebuild a map of the network that had been compromised (Greenberg, 2019: 194). One of Maersk's first decisions was to abandon its existing crisis management protocol, which had never included plans for such a comprehensive level of destruction, and instead rebuild its network from scratch, with the help of a consultancy firm that was given a blank cheque. The company got a lucky break when it discovered that a power outage in Nigeria had spared a single server that could be used to rebuild the entire network. Beginning with that one machine, within 10 days Maersk had re-installed 2,000 laptops and within a month had issued almost 50,000 clean machines to employees. Meanwhile, employees, reverting to pen and paper and taking orders via personal Gmail and WhatsApp accounts, managed to maintain 80% of container traffic (Crozier, 2018). Despite the high costs of the shock (between USD\$200 and 250 million), the company survived and its reputation was even enhanced by its willingness to share the lessons learned along the way. Its CEO stated that his new ambition was to use the incident "to get to a point where [Maersk's] ability to manage cyber security becomes a competitive advantage" (Crozier, 2018).

Maersk's experience epitomizes cyber-resilience in action: despite a level of preparation that was believed to be sufficient, an unexpected risk led to a situation that had never been encountered—or even imagined—and the company found its existing cybersecurity measures completely overwhelmed. Luck played a part in meeting this challenge but would have been insufficient if employees from all parts of the organization had not rallied and improvised solutions that kept operations running. External resources were also essential, as Maersk would probably have been unable to rebuild its network in such a short time without the expertise of the consulting firm it hired. The trust and goodwill Maersk had accumulated among its corporate peers were also important, as it was able to use their IT networks for new installs when the local supply of USB sticks ran out (Ritchie, 2019). Finally, this extreme shock triggered such a successful adaptation process that the organization was able to promote its newfound cybersecurity expertise at the World Economic Forum a year later (Olenick, 2018).

Despite examples like this, a long history of use in the fields of materials science, ecology, and psychology, and its recent extensive use as a strategy to address Anthropocene risks (Boin and van

Eeten, 2013), the concept of resilience remains peripheral in the literature on cyber-risks. When it does appear, it originates primarily in the field of computer science, where the main research questions are concerned with identifying the engineering features that can make cyber systems more robust and the metrics that can be used to evaluate their capacity to endure (Bodeau and Graubart, 2011; Linkov et al., 2013a). While these are important, a more holistic approach is needed to help understand what types of preparations, responses, recovery, and adaptation activities are needed to enhance an organization's cyber-resilience (The National Academies, 2012).

At its simplest, cyber-resilience can be defined as “the ability to continuously deliver the intended outcome despite adverse cyber events” (Björk et al., 2015: 312). This definition differentiates cyber-resilience from cybersecurity, whose main aim is to predict and prevent detrimental events – the case study described above could be depicted as a cybersecurity failure that ended up as a cyber-resilience success. Our work focuses on providing a broader organizational understanding of cyber-resilience. Qualitative data collected from a sample of cybersecurity professionals in the financial sector—a particularly exposed field – were used to uncover the meanings these professionals assign to the notion of cyber-resilience and the security measures that result. Resisting the temptation to delineate an elegant cyber-resilience framework that could be implemented mechanically by any organization—an attractive but, in our opinion, futile endeavour—our objectives are more modest. We have tried to benefit from the experience of those who craft and implement cyber-resilience practices in a very demanding business environment to learn what works, what does not, and the main technical and organizational constraints that must be overcome in the process.

We start with a quick overview of the existing literature on resilient organizations and how it applies to the handling of cyber-risks. We then describe our qualitative methodology and the sample of financial sector cybersecurity professionals who participated in this research. The next sections detail our examination of the three most salient dimensions associated with cyber-resilience practices in financial institutions: the sense attributed to cyber-risks, including their inherent uncertainty, the effectiveness of the organizational strategies used to prepare for and mitigate cyber-attacks, and the adaptive outcomes that result from these incidents.

Cyber-resilience, beyond engineering

The concept of resilience has been embraced by disaster management researchers in their attempts to understand how large-scale adverse events are handled by organizations (Manyena, 2006; Paton and Johnston, 2017). Two of the main insights generated by the resulting literature are related to different meanings of resilience and interactions between various levels of resilience. Resilience can be interpreted very differently by organizations with different levels of maturity in this area. For those at a basic level, it involves an ability to maintain the status quo and absorb the impact of disturbances, while more advanced organizations embrace an adaptive understanding of resilience that relies on self-organization and adoption of new practices that do not compromise structure or functions. A minority of organizations are able to leverage the transformative power of resilience to seize the new opportunities created by a changing

environment (Davidson et al., 2016). Studying resilience in complex systems requires mapping the cross-scale interactions produced by geographical, temporal, functional, and technological factors (Ansell et al., 2010).

To understand cyber systems like those financial institutions use to conduct their business, this multi-level approach examines how cyber-resilience requires building hardened and redundant infrastructures (hardware and software), accessing relevant information on which to base decisions, mobilizing the cognitive processes of various human actors, and making use of social capital inside and outside the organization (Linkov et al., 2013b). Given the technical and business interdependencies that link financial institutions, the cyber-resilience of a particular bank is tightly coupled to the cyber-resilience of the national and supranational financial systems in which it operates (Björk et al, 2015), connections that are the basis of financial regulators' concern about the ability of regulated entities to address cyber-risks. Cyber-resilience is both a state and the processes that lead to that state (Kaplan, 1999), so to understand it we need to consider the three types of work performed by institutions: operational work (implementation), regulatory work (oversight), and constitutional work (governance) (Gim et al., 2019).

The abundant literature on the operational work of cyber-resilience has identified five high-level dimensions that contribute to—or hinder—it (Dupont, 2019). First, cyber-resilience is dynamic: it can only be achieved through a permanent cyclical process of preparations, responses, recoveries, and adaptations that require a myriad of decisions. These decisions must be made according to operational tempos that range from minutes to months, as seen in the Maersk case study, and result in bifurcated outcomes of deliverance or destruction (Grossetti, 2009). Second, cyber-resilience is networked: modern organizations and their socio-technical systems are embedded in complex webs of interdependence that simultaneously enable and constrain them. This networked aspect of cyber-resilience and the multiplicity of operational ties it involves is particularly important for organizations, such as financial institutions, that may face transboundary crises that affect multiple jurisdictions and undermine various types of critical infrastructures (Ansell et al., 2010; Bossong and Wagner, 2017). Third, cyber-resilience involves practice: the capacity to manage dynamic events and mobilize the expertise and resources needed to maintain operations when faced with unexpected disruptions does not appear *ex nihilo*. Cyber-resilience cannot be improvised but improvisation skills can be developed through rehearsing crisis scenarios, which also helps the organization overcome startle and surprise effects (Bastien and Hostager, 1988; Staal, 2004). Fourth, cyber-resilience is adaptive both during and after a crisis. During a crisis, the flexibility required to tackle turbulence can be achieved through redundancy and diversity (Ansell et al., 2010; Bodeau and Graubart, 2011). After a crisis, cyber-resilient organizations are able to absorb the lessons learned to improve or transform their systems and procedures, making disasters a “catalyst for development” (Paton and Johnston, 2006: 8). Fifth, cyber-resilience is contested: the considerable investments it requires and the limited returns it provides in the short term collide with the usual business sensemaking frame, which favors efficiency and optimization in the pursuit of larger market shares and higher profits. The concept of sensemaking is useful to capture the collective cognitive and story-telling processes that flow through organizations and allow individuals and the groups to which they belong to structure the unknown and assign meaning to unexpected facts or events (Weick, 1995). The collective nature

of sensemaking means that a shared meaning must emerge from a diverse range of perspectives, including contested rationalities that must align to preserve the organization.

Despite a growing body of literature on the various conceptual aspects of cyber-resilience, there is still a very limited pool of empirical studies examining how organizations make sense of and implement cyber-resilience in practice (Maennel et al., 2017; Catota et al., 2018; Carias et al., 2019; Fujs et al., 2019). Such studies would help us “develop a stronger grasp on the relation between organizational characteristics, processes and outcomes” (Boin and van Eeten, 2013: 443). This article contributes to the empirical literature on cyber-resilience by examining the practices of cybersecurity professionals in financial institutions across five geographic areas.

Data and methods

Our study blends three qualitative methodologies to capture the experience of cybersecurity professionals who deal routinely with cyber-attacks in the financial sector. Forty-four respondents from 28 organizations were interviewed. A purposeful sampling approach (Patton, 2015) was adopted to achieve a diversity of views and experiences across five specific dimensions (geography, institutional type, institutional size, interviewee role, and interviewee experience) and to ensure that the networked aspect of cyber-resilience was adequately represented. The geographical diversity of the sample recognizes both the global nature of the cyber-resilience challenges faced by financial institutions and the local cultural or regulatory features that may foster different national practices. Respondents were interviewed in Canada, the US, the UK, the Netherlands, and France. Some organizations in each country had a very broad international exposure, operating in dozens of markets, while others maintained essentially a local footprint. The size of institutions for which the respondents worked also varied greatly – some of them have less than a billion USD\$ in annual revenue, while others’ profits can reach five to ten times that amount – leading to varying levels of resources and expertise available to implement cyber-resilience practices. The financial sector provides diverse services to retail and commercial customers and to take this into account the sample included cybersecurity professionals who work for banks, insurance companies, pension funds, and stock exchanges. The consulting and incident-response firms that provide cybersecurity services to financial firms and the regulators who oversee their activities were also interviewed. The positions of respondents ranged from Chief Information Security Officers (CISOs) and Chief Risk Officers (CROs) to Directors of Security Operations Centres (SOCs), Incident Response Teams (CSIRTs), and business continuity units; leaders of penetration-testing teams and red teams; and IT governance and security advisors. Experience in a cyber-risk management or regulation role ranged from less than a year to almost thirty years in the field.

Interviews were conducted between August 2018 and February 2020, and were done in person (25), by phone (18), or by email (1). Eleven respondents (25%) were female, a higher representation than the 11% that is the average for women in the cybersecurity workforce (Frost and Sullivan 2017). Interviews lasted for 57 minutes on average (range: 31 minutes to 1 ½ hours) and were recorded and transcribed for qualitative analysis, with the exception of three interviews that took place in public settings (café or restaurant) where the noise level was too high for

recording and handwritten notes were taken instead. The transcribed interviews were then imported into QSR International's NVivo 12, a qualitative analysis software package that facilitates the exploration, coding, and visualization of large quantities of unstructured data. Coding relied on theoretical categories identified in the previous section to see how they are translated—if at all—into everyday organizational practices. The coding process was also designed to pay particular attention to the tensions and challenges that might be associated with cyber-resilience practices and to the strategies that respondents deployed to negotiate these hurdles.

To facilitate the coding process, all interviews followed a similar script: respondents were first asked to explain how they defined cyber-resilience and then asked to recall the most severe cyber-attack they had experienced. These questions were asked at the beginning of the interview to elicit specific and concrete recollections of disruptive adverse events unique to the participant's organization and of how these events were managed. An indirect objective was to minimize participant reliance on generic statements or on highly publicized cases such as the one in this article's introduction, responses that are often used to deflect questions about a sensitive topic or one for which the organization has no response. The interview script then proceeded with questions about the technologies and procedures (including standards) used to foster cyber-resilience, the role played by public-private partnerships and external expertise, the organizational barriers to cyber-resilience, the impact of the human factor on cyber-resilience, and the regulatory aspects of cyber-resilience. A final open-ended question allowed respondents to identify any issues they thought had been overlooked.

While this study used interviews as its primary research material, we were also able to take notes and ask candid questions at a meeting in the summer of 2019 where the outcomes of a large international cyber-resilience exercise were reported to a dozen representatives from large multinational financial institutions and national regulators. Finally, three organizations shared with us, under strict confidentiality agreements, two complete cyber incident response plans, a series of ten post-incident reports, and a benchmarking report comparing the cyber crisis management models of ten multinational businesses from a variety of sectors (including finance and insurance). These documents provided useful contextual information for the study.

Sensemaking, cyber-risks, and the instability of decision-making processes

The first challenges cybersecurity professionals encounter in their attempts to design and implement cyber-resilience practices are the complexity of cyber-risks and the difficulty of making sense of them, of understanding what is happening. With well-known risks, established framings that make sense of events and identify response pathways can be quickly and easily deployed. With cyber-risks, where 'newness' abounds, framings need to be developed immediately in environments of uncertainty, crisis, and urgency. A crucial and necessary task for cybersecurity professionals is making sense of what has happened – Weick's "sensemaking" (Weick, 1995). The need for sensemaking imposes significant constraints on their ability to make appropriate decisions. In this section we explore the various barriers that impede the application of sensemaking processes to cyber-risks and constrain cyber-resilience practices. These constraints are discussed in terms of experiences both before and during the occurrence of severe adverse

events to reflect the distinction between decisions that prepare the organization to be cyber-resilient and those that enact this cyber-resilience.

The first barrier is related to the features that differentiate cyber-risks from more conventional forms of risk. Sensemaking processes are made more difficult by the dynamic nature of cyber-risks that are ‘manufactured’ by adversaries and for which there is “very little previous experience” (Giddens, 1999: 4). Adversaries innovate constantly, developing attack strategies and tools that have never been encountered before and for which there are no known defenses (Bilge and Dumitras, 2012; Ablon and Bogart, 2017). These so-called zero-day attacks introduce high levels of uncertainty that information-sharing arrangements between financial institutions, a form of distributed sensemaking, cannot alleviate.

“... that approach only works against things that have already happened to others; the new things that are coming along, the zero-day threats, the brand new virus that no one has seen yet, those are the things you have to watch for, that information sharing will never address because you have nothing to share because it hasn’t happened yet, and every day there are new things being invented.” (Canada 22)

The same respondent added that these sudden and destabilizing shifts emerge from an ocean of noisy data. His organization, for example, had had to deal with a trillion security events over the previous year, and the only way to handle such large amounts of alerts was to delegate sensemaking processes to artificial intelligence (Canada 22).

The dynamic nature of cyber-risks can destabilize sensemaking processes at different stages of an adverse event. Respondents recalled many cases where what had initially been identified as a fairly minor incident quickly escalated into a much more complex crisis that unfolded over many months. In one example, infection of an employee’s laptop by a malicious software, which would usually have been dealt with remotely in a few hours, led to activation of a crisis team when forensic analysis indicated that troves of email had been compromised. As well, this particular employee was the point of contact with multiple industry regulators and also organized the travel of the organization’s high-level management so had access to personal information such as passports, credit card numbers, etc. During a crisis, discoveries such as this can, and do, provoke sudden changes in the sensemaking process, which in turn can increase the probability of errors. Mindful of this pattern, one organization in our study had introduced an informal deferred decision-making approach to enable more thorough sensemaking assessments of a situation and avoid implementation of hasty measures that could prove counterproductive. Even when an incident has been resolved technically, its negative impact (such as the malicious use of stolen credentials or personal information) can linger for many months and require further sensemaking in a difficult and hostile environment.

Cyber-risks are often difficult to contain, generating risk-cascades (van Eeten et al., 2011) that increase the dynamic properties of cyber-risks and amplify a crisis. The move to cloud infrastructures provided by third parties exemplifies this challenge. The concentration of the cloud industry around three major providers (Amazon, Google, and Microsoft), which are not regulated

by the same organizations as their financial customers, introduces new forms of uncertainty in case of failure. US insurers (AIR, 2018) and legislators (Schroeder, 2019) have expressed concern and almost a quarter of participants mentioned that this shift to the cloud complicated their risk-management practices and even “made them blind” (Canada 16).

Another significant source of interference with the sensemaking process is the obfuscation of cyber-risks. The secrecy that often envelops the management of some incidents, the existence of ‘Shadow IT’ systems that are sometimes hidden from cybersecurity professionals (Hagenaars, 2019), and the loss of the expertise required to properly secure multiple stacks of ageing legacy systems all contribute to this obfuscation.

“If you have this big sprawling mixture of technology and legacy architecture and infrastructures that you’ve acquired over twenty-five to fifty years, depending on how long you’ve been in business, it can be really hard to wrap that in something that looks resilient, because it’s a leaky boat.” (Canada 18)

These distinctive features of cyber-risk can degrade the quality of sensemaking by making the severity of incidents harder to assess, their ramifications for the organization and its external partners harder to understand, and the level of response required harder to calibrate. These sensemaking blind spots are directly reflected in the quality of the response plans and ‘playbooks’, which function as broad sensemaking tools, that the financial industry has developed to manage adverse events.

Not all sensemaking challenges can be attributed to the external pressures of a fast-changing risk landscape. The second source of tension originates from the contested rationalities (or sensemaking frames) of business operational requirements and cyber-resilience. With digital technologies transforming financial institutions, the importance of using these new tools to optimize resources and maximize profits collides with a more cautious cyber-resilience approach in which innovation is slowed until proven safe. It also requires acknowledging that significant investments in redundancy, diversity, and training are necessary, even if they may not show immediate benefits. The decision to deploy redundant technologies often involves a contest of rationalities:

“As a general rule, ‘simple’ is easy to interact with, but ‘simple’ is also potentially not as resilient as ‘diverse’ and ‘complex’, but ‘diverse’ and ‘complex’ are more difficult to interact with, and so the questions become what are your business goals, what are the risks you face, and whether or not those pros and cons make sense in your business.” (Canada 23)

To resolve this tension, cybersecurity professionals implementing cyber-resilience practices inside their organization place a strong emphasis on communication. They are mindful of their users’ business needs, incorporate them into their risk management mandate, and are careful about communicating this mandate. Sometimes they even borrow sensemaking patterns from their business users to engage them in their cyber-resilience efforts.

“When you’re a bank, you’re making credit decisions all the time and there is a well-established model for measuring risk, how much risk are we accepting from a risk appetite. We’re trying to bring those practices that have evolved in banks from a credit risk perspective to cyber-risk and operational risk and so that’s where we’re going in terms of trying to calculate our risk on what we’re doing with our systems.” (Canada 22)

The pre-eminence of a business rationality temporarily cedes ground to a cyber-resilience rationality when a major crisis erupts. As many participants noted, nothing focuses the mind of CEOs and board members and increases their interest in cyber-resilience like a highly publicized data breach or cyberattack. They recalled how an occurrence of these disruptive events in their organization or in other financial institutions sparked a review of existing arrangements and unlocked significant investments that they had been unable to secure previously.

Finally, the third source of sensemaking tension originates with regulators, whose oversight activities generate geographic and temporal turbulences. Many respondents worked in financial institutions with branches in a large number of markets (sometimes more than fifty) that operate under a broad range of regulatory regimes. Some countries have adopted a principles-based approach to the regulation of cyber-risks, while others, such as the UK or the Netherlands, have been more prescriptive and have developed proactive testing strategies (CBEST in the UK and TIBER in the Netherlands) in which external ‘red teams’ mimic the types of attacks carried out by sophisticated actors (Hielkema and Kleijmeer, 2019). Financial institutions must consolidate these variations into their sensemaking processes to ensure they are compliant across the whole regulatory spectrum, introducing an additional level of complexity. The time available for sensemaking can also be decreased by some regulators’ requirement that the nature and scale of cyberattacks or data breaches be rapidly disclosed to the public even though the dynamic nature of cyber-risks and the technical complexity of digital infrastructures mean that assessment of an incident’s full impact may go through multiple iterations that alter how the crisis is understood. By forcing financial institutions to make their sensemaking processes transparent in a shorter timeframe, this regulatory strategy can lead to unexpected and detrimental outcomes.

“You know how in a lot of incidents that have gone public in the last number of years, you’ll get someone from the communications department speaking, saying within two or three days of an incident being announced that they’ve got it contained. Well the truth is, ninety percent of the time they have to come back in a few days or a week later and say “look, you know how we thought we had forty-thousand customer data records breached, oh shit, it’s four-hundred-thousand”. ... Because the fog of war means that half the time, you’re wrong, but don’t go out and say to your regulator or the public or your constituency, that you’ve got it fixed right? If you do that more than a couple of times, your trust and brand get destroyed.” (Canada 20)

The three sensemaking tensions outlined in this section reverberate across the plethora of decision-making processes activated by financial institutions’ exposure to cyber-risks. They significantly complicate the job of cybersecurity professionals, who are selected for their technical expertise or business acumen but may not be comfortable dealing with unpredictability,

uncertainty, ambiguity, and controversy. One respondent summed this up bluntly when he stated that these tensions provide fertile ground for “narrative fallacies that justify things that are not necessary” and allow “charlatans [to] proliferate to profit” (United Kingdom 2). In the next section, we explore the strategies and practices that financial institutions rely on to withstand cyber-shocks.

Organizational practices of cyber-resilience

The epigraph to this article highlights the value of planning over the plans themselves. That general approach was popular among participants, who often used the “muscle memory” analogy to convey the principles that guide their cyber-resilience practices. Mindful of the intrinsically unpredictable nature of cyber crises, they emphasized the development of general resources and practices that can be quickly adjusted to deal with unexpected adverse events.

This adjustment process generally starts with a comprehensive mapping of the critical functions a financial institution must recover in case of an extreme adverse event. Such mapping exercises are not new but in the past were more likely to focus on individual risks, making it more difficult to uncover interdependencies between critical services. This focus is changing in an environment where the complete loss of IT resources is a possibility, and where different teams must be ready to coordinate their efforts quickly to restore access to markets and resume services to customers. Mapping is not limited to internal processes but must also extend to third parties, which complicates matters when the latter are reluctant to share sensitive information (United States 1). The outcomes of these mappings are then combined with intelligence about the threat landscape to design scenarios of possible adverse events and create response playbooks.

The financial institution for which one of our participants worked maintains sixteen playbooks that are reviewed every quarter to see if new scenarios based on emerging modes of attacks are needed (Canada 22). Not all participants were so well prepared, however, and a few had just completed their first cyber-specific playbook or were in the process of developing it. Playbooks take time to develop and one participant explained that the creation of a playbook had involved several rounds of consultation and testing over almost a year to ensure that it captured the different perspectives, capacities, and methodologies of all the teams it was supposed to coordinate (Canada 4). Several respondents warned against an over-reliance on playbooks, which cannot possibly anticipate all the surprises encountered in real-life incidents. They highlighted that a cyber-resilient organization needs to be able to deviate from a playbook—sometimes radically—to adapt its response to unexpected conditions (Canada 1, Canada 32).

Two technical features usually associated with cyber-resilient systems are redundancy and diversity (Bodeau and Graubart, 2011; Zhang et al., 2016). Redundancy refers to the availability of multiple instances of a particular resource, while diversity references the existence of heterogenous resources that can be deployed to minimize exposure to a single type of risk. Both involve “the ability to quickly substitute technologies” (United States 2) and are often identified by practitioners as a “suspenders and belt approach”:

“if you don’t want your pants to fall down, you invest in two things to help make sure that doesn’t happen, in case one of them fails.” (Canada 22)

Both these approaches have proved effective in enhancing cyber-resilience. For example, participants noted that diversity had become an important feature in teams that manage cyber-crises (Canada 27, United Kingdom 2). Some organizations had built or were building multidisciplinary teams that drew on a wide array of backgrounds, perspectives, and expertise to ensure that their decisions did not overlook weak signals or discard unorthodox approaches because of groupthink (Janis, 1972). However, as several participants noted, these strategies can be difficult to manage because they introduce complexity and additional costs (Canada 21).

The importance of the human factor as a source of cyber-resilience was emphasized by the most experienced respondents, who often reminded us that people trump systems and procedures in dealing with an extreme cyberattack.

“People will save businesses in time of crises. If you train people, if you retain them, if you treat them well, you accumulate knowledge. And that knowledge in time of crisis will be crucial. We did have several quite severe incidents. And again, it was people who were at the front end, at the edge, saving the business. Not technology. Technology was useless.” (United Kingdom 2)

When asked to outline what personal traits were particularly useful in cyber-resilience, participants mentioned that the best performers in their incident response teams displayed higher-than-average curiosity, creativity, and flexibility. This gave them the ability to identify hidden patterns in large amounts of information, deviate from established procedures (or playbooks) when novel situations emerged, and quickly improvise previously unconsidered solutions. Without being reckless, they were comfortable with imperfect decision-making environments and were not prone to the “startle effect” that can lead to delay, panic, and even paralysis (Staal, 2004). They were good communicators who knew how to translate technical approaches so they could be understood by all in the organization and were able to explain the reasons behind inconvenient or drastic measures. They were also good listeners who could integrate multiple—and sometimes contradictory—perspectives into their own decisions. These results corroborate the findings of Chen et al. (2014), who conducted individual and team task analyses with three computer security incident response teams.

To avoid boredom and attrition for those on cyber risk teams, some participants have implemented cross-training programs that expose employees to the work of colleagues in different domains, broadening their capacity to identify and address emerging problems. Others use rotation systems, with employees occupying different positions within their team. Effective cyber-resilience professionals have a lot in common with jazz musicians who learn to create musical pieces from minimal structures in turbulent task environments where they must balance their individual skills and group coordination (Bastien and Hostager, 1988).

It would be misleading to suggest that a handful of naturally gifted operators are responsible for the overall cyber-resilience of an organization. Just like good jazz musicians, who develop their mastery over years of relentless practice, the performance of those in cyber-resilience teams is improved by focused training and incident rehearsals. Almost all participating organizations conduct simulations and tabletop exercises that try to recreate the conditions of a cyberattack as realistically as possible so that employees in a broad range of functions can familiarize themselves with existing playbooks and practice response and recovery protocols. The most mature financial institutions conduct up to half a dozen simulations per quarter at various levels (head office, specific business lines, regional branches). But these simulation activities are resource-intensive to design and to run, which explains why their quality varies greatly. They sometimes lack the level of detail that would allow employees to practice certain tasks in stressful situations, or they rely on scenarios that are not sufficiently challenging (Canada 24). They can also be seen as a distraction by senior decision-makers, who send delegates rather than attending in person, therefore defeating their purpose (United States 1, Canada 32).

Empirical work in the crisis management domain has shown that exercises alone are not sufficient antidotes to organizational pathologies that can undermine their benefits. For instance, the US Government's disorganized response to Hurricane Katrina in 2005 failed to take into account lessons learned during a major exercise undertaken in 2004 (Boin and McConnell, 2007). Despite the general consensus among respondents that training and simulation exercises were necessary to enhance the cyber-resilience of their organizations, the approaches advocated remained very conventional. No respondent mentioned the more focused evidence-based strategies to improve adaptability, problem-solving, communication, trust-building, and knowledge-sharing used by response teams in emergency medicine, the military, or the nuclear industry (Steinke et al., 2015).

We mentioned the importance of good communication in preparing for, responding to, recovering from, and adapting to cyberattacks. In the cyber-resilience context, the professionals we interviewed relied on dense internal and external organizational networks to improve the speed and effectiveness of communication flows. Despite the natural tendency in many financial institutions to segment expertise and require secrecy when crises unfold, many respondents highlighted the benefits of having developed bridging capital and weak ties throughout the organization to facilitate dealing with adverse events (Granovetter, 1973). For some, this meant embedding security workers inside business units to better understand their culture and technological constraints but also attempting to "build fundamental security into the business processes" (Canada 18). Other participants are establishing 'fusion centres' of various security units (fraud, cyber, physical, business continuity) to consolidate sensemaking and decision-making capacities. Awareness campaigns and cybersecurity 'ambassador programs' can also contribute to creating internal networks that can be activated in times of crisis. In another industry, Netflix has gone even further and launched a Reservist Program in which auxiliary crisis managers are trained across the organization to distribute and scale incident response expertise (Joshi, 2020).

Finally, external networks play a crucial role in organizational cyber-resilience. Financial institutions are embedded in a dense web of business partnerships and their sensemaking and incident response processes rely on the ability to collect information from outside the organization

quickly and to access 'surge capacities' while limiting bureaucratic or contractual frictions. Third parties, especially those providing IT services, need particular attention. Prompted by regulatory requirements, financial institutions are dedicating resources to assess the cyber-resilience of third parties and monitor how this impacts their own posture. However, as some respondents noted, these processes can expand to unsustainable levels: third parties have their own third parties, not always recognized before an incident, and modelling these risk cascades across organizations can quickly become extremely complex.

The main function of external networks remains the sharing of intelligence, best practices, and best thinking. One participant used the medical analogy of inoculation to describe the utility of sharing information across financial institutions, while acknowledging that this approach offered protection only against known threats. But many extolled information sharing as one of the most effective strategies to stop the contagion effect that can destabilize the financial system once attackers have found an industry-wide vulnerability.

“Networks of people who talk about what they’re experiencing, I think is very valuable, and in fact it’s sometimes more valuable than the consultants who come in and tell you stuff because, and I say this having been, given my prior history, essentially a consultant for a long period of time, the people who are out at the sharp end, sharing stories, are typically very open in the right setting and you learn more from that than you would do through a six-week consulting engagement and you’ll learn it faster.” (Canada 25)

The external networks conducive to effective information sharing blend informal and formal structures that can extend from small peer groups to large industry consortiums. One respondent estimated that the not-for-profit information-sharing initiatives in which his bank participated gave him access to threat indicators three and a half weeks earlier than the notifications he received from commercial feeds (Canada 31). Respondents insisted that fully benefitting from these external resources required trust as well as building and maintaining personal relationships over time, so that people have accumulated enough social capital to “call and ask for favours when they need to” (Canada 19).

Learning to adapt

In the proliferation of cybersecurity industry reports on this now trendy topic, cyber-resilience is often reduced to activities and processes that are undertaken as the immediate response to an incident (Dupont, 2019). But the ultimate goal of resilience is not survival until the next crisis but adaptation to a dynamic environment to reach a new state of equilibrium. Respondents discussed three different forms of adaptation associated with major adverse events.

The first form of adaptation is voluntary and reflects the learning that takes place after a major unexpected incident or after a poorly handled routine incident. As mentioned, highly publicized incidents such as the wave of Distributed Denial of Service Attacks against American banks in 2012, the Equifax breach in 2017, or the Capital One hack in 2019 sent shockwaves through the financial industry, highlighting the fragility of existing processes and leading to significant changes (Canada

5, Canada 16, Canada 17, United Kingdom 1, United States 2). Many smaller incidents that are never brought to the attention of the press as well as simulations that enact future-oriented scenarios also reveal the inadequacy of existing security measures and response processes. The lessons learned during these events by those involved in their mitigation are usually captured in post-incident reviews.

The review documents we were granted access to summarized the causes of the incident, its impacts on the organization and its customers, how it was resolved, what lessons were learned, and what adaptations were required as a result. But it was difficult to assess how these insights had been incorporated into the organization's cyber-resilience practices. Echoing this impression, a respondent regretted that there was no technology available to tap into the accumulated organizational memory that these reports contained, including a track record of the good and bad decisions that had been made and their outcomes (Canada 3). To ensure that all the data needed to produce such reviews are collected, including the most sensitive and embarrassing, a few respondents emphasized the need to create a safe environment for the employees at the origin of an incident. This "no-fault learning" approach was reiterated very publicly in one of the incidents described above.

"[name withheld], who is the Senior VP, even recorded a video to say that it is ok to make mistakes. We can make mistakes. What's not right is to keep making the same mistakes over and over again without correcting yourself, without thinking: Yes, I made a mistake, but what can I do to avoid it? And also, to realize, if I made a mistake in one system, in one way, can that mistake be reproduced elsewhere? So learn from our mistakes." (Canada 2)

The second form of adaptation is guided by cybersecurity standards and their cyber-resilience components. Sometimes defined as a "recipe for reality", standards have become ubiquitous in a complex world where technical and organizational infrastructures have to be coordinated on a global scale. They facilitate interactions between businesses by making explicit "the rules that others follow, or the range of categories from which they may choose" (Busch, 2011: 28). Two general sets of standards – ISO (International Organization for Standardization) and NIST (National Institute of Standards and Technology) – dominate the cybersecurity field, complemented by more specialized standards focused on a particular sector or function (Dupont, 2019).

Standards gradually incorporate lessons learned from past incidents and then help propagate best practices, raising the bar for everyone. But some respondents expressed doubts about the false sense of resilience that standards might introduce. Because of their complexity (often involving hundreds of criteria or controls), it is almost impossible for an organization to be fully compliant (Canada 19). Standards are also, by definition, very rigid and therefore ill-suited to help deal with the unknown (Canada 3). Their static nature can become a source of vulnerability since dynamic attackers can use them to calibrate their efforts: "Well, if that's your standard, all I have to do [to infiltrate your systems] is aim here. Above the standard" (Canada 7).

The third form of adaptation stems from the regulatory activity to which financial institutions are subjected. Respondents identified "a trend towards more regulation and more specific regulation"

(Canada 23), with certain jurisdictions becoming much more directive about cyber-resilience. Although a majority of participants preferred principle-based regulatory requirements out of concern that an excessively detailed and prescriptive approach would erode their flexibility, others explained how detailed regulations that mandated specific measures could accelerate adaptation. Even when financial institutions understand the value of technologies or processes that can enhance cyber-resilience, the costs associated with their deployment and the fear of being the only one to adopt them and losing customers to competitors that support customer experience rather than resilience act as powerful deterrents. Prescriptive regulations that force the whole industry to adapt can overcome this competitive barrier and lead to support for investments that would have been much more difficult to justify otherwise. Obviously, the level of this more intrusive regulatory approach remains a sensitive issue among stakeholders that want to avoid overreach or stifle innovation. The importance of avoiding regulatory capture by lobbyists and vendors that try to embed their products into norms is also a concern (France 1), as is the tendency for certain regulators to provide vague guidance that leads to interpretative uncertainty (United States 1, Canada 28).

Conclusion

This article provides a detailed overview of the current cyber-resilience practices implemented by financial institutions in five geographic areas, as well as the challenges they face. Although we initially expected to be able to identify different ways in which cyber-resilience was approached and enabled across countries, the limited size of our sample did not make this comparison possible. What became clear, however, was that even within the same jurisdiction or market (Canada, for example), it was impossible to design a one-size-fits-all template that organizations could adopt and implement to enhance their cyber-resilience. Cyber-resilience appears to be highly contextual and depends on a variety of unique factors, such as the history, size, business culture, international footprint, IT priorities, regulatory environment, and leadership style of each organization. There is no ideal cyber-resilience posture, only customized and tailored practices that can deliver improved levels of reliability and survivability when an organization confronts severe turbulences. This means that it is hard to provide perceptive analyses of an organization's cyber-resilience or to offer guidance on ways to improve it without a detailed understanding of its operational and strategic drivers and constraints.

By sharing some of the lessons learned—both in terms of successes and challenges—by cybersecurity professionals working in one of the sectors most exposed to cyber-risks, we have sought to demystify the practices of cyber-resilience, which are all too often shrouded in a veil of trendy buzzwords and glib normative agendas. Our particular focus has been to describe in concrete terms how cyber-resilience is embedded in a complex web of interactions that link technical systems, organizational processes, and human behaviors and is constrained by tensions in framing processes that lead to the prioritization of particular choices by making some actions thinkable and others inconceivable (Smith, 1987; Simpson et al., 2019).

The web of interactions involved and the tensions inherent in it, along with the particular context of competing adversaries, help explain why cyber-resilience cannot be reduced to dealing with

business continuity and disaster management. Conventional response models are designed to handle predictable and stable risks such as natural disasters. While the frequency and impact of such disasters can be expected to increase over the next decades as we have entered the Anthropocene, they are not, unlike cyber-risks, the result of actions by innovative and thinking adversaries. Our research also illuminated tensions that result from the contesting rationalities of business performance and institutional security. This quintessential dilemma was discussed by Holling in his foundational work on ecological resilience, where he cautioned that conditions favourable to short-term economic productivity (such as reduced diversity and redundancy, which allow economies of scale and resource optimization) may be detrimental to resilience and increase vulnerability (Holling, 1996: 38).

One source of framing and sensemaking tension, or of institutional blindness, that was not discussed by any of the professionals we interviewed arises from a range of individual and collective cognitive biases that interfere with disaster preparedness and crisis management. It came as a surprise that none of these highly experienced and battle-tested experts had considered that well-documented heuristics might lull their organizations into a false state of security. Ten heuristics seem particularly relevant in this context: the myopia bias (the tendency to focus on present benefits rather than future harms), the amnesia bias (the tendency to quickly forget the lessons of past disasters), the optimism bias (the tendency to minimize the impact an adverse event can have on us even while acknowledging it will affect others), the inertia bias (the tendency to remain passive when confronted with high levels of uncertainty), the simplification bias (the tendency to consider only convenient factors when confronted with complex risks), the herding bias (the tendency to align with the actions of others rather than rely on a more specific analysis of the situation), the familiarity bias (the tendency to rely on past actions as guides for behaviour), the consistency bias (the tendency to maintain an approach once an initial decision is made), the expert halo bias (the tendency to assess leaders' skills based on an overall positive impression rather than specific information), and the social facilitation bias (the tendency to take more risks when other people are involved) (McCammon, 2004; Meyer and Kunreuther, 2017). These ten biases represent significant 'unknown unknowns' for cybersecurity professionals who should therefore consider deploying mitigating remedies, ensuring that new pathways for realising resilience are explored and practiced safely.

This research project has attempted to draw on specific examples of highly disruptive cyber-shocks experienced by financial institutions to elicit practical insights on the benefits and limitations of the most common cyber-resilience measures. The wealth of data generated by this approach will contribute toward a more empirical understanding of cyber-resilience as a practice, a process, and an outcome. But our qualitative methodology can capture only a limited sample of the most memorable incidents experienced by each of the cybersecurity professionals who agreed to be interviewed. A more systematic approach based on information about tens—and possibly hundreds—of cyber-incidents, whose causes, responses and outcomes would be recorded using a set of predetermined criteria, would enable us to make stronger inferences about the efficiency and effectiveness of specific cyber-resilience measures. There is a long history of using methodologies such as the Critical Incident Technique to identify the most useful behaviours for solving practical problems (Flanagan, 1954). Translating these methodologies into the cyber-

resilience research domain would enable us to build a more robust knowledge base and avoid falling into the availability bias trap, where decisions are based on what can be recalled easily rather than through more thorough sensemaking processes.

References

- Ablon, L., & Bogart A. (2017). *Zero days, thousands of nights: The life and times of zero-day vulnerabilities and their exploits*. RAND Corporation. doi:10.7249/RR1751
- AIR. (2018). *Cloud down: Impacts on the US economy*. Lloyd's.
- Ansell, C., Boin, A., & Keller, A. (2010). Managing transboundary crises: Identifying the building blocks of an effective response system. *Journal of Contingencies and Crisis Management*, 18(4), 195-207. doi:10.1111/j.1468-5973.2010.00620.x
- Arquilla, J., & Ronfeldt, D. (2007). Cyberwar is coming! *Comparative Strategy*, 12(2), 141-165. doi:10.1080/01495939308402915
- Bastien, D. T., & Hostager, T. J. (1988). Jazz as a process of organizational innovation. *Communication Research*, 15(5), 582-602. doi: 10.1177/009365088015005005
- Bilge, L., & Dumitras, T. (2012). Before we knew it: An empirical study of zero-day attacks in the real world. In *CCS '12: Proceedings of the 2012 ACM conference on Computer and communications security*, 833-844. doi:10.1145/2382196.2382284
- Björk, Fredrik, Henkel, M, Stirna, Janis, and Zdravkovic, J. (2015), "Cyber resilience – Fundamentals for a definition", in Rocha, Alvaro, Correia, Anna Maria, Costanzo, Sandra, and Reis, Luis Paulo (eds.), *New contributions in information systems and technologies*, Springer, London, pp. 311-316. doi: 10.1007/978-3-319-16486-1_31
- Bodeau, D., & Graubart, R. (2011). *Cyber resiliency engineering framework*. The MITRE Corporation.
- Boin, A., & McConnell, A. (2007). Preparing for critical infrastructure breakdowns: The limits of crisis management and the need for resilience. *Journal of Contingencies and Crisis Management*, 15(1), 50-59.
- Boin, A., & van Eeten, M. (2013). The resilient organization: A critical appraisal. *Public Management Review*, 15(3), 429-445.
- Bossong, R., & Wagner, B. (2017). A typology of cybersecurity and public-private partnerships in the context of the EU. *Crime, Law and Social Change*, 67(3), 265-288.
- Bouveret, Antoine (2018), *Cyber risk for the financial sector: A framework for quantitative assessment*. *IMF Working Paper*, WP/18/143: 1-28.
- Busch, L. (2011). *Standards: Recipes for reality*. The MIT Press.

Carias, J. F., Labaka, L., Sarriegi, J. M., & Hernantes, J. (2019), Defining a cyber resilience investment strategy in an industrial internet of things context. *Sensors* 19(1), 1-16. doi: 10.3390/s19010138

Catota, F. E., Morgan, M. G., & Sicker, D. C. (2018), Cybersecurity incident response capabilities in the Ecuadorian financial sector. *Journal of Cybersecurity* 4(1), 1-20. doi: 10.1093/cybsec/tyy002

Chen, T. R., Shore, D. B., Zaccaro, S. J., Dalal, R. S., Tetrick, L. E., & Gorab, A. K. (2014). An organizational psychology perspective to examining computer security incident response teams. *IEEE Security & Privacy*, 12(5), 61-67. doi: 10.1109/MSP.2014.85

Crozier, R. (2018), Maersk had to reinstall all IT systems after NotPetya infection, *itnews*, 25 January, retrieved from <https://www.itnews.com.au/news/maersk-had-to-reinstall-all-it-systems-after-notpetya-infection-481815>.

Davidson, J. L., Jacobson, C., Lyth, A., Dedekorkut-Howes, A., Baldwin, C. L., Ellison, J. C., Holbrook, N. J., Howes, M. J., Serrao-Neumann, S., Singh-Peterson, L., & Smith, T. (2016). Interrogating resilience: Toward a typology to improve its operationalization. *Ecology and Society*, 21(2), 1-15. doi:10.5751/ES-08450-210227.

Dupont, B. (2019). The cyber-resilience of financial institutions: Significance and applicability. *Journal of Cybersecurity*, 5(1), 1-17. doi:10.1093/cybsec/tyz013

Eisenhower, D. E. (1958), Remarks at the National Defense Executive Reserve Conference – November 14, 1957, in *1957: containing the public messages, speeches, and statements of the president, January 1 to December 31, 1957*, Office of the Federal Register, National Archives and Records Service, Washington DC, pp. 817-820.

ESRB (2020, January 7), The General Board of the European Systemic Risk Board held its 36th regular meeting on 19 December 2019, *Press Release*, retrieved from <https://www.esrb.europa.eu/news/pr/date/2020/html/esrb.pr200107~29129d5701.en.html>.

Flanagan, J. C. (1954). The critical incident technique. *Psychological Bulletin*, 51(4), 327-358.

Frost & Sullivan (2017). *The 2017 global information security workforce study: Women in cybersecurity*. Frost & Sullivan.

Fujs, D., Mihelic, A., & Vhrovec, S. (2019, August 26–29). The power of interpretation: Qualitative methods in cybersecurity research. In *Proceedings of the 14th International Conference on Availability, Reliability and Security (ARES 2019)*, Canterbury. doi:10.1145/3339252.3341479

Giddens, A. (1999). Risk and responsibility. *The Modern Law Review*, 62(1), 1-10. doi:10.1111/1468-2230.00188

Gim, C., Miller, C. A., Hirt, P. W. (2019). The resilience work of institutions. *Environmental Science and Policy*, 97, 36-43.

Granovetter, M. (1973). The strength of weak ties. *The American Journal of Sociology*, 78(6), 1360-1380.

Greenberg, A. (2019). *Sandworm: A new era of cyberwar and the hunt for the Kremlin's most dangerous hackers*. Doubleday.

Grossetti, M. (2009). Imprévisibilités et irréversibilités : Les composantes des bifurcations. In M. Grossetti, M. Bessin, & C. Bidart (Eds.), *Bifurcations : Les sciences sociales face aux ruptures et à l'événement* (pp. 147-159).

Hagenaars, K. J. C. (2019). *An empirical study into how cyber security professionals deal with uncertainty in information security risks assessments*. Management of Technology, Technische Universiteit Delft.

Hielkema, P., & Kleijmeer, R. (2019). *Lessons learned and evolving practices of the TIBER framework for resilience testing in the Netherlands*. Carnegie Endowment for International Peace. Retrieved from https://carnegieendowment.org/files/WP_Hielkema_Kleijmeer_TIBER1.pdf.

Holling, C. S. (1996). Engineering resilience versus ecological resilience. In P. Schulze (Ed.), *Engineering within ecological constraints* (pp. 31-44).

Janis, I. L. (1972). *Victims of groupthink: A psychological study of foreign-policy decisions and fiascoes*. Houghton Mifflin.

Joshi, S. (2020, January 29). Reservist model: Distributed approach to scaling incident response. *Enigma Conference*. Retrieved from <https://www.usenix.org/conference/enigma2020/presentation/joshi>.

Kaplan, H. (1999). Toward an understanding of resilience: A critical review of definitions and models. In M. D. Glantz & J. L. Johnson (Eds.), *Resilience and development: Positive life adaptations* (pp. 17-83).

Linkov, I., Eisenberg, D., Plourde, K., Seager, T., Allen, J., & Kott, A. (2013a). Resilience metrics for cyber systems. *Environment Systems and Decisions*, 33(4), 471-476. doi:10.1007/s10669-013-9485-y

Linkov, I., Eisenberg, D. A., Bates, M. E., Chang, D., Convertino, M., Allen, J. H., Flynn, S. E., & Seager, T. P. (2013b). Measurable resilience for actionable policy. *Environmental Science & Technology*, 47(18), 10108-10110. doi:10.1021/es403443n

Maennel K., Ottis R., & Maennel O. (2017). Improving and measuring learning effectiveness at cyber defense exercises. In H. Lipmaa, A. Mitrokotsa, & R. Matulevičius (Eds.), *Secure IT Systems*. Springer (pp. 123-138). doi:https://doi.org/10.1007/978-3-319-70290-2_8

Manyena, S. B. (2006). The concept of resilience revisited. *Disasters*, 30(4), 433-450. doi:10.1111/j.0361-3666.2006.00331.x

McCammon, I. (2004). Heuristic traps in recreational avalanche accidents: Evidence and implications. *Avalanche News*, 68, 1-10.

Meyer, R., & Kunreuther, H. (2017). *The ostrich paradox: Why we underprepare for disasters*. Wharton Digital Press.

NCSC (2018, February 14), *Russian military 'almost certainly' responsible for destructive 2017 cyber attack*, National Cyber Security Centre, London, retrieved from <https://www.ncsc.gov.uk/news/russian-military-almost-certainly-responsible-destructive-2017-cyber-attack>.

Olenick, Doug (2018, January 26), NotPetya attack totally destroyed Maersk's computer network: Chairman, *SCMagazine*, retrieved from <https://www.scmagazine.com/home/security-news/ransomware/notpetya-attack-totally-destroyed-maersks-computer-network-chairman/>.

Paton, D., & Johnston, D. (2017). *Disaster resilience: An integrated approach*. Charles C. Thomas.

Patton, M. Q. (2015). Sampling, qualitative (purposeful). in G. Ritzer (Ed.), *The Blackwell encyclopedia of sociology*. John Wiley & Sons. doi:10.1002/9781405165518.wbeoss012.pub2

Rid, T. (2012). Cyber war will not take place. *Journal of Strategic Studies*, 35(1), 5-32. doi:10.1080/01402390.2011.608939

Rid, T., & Buchanan, B. (2015). 'Attributing cyber attacks'. *Journal of Strategic Studies*, 38(1-2), 4-37. doi:10.1080/01402390.2014.977382

Ritchie, R. (2019, August). Maersk: Springing back from a catastrophic cyber-attack. *I-CIO*, retrieved from <https://www.i-cio.com/management/insight/item/maersk-springing-back-from-a-catastrophic-cyber-attack>.

Schroeder, P. (2019). U.S. House lawmakers ask regulators to scrutinize bank cloud providers. *Reuters*, 23 August. Retrieved from <https://www.reuters.com/article/us-usa-congress-cloud/u-s-house-lawmakers-ask-regulators-to-scrutinize-bank-cloud-providers-idUSKCN1VD0Y4>.

Simpson, N., Shearing, C. D., & Dupont, B. (2019). Climate gating: A case study of emerging responses to Anthropocene risks. *Climate Risk Management*, 26, 1-10. doi:10.1016/j.crm.2019.100196

Smith, D. E. (1987). *The everyday world as problematic: A feminist sociology*. Northeastern University Press.

Staal, M. (2004). *Stress, cognition and human performance: A literature review and conceptual framework*. NASA Ames Research Center. Retrieved from <https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20060017835.pdf>.

Steinke, J., Bolunmez, B., Fletcher, L., Wang, V., Tomassetti, A. J, Repchick, K. M., Zaccaro, S. J., Dalal, R. S., & Tetrick, L. E. (2015). Improving cybersecurity incident response team effectiveness using teams-based research. *IEEE Security & Privacy*, 13(4), 20-29. doi: 10.1109/MSP.2015.71

The National Academies. (2012). *Disaster resilience: A national imperative*. The National Academies Press. doi:10.17226/13457

Van Eeten, M., Nieuwenhuijs, A., Luijff, E., Klaver, M., & Cruz, E. (2011). The state and the threat of cascading failures across critical infrastructures: The implications of empirical evidence from media incident reports. *Public Administration*, 89(2), 381-400.

Weick, K. E. (1995). *Sensemaking in organizations*. SAGE Publications.

Weick, K. E., & Sutcliffe, K. M. (2015), *Managing the unexpected: Sustained performance in a complex world – Third edition*. John Wiley & Sons.

Zhang, M., Wang, L., Jajodia, S., Singhal, A., and Albanese, M. (2016). Network diversity: A security metric for evaluating the resilience of networks against zero-day attacks. *IEEE Transactions on Information Forensics and Security*, 11(5), 1071-1086. doi:10.1109/TIFS.2016.2516916