CHAPTER

10

Cyber situational awareness issues and challenges

Ulrik Franke[a,b], *Annika Andreasson*[a], *Henrik Artman*[a,c], *Joel Brynielsson*[a,c], *Stefan Varga*[a,d], *and Niklas Vilhelm*[a,e]

[a]KTH Royal Institute of Technology, Stockholm, Sweden, [b]RISE Research Institutes of Sweden, Kista, Sweden, [c]FOI Swedish Defence Research Agency, Stockholm, Sweden, [d]Swedish Armed Forces Headquarters, Stockholm, Sweden, [e]Norwegian National Security Authority, Sandvika, Norway

1 Introduction

Modern society is full of information technology, underpinning the operations of corporations, governments, and nongovernmental organizations, as well as the personal and social lives of billions of individuals. While there are obvious benefits to these digital tools, their widespread use also means that they must be increasingly dependable. As a result, enterprises of all sorts need to make informed decisions that uphold the dependability of digital services. If the information stored in digital systems cannot be trusted, or if the systems themselves cannot be reached when needed, the benefits of digitalization will not materialize.

1.1 Cyber risk management

Clearly, there are numerous risks to networks and the information in them (see Cebula, Popeck, & Young, 2014 for a taxonomy). One class of risk is systems and technology failures, internal or external, that occur without any adversarial intent. Another class of risk is the threat from adversaries; threat actors who can reason rationally and adapt to different defensive strategies. It is important to consider both classes of risk, and in

https://doi.org/10.1016/B978-0-323-90570-1.00015-2

particular not to forget about nonadversarial risks just because there is much talk about deliberate attacks. Accidents can be equally damaging.

However, while it is easy to make the theoretical distinction between these two classes, it is not equally clear in practice, where adversaries have every reason to conceal themselves and their actions as mere accidents. Indeed, there is some evidence to suggest that defenders do not always care about this distinction (Varga, Brynielsson, & Franke, 2018, 2021), and that decision makers assess the probabilities of the two classes of threats as very similar (Franke & Wernberg, 2020).

Properly managing these risks typically involves many different actions. Some are technical in nature, and involve things like setting up and automating updates, backups, encryption, access privileges, and network monitoring. However, it is important to realize that this is not enough. Organizational measures such as training of employees (McCrohan, Engel, & Harvey, 2010) and exercises in cyber incident management (Maennel, Ottis, & Maennel, 2017) are also important. Human error is often the root cause of incidents, and technical measures rarely suffice if a legitimate user can be tricked into assisting in an attack (Krombholz, Hobel, Huber, & Weippl, 2015). Furthermore, senior management needs to work actively with identifying how operations depend upon technology, and document this in a strategy that is regularly updated (Dunbar, 2012; Dutta & McCrohan, 2002). It is important to realize that such risk management work critically depends upon understanding one's own organization and operations—it is not something that can be outsourced.

1.2 Cyber situational awareness

The brief outline above suggests that prudent cyber risk management encompasses a multitude of decisions. To make the best decisions, however, decision makers need to have an accurate situational awareness (SA)—colloquially, to *know what is going on*. Such SA with respect to information technology and networks—cyber situational awareness (CSA)—is the topic of this chapter.

Somewhat more formally, our point of departure is the definition of SA that stems from Endsley (1988) and her investigation of how aircraft pilots understand the situation. Using her definition, SA is "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" (Endsley, 1988, p. 792). Based on this definition, Endsley in later work defines three levels of SA: (i) perception, (ii) comprehension, and (iii) projection (Endsley, 1995). In this chapter, we adhere to the Endsley-inspired definition of CSA given in our previous work: "a subset of situational awareness, i.e., cyber situational awareness is the part of

situational awareness which concerns the 'cyber' environment" (Franke & Brynielsson, 2014, p. 20).

Much of the existing CSA literature is very technologically oriented, that is, focused on awareness of the state of the technical artifacts that constitute the network. In the literature review we published in 2014, we found comparatively much work on CSA for industrial control systems, and for general algorithms and information fusion in intrusion detection systems (IDSs). By contrast, there was considerably less research done in other areas such as information exchange, the risks of cyber deception, and of cyber battle damage assessments in military operations. Evancich et al. (2014) suggested that aspects such as operational awareness and threat awareness are also relevant for CSA.

1.3 Common operational picture and situational awareness

Since the two are often discussed in parallel and sometimes confused, it is useful to make the distinction between a common operational picture (COP) and SA. The former is an artifact—such as a shared big screen, a whiteboard, or a map with symbols—which aims to provide stakeholders with a "common picture" of what is going on (the concept of a COP has military origins; Hager, 1997). The latter is a mental state—awareness of what is going on. Of course, the COP is designed, more or less deliberately, to facilitate SA (CSA, in our case), but even though the two are closely related, they are conceptually different. This COP/CSA distinction is reminiscent of how Gutzwiller, Hunt, and Lange (2016) argued that there is a need to differentiate (the products of) data fusion (typically a COP) from actual CSA, where human cognition is a central component.

The simplest way in which the COP facilitates SA is that an individual just looks at the COP to gain SA. Here, the COP serves as the "information warehouse" (Copeland, 2008) where relevant information is stored and computed. However, particularly when several individuals are involved, the COP can also be perceived as a process (Wolbers & Boersma, 2013), where the meaning and interpretation of information is constantly (re)negotiated as stakeholders interact. From this perspective, the COP is a "trading zone" (Wolbers & Boersma, 2013), and all the interactions within this zone facilitate SA (CSA, in our case).

1.4 Outline

In the rest of this chapter, we further develop and elaborate the themes introduced above. A recurring theme is the necessity to include sociocognitive and organizational factors, in addition to technology, to achieve the most relevant and useful kind of CSA, and further that an important challenge to be addressed concerns adversarial behavior.

First, however, in Section 2, some of the technological building blocks necessary for CSA are discussed. Once these are in place, the scene is set for a look at socio-cognitive aspects in Section 3. Based on the previous two sections, we proceed to organizational issues in Section 4. Section 5 then investigates the prospects for reasoning about adversarial actions. As mentioned earlier, not all cyber incidents are adversarial (accidents and honest mistakes abound), but when there is an adversary, this adds an important dimension to CSA. In Section 6, some directions for future research are proposed, before Section 7 concludes this chapter.

2 The technological perspective

It is obvious that CSA is dependent on information from technical sensors. In this section, we review some of the technologies necessary to adequately understand what is going on in the cyber domain. How can IDSs, techniques for visualization of large datasets, and fusion of information from different sources contribute to CSA?

2.1 Nonadversarial cyber incidents, design, and resilience

The strategies to build secure and dependable IT systems have evolved over the years. In the infancy of computing, hardware failure was behind many outages and incidents. However, with improved manufacturing techniques and quality control in computer components, as well as reliability designs such as redundant components with failover, hardware failure no longer accounts for a significant share of downtime (Genadis, 1996).

Instead, for the last three or four decades, the major culprits of nonadversarial incidents have been software errors and mistakes in IT administration (Gray, 1985, 1990). For example, an investigation of failures in web applications found some 80% to be caused by software or human error. Similarly, Gartner found that people and process failures were behind approximately 80% of "mission-critical" downtime (Malik & Scott, 2010).

Avoiding nonadversarial cyber incidents, thus, has evolved from a matter of building more reliable hardware to having quality checks on the software being used and finding ways to make sure that the people managing it adhere to appropriate processes. This is not the place to provide full introductions to the fields of software testing (see, e.g., Bertolino, 2007; Whittaker, 2000) or IT service management (see, e.g., Galup, Dattero, Quan, & Conger, 2009; Iden & Eikebrokk, 2013), but we do note that adequate CSA must include a basic understanding of these issues—not least to be able to distinguish nonadversarial events caused by mistakes and process failures from adversarial attacks.

2.2 Adversarial cyber incidents and defense against them

As for nonadversarial incidents, the strategies to defend against attacks have evolved over time. Li, Huang, Wang, and Li (2019, p. 2) identified four paradigms. The first paradigm occurred before the 1960s. Here security was a design feature; it was assumed that a solid hardware and software architecture would prevent attacks. In the 1970s and 1980s, security assistance systems, such as, for example, IDSs, were added to the existing systems to provide protection from attacks. Later in the 1990s, more sophisticated approaches emerged. A common theme for these, arguably, was that they aimed to be proactive. Attacks were modeled using techniques such as state graphs and attack trees (Mauw & Oostdijk, 2005; Schneier, 1999), allowing preparations for defense in advance. The latest and present paradigm, according to Li et al. (2019), commenced after the turn of the millennium. This paradigm prescribes that a wider perspective than before has to be taken into account to improve defensive efforts. Li et al. (2019) argued that today it is necessary to try and predict future trends with techniques such as complex network theory (Wan, Cao, Chen, & Huang, 2017; Xu, 2014), game theory (Attiah, Chatterjee, & Zou, 2018), etc. These new requirements have emerged not because the probabilities of chance events have changed in a significant way, but rather because of the continuous improvement of adversarial (cyber) attack methods.

2.3 The expanding scope of cyber situational awareness

As illustrated earlier, strategies both for (i) resilience against nonadversarial incidents and (ii) security against adversarial attacks are ever-expanding in scope. It seems clear that the cyber defense community has abandoned the position that security could be solved by design only, once and for all. Instead, the scope of adequate CSA now also includes knowledge of attackers and the *modus* of attacks, as well as factors beyond the technical network itself.

For example, with increasing use of third-party cloud services, maintaining awareness of *internet connectivity* is of vital interest (Gunawi et al., 2016). Are there planned outages? Are there incidents in other places that may result in lower quality of service as the load rebalances (Omer, Nilchiani, & Mostashari, 2009)? Have necessary precautions been taken, and is there a way to insure against the residual risk that cannot be managed by technical means (Franke, 2018)?

Another example concerns the *insider threat*. To know whether an unauthorized employee is moving around at the physical premises should be of interest for cyber defenders, because such an individual could have gained physical access to the IT system and pose a (cyber) threat. Similarly, if a suspicious hardware device is seen in an office room, that is also

of interest. These two examples are mentioned by Vielberth, Menges, and Pernul (2019), and illustrate the point that CSA is not only about "cyber" events—things going on in the physical world can be just as important. In the longer term, vetting and recruitment processes of cybersecurity personnel to ensure that competent and trustworthy personnel with the required psychological profiles is hired, are also beneficial for cybersecurity (Chen et al., 2014).

2.4 Intrusion detection systems

As mentioned earlier, the IDS is one the most important technical tools used for CSA with respect to adversarial attacks. An IDS collects data from networks or files on devices, and then determines if this data contains any indicators of compromise. Broadly, IDSs can be classified into two categories: (i) signature-based systems and (ii) anomaly-based systems (Khraisat, Gondal, Vamplew, & Kamruzzaman, 2019). The first category uses signatures of already known threats to detect them when encountered again. The second category instead looks for anomalies that may indicate that something is wrong in the network as a result of an attack. While different in principle, most commercial systems are a combination of both. Furthermore, while it might naïvely be assumed that signature-based systems are unable to detect previously unencountered threats, this is not the case. On the contrary, it has been shown that a rule-based IDS can detect attacks not described by its set of rules, though its detection rate is lower than for known attacks (Holm, 2014).

To successfully detect an incident the IDS depends on capturing the correct data and interpreting it correctly. Current approaches to detecting intrusions can be summarized as the use of (i) known pattern templates, (ii) threatening behavior templates, (iii) traffic analysis, (iv) statistical anomaly detection, and (v) state-based detection (Bass, 2000). Decisions about where data should be collected and at what rate are of critical importance. Adding more collection points in the infrastructure and more frequent data captures increase resource requirements. However, in real-world implementations trade-offs have to be made with respect to data collection (Werlinger, Hawkey, Muldner, Jaferian, & Beznosov, 2008).

A significant challenge to the use of an IDS is its detection characteristics: the rate of false positives (threats that are reported, but not real) and false negatives (threats that are real, but not reported). Too many reports might overwhelm system administrators, but too few might not provide adequate information. The sensitivity to which indicators should trigger a report is a balance between these two diverging metrics (Sommestad & Franke, 2015). Improving these rates is, therefore, an active research area (Goeschel, 2016; Spathoulas & Katsikas, 2010). However, it is also noteworthy that even though an IDS is a technical tool, it matters who uses it: competent system

administrators are important for IDS usage to be effective (Sommestad & Hunstad, 2013). This serves as a cautionary reminder about the importance of including human factors when designing and using technical solutions for CSA.

To generate the best possible assessment of whether attacks are ongoing, an IDS has to combine many different kinds of data. However, combining data and information from heterogeneous origins and producing reports useful to human decision makers is challenging (Bass, 2000). A collection of low-level responses from security mechanisms and other systems is not easily interpreted in isolation and typically needs to be put into a larger context and be analyzed to provide reliable information about the system. Still, inferences about the systems can be assumed to be correct to a greater degree of certainty if supported by more evidence. Conversely, errors are more reliably spotted if information on the same issue is gathered from multiple sources.

2.5 Intrusion detection systems and explainable artificial intelligence

With recent advances in artificial intelligence and in particular deep neural networks (DNNs), it is not surprising that these techniques have been applied to IDSs (Kang & Kang, 2016; Kim, Shin, Jo, & Kim, 2017; Roy, Mallik, Gulati, Obaidat, & Krishna, 2017). Since manual analysis of indicators of compromise tends to be very labor-intensive, the possibility to automate such analysis using AI methods is very attractive.

However, DNNs typically use a huge parametric space with hundreds of layers and millions of parameters and are thus for practical purposes considered "black boxes" which are very difficult to interpret for humans (Arrieta et al., 2020). This lack of transparency has led to an intense discussion—in the scientific community and in society at large—about the risk that black box AI systems will make biased or erroneous decisions in the absence of meaningful human supervision. For example, many AI systems have been found to exhibit biases disadvantaging, for example, poorer people and those from minorities (Nature, 2016).

For DNN-based IDSs to become more useful and trusted to make critical decisions autonomously, the explainability issues mentioned earlier need to be addressed. Within the field of explainable AI (XAI), methods are developed to make AI (and in particular DNN) systems less opaque. Review articles include Guidotti et al. (2018) and Du, Liu, and Hu (2019). Though many questions remain unsolved, advances within XAI are widely acknowledged as a critical feature for the practical deployment of AI models (Arrieta et al., 2020). Such research should also include studies of explainability from the perspective of human-computer interaction (Abdul, Vermeulen, Wang, Lim, & Kankanhalli, 2018).

2.6 Network-centric and domain-centric cyber situational awareness

It seems intuitively plausible that a computer network that is set up with state-of-the-art cyber defense software fully loaded with relevant and timely IDS signatures and the likes, will fare better than a system that lacks up-to-date threat information. Similarly, it also seems likely that a network defended by individuals who possess a deep understanding of the techniques, tactics and procedures of potential adversaries, along with their underlying motivations and rationales (Brynielsson, Franke, Tariq, & Varga, 2016), are able to perform a better job in defending their network than people who do not know anything about these things. Unfortunately, to the best of our knowledge, neither of these two propositions have been rigorously examined through research efforts. We note that the former question, which comprises the "on-the-network" fight (Borum, Felker, Kern, Dennesen, & Feyes, 2015), draws more academic interest than the latter more tactical and strategic dimensions of cyber intelligence (Borum et al., 2015).

Cyber threat intelligence (CTI) is often discussed with the tacit assumption that the "use case" is cyber defense from the context of cyber analysts. However, cyber defense tasks can be performed at multiple managerial levels: ranging from the operational all the way up to the strategic level. Therefore, CTI should be thought of in terms of multiple levels as well (Mattern, Felker, Borum, & Bamford, 2014). Cyber analysts need to understand the human side of the threat, and not only the technical manifestations of it (Mattern et al., 2014). This means that analysts need to understand underlying motivations and rationales for people who are behind attacks on a strategic level, but also the role of humans as offensive cyber operators on the tactical and operational levels. In line with this thinking, Zheng and Lewis (2015) assert that there are two basic categories of cyber threat information (to share), first: technical threat indicators (e.g., IP addresses, specific strings of data, file hashes, exploit toolkits or payloads, adversary tactics, techniques and procedures), and second: contextual threat intelligence (e.g., exploit targets, exfiltrated content, incident details, and specific courses of action). The technical indicators provide direct descriptive information related to the actual breach, while the contextual indicators provide more indirect information such as related effects and other pieces of explanatory information. In line with this distinction, we propose that it is useful to distinguish between *network-centric* and *domain-centric* CSA, as outlined in Table 1.

The network-centric category mainly relates to awareness of technological (network) level data and information. Data and information at this level include low-level indicators of compromise, such as known malware hash values and command and control servers, IP addresses, and so on

TABLE 1 Network-centric and domain-centric CSA.

	Network	Domain
Focus	Technical	Other (organization, mission)
Scope	Narrow (computers and network)	Wide
Automation	Yes	No
Analysis	Fusion	Aggregation, semantic meaning
Human intervention required	No (yes)	Yes
Timescale	Short term	Long term

(Zheng & Lewis, 2015). Furthermore, Brown, Gommers, and Serrano (2015) point out that related "network"-level intelligence management processes probably can be fully automated, as discussed in the previous section. Network-centric awareness may lead to immediate corrective actions, such as the patching of vulnerabilities or other "stop the bleeding"-kind of activities that will improve the level of cybersecurity in the short term (Ahmad, Desouza, Maynard, Naseer, & Baskerville, 2020).

By contrast, the domain-centric category relates to awareness of issues in the organizational and threat dimensions, including information or intelligence about more complex matters such as, for example, descriptions of threat actors and their motivations, modes of attack, and other characteristics. Brown et al. (2015) also mention information about the use of different malware families *over time*. A deeper understanding of this kind of information can result in more profound organizational changes, such as amendments of current security strategies, processes, and workflows, that will lead to improvements of the overall cybersecurity posture in the long term.

To further clarify this distinction, consider how different types of specific information can be assigned to one or the other of the two categories. We do this by using STIX, the *Structured Threat Information eXpression* language and serialization format, since it is the de facto standard for exchange of cyber threat intelligence (Sauerwein, Sillaber, Mussmann, & Breu, 2017). In STIX, we find a taxonomy of relevant information elements for cyber defense in the so-called STIX domain objects (SDOs).

Hence, we propose that network-centric CSA should correspond to the following example SDOs: *indicator*, *observed data*, *tool*, and *vulnerability*. These categories are formed by what we may call atomic information elements. Here we mean atomic in the sense of foundational and

minuscule. Atomic information (elements) should leave no or little room for interpretation—it is either right or wrong. Further, our chosen atom metaphor is appropriate because it also alludes to the possibility of combining multiple atoms to infer a high-level understanding of the phenomena under investigation.

The domain-centric CSA can be built upon the analysis of atomic information elements and higher-order pieces of information or intelligence. Corresponding STIX SDO examples are: *campaign*, *course of action*, and *threat actor*. To make sense of this type of information, richer descriptions are required. At the moment, these SDOs need to be processed and acted on by humans who can understand, and discuss, their meaning. This sets the scene for the discussion about socio-cognitive aspects in the next section.

3 The socio-cognitive perspective

From the previous section, we are now acquainted with some of the cyber information available to decision makers. But how do human decision makers process and combine this information? How do we reach CSA when working in teams? What pitfalls need to be avoided?

Endsley's definition of SA and, by inheritance, our definition of CSA as given in Section 1 is individualistic. There is a world, the individual can pay attention to or fail to pay attention to it, and (C)SA is in some sense about mirroring the world within the individual mind. However, much work with complex problems and systems, with a great emphasis on security and reliability that requires significant and fast action, is conducted in teams (Artman, 2000; Artman & Wærn, 1999). This *teamwork* introduces an additional and important perspective. While some aspects of teamwork are addressed in Section 4, we start this exposition of the socio-cognitive perspective with some considerations about team SA (Demir, McNeese, & Cooke, 2017), before considering the implications of the cyber environment for (i) perception, (ii) comprehension, and (iii) projection into the future.

3.1 Team cyber situational awareness

In most teams working to manage cyber incidents, responsibilities are distributed for different, albeit overlapping, tasks which must somehow be coordinated if problems are to be solved. Coordination in such cases is usually a continuous process where different team members tell each other what they are paying attention to or looking for, and how they understand the situation. In these communications between different members, a shared model of understanding of the course of events, and possibly also explanatory models of why they arose, is also created (Artman, Brynielsson, Johansson, & Trnka, 2011).

The physical world is governed by the laws of physics. Based on examining physical objects and their relationship, one can create fairly reliable models and causal relationships. Social systems are not equally amenable to description by laws, because they are based on intangible relationships such as trust, confidence, experience, social hierarchies, knowledge, and the communication of these between people. In our context, this highlights the difficulty of understanding and communicating about a specific attack and the intentions behind it (Artman et al., 2011). Since you can never know for sure another person's intentions with specific actions, you have to make assumptions. Such assumptions may well be based on experience and knowledge, but are nevertheless based on detached prejudices and hypotheses.

Such a shared model—created for the moment and in order to understand a complex process—risks becoming so influential that it in itself becomes an indisputable reality for the group. This may lead the group to look more for evidence confirming what is (thought to be) known than for contradicting evidence. For individuals, this is known as *confirmation bias* (Nickerson, 1998; Tversky & Kahneman, 1974), whereas the corresponding group phenomenon is sometimes called *groupthink*. This concept has famously been used to explain decisions made during the Bay of Pigs invasion, the Vietnam war, and the Watergate scandal (Kramer, 1998; Raven, 1998).

3.2 Perception

Starting from Endsley's definition, what does it mean to translate it to the cyber environment, as we do in this chapter? Endsley's SA concept is developed for aviation. In aviation (and more generally, in most physical environments), perception—of views, sounds, vibrations, etc.—cannot be avoided. In the cockpit, the pilot can see what is happening in the physical space either through the cockpit window or through instruments, which mediate elements in the environment beyond visual range. Auditory aspects are also important to be able to assess the state of the aircraft. Other sensory input such as G-forces also play a significant role.

This differs from the cyber environment, where perception only takes place once appropriate sensors—firewalls, IDSs, vigilant users—are in place. Log file analysis systems that warn of unusual and critical events can, perhaps, be likened to a radar system, but the systems must be designed to actively search for unusual elements or events. Radar systems search for physical materials that can be distinguished from other physical materials. In the cyber world, however, all elements and all events are of the same immaterial nature. What remains are intangible processes that must differ from the normal state of the system one is trying to protect. Anyone who tries to attack and break into the digital system knows this,

of course, and can continuously delete any digital tracks to avoid detection. This cleaning of digital tracks also differs greatly from how a physical element can be masked. Furthermore, an attacker can introduce false tracks to confuse. The attacker can impersonate a less competent actor or, conversely; a competent person who is trying to attack a particular part of the system may act so as to hide that he/she is "actually" trying to intrude on another part of the system, which may also be less defended. Thus, while SA in the physical environment is *perceptually driven*, CSA is rather *search-driven*.

In aviation, these elements typically appear some time before there is an actual crisis—and the pilot will have at least some time to act and react—while elements and events in cyberspace might appear only at the moment of imminent attack. At least, this is the case for network-centric CSA as discussed in Section 2; it may be less so for domain-centric CSA. In the cyber world, the constraints of time and space are less binding. For example, a person can be on the other side of the globe and attack a system without any significant time constraints.

3.3 Comprehension

The next level in Endsley's model is about understanding what the different elements or events mean. The discussion above has already touched on meanings and meaning because in cyberattacks there is always an acting person and this person has intentions with his/her actions. However, the intention may vary; some attack a system or organization for criminal reasons, others do it to learn and test their skills, and some do it to warn the organization about loopholes in the system. There are certainly many other purposes for a cyberattack as well (Tariq, Brynielsson, & Artman, 2012). As we have already mentioned, adversarial attacks mean that as a defender of the system or organization, one must have an understanding of who the opponent is, what intentions they may have, and how many and well-organized they are. Events and elements that can be traced in log files or during ongoing attacks contain weak signals to be used for the purpose of determining the intentions and goals of the opponent. As we have already pointed out, this means that CSA must to a large extent rely on creating plausible hypotheses and constantly examining the course of events.

3.4 Projection into the future

Such hypotheses in turn govern how one understands future events and can act to prevent these from happening. A team working to counter future potential events is well advised not only to understand the intentions of the

opponent, but also its skills and capacity to carry out attacks. One must therefore to a large extent preempt the opponent's future activities and block the possibilities for further intrusion. In the physical world, there is usually a certain time difference between starting an activity until you can complete it. In the cyber world, this time difference is minimized, which makes it difficult to take the time to reflect on a potential course of events—there is, hence, a risk of "extinguishing ongoing fires" rather than preventing the fire from spreading.

3.5 The challenge of a complex nonphysical environment

When trying to understand complex dynamic systems where there are many interrelated dependencies—often in chains of events and where it is not obvious how elements, subsystems, and individual systems relate to each other—there is a risk that we create simplified mental models that work well in the short run, but have unknown and unintended adverse effects in the long run. Among natural systems, the climate is an obvious example, where effects occur in long causal chains which are difficult to fully grasp. Human use of raw materials releases carbon dioxide into the atmosphere, which in turn increases the temperature, which in turn causes animals and plants to migrate, etc. Among social systems, the economy is a good example, where the possibility or impossibility of central planning was a hotly contested topic in the early 20th century (and revisited time and again since), and many arguments focus precisely on the impossibility for any single individual to gather all the data and properly understand all the long causal chains involved in economic transactions (Cottrell & Cockshott, 1993; Hayek, 1982; Yeager, 1994). In such cases, there is a risk that decision makers and analysts create such a simplified worldview that even though a single good is closely monitored and achieved, other elements are ignored, to the detriment of other important desiderata. This phenomenon is known as *encapsulation* (Dörner & Schaub, 1994).

Discussing CSA and the intangible cyber environment, the risk of encapsulation must be considered: do analysts and decision makers have oversimplified conceptions of what is going on, what they themselves try to achieve, and what the attackers are doing? Such oversimplifications can also be entrenched by groupthink.

4 The organizational perspective

Different stakeholders need different forms of CSA. For example, an operator in a security operations center (SOC) has a different time perspective than the CEO. But sometimes the operator might need to escalate a

question quickly to the CEO. How should an organization be designed to facilitate the kind of information flows and mandates that enable making correct and timely decisions?

4.1 Operational level

At first glance, CSA usually seems to be discussed in relation to analysts in a SOC. It is often in the SOC that adversarial cyberattacks are detected, responded to, and recovered from by the SOC staff. But SOC analysts are not the only members of an organization with a need for CSA. Understanding the complex inner workings of the SOC is, indeed, a prioritized undertaking, but the cybersecurity efforts of an organization should not stop there. One way of addressing CSA in organizations is to define the CSA requirements of different professional security roles, as proposed by, among others, Gutzwiller, Dykstra, and Payne (2020). One idea suggested is to use the *Workforce Framework for Cybersecurity (NICE Framework)*—developed by the US National Institute of Standards and Technology (NIST)—as a starting point (Petersen, Santos, Smith, Wetzel, & Witte, 2020). The framework is an effort to establish what different professional cybersecurity roles entail, without being dependent on where the role is situated. Each role has skills and tasks associated with it. This would result in generic CSA requirements that organizations could use as a starting point for establishing what information is needed for the work performed by different cybersecurity operators within the organization.

4.2 Tactical and strategic level

However, in an organization there are other decision makers in need of CSA than the immediate cybersecurity workforce. When an organization is under attack, information about the present situation often needs to reach higher up in the decision chain, as the event could pose a potential threat to ongoing operations and even the enterprise as a whole. Goals, timelines, and CSA requirements vary with the roles in the organizational hierarchy. The distinction between network-centric and domain-centric CSA as illustrated in Table 1 can be seen as an example of this. As pointed out by McKenna, Staheli, and Meyer (2015), organization members like the SOC analyst, SOC manager, CIO, and CEO, all need CSA to a varying degree. While under attack, the core issues and goals of the SOC analyst are quite different from the core issues and goals of the CEO. While one focuses on what is going on in the network, the other focuses on what that might mean in the wider context of future operations. In this linear organization example, information flows up in the hierarchy and decisions flow down.

4.3 Crisis management

It should not be forgotten, though, that there are roles not directly involved in the "vertical" line organization or that are part of the NICE framework roles, that nevertheless have a need for CSA. As an example, in the face of larger cyber events, special crisis management organizations are often activated, involving members from several different parts of an organization. While cybersecurity employees are involved in the immediate incident response, crisis management teams are simultaneously working on issues such as business continuity, stakeholder communications, and media relations. These teams, to a varying degree, also need CSA to be able to make the best decisions in going forward. Organizations need to ensure that each team has the CSA required to fulfill their goals, which is often complicated by the conflicting need to restrict regular users' access to sensitive information due to security concerns (Tariq, Brynielsson, & Artman, 2014). Aligned with the goals, the operators should have the mandate to make the decisions required to fulfill those goals.

4.4 Methods for research and organizational design

The method proposed by Endsley and Jones (2011) to capture SA requirements, is a goal-directed task analysis (GDTA). In a GDTA, the goals, decisions, and information requirements of different operators are identified. From there, an organization could design information sharing infrastructures to aid the CSA of different teams and operators with a need to know "what is going on" in the cyber domain. This could indeed be a way forward for organizations to determine where there are shared information requirements, and establish routines and ways of working where information reaches all the decision makers who need it. This reduces the risk that decision makers needing the same information collect it from sources of different quality, unbeknownst of each other.

5 Reasoning about adversarial behavior

Not all cyber incidents are adversarial, but the line can be difficult to draw—partly because an adversary might want to disguise attacks as mere accidental outages. Previous research suggests that in some circumstances decision makers do not take adversarial behavior into account to the extent that would be prudent (Varga et al., 2018, 2021). How can CSA be improved by reasoning about adversarial behavior to better identify things like diversions?

5.1 Threat modeling

There are numerous cyber threat models that aim to inform about cyber threats. Adam Shostack (2014) lists a few in his book *Threat modeling: Designing for security*. A few examples include the Threat Agent Risk Assessment (TARA) framework from the Intel corporation (Rosenquist, 2009) that lists 21 different threat agents; the OCTAVE Allegro information security risk assessment methodology that emphasizes the use of threat scenarios (Caralli, Stevens, Young, & Wilson, 2007); and the Military Activities and Cyber Effects (MACE) taxonomy that lists eight different adversary types (Bernier, 2013). The adversary types are described in terms of skill levels, maliciousness, motivation, and methods used. There are plenty of other models available.

These models generally draw a picture of various attackers based on certain characteristics. The idea is that conclusions about the threat against one's own organization should be drawn and put to use. Precisely *how* such inferences should be made and transformed into action, however, is far from trivial. This aspect does not appear to be extensively discussed in the literature. The term "threat actors" is mentioned a total of six times in the book *Cyber threat intelligence* (Dehghantanha, Conti, & Dargahi, 2018), but its relevance and meaning are not elaborated on beyond the fact that threat actors exist. Nor is the specifics of threat actors a subject for treatment in the book *Information security practices: Emerging threats and perspectives* (Traoré, Awad, & Woungang, 2017). Furthermore, in an extensive analysis of information security data sources by Sauerwein, Pekaric, Felderer, and Breu (2019), threat actor information was not mentioned at all. In conclusion, there appears to be a strong emphasis on information about vulnerabilities and other threats rather than on details on threat actor information in the contemporary literature about cyber threat information and intelligence. It seems that there is a sense of dutiful obligation to mention the presence of threat actors, but then to leave it at that.

5.2 Usefulness of threat actor information

Knowledge about potential adversaries' goals and motives, for example, their target selection priorities, as well as details about how they conduct their cyber operations, will surely help cyber defenders to prioritize their defensive efforts. A corporation, for example, with a range of products or services that almost certainly is not of interest for a nation state, might with some certainty assume that they will not have to dimension their cybersecurity to withstand that particular threat actor. Furthermore, detailed threat actor knowledge might help defenders to predict adversarial actions in sustained cyberattack campaigns.

To verify and indicate whether practitioners actually value and use threat actor information, we interviewed 10 experts about whether they perceived a need for threat actor information or not in a team cyber defense exercise (see the "Appendix" section). The experts *did not* unanimously ask for information about threat actors. Two respondents pointed out that atomic information (e.g., IP addresses) was considered to be enough. One respondent *explicitly declined* to receive information about threat actor identities. The remaining seven respondents expressed a wish to receive cyber threat actor information, although sometimes in vague terms. They generally asked for adversary goals and motives, as well as for details about their *modus operandi*. Moreover, two respondents raised the point that threat actor information is probably more useful at the strategic managerial level (for the purpose of determining threat sources, e.g., attribution), than at the technical level (for hands-on threat detection and mitigation). This is in line with the observation in Section 4 that different stakeholders need different information depending on their roles.

5.3 Cyber threat intelligence

Inquiries into how cyber threat information can and should be selected, processed and shared have emerged as a research field called cyber threat intelligence, CTI (Mavroeidis & Bromander, 2017; Sauerwein et al., 2017; Shin & Lowry, 2020; Tounsi & Rais, 2018; Wagner, Mahbub, Palomar, & Abdallah, 2019). CTI at large, hence, aims to design (cyber) information processing capabilities that help decision makers to leverage information in order to enable sensible cyber defense-related decisions to ultimately remedy the fundamentally disadvantageous position of the cyber defender (Mohaisen, Al-Ibrahim, Kamhoua, Kwiat, & Njilla, 2017). Burger, Goodman, Kampanakis, and Zhu (2014) propose a taxonomy for cyber threat information that encompasses the full scope of intelligence information exchange between stakeholders. The taxonomy has five layers that contain information elements ranging from straightforward to more complex. The layers are *transport*, *session*, *indicators*, *intelligence*, and *5W1H*. At one end of the spectrum, the lowermost transport-level involves information about the movement of bytes, for example, data streams that represent the cyber threat intelligence between enterprises. At the other end, we find the uppermost 5W1H-layer. Information here, which is fed from the underlying layers, aims to answer questions such as *who*, *what*, *when*, *where*, *why*, and *how*? Hence, it seeks to answer the overarching question of attribution: who or what organization is responsible for the threat?

5.4 Game-theoretic approaches

Based on what has been presented earlier in the chapter, this section serves to provide a basis for what would be required for an organization to be able to gain full understanding of the cyber threat, with a special emphasis on adversarial behavior, and how such a model can be used as the underlying model to be used for obtaining and sustaining CSA with respect to adversaries.

In Sections 2.1 and 2.2, a distinction was made between nonadversarial and adversarial cyber incidents, where nonadversarial incidents refer to different types of IT-related system errors while adversarial attacks concern some form of willful-thinking opponent being behind the attack. However, as discussed in Section 2.1, it is not easy to distinguish between these two cases because a willful-thinking opponent will typically be interested in hiding his/her intentions by, for example, masking them in terms of ordinary system outages or other shortcomings that can be related to the organization's ordinary information technology architecture. There is thus a need to be able to maintain and reason about the situation with respect to two conceivable, but widely differing, models where one model is based on technical problems and the other on a rational-thinking opponent.

Section 2.3 further raised the issue of an attack's "modus," which exemplifies that adversarial attackers can follow widely differing modus operandi: is, for example, a single bank attack "solely" about stealing money, or is it part of a more far-reaching operation in which a highly capable state actor performs the attack as part of a larger attack with the aim to destabilize the country's payment system? This rhetorical question points out that it is not enough to just distinguish between nonadversarial and adversarial cyber threats, but that the question is more difficult: responding to a state-sponsored cyber operation is very different from dealing with a "script kiddie." Cyber operations undertaken by foreign powers are typically characterized by their focus, longevity, (close to) unlimited resources, and professional organization with specialists covering each field (target identification, infrastructure, operators, support, etc.).

Thus far, in our quest for modeling the cyber threat, an increasingly complex three-level modeling scale can be discerned along the lines of: (i) systems and technology failures, (ii) threats from reasoning, but financially restricted, adversaries, and (iii) threats from state actors. However, other dimensions also need to be addressed. Section 2.3 also discusses the insider threat, which forms the basis for a completely different type of rational-thinking opponent. Such an adversary can cause great harm even without significant IT skills, and the insider probably also has a completely different set of core values compared to the external adversaries discussed previously. Hence, this calls for a different model than the three-level model anticipated above.

The depicted situation is characterized by the fact that it is game-theoretical in that it largely deals with "reasoning about reasoning." Furthermore, there is uncertainty as to which model, or "game" according to game theory jargon, actually applies, for example, the "insider model" and the "technology failure model" are very different, and we would like to incorporate the fact that we are uncertain about which model really applies. Moreover, even if an enduring state-sponsored attack is unlikely, we would still like to model that possibility and update the likelihood as new intelligence arrives.

In game theory, adhering to established notation (Camerer, 2003; Myerson, 1991; Schelling, 1960; von Neumann & Morgenstern, 1944), the model is called a *game*, consisting of "the *strategies* each of several *players* have, with precise rules for the order in which players choose strategies, the information they have when they choose, and how they rate the desirability, or *utility*, of resulting outcomes" (Brynielsson, 2006, p. 47). In our context, the players consist of (i) our own organization and (ii) the adversary that we would like to reason about given the intelligence at hand; our own strategies consist of the different means to defend ourselves that we have at our disposal; the adversarial strategies consist of the possible attack vectors that our organization exposes; and the utilities denote the payoffs that the players experience given certain outcomes, that is, payoff values need to be defined for each combination of outcomes, and therefore the payoff values in a sense define the whole game model.

Now, as highlighted above, in our context the game models for different attackers are likely to be very different in terms of game rules, strategies, and payoffs, that is, the strategies available to an insider will be totally different from the strategies available to a state-sponsored cyber operation, and the utilities, that is, how these different attackers value different outcomes, will also be very different. Assuming a Bayesian prior probability distribution representing what players believe about other players, the previously described game model can be extended into a so-called *Bayesian game*, as described by Harsanyi (1967–1968). Using such a game—conceptually a probability distribution over the relevant game models—it is possible to model the envisioned situation where the involved actors not only reason about the opponent's reasoning, but also models their reasoning about its own and the other actors' reasoning about who the opponent actually is in terms of the rules of the game and the opponent's capabilities, beliefs, desires, and intentions, thereby making it possible to model the full complexity of the cyber threat.

So far in the discussion, we have defined the model top-down, but the bottom-up perspective is just as important. As already noted in passing, the utility/payoff values define the whole game and are equally important to consider. These values must be related to the information assets and information technology that make up the organization. This gives rise

to the threat intelligence needed to feed the model. As mentioned in Section 2.6, there are a number of different types of IT sensors that give rise to both technical threat indicators and contextual threat intelligence. Based on our previous work, this intelligence ought to be used as input to feed the envisioned Bayesian game with utility values in a bottom-up fashion (Brynielsson & Arnborg, 2006).

6 Research directions

While much is known about CSA, there are also many open questions. In particular, previous research has a technological focus, while the cognitive and organizational areas have received comparatively less attention. In the following, we list a few directions where we believe valuable research contributions ought to be made.

First, it would be valuable to reassess the applicability of current SA models for the cyber environment. As discussed in Section 3, current models were developed and adapted to solve problems within the physical realm, but the cyber domain is different. Current models require distinctly defined "situations." Such situations are hard to delimit in the cyber environment. Furthermore, current models have poor applicability for network-centric CSA: the timescale, with, for example, "computational speeds," cannot always be timely comprehended, much less acted upon, by humans. Models may not be optimal for domain-centric CSA either. Here, the timescale may instead be too long, and the "situations" become hard to delimit because of this. It might be that ordinary risk management practices that aim to reduce or mitigate risks for operations can be adapted for this purpose.

Second, it would be interesting to study the team aspect of CSA further. As argued in Section 3, analyzing how CSA is obtained in a command and control center means that one must to a great extent understand not only the digital but also the social interaction. Studying distributed CSA, where different team members with different skills and different tasks collaborate towards a common goal, requires that an event can be followed from the start, monitoring how it is perceived, and how understanding is jointly created. There are few studies that show how these processes work and how information is transformed over time.

Third, there is a need to develop CSA measurement techniques and rigorous experimental designs. For CSA research to become an evidence-based research area based on how people actually act, a number of practices and principles need to be developed. Descriptive studies are required to be able to understand how the social processes are formed based on existing information and how they affect how one assumes a position for counterattack. Because such contexts are highly sensitive

and secretive, access may be difficult to obtain for researchers, but such studies would still have great value. Furthermore, experimental systems where teams work against each other under realistic but controlled circumstances are required, as suggested previously (Brynielsson et al., 2016).

Fourth, the existing literature lacks good empirical investigations of whether good CSA actually improves the overall cybersecurity stance or not. While it is very plausible that better decisions are made if informed by better CSA, such a simplistic wording hides details that might matter, for example, how much does long-term, domain-centric, CSA contribute to short-term incident management? And if time is short, as is often the case, are there some indicators or pieces of evidence that are more important—give a higher return on the invested time and mental resources—than others? Empirically based studies of such questions would have great practical value.

Fifth, and finally, the game-theoretic model described in Section 5.4 provides a promising framework for studying the multiple uncertainties that cyber situations entail. It would be valuable to develop it further; first theoretically and then empirically to test its applicability on real problems.

7 Conclusions

This chapter has argued that CSA is best understood by combining three complementary points of view: the technological, the socio-cognitive, and the organizational perspectives. Much of the existing literature focuses mostly on the first of these, perhaps assuming that issues related to cognition and organization are best handled through technical solutions. However, even if this would be true, such technical solutions cannot be built in the first place if the socio-cognitive and organizational perspectives are ignored when requirements are set and prioritized. The reasoning set forth in Sections 2–4 gives ample evidence for why an organization aiming to achieve adequate CSA needs to consider all three perspectives from the very start of the life cycle of any technical support systems.

Cyber incidents can be nonadversarial mistakes and errors, or adversarial attacks. Oftentimes, articles in the literature fail to acknowledge this duality, immediately leaping for one or the other strand. However, both are important, and neither should a priori be ignored—especially since a capable adversary might well hide behind the semblance of inconspicuous errors. Throughout the chapter, we have argued for the importance of maintaining both perspectives.

That said, the presence of an adversary poses particular challenges. "Thinking about thinking" is the title of the first chapter in Heuer's

(1999) seminal work on intelligence analysis—a mandatory must-read throughout the intelligence community. Reasoning about and maintaining an understanding of a thinking adversary have always been central to intelligence, and within the cyber domain, this is more important than ever: in today's online cyber world, it is harder than ever to form a reasonable hypothesis of what is going on, that is, to maintain a high level of CSA. It is also evident from our previous work (Varga et al., 2018, 2021) that this capability is both needed and sought for. To design such a model similar to what has previously been proposed for the intelligence domain (Brynielsson, Horndahl, Kaati, Mårtenson, & Svenson, 2009) and make it part of the cyber analyst's toolbox, is an important research undertaking on our current research agenda.

Acknowledgments

We would like to thank Gazmend Huskaj, Swedish Defence University, for the joint effort in conducting interviews with Locked Shields 2019 cyber defense exercise participants.

Appendix A Interview methodology

To determine the need for knowledge about adversaries for cyber defense analysts, we interviewed 10 Swedish expert participants from the cyber defense exercise Locked Shields 2019. The interviews were held between 1 and 2 months after the exercise. The exact question asked was: *Do you have any use for information about the threat actors themselves?*[a] The concept "threat actor," was neither defined specifically, nor discussed prior to the question. Two interviews were carried out while having two respondents from the same organization in the same room simultaneously. The interview protocol contained several other questions that were not related to the research presented in this chapter.

The annual cyber defense exercise Locked Shields is the world's largest unclassified defensive exercise. It is hosted by the NATO Cooperative Cyber Defence Centre of Excellence. The training audience consists of several so-called blue teams that are tasked to defend critical infrastructure. The exercise is also a competition, and the blue teams are awarded points depending on their performance. The Swedish blue team in 2019 was composed of around 60 cybersecurity experts. They represented both governmental agencies (70%) and commercial companies (30%). The team was highly motivated, and finished in third place out of 24 international competing teams.

[a]Authors' translation from Swedish: *Har du nytta av information om hotaktörer i sig?*

The 10 interviewees all either had leadership positions, or were designated subject matter experts according to the following:

- One was the team leader and one was the deputy team leader of the Swedish blue team.
- One was the chief of staff.
- Four people were group leaders for various functional groups (e.g., Windows, Linux, etc.), which each consisted of a number of personnel.
- One was a staff member responsible for intelligence issues.
- One was a staff member who performed threat-hunting and hash-cracking (via a GPU cluster). This individual was also responsible for real-world logistical support and general technical support.
- One was a technical expert on the domain name system (DNS) and the Apple Mac operating system computing environment.

A.1 Interview results

A synopsis of the answers from the 10 interviewees to the question *Do you have any use for information about the threat actors themselves?* follows:

Respondent 1: Yes, it was useful to hear from the red team after the exercise. It would be good to get the [adversary's] overarching goals and objectives, to understand what they are going to focus on. It would be good to understand their technical capabilities. If they use publicly known code, it would be useful to know which ones. It would even be useful if they use hitherto unknown code. As much threat actor information as possible would have been good, like: "This is how they code their trojans!"; "This is how they are setting up their command and control servers!"; and "This is how they hide themselves!" But, there would probably be issues of secrecy in a real scenario. I don't know how, and if, an intelligence service would share such information, and whether they plan to do so or not…

Respondent 2: It was a game, so that was an artificial situation. Maybe in real life… I would like to know what they are after, and what methods they use. In this case we were given a set of threat actors, and we knew roughly what we could expect. Then we identified additional threat actors. Some might have a use for it [threat actor information]. I didn't.

Respondent 3: I did have use for cyber threat actor information, especially TTPs [tactics, techniques, and procedures]. What are the attackers after? If you know the doctrine of the attackers, e.g., if they employ *maskirovka* [deception], then you can interpret some of the observed attacks as potential smoke screens to protect even more nefarious attacks. We could then prioritize our assets to counter this type of threat. Another thing is to pay attention to what is *not* done at a

specific time; these kind of observations might also reveal something about the motivations of the attackers.

Respondent 4: I did not have any use for that kind of information. We only focused on attackers' IPs, and thought that it would be an easy thing for them to change attack platforms, while remaining hostile. But, it would have been good to have an appreciation of the threat actor's "modus," in order to adapt our behavior.

Respondent 5: We did not get too much information about threat actors or their capabilities [in the exercise]. We had to make some assumptions. We still gathered that an operating principle for the adversary was to distract and have us concentrate our resources on nonessential tasks. The same question at my real job: Yes, information about capabilities, tactics and tools would be of interest. It would be good to know whether an attacker can or cannot perform certain things. It would be good to know their capabilities as well as their level of knowledge, e.g., if and how they can acquire privileges [in accounts] to execute certain things. How they attack us; how they perform lateral movement; where they get a foothold, and understand their perspectives. Where will they find the weakest part of our defense? We could use a bit of reverse thinking here: What would we go after if we were to be the attacker? [In the exercise] we managed to successfully predict what accounts the attackers would attack. So, if you have a good understanding of what the attackers want, it is easier to protect yourself.

Respondent 6: You had to see what was happening to be able to predict what was going to happen next. You had to know what the attacker was after. We would need [intelligence] analyses, and not only "raw data." It is not helpful to only be informed about events; we also need to understand their consequences.

Respondent 7: We had performed an analysis based on earlier exercises. We assumed that the same threat actor was in play, but with partly new "weapons." But what were they after? Threat actor information was not the most important thing for us [at the technical level]; it is more useful for the strategic level. On the technical level, we did not miss so much about the actors.

Respondent 8: We were not good at that! We did not get good info on the threat actors and their capabilities. We made assumptions. Some attacks were done to tie up our resources. It would be good to have info about goals and capabilities of threat actors so we could prioritize our own resources. We also found out that some users were cooperating with the red team. The actual identity of the threat actors is not important.

Respondent 9: The threat actors were not distinct. What will they do (given their capabilities)? Can we counter that? Maybe not that

important at the "technical" level? Attribution is important on the political level, though.

Respondent 10: No, that type of information is not interesting in this scenario. At my normal work, where I know my own system, it would be of use so that I could focus my attention to look out for certain things, and to be proactive. But I would assert that threat actor information is of no use for an ordinary company. I think a very mature organization is required in order to handle such threat information in a good way.

A.2 Interview validity and reliability

The potential pool of interviewees was the total number of Swedish participants in the exercise ($N = 63$). The selection criteria for interviewees were that they had to: (i) have some senior position in the exercise, (ii) have adequate expertise and experience in the field, and (iii) be available for interviews. All interviewees were drawn from the Stockholm region geographical area. Further, all respondents were working with practical network defense measures, that is, they were not working with managerial and leadership issues other than to a small extent.

We did not clarify the meaning of the term "threat actor" in the interviews. The received responses clearly showed how respondents interpreted its meaning in different ways. Rather, several respondents indicated that they immediately assumed that we were asking about (adversarial) tactics, techniques, and procedures, TTPs.

Guest, Bunce, and Johnson (2006) claimed that there virtually did not exist any guidelines for determining nonprobabilistic sample sizes. They then investigated the property of data saturation, that is, how various themes accumulate in a series of interviews in homogeneous groups. They found that saturation occurred within the first 12 interviews, that is, that the number of emergent themes were saturated and that the information value added in the consecutive interviews was negligible. Furthermore, they identified that elements of emerging "metathemes" could be discerned as early as in the sixth interview.

References

Abdul, A., Vermeulen, J., Wang, D., Lim, B. Y., & Kankanhalli, M. (2018). Trends and trajectories for explainable, accountable and intelligible systems: An HCI research agenda. In *Proceedings of the 2018 CHI conference on human factors in computing systems (CHI 2018)* (pp. 1–18). New York, NY: ACM. https://doi.org/10.1145/3173574.3174156.

Ahmad, A., Desouza, K. C., Maynard, S. B., Naseer, H., & Baskerville, R. L. (2020). How integration of cyber security management and incident response enables organizational learning. *Journal of the Association for Information Science and Technology*, *71*(8), 939–953. https://doi.org/10.1002/asi.24311.

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., … Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, *58*, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012.

Artman, H. (2000). Team situation assessment and information distribution. *Ergonomics*, *43*(8), 1111–1128. https://doi.org/10.1080/00140130050084905.

Artman, H., Brynielsson, J., Johansson, B. J. E., & Trnka, J. (2011). Dialogical emergency management and strategic awareness in emergency communication. In *Proceedings of the eighth international conference on information systems for crisis response and management (ISCRAM 2011)* (pp. 1–9). http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-61800.

Artman, H., & Wærn, Y. (1999). Distributed cognition in an emergency co-ordination center. *Cognition, Technology & Work*, *1*(4), 237–246. https://doi.org/10.1007/s101110050020.

Attiah, A., Chatterjee, M., & Zou, C. C. (2018). A game theoretic approach to model cyber attack and defense strategies. In *Proceedings of the 2018 IEEE international conference on communications (ICC 2018)* (pp. 1–7). Piscataway, NJ: IEEE. https://doi.org/10.1109/ICC.2018.8422719.

Bass, T. (2000). Intrusion detection systems and multisensor data fusion. *Communications of the ACM*, *43*(4), 99–105. https://doi.org/10.1145/332051.332079.

Bernier, M. (2013). *Military activities and cyber effects (MACE) taxonomy.* Technical Memorandum DRDC CORA TM 2013-226. Ottawa, Canada: Centre for Operational Research and Analysis, Defence Research and Development Canada. https://cradpdf.drdc-rddc.gc.ca/PDFS/unc218/p803340_A1b.pdf.

Bertolino, A. (2007). Software testing research: Achievements, challenges, dreams. In *Proceedings of future of software engineering (FoSE 2007)* (pp. 85–103). Piscataway, NJ: IEEE. https://doi.org/10.1109/FOSE.2007.25.

Borum, R., Felker, J., Kern, S., Dennesen, K., & Feyes, T. (2015). Strategic cyber intelligence. *Information and Computer Security*, *23*(3), 317–332. https://doi.org/10.1108/ICS-09-2014-0064.

Brown, S., Gommers, J., & Serrano, O. (2015). From cyber security information sharing to threat management. In *Proceedings of the 2015 ACM workshop on information sharing and collaborative security (WISCS 2015)* (pp. 43–49). New York, NY: ACM. https://doi.org/10.1145/2808128.2808133.

Brynielsson, J. (2006). *A gaming perspective on command and control* (Doctoral dissertation). Stockholm, Sweden: Royal Institute of Technology. http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-4029.

Brynielsson, J., & Arnborg, S. (2006). An information fusion game component. *Journal of Advances in Information Fusion*, *1*(2), 108–121. http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-24204.

Brynielsson, J., Franke, U., Tariq, M. A., & Varga, S. (2016). Using cyber defense exercises to obtain additional data for attacker profiling. In *Proceedings of the 2016 IEEE conference on intelligence and security informatics (ISI 2016)* (pp. 37–42). Piscataway, NJ: IEEE. https://doi.org/10.1109/ISI.2016.7745440.

Brynielsson, J., Franke, U., & Varga, S. (2016). Cyber situational awareness testing. In B. Akhgar, & B. Brewster (Eds.), *Combatting cybercrime and cyberterrorism: Challenges, trends and priorities* (pp. 209–233). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-38930-1_12.

Brynielsson, J., Horndahl, A., Kaati, L., Mårtenson, C., & Svenson, P. (2009). Development of computerized support tools for intelligence work. In *Proceedings of the 14th international command and control research and technology symposium (14th ICCRTS)*. Washington, DC: US Department of Defense CCRP. http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-89594.

Burger, E. W., Goodman, M. D., Kampanakis, P., & Zhu, K. A. (2014). Taxonomy model for cyber threat intelligence information exchange technologies. In *Proceedings of the 2014 ACM workshop on information sharing and collaborative security (WISCS 2014)* (pp. 51–60). New York, NY: ACM. https://doi.org/10.1145/2663876.2663883.

Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.

Caralli, R. A., Stevens, J. F., Young, L. R., & Wilson, W. R. (2007). *Introducing OCTAVE Allegro: Improving the information security risk assessment process*. Technical Report CMU/SEI-2007-TR-012. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University. https://doi.org/10.1184/R1/6574790.v1.

Cebula, J. J., Popeck, M. E., & Young, L. R. (2014). *A taxonomy of operational cyber security risks version 2*. Technical Note CMU/SEI-2014-TN-006. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University. https://doi.org/10.1184/R1/6571784.v1.

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. https://doi.org/10.1109/MSP.2014.85.

Copeland, J. (2008). *Emergency response: Unity of effort through a common operational picture*. Strategy Research Project. Carlisle, PA: US Army War College. https://apps.dtic.mil/sti/citations/ADA479583.

Cottrell, A., & Cockshott, W. P. (1993). Calculation, complexity and planning: The socialist calculation debate once again. *Review of Political Economy*, *5*(1), 73–112. https://doi.org/10.1080/09538259300000005.

Dehghantanha, A., Conti, M., & Dargahi, T. (Eds.). (2018). *Cyber threat intelligence*. Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-73951-9.

Demir, M., McNeese, N. J., & Cooke, N. J. (2017). Team situation awareness within the context of human-autonomy teaming. *Cognitive Systems Research*, *46*, 3–12. https://doi.org/10.1016/j.cogsys.2016.11.003.

Dörner, D., & Schaub, H. (1994). Errors in planning and decision-making and the nature of human information processing. *Applied Psychology*, *43*(4), 433–453. https://doi.org/10.1111/j.1464-0597.1994.tb00839.x.

Du, M., Liu, N., & Hu, X. (2019). Techniques for interpretable machine learning. *Communications of the ACM*, *63*(1), 68–77. https://doi.org/10.1145/3359786.

Dunbar, T. (2012). The first steps to managing cyber-risk. *Risk Management*, *59*(8). http://www.rmmagazine.com/articles/article/2012/10/05/-The-First-Steps-to-Managing-Cyber-Risk-.

Dutta, A., & McCrohan, K. (2002). Management's role in information security in a cyber economy. *California Management Review*, *45*(1), 67–87. https://doi.org/10.2307/41166154.

Endsley, M. R. (1988). Situation awareness global assessment technique (SAGAT). In *Proceedings of the IEEE 1988 national aerospace and electronics conference (NAECON 1988): Vol. 3* (pp. 789–795). Piscataway, NJ: IEEE. https://doi.org/10.1109/NAECON.1988.195097.

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, *37*(1), 32–64. https://doi.org/10.1518/001872095779049543.

Endsley, M. R., & Jones, D. G. (2011). *Designing for situation awareness: An approach to user-centered design*. Boca Raton, FL: CRC Press. https://doi.org/10.1201/b11371.

Evancich, N., Lu, Z., Li, J., Cheng, Y., Tuttle, J., & Xie, P. (2014). Network-wide awareness. In A. Kott, C. Wang, & R. F. Erbacher (Eds.), *Cyber defense and situational awareness* (pp. 63–91). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-11391-3_5.

Franke, U. (2018). Cyber insurance against electronic payment service outages. In *Proceedings of the 14th international workshop on security and trust management (STM 2018)* (pp. 73–84). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-030-01141-3_5.

Franke, U., & Brynielsson, J. (2014). Cyber situational awareness: A systematic review of the literature. *Computers & Security*, *46*, 18–31. https://doi.org/10.1016/j.cose.2014.06.008.

Franke, U., & Wernberg, J. (2020). A survey of cyber security in the Swedish manufacturing industry. In *Proceedings of the 2020 international conference on cyber situational awareness, data analytics and assessment (CyberSA 2020)* (pp. 1–8). Piscataway, NJ: IEEE. https://doi.org/10.1109/CyberSA49311.2020.9139673.

Galup, S. D., Dattero, R., Quan, J. J., & Conger, S. (2009). An overview of IT service management. *Communications of the ACM*, *52*(5), 124–127. https://doi.org/10.1145/1506409.1506439.

Genadis, T. C. (1996). A cost optimization model for determining optimal burn-in times at the module/system level of an electronic product. *International Journal of Quality & Reliability Management*, *13*(9), 61–74. https://doi.org/10.1108/02656719610150623.

Goeschel, K. (2016). Reducing false positives in intrusion detection systems using data-mining techniques utilizing support vector machines, decision trees, and naive Bayes for off-line analysis. In *IEEE SoutheastCon 2016* (pp. 1–6). Piscataway, NJ: IEEE. https://doi.org/10.1109/SECON.2016.7506774.

Gray, J. (1985). *Why do computers stop and what can be done about it?* Technical Report 85.7. Cupertino, CA: Tandem Computers, Inc.

Gray, J. (1990). A census of Tandem system availability between 1985 and 1990. *IEEE Transactions on Reliability*, *39*(4), 409–418. https://doi.org/10.1109/24.58719.

Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, *18*(1), 59–82. https://doi.org/10.1177/1525822X05279903.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, *51*(5), 1–42. https://doi.org/10.1145/3236009.

Gunawi, H. S., Hao, M., Suminto, R. O., Laksono, A., Satria, A. D., Adityatama, J., & Eliazar, K. J. (2016). Why does the cloud stop computing? Lessons from hundreds of service outages. In *Proceedings of the seventh ACM symposium on cloud computing (SoCC 2016)* (pp. 1–16). New York, NY: ACM. https://doi.org/10.1145/2987550.2987583.

Gutzwiller, R. S., Dykstra, J., & Payne, B. (2020). Gaps and opportunities in situational awareness for cybersecurity. *Digital Threats: Research and Practice*, *1*(3), 1–6. https://doi.org/10.1145/3384471.

Gutzwiller, R. S., Hunt, S. M., & Lange, D. S. (2016). A task analysis toward characterizing cyber-cognitive situation awareness (CCSA) in cyber defense analysts. In *Proceedings of the 2016 IEEE international multi-disciplinary conference on cognitive methods in situation awareness and decision support (CogSIMA 2016)* (pp. 14–20). Piscataway, NJ: IEEE. https://doi.org/10.1109/COGSIMA.2016.7497780.

Hager, R. S. (1997). *Current and future efforts to vary the level of detail for the common operational picture* (Master's thesis). Monterey, CA: Naval Postgraduate School. https://apps.dtic.mil/sti/citations/ADA341674.

Harsanyi, J. C. (1967–1968). Games with incomplete information played by "Bayesian" players. *Management Science*, *14*(3, 5, 7), 159–182, 320–334, 486–502. https://doi.org/10.1287/mnsc.14.3.159, https://doi.org/10.1287/mnsc.14.5.320, https://doi.org/10.1287/mnsc.14.7.486.

Hayek, F. A. (1982). Two pages of fiction: The impossibility of socialist calculation. *Economic Affairs*, *2*(3), 135–142. https://doi.org/10.1111/j.1468-0270.1982.tb01416.x.

Heuer, R. J., Jr. (1999). *Psychology of intelligence analysis*. Washington, DC: Center for the Study of Intelligence, Central Intelligence Agency.

Holm, H. (2014). Signature based intrusion detection for zero-day attacks: (Not) a closed chapter? In *Proceedings of the 2014 47th annual Hawaii international conference on system sciences (HICSS 2014)* (pp. 4895–4904). Piscataway, NJ: IEEE. https://doi.org/10.1109/HICSS.2014.600.

Iden, J., & Eikebrokk, T. R. (2013). Implementing IT service management: A systematic literature review. *International Journal of Information Management*, *33*(3), 512–523. https://doi.org/10.1016/j.ijinfomgt.2013.01.004.

Kang, M.-J., & Kang, J.-W. (2016). Intrusion detection system using deep neural network for in-vehicle network security. *PLoS ONE*, *11*(6), 1–17. https://doi.org/10.1371/journal.pone.0155781.

Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. *Cybersecurity*, *2*(1), 1–22. https://doi.org/10.1186/s42400-019-0038-7.

Kim, J., Shin, N., Jo, S. Y., & Kim, S. H. (2017). Method of intrusion detection using deep neural network. In *Proceedings of the 2017 IEEE international conference on big data and smart computing (BigComp 2017)* (pp. 313–316). Piscataway, NJ: IEEE. https://doi.org/10.1109/BIGCOMP.2017.7881684.

Kramer, R. M. (1998). Revisiting the Bay of Pigs and Vietnam decisions 25 years later: How well has the groupthink hypothesis stood the test of time? *Organizational Behavior and Human Decision Processes*, *73*(2–3), 236–271. https://doi.org/10.1006/obhd.1998.2762.

Krombholz, K., Hobel, H., Huber, M., & Weippl, E. (2015). Advanced social engineering attacks. *Journal of Information Security and Applications*, *22*, 113–122. https://doi.org/10.1016/j.jisa.2014.09.005.

Li, Y., Huang, G.-q., Wang, C.-z., & Li, Y.-c. (2019). Analysis framework of network security situational awareness and comparison of implementation methods. *EURASIP Journal on Wireless Communications and Networking*, *2019*(205), 1–32. https://doi.org/10.1186/s13638-019-1506-1.

Maennel, K., Ottis, R., & Maennel, O. (2017). Improving and measuring learning effectiveness at cyber defense exercises. In *Proceedings of the 22nd Nordic conference on secure IT systems (NordSec 2017)* (pp. 123–138). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-70290-2_8.

Malik, B., & Scott, D. (2010). *How to calculate the cost of continuously available IT services*. Gartner Research Paper G00209007. Stamford, CT: Gartner, Inc.

Mattern, T., Felker, J., Borum, R., & Bamford, G. (2014). Operational levels of cyber intelligence. *International Journal of Intelligence and CounterIntelligence*, *27*(4), 702–719. https://doi.org/10.1080/08850607.2014.924811.

Mauw, S., & Oostdijk, M. (2005). Foundations of attack trees. In *Proceedings of the eighth international conference on information security and cryptology (ICISC 2005)* (pp. 186–198). Berlin/Heidelberg, Germany: Springer. https://doi.org/10.1007/11734727_17.

Mavroeidis, V., & Bromander, S. (2017). Cyber threat intelligence model: An evaluation of taxonomies, sharing standards, and ontologies within cyber threat intelligence. In *Proceedings of the 2017 European intelligence and security informatics conference (EISIC 2017)* (pp. 91–98). Piscataway, NJ: IEEE. https://doi.org/10.1109/EISIC.2017.20.

McCrohan, K. F., Engel, K., & Harvey, J. W. (2010). Influence of awareness and training on cyber security. *Journal of Internet Commerce*, *9*(1), 23–41. https://doi.org/10.1080/15332861.2010.487415.

McKenna, S., Staheli, D., & Meyer, M. (2015). Unlocking user-centered design methods for building cyber security visualizations. In *Proceedings of the 2015 IEEE symposium on visualization for cyber security (VizSec 2015)* (pp. 1–8). Piscataway, NJ: IEEE. https://doi.org/10.1109/VIZSEC.2015.7312771.

Mohaisen, A., Al-Ibrahim, O., Kamhoua, C., Kwiat, K., & Njilla, L. (2017). Rethinking information sharing for threat intelligence. In *Proceedings of the fifth ACM/IEEE workshop on hot topics in web systems and technologies (HotWeb 2017)* (pp. 1–7). New York, NY: ACM. https://doi.org/10.1145/3132465.3132468.

Myerson, R. B. (1991). *Game theory: Analysis of conflict*. Cambridge, MA: Harvard University Press.

Nature. (2016). More accountability for big-data algorithms. *Nature*, *537*(7621), 449. https://doi.org/10.1038/537449a.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, *2*(2), 175–220. https://doi.org/10.1037/1089-2680.2.2.175.

Omer, M., Nilchiani, R., & Mostashari, A. (2009). Measuring the resilience of the global internet infrastructure system. In *Proceedings of the 2009 third annual IEEE international systems*

conference (SysCon 2009) (pp. 156–162). Piscataway, NJ: IEEE. https://doi.org/10.1109/SYSTEMS.2009.4815790.

Petersen, R., Santos, D., Smith, M. C., Wetzel, K. A., & Witte, G. (2020). *Workforce framework for cybersecurity (NICE framework)*. NIST Special Publication 800-181 Revision 1. National Institute of Standards and Technology, US Department of Commerce. https://doi.org/10.6028/NIST.SP.800-181r1.

Raven, B. H. (1998). Groupthink, Bay of Pigs, and Watergate reconsidered. *Organizational Behavior and Human Decision Processes*, *73*(2–3), 352–361. https://doi.org/10.1006/obhd.1998.2766.

Rosenquist, M. (2009). *Prioritizing information security risks with threat agent risk assessment*. IT@Intel White Paper. Intel Corporation.

Roy, S. S., Mallik, A., Gulati, R., Obaidat, M. S., & Krishna, P. V. (2017). A deep learning based artificial neural network approach for intrusion detection. In *Proceedings of the third international conference on mathematics and computing (ICMC 2017)* (pp. 44–53). Singapore: Springer. https://doi.org/10.1007/978-981-10-4642-1_5.

Sauerwein, C., Pekaric, I., Felderer, M., & Breu, R. (2019). An analysis and classification of public information security data sources used in research and practice. *Computers & Security*, *82*, 140–155. https://doi.org/10.1016/j.cose.2018.12.011.

Sauerwein, C., Sillaber, C., Mussmann, A., & Breu, R. (2017). Threat intelligence sharing platforms: An exploratory study of software vendors and research perspectives. In *Proceedings of the 13th international conference on wirtschaftsinformatik (WI 2017)* (pp. 837–851). Atlanta, GA: Association for Information Systems (AIS).

Schelling, T. C. (1960). *The strategy of conflict*. Cambridge, MA: Harvard University Press.

Schneier, B. (1999). Attack trees. *Dr. Dobb's Journal*, *24*(12), 21–29.

Shin, B., & Lowry, P. B. (2020). A review and theoretical explanation of the 'Cyberthreat-Intelligence (CTI) capability' that needs to be fostered in information security practitioners and how this can be accomplished. *Computers & Security*, *92*, 1–16. https://doi.org/10.1016/j.cose.2020.101761.

Shostack, A. (2014). *Threat modeling: Designing for security*. Indianapolis, IN: Wiley.

Sommestad, T., & Franke, U. (2015). A test of intrusion alert filtering based on network information. *Security and Communication Networks*, *8*, 2291–2301. https://doi.org/10.1002/sec.1173.

Sommestad, T., & Hunstad, A. (2013). Intrusion detection and the role of the system administrator. *Information Management & Computer Security*, *21*(1), 30–40. https://doi.org/10.1108/09685221311314400.

Spathoulas, G. P., & Katsikas, S. K. (2010). Reducing false positives in intrusion detection systems. *Computers & Security*, *29*(1), 35–44. https://doi.org/10.1016/j.cose.2009.07.008.

Tariq, M. A., Brynielsson, J., & Artman, H. (2012). Framing the attacker in organized cybercrime. In *Proceedings of the 2012 European intelligence and security informatics conference (EISIC 2012)* (pp. 30–37). Piscataway, NJ: IEEE. https://doi.org/10.1109/EISIC.2012.48.

Tariq, M. A., Brynielsson, J., & Artman, H. (2014). The security awareness paradox: A case study. In *Proceedings of the 2014 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM 2014)* (pp. 704–711). Piscataway, NJ: IEEE. https://doi.org/10.1109/ASONAM.2014.6921663.

Tounsi, W., & Rais, H. (2018). A survey on technical threat intelligence in the age of sophisticated cyber attacks. *Computers & Security*, *72*, 212–233. https://doi.org/10.1016/j.cose.2017.09.001.

Traoré, I., Awad, A., & Woungang, I. (Eds.). (2017). *Information security practices: Emerging threats and perspectives*. Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-48947-6.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124.

Varga, S., Brynielsson, J., & Franke, U. (2018). Information requirements for national level cyber situational awareness. In *Proceedings of the 2018 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM 2018)* (pp. 774–781). Piscataway, NJ: IEEE. https://doi.org/10.1109/ASONAM.2018.8508410.

Varga, S., Brynielsson, J., & Franke, U. (2021). Cyber-threat perception and risk management in the Swedish financial sector. *Computers & Security*, *105*, 1–18. https://doi.org/10.1016/j.cose.2021.102239.

Vielberth, M., Menges, F., & Pernul, G. (2019). Human-as-a-security-sensor for harvesting threat intelligence. *Cybersecurity*, *2*(1), 1–15. https://doi.org/10.1186/s42400-019-0040-0.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

Wagner, T. D., Mahbub, K., Palomar, E., & Abdallah, A. E. (2019). Cyber threat intelligence sharing: Survey and research directions. *Computers & Security*, *87*, 1–13. https://doi.org/10.1016/j.cose.2019.101589.

Wan, Y., Cao, J., Chen, G., & Huang, W. (2017). Distributed observer-based cyber-security control of complex dynamical networks. *IEEE Transactions on Circuits and Systems I: Regular Papers*, *64*(11), 2966–2975. https://doi.org/10.1109/TCSI.2017.2708113.

Werlinger, R., Hawkey, K., Muldner, K., Jaferian, P., & Beznosov, K. (2008). The challenges of using an intrusion detection system: Is it worth the effort? In *Proceedings of the fourth symposium on usable privacy and security (SOUPS 2008)* (pp. 107–118). New York, NY: ACM. https://doi.org/10.1145/1408664.1408679.

Whittaker, J. A. (2000). What is software testing? And why is it so hard? *IEEE Software*, *17*(1), 70–79. https://doi.org/10.1109/52.819971.

Wolbers, J., & Boersma, K. (2013). The common operational picture as collective sensemaking. *Journal of Contingencies and Crisis Management*, *21*(4), 186–199. https://doi.org/10.1111/1468-5973.12027.

Xu, S. (2014). Cybersecurity dynamics. In *Proceedings of the 2014 symposium and bootcamp on the science of security (HoTSoS 2014)* (pp. 1–2). New York, NY: ACM. https://doi.org/10.1145/2600176.2600190.

Yeager, L. B. (1994). Mises and Hayek on calculation and knowledge. *The Review of Austrian Economics*, *7*(2), 93–109. https://doi.org/10.1007/BF01101944.

Zheng, D. E., & Lewis, J. A. (2015). *Cyber threat information sharing: Recommendations for congress and the administration*. A Report of the CSIS Strategic Technologies Program. Washington, DC: Center for Strategic & International Studies. https://www.csis.org/analysis/cyber-threat-information-sharing.