

Re-thinking Decision-Making in Cybersecurity: Leveraging Cognitive Heuristics in Situations of Uncertainty

Thierry A. Schaltegger, Benjamin Ambuehl, Kurt A. Ackermann, Nico Ebert
ZHAW School of Management and Law, Switzerland
firstname.lastname@zhaw.ch

Abstract

The prevailing consensus in cybersecurity is that individuals' insecure behavior due to inadequate decision-making is a primary source of cyber incidents. The conclusion of this assumption is to enforce desired behavior via extensive security policies and suppress individuals' intuitions or rules of thumb (cognitive heuristics) when dealing with critical situations. This position paper aims to change the way we look at these cognitive heuristics in cybersecurity. We argue that heuristics can be particularly useful in uncertain environments such as cybersecurity. Based on successful examples from other domains, we propose that heuristic decision-making should also be used to combat cyber threats. Lastly, we give an outlook on where such heuristics could be beneficial in cybersecurity (e.g., phishing detection or incident response) and how they can be found or created.

Keywords: cybersecurity decision-making, cognitive heuristics, intuition, uncertainty, ecological rationality

1. Introduction

Following the logic that humans are the “weakest link in cybersecurity” (Schneier, 2000), employees have to comply with an increasing number of constraining security policies and regulations (Fuster & Jasmontaite, 2020) and must attend regular training with questionable efficacy to keep the company secure (Herley, 2009). Considering the limited effectiveness of this approach, it is pressing to question its underlying assumptions (Alsharnouby et al., 2015; Butavicius et al., 2022).

Human cognition, e.g., judgment and decision-making, is typically perceived as a weakness in cybersecurity (Zimmermann & Renaud, 2019). These shortcomings are often explained through people's reliance on cognitive heuristics – mental shortcuts and simple rules of thumb -which may offer an intuitive explanation for many things that go wrong in cybersecurity (Tsohou et al., 2015). Generally speaking, heuristics are defined as decision processes that do not consider all available information but focus

forms. One example is the availability heuristic, where one gauges the frequency of an event by the ease of recalling it (Tversky & Kahneman, 1974). This may work most of the time but fails if the event in question is overrepresented in the media or otherwise easier to recall, leading one to overestimate its frequency. In the case of cybersecurity, it is often assumed that people fall for phishing emails because they rely on unreliable or misleading cues instead of consciously and systematically processing every single email (Frauenstein & Flowerday, 2020) or because their biases lead them to underestimate threats or overestimate their abilities to detect malicious emails (Kwak et al., 2020). The solution following this assumption is to establish strict guidelines for the use of digital tools and train employees to suppress heuristic processing, and systematically process every single email (Butavicius et al., 2022).

However, this approach to combat cybersecurity risks faces two problems. First, while strict rules and training make sense in a stable environment, they are insufficient in an unstable and dynamic environment (Donaldson, 2001). Therefore, they need to permanently change if one wants to account for every possible incident. This problem is exacerbated in the domain of cybersecurity, where one has not only to deal with random incidents but with motivated opponents that react to one's security efforts and actively try to circumvent them. The attack strategies used by malicious actors are dynamic and change on an almost weekly basis (Sundaramurthy et al., 2017). Effective tools to counter these attacks need to be highly adaptive to be able to cope with these evolving threats.

The second problem lies in the disregard for economic aspects. Every rule and every additional hour of training brings direct and indirect costs that must be accounted for and conflict with other goals of the organization and its employees (e.g., productivity, autonomy). While, in theory, one may achieve complete security through a near infinite amount of rules and regulations that are continuously updated to preemptively defend against every possible attack (at least temporarily), this would incur costs that most likely supersede the damage of an intrusion. Only a few studies mention this trade-off and account for the economic dimensions of security training (Canfield & Fichtloff, 2018; Herley, 2009). Considering that

OPEN ACCESS. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

research in cybersecurity aims to deliver practical answers, researchers ought to consider the economic aspects of their proposed measures to increase security.

Heuristic decision-making may be the answer to both aforementioned problems. Heuristics are adaptive and allow decision-makers to better cope with an evolving threat landscape than more elaborate but strict guidelines, as they are less specific and focus on more general characteristics. There is an argument to be made that - in situations of uncertainty - heuristics are even more accurate than more complex strategies (Czerlinski et al., 1999). Additionally, they are economical as they rely on only a few cues to make fast decisions and are easy to acquire (Gigerenzer & Brighton, 2009; Tversky & Kahneman, 1974).

With this work, we want to make the following argument:

- I) In uncertain environments, simple heuristics are more economical and sometimes even more accurate than more complex strategies
- II) Cybersecurity is an uncertain environment
- III) Heuristics should be embraced rather than avoided

We will first give a brief overview of how heuristics shape our decision-making. This overview will be substantiated with real-world applications where heuristic decision-making proved advantageous. Lastly, we will hypothesize where heuristics are likely to be advantageous in dealing with different cybersecurity problems.

2. Cognitive heuristics & biases

A heuristic is a decision strategy that ignores part of the available information to reach faster conclusions than more resource-intensive and complex processes like analytical reasoning (Gigerenzer & Brighton, 2009). Heuristics can be tacit in the form of intuitions and gut feelings or explicit as simple rules of thumb or decision trees. The prevailing view in psychology and behavioral economics is that the human mind relies on the use of simple heuristics to cope with its limited resources for dealing with the world's complexities, leading to *bounded rationality* (Simon, 1955). While the concept of heuristic problem solving has been introduced earlier (Pólya, 1954; Simon & Newell, 1958), Tversky and Kahneman managed to bring this discourse into the mainstream in the 1970s with their groundbreaking research on cognitive heuristics and biases (Tversky & Kahneman, 1974). Their work challenged longstanding assumptions about human decision-making and rationality and shaped the discourse in psychology and economics for the following decades (MacCoun, 2002). While their

heuristics and biases program acknowledges the general effectiveness of heuristics, especially where quick decisions under uncertainty are required, it focuses on cases where they *fail* to lead to accurate outcomes and conceptualizes them as the main reason for *deviations from rational behavior*. Tversky and Kahneman propose that heuristics lead to distorted perceptions and suboptimal decisions, further referred to as *biases*. It is argued that biases would lead to systematic and predictable errors and, therefore, to bad choices, with adverse outcomes for health, wealth, and happiness (Thaler & Sunstein, 2008). Much of the subsequent scientific and popular literature adopted the negative framing of heuristics as irrational while ignoring their adaptive value. The rise to prominence of Kahneman and Tversky's work spurred a hunt for new, actual, or imagined biases that is still ongoing (Gigerenzer, 2018).

The most prominent opponent of this grim view of human decision-making capabilities is Gerd Gigerenzer, who argues for a more positive evaluation of heuristics and intuitive decision-making (Gigerenzer & Brighton, 2009). His main contention lies in the widely used definition of rationality, usually conceptualized as purely mathematical utility maximization, where the optimal choice is derived from a complex statistical model incorporating all available information. While in the common experiments to test decision-making, all relevant outcomes and their corresponding probabilities are known, this is rarely the case outside the laboratory. Decision-making scenarios in the real world are often much more complex and uncertain, such that decision-makers must deal with unreliable or incomplete information. Therefore, human decision-making is measured against an unrealistic benchmark of rationality that is hardly ever observed in the real world and focuses only on maximizing the expected utility of the outcome, ignoring search and processing costs. In the real world, searching for information costs time and/or money, leading to the diminishing utility of new information. Rationality should not be seen as something purely logical but *ecological*: adapted to a specific environment (Todd & Gigerenzer, 2012). To assess *ecological rationality*, the right question to ask is: "How well does a decision-making process perform *in a given environment*?"

Gigerenzer even argues that heuristics do not merely make a tradeoff between information and its associated costs and accuracy but that they can lead to equivalent or even more accurate inferences than other strategies in the right environment, highlighting their ecological rationality (Gigerenzer & Brighton, 2009). This claim has been substantiated through a comprehensive empirical investigation. For example, the effectiveness of heuristic processing, which relied on a single cue, has been extensively compared with complex models that incorporated all the available information. The comparison was conducted across 2035

different environments, varying from predicting fish fertility to fuel consumption and high school dropout rates, and a simple heuristic achieved the highest accuracy (Czerlinski et al., 1999).

These findings demonstrate that heuristics are not necessarily always the second-best solution, only employed because of cognitive shortcomings, and that more information, time, and computational power is not always better.

2.1. Fast & frugal heuristics

A distinct class of simple heuristics are fast-and-frugal decision trees (FFT) (L. Martignon et al., 2005, 2008). These heuristics are simple decision aids in the form of a binary decision tree that facilitate efficient and accurate human decisions with limited information. These decision trees are fast because they can be applied quickly as they rely on only a few relevant cues without requiring any computation, and they are frugal because they ignore most of the available information and focus only on the variables with the highest discriminatory power. Because FFTs are simple and rely on only a few cues, they are particularly useful in high-stress environments where limited information is available and where speed is of high importance. In contrast to complex decision algorithms, such as regression, FFTs can be easily understood and learned by humans and allow people to make quick and accurate decisions on the fly without the need for statistical training or technical tools. A fast and frugal heuristic is regularly applied by health emergency first responders to triage patients. Unconscious patients that are still breathing are placed in a recovery position for further assessment. In the absence of breathing, cardiopulmonary resuscitation is initiated immediately.

There are two approaches to creating FFTs: relying on data or on human expertise. The creation of such a decision tree relies on the identification of core variables that predict an outcome, are easy to measure, and are unlikely to change. This task needs a deep understanding of the problem and can be done either with the support of experts in the particular domain or with the help of supervised learning algorithms to find the variables with the highest predictive power (Phillips et al., 2017). The second approach needs a reliable database containing all the relevant variables to draw inferences. How realistic this is relies heavily on the specific context, the structure and the complexity of the problem and the data availability. If this data is not available, or the observable variables are not easily quantifiable, one may draw on the experience of experts to devise an FFT. Their simplicity not only assures transparency but makes them easy to understand and implement for non-experts, making them an excellent tool for knowledge transfer.

2.2. Where are heuristics successfully applied?

Successful heuristics exploit the structure of the environment they are adapted to (Todd, 2007). In what environment are heuristics ecologically rational, and when should we rather rely on complex models or intricate rules to guide one's decisions? Table 1 presents an overview of various different domains that embody uncertain environments where the successful application of heuristics has been demonstrated.

Even in the highly regulated and standardized medical field, heuristics have been shown to perform well in various applications. Simple heuristics have been shown to be both fast and accurate in diagnosing depression (Jenny et al., 2013) and categorizing heart and other diseases (Green & Mehr, 1997; L. F. Martignon et al., 2012; Phillips et al., 2017). These insights led to the proposition that medical education should embrace heuristics rather than train them away (Feufel & Flach, 2019) and that more clinical practice guidelines should be transformed into FFTs (Djulgovic et al., 2018).

A good, although counterintuitive, example of an environment that is better described by uncertainty than by risk is the financial system. The increasing complexity of this system invited the use of increasingly complex modeling and risk management strategies to understand and control it. Neither of which was very successful in predicting future developments, looking back at the last financial crisis or year-ahead predictions of the euro-dollar exchange rate (Gigerenzer, 2014). Simple heuristics have been shown to outperform complex modeling approaches for calculating banks' capital requirements and in predicting the failure of individual banks during the global financial crisis (Aikman et al., 2021).

Some other environments where heuristics and simple models have demonstrated their potential are: predicting recidivism in criminals (Dressel & Farid, 2018), predicting bail decisions of judges (Dhimi, 2003), detecting deception (Verschuere et al., 2023), modeling climate (Halide & Ridd, 2008), forecasting in the domains of weather, sports, crime, and business (Goldstein & Gigerenzer, 2009), and even reducing civilian casualties in military operations (Keller & Katsikopoulos, 2016). A great example to demonstrate the need to thoroughly analyze the environment in which they are applied is documented in the management literature. Simple heuristics performed better than more complex approaches in predicting individual customer behavior, but stochastic models delivered the best insights on the aggregate level (Wübben & Wangenheim, 2008). Similar findings are reported for car pricing, where a complex model best describes the aggregate market, but the use of heuristic pricing generated the most profits in an uncertain market (Artinger & Gigerenzer, 2016). These findings demonstrate that heuristics perform well in various

Table 1. Examples of the successful application of simple heuristics

Domain	Application of Cognitive Heuristic	Outcomes / Benefits
Medicine	• Triaging (Super, 1984)	Rapid assessment of victims in mass casualty incidents relying on a few simple cues.
	• Depression diagnostics (Jenny et al., 2013)	Simple heuristic was competitive with complex model while relying mostly on a single cue.
	• Medical prescription (Fischer et al., 2002)	Fast-and-frugal trees (FFT) performed close to the complex scoring system while being faster and more transparent.
	• Care unit admission (Wegwarth et al., 2009)	FFT is both faster and more accurate than complex tools and doctors' intuition.
	• Diagnosing different diseases (Phillips et al., 2017)	FFT was as or more accurate with less information than different complex models.
Finances	• Predicting bank failure (Aikman et al., 2021)	FFT outperforms logistic regression, especially with limited information.
	• Investing decisions of venture capitalists (Woike et al., 2015)	The simple strategy was competitive with complex ones while needing less information.
	• Investing: simple heuristic vs. different complex models (DeMiguel et al., 2009)	No model consistently outperformed a simple heuristic across different datasets.
	• Investment with simple heuristic vs. market and different portfolios (Borges et al., 1999)	Heuristic outperformed the stock market, individual traders, and big funds.
Legal	• Predicting recidivism (Dressel & Farid, 2018)	Not a heuristic in the strict sense, but a model with two cues achieved the same accuracy as one with 137 cues.
Business	• Customer base analysis (Wübben & Wangenheim, 2008)	Statistics performed well for aggregate analysis, but heuristics predicted individual customer purchasing behavior better.
Military	• Detecting unexploded ordnance (Fernandez et al., 2010)	Heuristic outperformed or matched regression and machine learning with a 100% hit rate and a 3% false alarm rate.
	• FFT to classify hostility of approaching cars (Keller & Katsikopoulos, 2016)	Potential reduction of civilian casualties in military operations by 60%.
	• FFT to support decision-making in combat situations (Banks et al., 2020)	FFT was used to transfer knowledge from senior to junior officers and reduced mental load in stressful situations.
Deception detection	• Detecting deception in writing, video, or live interviews (Verschuere et al., 2023)	Reliance on a single cue led to lie detection well above chance (59-79%).

environments but that we cannot apply them blindly and have to evaluate their validity for the different environments we want to employ them.

3. The potential of cognitive heuristics in cybersecurity

Many of the environments where heuristics succeeded share similarities with cybersecurity regarding their unpredictability and the role that time pressure plays in resolving incidents. In this chapter, we want to emphasize these similarities and look at specific examples from cybersecurity where cognitive heuristics may be beneficial.

3.1. Why leverage heuristics in cybersecurity?

In computer science, heuristic algorithms allow us to deal with processing constraints in computational tasks when a problem is too complex or too large to be solved optimally in a reasonable timeframe (Rothlauf, 2011). Similar to cognitive heuristics, these heuristic algorithms are not guaranteed to find the best solution,

they usually provide a satisfactory solution in environments that cannot be comprehensively mapped out and retain an element of uncertainty.

The concept of *cognitive* heuristics in decision-making has later been introduced to computer science, and specifically cybersecurity (e.g., Pfleeger & Caputo, 2012). However, the cybersecurity literature seems to have adopted the mainly negative frame from modern behavioral economics, where heuristics are associated with error-prone “mental software,” leading to systematic blunders and biases (Thaler & Sunstein, 2008). Whereas Kahneman and Tversky emphasized the utility of heuristics, the subsequent (cybersecurity) literature largely omitted their advantages, and heuristics are now almost exclusively seen as a weakness and a primary cause of human error in cybersecurity (Frauenstein & Flowerday, 2020; Kwak et al., 2020; Petrič & Roer, 2022; Tsohou et al., 2015). In some cases, heuristics serve as a convenient post-hoc explanation for sub-optimal decisions regarding cybersecurity (Gavett et al., 2017). Therefore, popular advice for reducing human errors is to “focus on strategies that inhibit the use of heuristics” (Butavicius et al., 2022) or “stimulate users

to process systematically” to save users from phishing emails or other digital dangers (Petrič & Roer, 2022). While not all of the mentioned cybersecurity examples paint heuristics in such a grim light, they are mostly seen as a hindrance to making good decisions and disregard their potential benefits.

The research stream of usable security analyzes human decision-making in an attempt to tailor applications to those constraints instead of solely blaming the human as the weakest link (Wash, 2010). Our work goes a step further to explore where human decision-making strategies could be used as an asset for security.

Cyberspace is a complex world with limited or ambiguous information and where probabilities are uncertain (Gomez & Villar, 2018). Emerging technologies and their impact are hard, sometimes even impossible, to foresee and create uncertainty (Lewallen, 2021). The behavior of both the attacker and the defender in cyberattacks is hard to predict, and so are the payoffs of their respective strategies (Sinha et al., 2015). This uncertainty negatively impacts the accuracy of risk assessments (Fielder et al., 2018), which in turn creates uncertainty in calculating the optimal investment in IT security (Paul & Wang, 2019). The downstream effect of these numerous uncertainties is further uncertainty in the choice of solutions to defend oneself against all the unforeseeable threats and their unpredictable consequences (Renaud & Weir, 2016). We therefore argue that the use of cognitive heuristics is a promising but neglected approach to combat cyber threats.

3.2. Potential domains for heuristics in cybersecurity

Table 2 gives a short overview of the problems the different domains of cybersecurity face and describe the role heuristics could play in alleviating them. This

list is a deductive selection of cybersecurity challenges characterized by high uncertainty and/or high time-pressure.

The most obvious application is the detection of *phishing emails*, as the literature already mentions cognitive heuristics, but mostly as the reason one falls victim to them. Furthermore, heuristics are already in use to aid detection, but only in the form of software algorithms (Khonji et al., 2013). Heuristics used by humans may incorporate context factors that are not accessible to the computer. One prominent approach to reduce phishing susceptibility is to make every user an expert in this domain, for example, by teaching them every possible strategy and giving them an extensive list of rules to follow (Canham et al., 2020). The outcomes of this approach are less than stellar, and studies have yet to demonstrate that the effect of this training persists for more than a few weeks (Lain et al., 2022). Another approach is to raise awareness about the threat and to give users a handful of rules like “always inspect the sender” and “hover over every link and look them up online to see if they are real” (Downs et al., 2006). While this approach seems quite frugal, it stops being so if this has to be done for every single email and, more importantly, is relatively easy to circumvent if one knows of this strategy, for example, by spoofing email addresses and websites. Both of these approaches aim to change the behavior of the user by increasing their knowledge.

Instead, cybersecurity trainings based on heuristics would address the decision-making process and aim to change the way people reason before making a decision. Here we see potential for FFTs which could be, if well-constructed, both faster than the extensive training and more robust than the rules that are currently in use. In the case of phishing emails, we could exploit the fact that there always needs to be a stage where the attacker has to “close the deal” in the form of making a request to share confidential

Table 2. Potential use cases of cognitive heuristics in cybersecurity

Use case	What is the problem?	How could a heuristic help?
Phishing	The potential victim is confronted with deception and time pressure. Established rules to detect phishing are often complex and time intensive (e.g., patterns of fake URLs).	The use of the correct heuristic relying on a few cues could help employees to quickly recognize suspicious emails for a more thorough analysis.
Threat detection / Incident response	Decisions have to be made under time pressure in an uncertain environment, as attack vectors are changing regularly. Established incident-handling processes are complex, making them hard to learn and use.	Converting complex incident handling guidelines to simple heuristics could support knowledge transfer and make them easier to apply in stressful situations by different stakeholders.
Risk assessment	Many variables to estimate cyber risks cannot be objectively measured and have to be approximated relying on historical data, making it bad at predicting future events. The process is complex and time-consuming.	The reduction of model variables may leverage the advantage of heuristic decision-making and allow for better predictions while saving time and resources. This can help both risk analysts and executives that need to decide how to manage risks.
Vulnerability analysis	A full analysis of the attack surface of an organization is time-consuming and expensive, increasing with the size of the organization.	Heuristics could be used by different stakeholders to make the analysis process simpler and to choose what to prioritize in the response to this assessment.
Security policies & guidelines	Security policies and guidelines are often complex, making them time-consuming to implement, especially for smaller companies.	Policies and guidelines in the form of simple heuristics could streamline their implementation and make them accessible, especially for smaller companies.
IT-Investment decisions	The benefits of investments are uncertain, and the return on investment is hard to predict.	Heuristics in the form of rules of thumb could support investment decisions regarding prioritization and the investment.

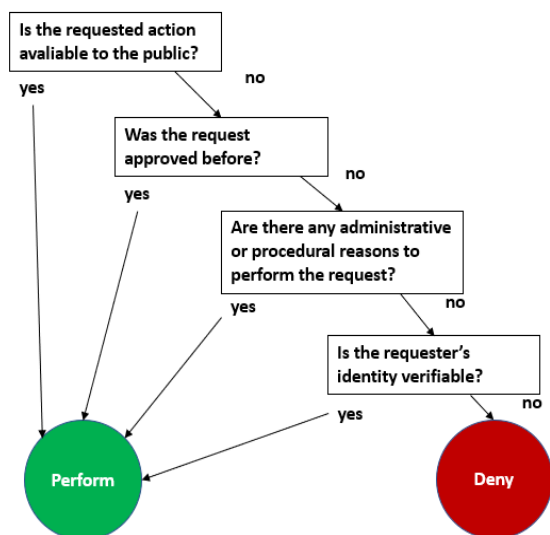


Figure 1. Illustration of a hypothetical fast and frugal tree derived from the social engineering detection model (Mouton et al., 2017).

information, approve an invoice, download malicious software, etc. This approach is insofar more adaptive than strict checklists, as it can be extended to different social engineering contexts (e.g., email, SMS, voice, etc.) and can better react to new angles of attack due to focusing on general actions instead of specific strategies.

Studies have already mentioned that sometimes a bad gut feeling is a good indicator of malicious emails (Williams et al., 2018) and that improving heuristics should be the first step to combat phishing threats (Vishwanath et al., 2018), but they did so without further elaborating what makes a good heuristic. Specifically, the work on constructing FFTs (Phillips et al., 2017) offers a good starting point to explore the potential of cognitive heuristics for phishing prevention. Figure 1 illustrates how a social engineering FFT for employees could look like. This FFT is based on the social engineering detection model and should merely act as an example with no judgement of its actual efficacy (Mouton et al., 2017). On receiving a request for an action (through any channel), one proceeds through each step of the decision tree in order. A positive answer to any of the questions leads to immediate termination of the decision algorithm, and the request is granted. If each of the answers is negative, the request is denied. In the first step, the recipient of a request (irrespective of the channel) checks if the requested action is available to the public or only to a specific group (e.g., employees). If not, one proceeds to the second step of the decision tree and checks if the request was approved beforehand by a higher instance. If that is not the case, one checks if there are any administrative or procedural rules in place to perform that action. If the answer is negative, the requester's identity is verified, and the decision process is terminated with either a positive or negative answer to the request.

Other fields of application that should be explored are *threat detection* and *incident response*. These have quite some resemblances with the medical field, especially emergency response, as medical personnel must make accurate categorizations with little information and must choose the correct response under time pressure. Typically during a cyberattack, decision-makers have at most 1 minute to detect it, 10 minutes to understand it, and 60 minutes to contain it before it spreads throughout the system (Pool, 2020). Different approaches to handle these tasks are in use at the moment. For example, machine learning models have been proposed to detect hostile traffic, but these often fail to generalize to new environments (Gehri et al., 2023). Here, heuristic approaches may help to alleviate this problem.

While some incident response teams have elaborate playbooks that guide every step of the response process, others provide less guidance and mostly rely on the prior knowledge of the different responders. Relying on a strict playbook may not be adaptive enough to cope with the changing threat landscape while being potentially inefficient. The use of heuristics, in this case, may allow for greater adaptability to emerging threats while being much more economical and easier to use and learn. Therefore, simple heuristics could also help to reduce staff burnout, which is a typical phenomenon in computer security incident response teams (e.g., Killcrece et al., 2003).

The military demonstrated their use for transferring expert knowledge (Banks et al., 2020), which should make them an excellent tool for training new cybersecurity personnel. The triaging heuristic (Super, 1984) could offer inspiration to construct a similar heuristic for cyber emergencies so that incident response teams can quickly assess the situation and initiate a suitable response. If FFTs are good enough to save lives in the hands of first responders, we should consider their potential in the hands of a cyber incident response team.

Another domain that could be improved by simple heuristics is *risk assessment*. Risk assessment entails filling out complex risk matrices while relying on rough estimations of the expected probability of breaches. These risk scores have not yet proven useful, and some even argue that they are harmful by providing a false sense of security (Hubbard & Seiersen, 2023). It was found that more data points improved confidence but actually worsened the prediction quality. These findings are an argument for reducing the complexity of risk models. The algorithms for choosing the right cues for FFTs could offer guidance in choosing the correct variables to improve risk assessments. Furthermore, smaller firms that do not have resources to spare could make good use of heuristics to get a rough estimation of the risks they are exposed to.

Vulnerability analysis is a closely related task that could likewise benefit from the use of heuristics. Both the assessment of vulnerabilities and their resolution is an effortful and time-intensive tasks. The industry standard vulnerability assessment tools employ complex scoring schemes and require extensive domain knowledge (e.g., CVSS). Furthermore, the results of this scoring system are detailed, but their complexity makes them difficult to use for prioritizing weaknesses (David Seidl & Mike Chapple, 2022). Heuristic guidelines could aid us in interpreting the results of such an assessment and help us in prioritizing the parts of the infrastructure we need to assess first and which steps we should take first to protect the most vulnerable systems from harm. A more comprehensive approach would be to break the scoring process itself down into simple heuristics. In this case, a good heuristic could improve the efficiency of the vulnerability analysis while being transparent and easy to understand so that less expert knowledge would be needed, and new employees could quickly be trained for this task.

Implementation or even replacement of **cybersecurity guidelines** (e.g., NIST Cybersecurity Framework) is another domain where the use of heuristics may be advised. The implementation of such guidelines is limited by their complexity and cost and by a lack of knowledge and missing skillsets in the implementing staff (Yvon, 2020). Furthermore, these guidelines are still high-level and leave plenty of room for individual judgement and decision-making. Heuristics may bridge the gap between general rules and their specific application. Breaking down (parts of) complex guidelines into simpler heuristics may, on the one hand, aid the knowledge transfer to train new staff for this task and, on the other hand, reduce their cost and time commitment, making it more worthwhile for smaller businesses that do not have the necessary resources to implement the full framework.

IT investment decisions are difficult endeavors that often rely on complex statistical models to predict the potential damage caused by future attacks and the return on investment of implementing defensive measures (Cavusoglu et al., 2004). What we observed for classical investment decisions may hold true in the domain of cybersecurity insofar as a heuristic allows us to choose the optimal strategy in such an uncertain environment.

3.3. Potential pitfalls of heuristics in cybersecurity

While heuristics are a promising solution for many problems one may encounter in the field of cybersecurity, they are not a silver bullet. First, a heuristic has to be applied in the right environment to perform optimally. Application in the wrong environment leads to inadequate perception and reasoning, manifested as systematic judgment errors.

Therefore, a heuristic that works in one environment should not be carelessly applied in another environment without gaining a deeper understanding first. Second, heuristics are only as good as their user. Studies have shown that there are individual differences in the application of even simple heuristics, which are explained by intelligence and other, sometimes immutable, psychological factors (Bröder, 2003). Even with a heuristic that is theoretically error-free, there will be cases where the practical application fails and compromises security. Therefore, heuristics are not a substitute for contingency planning. Third, heuristics can turn bad in hostile environments where other agents can discern the cues that trigger specific responses and may alter their behavior strategically (Evans & Stankovich, 2013). Adaptive adversaries may learn to circumvent popular heuristics, leading to yet another continuous arms race between attackers and defenders. Being aware of this, our response should be to analyze the environment in which a heuristic is employed and use this knowledge to either create better heuristics that cannot be exploited easily or abandon them in this specific environment if that cannot be achieved. Fourth, in cases where there is such an abundance of data that the whole decision space can be accurately mapped and complex algorithms dominate simple ones, one should ask the question if the human component is even needed or if that task can be entirely delegated to a machine.

Future research needs to uncover which specific cybersecurity contexts are suitable for the use of heuristics, who can benefit from them, and how to make them robust against adaptive adversaries.

4. Conclusion

We have elaborated on what heuristics are, their advantages, and where they are already successfully used. Furthermore, we have pointed out the similarities (e.g., uncertainty and time pressure) between these environments and cybersecurity to argue that the use of heuristics may be a good alternative to the established tools to combat cyber risks (e.g., in phishing detection or incident response).

Chesterton's fence, ironically in itself a heuristic, should act as a guide on how to proceed with heuristics in cybersecurity (Chesterton, 1929):

There exists, in such a case, a certain institution or law; let us say, for the sake of simplicity, a fence or gate erected across a road. The more modern type of reformer goes gaily up to it and says, "I don't see the use of this; let us clear it away." To which the more intelligent type of reformer will do well to answer: "If you don't see the use of it, I certainly won't let you clear it away. Go away and think. Then, when you can come back and tell me that you do see the use of it, I may allow you to destroy it."

This quote seems quite applicable to the way heuristics are dealt with now in cybersecurity. While heuristics are rightfully criticized in the cases where they lead to errors, the cases where they work well have been widely omitted in the cybersecurity literature. In our view, misunderstandings concerning the concept of heuristics and their context dependency led to their hasty dismissal in an environment where heuristic decision-making could, in fact, be beneficial.

Instead of abandoning heuristics, further empirical research is needed to uncover the specific environments where heuristics are advantageous for cybersecurity and where we should stick to complex models and comprehensive rules.

Acknowledgment

The study was financed with a DIZH grant by the Canton of Zurich and a grant of the Swiss National Science Foundation (No. 207550).

5. References

- Aikman, D., Galesic, M., Gigerenzer, G., Kapadia, S., Katsikopoulos, K., Kothiyal, A., Murphy, E., & Neumann, T. (2021). Taking uncertainty seriously: Simplicity versus complexity in financial regulation. *Industrial and Corporate Change*, 30(2), 317–345. <https://doi.org/10.1093/icc/dtaa024>
- Alsharnouby, M., Alaca, F., & Chiasson, S. (2015). Why phishing still works: User strategies for combating phishing attacks. *Int. J. Human-Computer Studies*, 82, 69–82. <https://doi.org/10.1016/j.ijhcs.2015.05.005>
- Artinger, F. M., & Gigerenzer, G. (2016). *Heuristic Pricing in an Uncertain Market: Ecological and Constructivist Rationality*. <https://dx.doi.org/10.2139/ssrn.2938702>
- Banks, A. P., Gamblin, D. M., & Hutchinson, H. (2020). Training fast and frugal heuristics in military decision making. *Applied Cognitive Psychology*, 34(3), 699–709. <https://doi.org/10.1002/acp.3658>
- Borges, B., Goldstein, D. G., Ortmann, A., & Gigerenzer, G. (1999). Can ignorance beat the stock market. In *Simple Heuristics That Make Us Smart* (pp. 59–72). Oxford University Press.
- Bröder, A. (2003). Decision making with the “adaptive toolbox”: Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(4), 611–625. <https://doi.org/10.1037/0278-7393.29.4.611>
- Butavicius, M., Taib, R., & Han, S. J. (2022). Why people keep falling for phishing scams: The effects of time pressure and deception cues on the detection of phishing emails. *Computers & Security*, 123. <https://doi.org/10.1016/j.cose.2022.102937>
- Canfield, C. I., & Fischhoff, B. (2018). Setting Priorities in Behavioral Interventions: An Application to Reducing Phishing Risk. *Risk Analysis*, 38(4), 826–838. <https://doi.org/10.1111/risa.12917>
- Canham, M., Posey, C., & Bockelman, P. S. (2020). Confronting Information Security’s Elephant, the Unintentional Insider Threat. In D. D. Schmorrow & C. M. Fidopiastis (Eds.), *Augmented Cognition. Human Cognition and Behavior* (Vol. 12197, pp. 316–334). Springer International Publishing. https://doi.org/10.1007/978-3-030-50439-7_22
- Cavusoglu, H., Mishra, B., & Raghunathan, S. (2004). A model for evaluating IT security investments. *Communications of the ACM*, 47(7), 87–92. <https://doi.org/10.1145/1005817.1005828>
- Chesterton, G. K. (1929). *The thing*. Aeterna Press.
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In *Simple heuristics that make us smart* (pp. 97–118). Oxford University Press. <https://doi.org/10.1002/acp.843>
- David Seidl & Mike Chapple. (2022). Analyzing Vulnerability Scans. In *ComptIA PenTest+ Study Guide: Exam PT0-002* (pp. 151–193). Wiley Data and Cybersecurity.
- DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal Versus Naive Diversification: How Inefficient is the 1/ N Portfolio Strategy? *Review of Financial Studies*, 22(5), 1915–1953. <https://doi.org/10.1093/rfs/hhm075>
- Dhami, M. K. (2003). Psychological Models of Professional Decision Making. *Psychological Science*, 14(2), 175–180.
- Djulgovic, B., Hozo, I., & Dale, W. (2018). Transforming clinical practice guidelines and clinical pathways into fast-and-frugal decision trees to improve clinical care strategies. *Journal of Evaluation in Clinical Practice*, 24(5), 1247–1254. <https://doi.org/10.1111/jep.12895>
- Donaldson, L. (2001). Structural Contingency Theory. In *International Encyclopedia of the Social & Behavioral Sciences* (pp. 15208–15210). Elsevier. <https://doi.org/10.1016/b0-08-043076-7/04214-5>
- Downs, J. S., Holbrook, M. B., & Cranor, L. F. (2006). Decision strategies and susceptibility to phishing. *Proceedings of the Second Symposium on Usable Privacy and Security - SOUPS '06*, 79. <https://doi.org/10.1145/1143120.1143131>
- Dressel, J., & Farid, H. (2018). The accuracy, fairness, and limits of predicting recidivism. *Science Advances*, 4(1). <https://doi.org/10.1126/sciadv.aao5580>
- Evans, J. St. B. T., & Stanovich, K. (2013). Dual-Process Theories of Higher Cognition: Advancing the Debate. *Perspectives on Psychological Science*, 8(3), 223–241.
- Fernandez, J. P., Katsikopoulos, K. V., & Shubitizde, F. (2010). *Simple geometric heuristics for the detection of unexploded ordnance* [Unpublished Manuscript].
- Feufel, M. A., & Flach, J. M. (2019). Medical education should teach heuristics rather than train them away. *Medical Education*, 53(4), 334–344. <https://doi.org/10.1111/medu.13789>
- Fielder, A., König, S., Panaousis, E., Schauer, S., & Rass, S. (2018). Risk Assessment Uncertainties in Cybersecurity Investments. *Games*, 9(2), 34. <https://doi.org/10.3390/g9020034>

- Fischer, J. E., Steiner, F., Zucol, F., Berger, C., Martignon, L., Bossart, W., Altwegg, M., & Nadal, D. (2002). Use of Simple Heuristics to Target Macrolide Prescription in Children With Community-Acquired Pneumonia. *Archives of Pediatrics & Adolescent Medicine*, *156*(10), 1005–1008. <https://doi.org/10.1001/archpedi.156.10.1005>
- Frauenstein, E. D., & Flowerday, S. (2020). Susceptibility to phishing on social network sites: A personality information processing model. *Computers & Security*, *94*. <https://doi.org/10.1016/j.cose.2020.101862>
- Fuster, G. G., & Jasmontaite, L. (2020). Cybersecurity Regulation in the European Union: The Digital, the Critical and Fundamental Rights. In M. Christen, B. Gordijn, & M. Loi (Eds.), *The Ethics of Cybersecurity* (Vol. 21, pp. 97–115). Springer International Publishing. https://doi.org/10.1007/978-3-030-29053-5_5
- Gavett, B. E., Zhao, R., John, S. E., Bussell, C. A., Roberts, J. R., & Yue, C. (2017). Phishing suspiciousness in older and younger adults: The role of executive functioning. *PLoS ONE*, *12*(2), e0171620. <https://doi.org/10.1371/journal.pone.0171620>
- Gehri, L., Hulliger, D., Meier, R., & Lenders, V. (2023). *Towards Generalizing Machine Learning Models to Detect Command and Control Attack Traffic*. 2023 15th International Conference on Cyber Conflict, Tallinn.
- Gigerenzer, G. (2014). *Risk Savvy: How to Make Good Decisions*. Penguin.
- Gigerenzer, G. (2018). The Bias Bias in Behavioral Economics. *Review of Behavioral Economics*, *5*(3–4), 303–336. <https://doi.org/10.1561/105.00000092>
- Gigerenzer, G., & Brighton, H. (2009). Homo Heuristicus: Why Biased Minds Make Better Inferences. *Topics in Cognitive Science*, *1*(1), 107–143. <https://doi.org/10.1111/j.1756-8765.2008.01006.x>
- Goldstein, D. G., & Gigerenzer, G. (2009). Fast and frugal forecasting. *International Journal of Forecasting*, *25*, 760–772.
- Gomez, M. A., & Villar, E. B. (2018). Fear, Uncertainty, and Dread: Cognitive Heuristics and Cyber Threats. *Politics and Governance*, *6*(2), 61–72. <https://doi.org/10.17645/pag.v6i2.1279>
- Green, L., & Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit? *The Journal of Family Practice*, *45*(3), 219–226.
- Halide, H., & Ridd, P. (2008). Complicated ENSO models do not significantly outperform very simple ENSO models. *International Journal of Climatology*, *28*(2), 219–233. <https://doi.org/10.1002/joc.1519>
- Herley, C. (2009). So long, and no thanks for the externalities: The rational rejection of security advice by users. *Proceedings of the 2009 Workshop on New Security Paradigms Workshop*, 133–144. <https://doi.org/10.1145/1719030.1719050>
- Hubbard, D. W., & Seiersen, R. (2023). *How to measure anything in cybersecurity risk*. John Wiley & Sons.
- Jenny, M. A., Pachur, T., Williams, S. L., Becker, E., & Margraf, J. (2013). Simple rules for detecting depression. *Journal of Applied Research in Memory and Cognition*, *2*, 149–157. <https://doi.org/10.1037/h0101797>
- Keller, N., & Katsikopoulos, K. (2016). On the role of psychological heuristics in operational research; and a demonstration in military stability operations. *European Journal of Operational Research*, 1063–1073. <https://doi.org/10.1016/j.ejor.2015.07.023>
- Khonji, M., Iraqi, Y., & Jones, A. (2013). Phishing Detection: A Literature Survey. *IEEE Communications Surveys & Tutorials*, *15*(4), 2091–2121. <https://doi.org/10.1109/SURV.2013.032213.00009>
- Killcrece, G., Kossakowski, K.-P., Ruefle, R., & Zajicek, M. (2003). *State of the Practice of Computer Security Incident Response Teams (CSIRTs)*: Defense Technical Information Center. <https://doi.org/10.21236/ADA421664>
- Kwak, Y., Lee, S., Damiano, A., & Vishwanath, A. (2020). Why do users not report spear phishing emails? *Telematics and Informatics*, *48*, 101343. <https://doi.org/10.1016/j.tele.2020.101343>
- Lain, D., Kostiaainen, K., & Čapkun, S. (2022). Phishing in Organizations: Findings from a Large-Scale and Long-Term Study. *2022 IEEE Symposium on Security and Privacy (SP)*, 842–859. <https://doi.org/10.1109/SP46214.2022.9833766>
- Lewallen, J. (2021). Emerging technologies and problem definition uncertainty: The case of cybersecurity. *Regulation & Governance*, *15*(4), 1035–1052. <https://doi.org/10.1111/rego.12341>
- MacCoun, R. (2002). Why a Psychologist Won the Nobel Prize in Economics. *American Psychological Society OBSERVER*, *15*(10).
- Martignon, L. F., Katsikopoulos, K. V., & Woike, J. K. (2012). Naïve, Fast, and Frugal Trees for Classification. In P. M. Todd & G. Gigerenzer (Eds.), *Ecological Rationality: Intelligence in the World* (pp. 360–378). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195315448.003.0106>
- Martignon, L., Katsikopoulos, K. V., & Woike, J. K. (2008). Categorization with limited resources: A family of simple heuristics. *Journal of Mathematical Psychology*, *52*, 352–361. <https://doi.org/10.1016/j.jmp.2008.04.003>
- Martignon, L., Vitouch, O., Takezawa, M., & Forster, M. R. (2005). Naive and Yet Enlightened: From Natural Frequencies to Fast and Frugal Decision Trees. In D. Hardman & L. Macchi (Eds.), *Thinking: Psychological Perspectives on Reasoning, Judgment and Decision Making* (pp. 189–211). John Wiley & Sons, Ltd. <https://doi.org/10.1002/047001332X.ch10>
- Mouton, F., Teixeira, M., & Meyer, T. (2017). Benchmarking a mobile implementation of the social engineering prevention training tool. *2017 Information Security for South Africa (ISSA)*, 106–116. <https://doi.org/10.1109/ISSA.2017.8251782>
- Paul, J. A., & Wang, X. (Jocelyn). (2019). Socially optimal IT investment for cybersecurity. *Decision* Page 4742

- Support Systems*, 122, 113069.
<https://doi.org/10.1016/j.dss.2019.05.009>
- Petrič, G., & Roer, K. (2022). The impact of formal and informal organizational norms on susceptibility to phishing: Combining survey and field experiment data. *Telematics and Informatics*, 67, 101766.
<https://doi.org/10.1016/j.tele.2021.101766>
- Pfleeger, S. L., & Caputo, D. D. (2012). Leveraging behavioral science to mitigate cyber security risk. *Computers & Security*, 31(4), 597–611.
<https://doi.org/10.1016/j.cose.2011.12.010>
- Phillips, N. D., Neth, H., Woike, J. K., & Gaissmaier, W. (2017). FFTrees: A toolbox to create, visualize, and evaluate fast-and-frugal decision trees. *Judgment and Decision Making*, 12(4), 344–368.
<https://doi.org/10.1017/S1930297500006239>
- Pólya, G. (1954). *Induction and Analogy in Mathematics: Vol. 1 of Mathematics and Plausible Reasoning*. Princeton University Press.
- Pool, R. (2020). 1-10-60: Measuring the speed of incident response. *Cyber Security: A Peer-Reviewed Journal*, 3(4), 308–314.
- Renaud, K., & Weir, G. R. S. (2016). Cybersecurity and the Unbearability of Uncertainty. *2016 Cybersecurity and Cyberforensics Conference (CCC)*, 137–143.
<https://doi.org/10.1109/CCC.2016.29>
- Rothlauf, F. (2011). *Design of Modern Heuristics: Principles and Application*. Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-540-72962-4>
- Schneier, B. (2000). *Schneier on Security—The Process of Security*.
https://www.schneier.com/essays/archives/2000/04/the_process_of_secur.html
- Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99–118. <https://doi.org/10.2307/1884852>
- Simon, H. A., & Newell, A. (1958). Heuristic Problem Solving: The Next Advance in Operations Research. *Operations Research*, 6(1), 1–10.
<https://doi.org/10.1287/opre.6.1.1>
- Sinha, A., Nguyen, T. H., Kar, D., Brown, M., Tambe, M., & Jiang, A. X. (2015). From physical security to cybersecurity. *Journal of Cybersecurity*, 1(1), 19–35. <https://doi.org/10.1093/cybsec/tyv007>
- Sundaramurthy, S. C., Wesch, M., Ou, X., McHugh, J., Rajagopalan, S. R., & Bardas, A. G. (2017). Humans Are Dynamic—Our Tools Should Be Too. *IEEE Internet Computing*, 21(3), 40–46.
<https://doi.org/10.1109/MIC.2017.52>
- Super, G. (1984). START: A triage training module. *Newport Beach, CA: Hoag Memorial Hospital Presbyterian*.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. (pp. x, 293). Yale University Press.
- Todd, P. M. (2007). How much information do we need? *European Journal of Operational Research*, 177, 1317–1332.
<https://doi.org/10.1016/j.ejor.2005.04.005>
- Todd, P. M., & Gigerenzer, G. (Eds.). (2012). *Ecological rationality: Intelligence in the world*. Oxford University Press.
- Tsohou, A., Karyda, M., & Kokolakis, S. (2015). Analyzing the role of cognitive and cultural biases in the internalization of information security policies: Recommendations for information security awareness programs. *Computers & Security*, 52, 128–141.
<https://doi.org/10.1016/j.cose.2015.04.006>
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science*, 185(4157), 1124–1131.
- Verschuere, B., Lin, C.-C., Huismann, S., Kleinberg, B., Willemse, M., Mei, E. C. J., van Goor, T., Löwy, L. H. S., Appiah, O. K., & Meijer, E. (2023). The use-the-best heuristic facilitates deception detection. *Nature Human Behaviour*.
<https://doi.org/10.1038/s41562-023-01556-2>
- Vishwanath, A., Harrison, B., & Ng, Y. J. (2018). Suspicion, Cognition, and Automaticity Model of Phishing Susceptibility. *Communication Research*, 45(8), 1146–1166.
<https://doi.org/10.1177/0093650215627483>
- Wash, R. (2010). Folk models of home computer security. *Proceedings of the Sixth Symposium on Usable Privacy and Security*, 1–16.
<https://doi.org/10.1145/1837110.1837125>
- Wegwarth, O., Gaissmaier, W., & Gigerenzer, G. (2009). Smart strategies for doctors and doctors-in-training: Heuristics in medicine. *Medical Education*, 43(8), 721–728.
<https://doi.org/10.1111/j.1365-2923.2009.03359.x>
- Williams, E. J., Hinds, J., & Joinson, A. N. (2018). Exploring susceptibility to phishing in the workplace. *International Journal of Human-Computer Studies*, 120, 1–13.
<https://doi.org/10.1016/j.ijhcs.2018.06.004>
- Woike, J. K., Hoffrage, U., & Petty, J. S. (2015). Picking profitable investments: The success of equal weighting in simulated venture capitalist decision making. *Journal of Business Research*, 68(8), 1705–1716.
<https://doi.org/10.1016/j.jbusres.2015.03.030>
- Wübben, M., & Wangenheim, F. v. (2008). Instant Customer Base Analysis: Managerial Heuristics Often “Get it Right.” *Journal of Marketing*, 72(3), 82–93.
<https://doi.org/10.1509/jmkg.72.3.082>
- Yvon, T. (2020). *Exploring factors limiting implementation of the national institute of standards and technology cybersecurity framework*.
- Zimmermann, V., & Renaud, K. (2019). Moving from a ‘human-as-problem’ to a ‘human-as-solution’ cybersecurity mindset. *International Journal of Human-Computer Studies*, 131, 169–187.
<https://doi.org/10.1016/j.ijhcs.2019.05.005>