



DECISION MAKING IN CYBERSECURITY

By: Ricardo Villadiego

TABLE OF CONTENTS

Executive Summary	-----	03
How We Make Decisions	-----	04
Complexity in Cybersecurity Decision Making	-----	05
The Prevalence of Poor Decision Making in Cybersecurity	-----	06
The Psychology of Error: Biases in Our Perception of Security	-----	07
The Cost of Insecurity	-----	09
Better Data Makes Better Decisions	-----	11
Conclusions	-----	12



EXECUTIVE SUMMARY

Data breaches have become a fact of life. However, a closer look at the most egregious breaches shows that better decision making could have prevented these attacks, or at least mitigated their effects. Responsible cybersecurity managers would do well to look at these examples and ask themselves “What would I have done better?”

In this paper, we consider two types of decision making required in cybersecurity.

1. Event-based decision making. This open-loop approach is faster, tactical, and necessary for day-to-day choices.
2. Risk-based decision making required for strategic investments. Functions such as detection, prevention, and response need to be invested in, and it is difficult to determine their interactions, delays, and the effect they have in reducing the cost of a breach. This approach considers cybersecurity architecture as an interconnected system whose output needs to be measured in order to close the feedback loop.

Both decision-making frameworks require being armed with the best information. The facts help to overcome biases and measure the value of security (and insecurity) accurately. This allows making well-informed decisions regarding the importance of security. Furthermore, information that sets a compromise baseline will allow us to measure the effectiveness of investments. In the event of a breach, having good information at one's fingertips makes for better reactive decision making. In order to deliver on its potential, this information needs to:

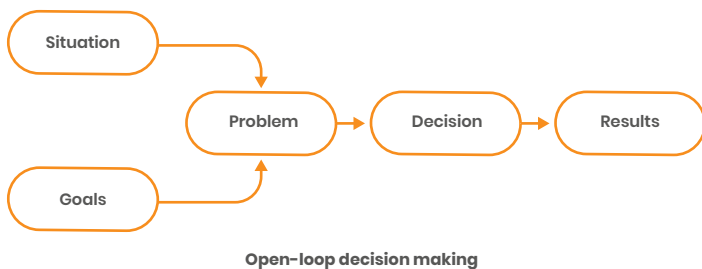
- Be timely and up-to-date
- Be of consistent quality
- Offer greater visibility
- Support taking action

HOW HUMANS MAKE DECISIONS

Human minds, though amazing in their own right, are unable to adequately encompass the complexity of real-world systems, according to cognitive psychologist Herbert A. Simon's concept of Bounded Irrationality^[1]. We frequently resort to reasoning shortcuts and other mental biases that lead us to adopt 'satisfying' solutions, rather than optimal ones.

Event-based Decision Making

According to renowned author, John D. Sterman^[2] "Where the world is dynamic, evolving, and interconnected, we tend to make decisions using mental models that are static, narrow, and reductionist." Sterman goes on to state that we tend to interpret experience through a series of open-loop events, where problems resulting from disparities between our goals and situation, require decisions that lead to results. In this paradigm, the decision-maker acts as a 'satisficer' and does not consider the inter-connectivity and feedback from real-world, dynamic systems.



We have evolved to make decisions in this manner in order to continue functioning under stringent limitations. Such constraints include having limited time to make decisions, too much information to process, not enough meaning from the information, and fallible memories for retaining it all. Cognitive psychologist Daniel Kahneman calls this thinking fast^[3]— when we are not able to think slowly. While these "short cuts" may lead to cognitive biases—more on that later—they are crucial for decision-making efficiency.

The average cybersecurity team could make hundreds of these reactive decisions in a day. Every alert is an opportunity for a decision and there simply isn't time to think 'slowly' at every juncture. The trick is knowing if the decision calls for some slow thinking.

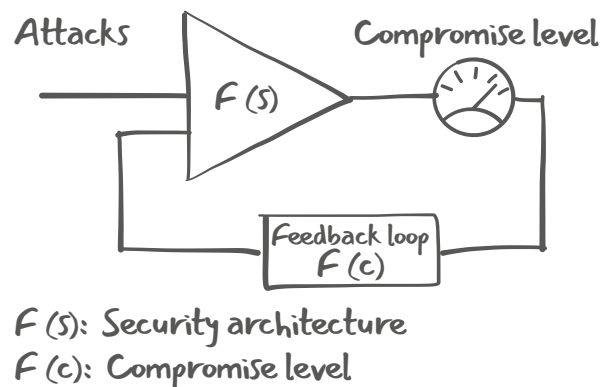
Proactive Investment Strategies

In *Managerial Perspectives on Risk and Risk Taking*^[4], James G. March and Zur Shapira state that "in conventional decision theory formulations, choice involves a trade-off between risk and expected return." Therefore, rational decision-makers invest in cybersecurity when the investment will yield a positive return, or rather when the cost of the investment is less than the potentially catastrophic loss it prevents. Indeed, the greatest responsibility of any modern CISO considers just this: how to invest budgets and resources in a way that most effectively reduces breaches and their consequences.

"THE FEEDBACK LOOP NEEDS TO BE CLOSED IN ORDER TO MEASURE THE EFFECTS OF THE CHANGES TO THE SYSTEM".

In any cybersecurity architecture, investments must be made into prevention, detection, and response systems, all of which have an influence on the other. This means that the feedback loop needs to be closed in order to measure the effects

of the changes to the system. Traditionally this has been difficult to achieve since organizations have not had the ability to measure compromise as an output of the system.



Lumu closes the feedback loop in cybersecurity by measuring compromise

There is a time and place for each of these decision making paradigms. The latter type of decision is strategic. They consider having the right tools and that enough resources are available when they are needed. The former is eminently tactical, covering how to employ such tools and resources to minimize damage. The important part is understanding when each type of decision making is called for.

COMPLEXITY IN CYBERSECURITY DECISION MAKING

OUR INABILITY TO MEASURE THE COST OF A HYPOTHETICAL INCIDENT COMBINED WITH UNCERTAINTY ABOUT ATTACK FREQUENCY CREATES AN IMPRESSION OF THE "EXPECTED COST" OF INSECURITY THAT DOES NOT BEAR RESEMBLANCE TO REALITY.

In a research paper^[5], issued by the Cybersecurity Interdisciplinary Systems Laboratory at MIT Sloan, researchers attempted to determine why poor decision making was so prevalent in cybersecurity. The researchers ran a cybersecurity simulation game that mimicked the complex systems—including prevention, detection, and response—needed in a

modern enterprise's cybersecurity program. Players had to choose how to invest in these processes, in order to protect against attacks and ultimately protect their enterprise's bottom line. Two groups of players were invited to play the simulation game. One group consisted of cybersecurity professionals, the other of inexperienced players.

The study found that both groups struggled in making effective decisions, but over multiple iterations, both groups managed to improve their scores. There were two major sources of complexity that needed to be overcome:

Uncertainty Concerning Cyber Incidents

The study considers that uncertainty surrounding the cost of an incident hampers decision making. In cases where

deterrents are successful, it can be difficult to measure the cost of a hypothetical cyber incident. Operators may also underestimate the frequency of attacks. The combination of these two factors creates an impression of the "expected cost" of insecurity that does not bear resemblance to reality. Even in cases where security operators have a good grasp on the expected cost of a breach, biases might cause operators to act irrationally—see the section *The Psychology of Error: Biases in Our Perception of Security* for more on this topic.

Delays in Complex Systems

The study looks at how investments in prevention, detection, or response can take time to have observable effects once implemented. Additionally, each investment needs time for implementation, and operators need time for training and overcoming learning curves. In a reactive decision-making paradigm, the development of cybersecurity capabilities only after the detection of an attack, means the organization's information systems will not properly recover in time and will remain vulnerable. A closed-loop decision-making process fares better, but the delays in feedback would mean that constant adjustment and measurement of the system would be needed to reach an optimal state.

The MIT paper showed that exposure to such large-scale breach events and their management improves overall cybersecurity decision making. However, waiting for such events to occur is a very expensive way to learn how to deal with them.

THE PREVALENCE OF POOR DECISION MAKING IN CYBERSECURITY

Equifax – a Compendium of Errors

The poster-child of data breaches' first example of poor decision making was a lack of preventative maintenance. Hackers made use of a widely-known vulnerability (that had been reported only 3 days earlier) in their complaints portal to gain initial access. If the vulnerability had only been promptly patched, there would not have been a breach.

The attackers' second move—moving laterally while escalating privileges—was also made easier by a lack of preventative measures. If Equifax had chosen to invest in the proper segmentation of systems, the attack would have been more easily limited to their customer complaint platform.

The attackers were able to have access to Equifax's databases for 76 days^[6]. At that time, they had reportedly not renewed an encryption license. Therefore, the encrypted personal information of approximately half of all Americans was able to pass through their HTTPS interception without being inspected. Only when the encryption had been updated—ten months late—did full network visibility resume, and was the attack detected.

Once the attack was discovered, Equifax's response showed terrible event-based reactive decision making. They delayed publicizing the breach for a month, when transparency in such events is the best policy. During that time little was done in terms of mitigating its effect on the American people, although several executives sold stock in the company—one being convicted for insider trading.

Capital One

In early 2019, an attacker exposed a vulnerability in Capital One's cloud integration in order to steal the credentials from over 100 million credit applications. The attacker executed a Server Side Request Forger^[7] to trick a misconfigured web

application firewall into relaying information including current credentials. This type of vulnerability had been known for years, but required specialized knowledge related to Amazon Web Services' Identity and Access Management as well as EC2 to identify and fix. Ultimately, a lack of investment in these in-demand cybersecurity skills led to a vulnerability that could have easily been avoided.

Marriott – the Breach that Lasted 4 Years

On November 30th, 2018, Marriott Hotels announced a breach^[8] that had been detected on September 8th. The breach affected the network of a chain of hotels—Starwood—that Marriott had purchased in 2016. It soon became apparent that Starwood had been breached in 2014 and remained compromised for 4 years. The attack exposed over 500 million customer records including passwords and credit card details. The breach was typical of a phishing attack that installed a Remote Access Trojan and a password sniffer in order to gain access and administrator privileges.

The most worrying aspect of the Marriott breach is that the compromise was allowed to persist for 4 years. This reveals that a key cybersecurity rule was not followed: assume you are compromised and prove otherwise^[9]. It also highlighted the importance of IT and security due diligence in the event of mergers and acquisitions. As the proprietor of Starwood, Marriott laid off most of their corporate staff, including IT and security staff. The new reservation system was not ready to manage the hundreds of newly acquired hotels, so the old understaffed and malware-maligned system was allowed to continue serving customers until the breach was discovered two years later.

Marriott's response^[10] to the breach caused further problems by using a wide range of email domains and websites, some of which lacked HTTPS certification. This led to a variety of phishing attacks imitating Marriott in the wake of the breach.

THE PSYCHOLOGY OF ERROR: BIASES IN OUR PERCEPTION OF SECURITY

Security comes at a cost, whether it is in the form of a loss of money, convenience, or opportunities. For example, locking your front door means trading increased security at the cost of a minor inconvenience. We all have an instinctive understanding that a trade-off needs to be made.

We have developed the ability to make these cost-benefit decisions quickly through cognitive biases: shortcuts that go around our limitations in time, memory, meaning, and dealing with excessive information. As Bruce Schneier says in his TEDtalk^[1], “We are highly optimized for risk decisions that are endemic to small family groups in the East-African highlands in 100,000 BC.” Our instincts inform our perception of security. Unfortunately, as the following examples illustrate, how we perceive security can differ greatly from its reality.

We Exaggerate Rare Risks

Many people fear flying even though it is safer than driving a car^[2]. This is because we tend to underestimate common risks. News stories of flaming airplane wreckage feature prominently in peoples’ association of flying. Part of the problem is that it is precisely the rarity of these events that make them newsworthy. However, the more attention is devoted to these events in news headlines, the larger the risk seems to us. In fact, air travel has become progressively safer over the years.

Comparing the numbers of commercial air disasters with the numbers of data breaches reveals increasing security in air travel, and decreasing security for personal data. Yet ‘having your data stolen’ is not a fear that people hold, despite

DEPLOYING THE LATEST TECHNOLOGY CAN MAKE SECURITY TEAMS FEEL LIKE THEY ARE IN CONTROL, UNDERESTIMATING ADVERSARIES AND THEIR ABILITY TO GET AROUND THESE MEASURES.

breaches becoming so commonplace that they rarely make the front page. Given the number of known records breached—and allowing for some unknown breaches—every one of us has had our private data breached multiple times.

The Unknown Is Feared More than the Familiar

We tend to trust people or things we know rather than those we do not know. System administrators do not patch known vulnerabilities for fear of introducing instability in their systems. Additionally, adopting new technologies is delayed in preference for more familiar legacy technologies. It is for this simple reason that phishers target users with emails that imitate trusted senders.

Personified Risks Are Given Priority Over Anonymous Risks

We struggle to accept risks when they are just abstractions. This is the reason why faceless attack groups—as well as hurricanes—are given names. It becomes more urgent when you know that Samurai Panda or APT4 is after you than some obscure Chinese officer.

“YET ‘HAVING YOUR DATA STOLEN’ IS NOT A FEAR THAT PEOPLE HOLD, DESPITE BREACHES BECOMING SO COMMONPLACE THAT THEY RARELY MAKE THE FRONT PAGE”.

We Underestimate Risks in Situations Where We Feel in Control

When we willingly adopt a risk posture, we tend to underestimate it. People feel in control when they have just deployed a new firewall, some magical virtualization technology,

or even a visibility solution. A CISO may think they are in control because they just deployed the latest state-of-the-art EDR. This can lead to seriously underestimating adversaries and their ability to get around these measures.

We Misjudge Objects When We Have Poor Visibility

In behavioral psychology, it has been found that people with poor vision tend to think that objects are farther away than they really are. The same happens when security operators have poor visibility into the compromises in their network infrastructure. In these cases, it is assumed that the risk of compromise is more remote than it actually is.

Overcoming Biases

$$R_s = C(P_0 - P_s)$$

R_s: Return of investment of a given solution *s*

C: Cost of a breach for my organization

P₀: Probability of a breach in a given time frame, with the current posture

P_s: Probability of a breach in the same time frame, adopting the solution

How can we align our perception of security with its reality? How do we know if the proper amount is being spent on security and spent effectively? It's important to realize that we are all susceptible to biases. However, the first true step towards achieving this is arming ourselves with the facts—and keeping these facts updated.

The first fact that needs clarity is the cost of insecurity. A clear understanding of the cost of a breach forms one part of the equation that tells you if your security trade-off is balanced. This used to be a difficult number to quantify, but each year brings better reporting^[3] that helps you understand the consequences for your industry, company size, and geographic region.

The second critical fact is your business' individual risk of a breach. Lumu's Continuous Compromise Assessment was developed to determine your organization's real-time factual level of compromise. The result of this process is a baseline for your cybersecurity architecture. This metric informs those big strategic decisions like "Are my security tools delivering on their promises?" and "Where do I need further investment?"

THE COST OF INSECURITY

As we have stated before, investment decisions are transactional. An investment has to be justified by its return. In cybersecurity, the return is the costs associated with the breach that is avoided by the investment. It has been noted that “difficulties in measuring the costs and benefits of information security investments cloud the vision of the rational decision-maker.”^[14] However, with each year better information regarding the cost and frequency of breaches becomes available through a range of reputable resources. It has become pivotal for cybersecurity operators to acquaint themselves with the real cost of insecurity in order to make an informed decision.

IT IS PIVOTAL FOR CYBERSECURITY OPERATORS TO ACQUAINT THEMSELVES WITH THE REAL COST OF INSECURITY IN ORDER TO MAKE AN INFORMED DECISION.

What Motivates Attackers?

Cybercrime is big business. A report by Atlas VPN^[15] estimated that cybercrime generates \$1.5 trillion annually. The largest component—\$860 billion—of this total comes from illegal online trading. The selling of trade secrets and intellectual property theft accounts for another \$500 billion. Trading stolen data—anything from credit cards to birthdates—generates another \$160 billion. A further \$1.6 billion is made by selling crimeware or Crimeware-as-a-Service. While individual Ransomware attacks provide great returns for threat actors and cause extensive damage, it ‘only’ accounts for \$1 billion of the total revenues of cybercrime.

State actors are driven by more than profit motive. These might conjure up images of strategic attacks like those we have seen carried out against nuclear centrifuges or election meddling. However, private citizens are also at risk. The Equifax breach that exposed the personal data of nearly half of all

Americans were believed to have been carried out by Chinese spies for the purposes of espionage.

How Are Attackers getting In?

It should be no surprise that as in previous years, the most common method of entry for breaches was hacking/intrusion^[16]. This category, accounting for 39% of all breaches, includes breaches through phishing, ransomware/malware, and skimming. The second-largest category, unauthorized access (37%) continued its growth trend from 2018, largely due to the increased prevalence of credential stuffing. The remaining 24% of compromises resulted from employee negligence, accidental exposure, data on the move, physical theft, and insider theft.

How Long Are They Avoiding Detection?

The average time to detect a compromise increased to 207 days in 2020^[17]. A further 73 days were required to contain these threats. Interestingly, these figures varied greatly depending on their region or industry. For example, German organizations required 160 days to identify and contain compromises, compared to 380 days in Brazil. Financial and banking organizations performed somewhat better than most, requiring 233 days while healthcare providers performed worst, requiring 329 days.

What Are They Getting Out?

The number of breaches increased in 2019 and so did the number of records exposed. In total, 870 million records were exposed, of which 165 million are considered to be ‘sensitive records’. Financial institutions were attackers’ main source of sensitive records, accounting for 101 million exposed records.

The Impact And Cost

Cybersecurity spending has increased by 44% since 2014, and yet we continue to see an increase in the number of breaches and records exposed. In 2019 the number of breaches increased by 17%. The impact of each breach also increased, especially in the USA, where the average cost of a breach amounted to \$8.64 million, more than double the global average.

From the data, it is clear that no industry is safe from breaches. Even the industries that were fastest to detect and contain compromises were still unacceptably slow. Industries subject to the most stringent regulations are failing to protect sensitive

data. Complying with the minimum demands of regulators or comparing yourself with industries that are faring worse, is far from enough.

Despite the direct correlation between dwell time and ransomware attacks, the time required for compromise detection is only increasing. Threat actors are constantly evolving their tactics, techniques, and procedures to ensure better deliverability. There needs to be a tactical and mindset change if strategists and operators are going to be able to turn around the hard reality our industry is up against.

BETTER DATA MAKES BETTER DECISIONS

Whether making quick tactical decisions or longer-term strategic ones, acting upon good information always aids the process. Let's look at some of the qualities this information needs.

Timely and Up-to-Date

Being able to make decisions quickly requires access to the newest information. Lacking information can lead to uncertainty and delays. Delays, in turn, can lead to growing doubts and more ineffective decision making.

Consistent Quality

Comprehensiveness should not come at the expense of quality. An example would be the prevalence of false alarms. Low-quality alerts cause alert fatigue and security operators to ignore alerts, as in the case of the boy who cried wolf. Alerts can only achieve certainty in response to known attacks with documented techniques and assets. Novel attacks will have to be represented by anomalies that require investigation. However, the investigative burden can be eased and alert fatigue lessened by improving the orchestration between alerts and investigating teams, and by providing contextual information.

Greater Visibility

As with poor eyesight, poor network visibility leads to errors in judgment. Greater network visibility helps to understand the main output of a cybersecurity system: its level of compromise. This level of compromise is crucial feedback information that can inform where additional investment is necessary in the system and tell you if investments are performing according to their promise.

Support Taking Action

Having too many options exacerbates delays in decision making. We frequently spend a lot of time trying to choose the best option. Paradoxically, it can be best to make a good choice, and then commit the resources that it needs to become a great choice in retrospect. However, to do so requires that the initial choice was made based on accurate intelligence and that the necessary resources are available for its follow up.

CONCLUSIONS

Attackers continue to successfully breach networks and avoid detection largely due to human error. Only being human, we are all susceptible to biases that can hamper effective decision making in cybersecurity. When making short-term decisions that deal with specific events, it is necessary to arm ourselves with information so that 'thinking fast' does not hamper thinking effectively. This information must be timely and reliable while offering comprehensive visibility and supporting taking action.

In the case of investment decisions, it is important to consider the effect of each investment on the system as a whole. In the cybersecurity industry, this is made more difficult by the fact that the interlinked functionalities and resources create a highly complex system, as well as significant delays between a feature's implementation and its visible effects. For this reason, it is critical to measure systems' output: compromises. This information can be used to close the feedback loop in cybersecurity, allowing you to determine the effectiveness of the cybersecurity system and its components.

Lumu's Continuous Compromise Assessment closes the feedback loop for cybersecurity systems by continuously and intentionally measuring this level of compromise. Closing the feedback loop allows knowing if investments need to be made and if existing functionalities are underperforming. Additionally, Lumu delivers detected compromises with contextual information, for swift and effective remediation.

REFERENCES

- [1] Simon, Herbert (1957). A Behavioral Model of Rational Choice, in Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting. New York: Wiley.
- [2] Sterman, John D. (2000). Business Dynamics: Systems thinking and modeling for a complex world. McGraw Hill. https://en.wikipedia.org/wiki/Business_dynamics
- [3] Kahneman, Daniel (2011). Thinking, Fast and Slow. London: Penguin Books. <https://archive.org/details/thinkingfastslow0000kahn/page/14>
- [4] March, James G.; Shapira, Zur (1987). Managerial Perspectives on Risk and Risk Taking. Management Science. <https://semanticscholar.org/paper/bed01b90c5f0b03cd60b20b8ace2ba56c1b2f942>
- [5] Jalali MS, Siegel M, Madnick S. Decision-Making and Biases in Cybersecurity Capability Development: Evidence from a Simulation Game Experiment. Journal of Strategic Information Systems 2018. <https://scholar.harvard.edu/jalali/publications/decision-making-and-biases-cybersecurity-capability-development-evidence>
- [6] Tech Crunch: Equifax breach was 'entirely preventable' had it used basic security measures, says House report. <https://techcrunch.com/2018/12/10/equifax-breach-preventable-house-oversight-report/>
- [7] Hackerone: Server-Side Request Forgery (SSRF) <https://www.hackerone.com/blog-How-To-Server-Side-Request-Forgery-SSRF>
- [8] CSO Online: Marriott data breach FAQ: How did it happen and what was the impact?. <https://www.csoonline.com/article/3441220/marriott-data-breach-faq-how-did-it-happen-and-what-was-the-impact.html>
- [9] Krebs on Security: What the Marriott Breach Says About Security. <https://krebsonsecurity.com/2018/12/what-the-marriott-breach-says-about-security/>
- [10] TechCrunch: Marriott's breach response is so bad, security experts are filling in the gaps — at their own expense. <https://techcrunch.com/2018/12/03/marriott-data-breach-response-risk-phishing/>
- [11] Schneier, B: (2010, October) The Security Mirage , Retrieved from https://www.ted.com/talks/bruce_schneier_the_security_mirage.
- [12] Forbes, Ricardo Villadiego: Aiming For The Sky: Cybersecurity Lessons From The Airline Industry <https://www.forbes.com/sites/forbestechcouncil/2020/01/09/aiming-for-the-sky-cybersecurity-lessons-from-the-airline-industry/#4533d67376b9>
- [13] Lumu Technologies: 2020 Compromise Flashcard, 2020 Ransomware Flashcard. <https://lumu.io/resources/2020-compromise-flashcard/>
<https://lumu.io/resources/2020-ransomware-flashcard/>
- [14] Chai et Al, 2011: Firms' information security investment decisions: stock market evidence of investors' behavior
- [15] AtlasVPN: Cybercrime annual revenue is 3 times bigger than Walmart's. <https://atlasvpn.com/blog/cybercrime-annual-revenue-is-3-times-bigger-than-walmarts>
- [16] ID Theft Centre: 2019 End-of-Year Data Breach Report. <https://www.idtheftcenter.org/2019-end-of-year-data-breach-report-download/>
- [17] IBM: Cost of a Data Breach Report 2020. <https://www.ibm.com/security/data-breach>



**Illuminating threats
and adversaries**

www.lumu.io