CYBERSECURITY MYTHS *and* MISCONCEPTIONS

Avoiding the Hazards and Pitfalls that Derail Us

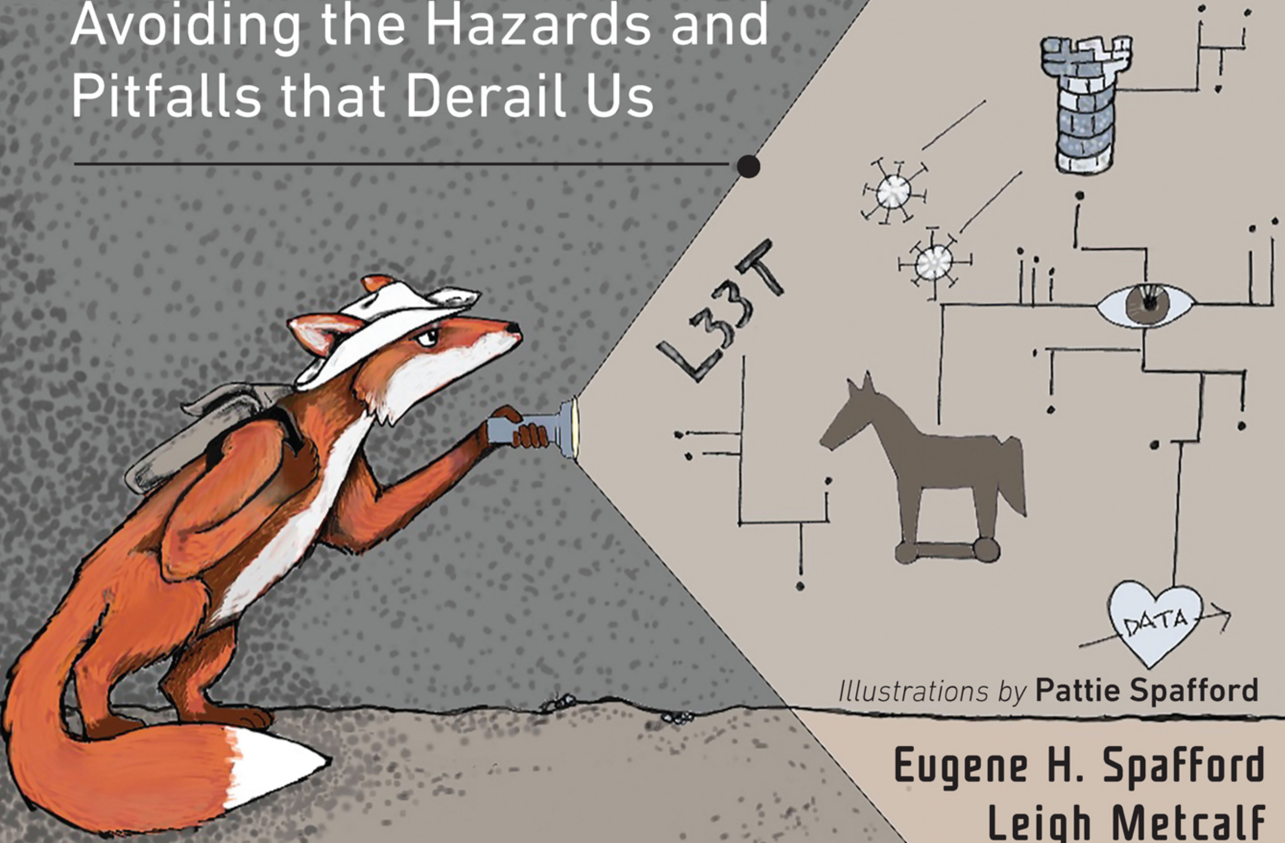

Illustrations by Pattie Spafford

Eugene H. Spafford
Leigh Metcalf
Josiah Dykstra

Foreword by Vint Cerf

“Many security leaders are traditionally in charge of correcting misconceptions just as much as they are in charge of building up solid security practices. We have plenty of resources on practices—but this book is the crucial guide to that essential myth busting.”

—Phil Venables
CISO, Google Cloud

“I’m writing this on my phone, over Wi-Fi, in an airplane on my way to Black Hat, one of the world’s largest security conferences. The fact that I’m able to do this at all shows how much we’ve really learned about cybersecurity over the decades. Now it’s all collected in one place for everyone to share. Thank the wise authors, and most importantly: GET OFF THEIR LAWN.”

—Wendy Nather
Head of Advisory CISOs, Cisco

“This book is astounding. A true tour de force—which I have never said about any other book. Inverting the viewpoint is a stroke of genius. This is going to be on my grabbable-at-any-time shelf. What I learned, recalled, and was refreshed on with technically astute agnosticism cannot be measured; just appreciated as a profound historical compilation of security practice and theory. Bravo!”

—Winn Schwartau
Founder and Chief Visionary Officer, The Security Awareness Company

“I am happy to endorse the central idea of this book—that cybersecurity is rife with myths that are themselves part of the problem. The brain wants to understand, the world grows ever more complicated, and the sum of the two is myth-making. As the authors say, even if some understanding is true at some time, with enough change what was true becomes a myth soon enough. As such, an acquired immunity to myths is a valuable skill for the cybersecurity practitioner if no other. The paramount goal of all security engineering is No Silent Failure, but myths perpetuate if not create silent failure. Why? Because a state of security is the absence of unmitigable surprise and you cannot mitigate what you don’t know is going on. Myths blind us to reality. Ignorance of them is not bliss. This book is a vaccine.”

—Dan Geer
CISO, In-Q-Tel

“This is a fun read for all levels. I like their rapid fire delivery and the general light they cast on so many diverse myths. This book will change the cybersecurity industry for the better.”

—Michael Sikorski
Author of *Practical Malware Analysis* & CTO, Unit 42 at Palo Alto Networks

A ***logical fallacy*** is an error in reasoning within an argument when trying to explain something or persuade someone. Fallacies make arguments less effective and compelling. Logic has been studied and taught for millennia, and most fallacies are thus well understood. As a result, some have fancy Latin names, such as the *ad hominem* fallacy. Others are better known by more colloquial names, such as *The Gambler's Fallacy*. Many of these fallacies result from not comprehending underlying probabilities or logic. People use some fallacies to win arguments. They might even know what they present is fallacious but use the argument unethically. Understanding these mental pitfalls—and avoiding them—can help make us better professionals, better communicators, and better critics.

The fact that humans are susceptible to these mental errors might seem like a human design flaw. The complex and information-rich environment around us has changed more quickly than our biology has evolved. A mental shortcut that worked historically in "most circumstances" might fall short in today's world.[3]

Humans also sometimes misunderstand concepts. We have opportunities to learn, either through instruction or self-study. In that process, we might be taught incorrect material (recall the "you lose the most heat through your head" trope in the Introduction), or we might not have the foundation to grasp the nuances. There are also occasions when we are exposed to false information presented as truth, and we learn the wrong facts! There might also be things we simply do not know about. The result is we might make future decisions based on an incorrect understanding of the world.

This chapter presents 14 common logical fallacies and misunderstandings in cybersecurity, from correlation and causation to the sunk cost fallacy. While there are dozens of others, this discussion will help you recognize some specific pitfalls and perhaps understand logical fallacies more generally. See the Further Reading resources at the end of the chapter to learn about others.

The False Cause Fallacy: Correlation Is Causation

Correlation is a statistical artifact between numbers often assumed to mean more than "only a number." However, this relationship is often nothing more than a coincidental, statistical artifact without intrinsic meaning. Further, in-depth explanation of this fallacy is presented in Chapter 14, "Lies, Damn Lies, and Statistics" in the "Correlation Implies Causation" section.

Imagine that an alert triggers for a cloud-based server, and the SIEM shows the logs in Table 4.1. What happened? A novice infosec professional could tell this appears to be a password guessing attack. There were multiple failed logins in a short time originating from the same source and then a successful login. Password guessing attacks do generate logs like this. It could have been a password guessing attack, or it could have been a frustrated user who accidentally left the caps lock on and kept trying until they realized that.

[3] For more, see Haselton, Martie G., Nettle, Daniel, and Murray, Damian R., "The Evolution of Cognitive Bias," *The Handbook of Evolutionary Psychology* (2015): 1–20.

TABLE 4.1 Sample login logs showing correlation.

Date	User	Host	Login Status	Source Address
5/1/2022, 7:01:30 AM	user01	vm01	Failure	203.0.113.6
5/1/2022, 7:01:31 AM	user01	vm01	Failure	203.0.113.6
5/1/2022, 7:01:32 AM	user01	vm01	Failure	203.0.113.6
5/1/2022, 7:01:33 AM	user01	vm01	Failure	203.0.113.6
5/1/2022, 7:01:34 AM	user01	vm01	Failure	203.0.113.6
5/1/2022, 7:01:35 AM	user01	vm01	Failure	203.0.113.6
5/1/2022, 7:01:36 AM	user01	vm01	Failure	203.0.113.6
5/1/2022, 7:01:37 AM	user01	vm01	Failure	203.0.113.6
5/1/2022, 7:01:38 AM	user01	vm01	Success	203.0.113.6

Importantly, a log such as this does *not* alone tell us the cause of the events, only that they occurred.[4] There are plausible alternatives for the cause of these login failures. It could have been the legitimate user making mistakes. Uncertainty does not mean the logs are useless; they might be the key indicator of a compromised account. Regardless, we should be honest that the logs *suggest* a password guessing attack, and we cannot say for sure (without other data) that it was. In the end, such logs offer a helpful hint, not a definitive cause.

The False Cause Fallacy is not new. In ancient Rome, sacred chickens were consulted about major undertakings. A widely known story was how Publius Claudius Pulcher (consul, 249 BC) consulted the chickens before a planned naval battle in the First Punic War. The chickens were not eating their feed, which was viewed as a bad omen. Claudius, enraged, said something (in Latin, undoubtedly) along the lines of “If they will not eat, let them drink!” and had them thrown into the sea. The battle went badly for Rome, resulting in the loss of the fleet. People of the time saw the correlation and assumed causation was involved: Pulcher narrowly escaped with his life. Few people today base major decisions on chickens.[5]

Correlation does not mean causation, but it can give a nudge that the items should be examined more closely to see if there is a functional relationship.

[4] Somtimes, logs do include causation, such as the event that triggered an error.

[5] Though some chickens have made better decisions than people. https://money.cnn.com/1996/09/27/personalfinance/yomo_ worst/

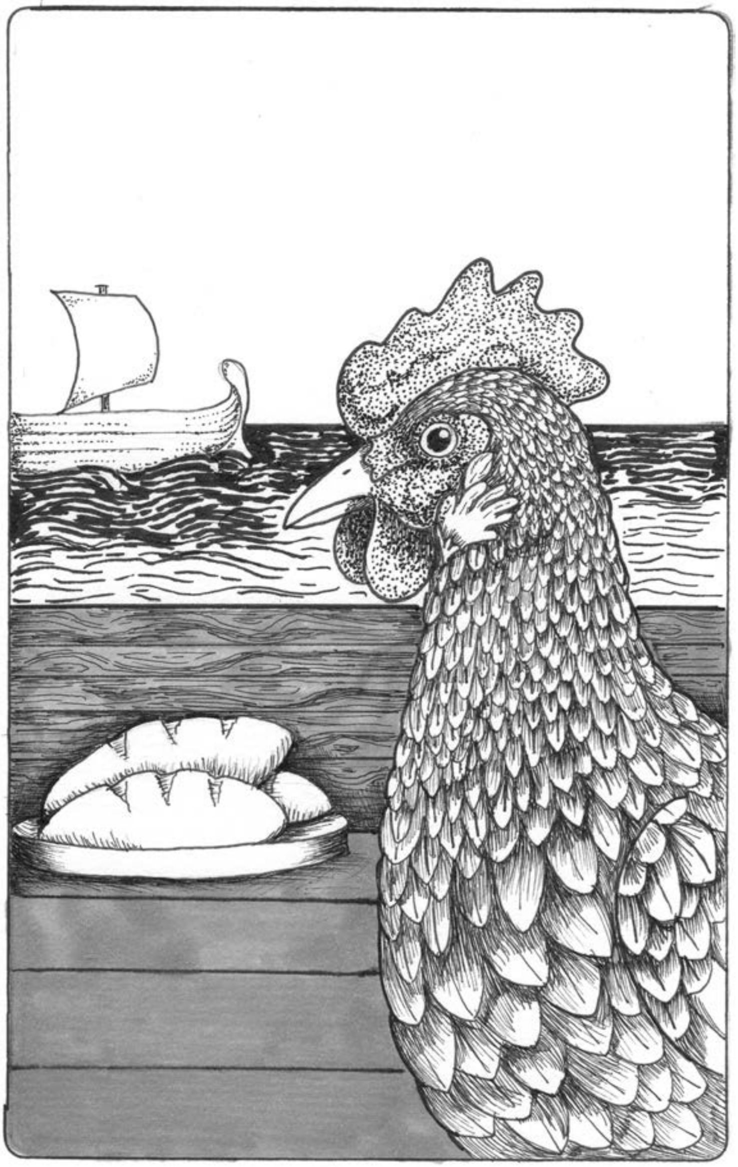

FIGURE 4.1 Killing chickens will not cause the battle to go better.

The "Magic Bullet" and Other Errors in Cause and Effect

The concept of a *magic bullet* (sometimes *silver bullet*) creeps into cybersecurity now and then. It suggests a highly effective—even perfect—single solution to a specific problem. The phrase is commonly "there is no magic bullet in cybersecurity." But that's not the end of the story.[a]

As we will continue to discuss, cybersecurity is complex, and no single tool or solution exists that can fully mitigate security challenges or deliver perfect outcomes. In that sense, people generally understand that there is no single security appliance or piece of software for complete cybersecurity.

A related pitfall exists in the investigation and causation of cyber incidents. How did the breach happen, for example? Humans seem to be wired to latch on to simple explanations. "The breach resulted from a password guessing attack," we say. That's a direct and understandable cause.

It gets worse. Once we spot the simple cause that feels right, we stop looking for other reasons. Breaches can have *many* causes, including insufficient policies, lack of multifactor authentication, allowing too many password attempts, misconfigured system configurations, and on and on. There can be many causes, not only the first one that sticks firmly in our minds. Furthermore, attackers increasingly use exploit chains where multiple vulnerabilities are combined to achieve the desired outcome. So, it's no longer one cause but many.

Finally, be aware of a related bias known as the *fundamental attribution error*. This error surfaces when evaluating a situation where a human is involved. Research shows that we tend to over-emphasize personality-based explanations—Morgan clicked on the phishing link because of laziness—rather than situational and environmental factors. Maybe Morgan is super busy or undertrained, or the spam detector is broken. Again, be careful about tunnel vision and losing sight of other causes.

We will talk a lot more about troubles with other analogies in Chapter 8, "Pitfalls of Analogies and Abstractions."

[a] To the best of our knowledge, silver bullets are needed only against werewolves. Our advice is to avoid investing in precious metal ammunition. If you find you are being attacked by lycanthropes, summon the Long Ranger as an outside consultant before the next full moon. Also, contact the authors: We have questions.

Absence of Evidence Is Evidence of Absence

If someone does not see an intruder in their apartment, does that mean one is not hiding there? If their doctor has not told them that they have cancer, does that mean that they definitely do not have cancer? If their security tools do not tell them there is an attacker, does that mean there isn't an attacker present?

Unfortunately, the answer to each of these questions is "no." An attacker might be using a method that the security appliance does not detect. It's possible that the attacker was in place before the appliance was installed and was not detected. It's possible that the appliance itself is defective or has a faulty configuration. It's possible that the attacker is clever and knows how to fool the appliance. It might even be the case that the attacker got in, did something nefarious, and left before the appliance could detect anything. Those are not all the possibilities of what could happen, but the same underlying reality is that because there is no evidence that it's happening does not mean it is not happening (or happened).

Even with all of these considerations, if the appliance indicates there is a problem, most likely there is a problem. There is evidence that the problem might exist (although false alarms happen). No evidence of a problem does not mean there is no problem.

This fallacy shows up in many aspects of our lives. Only 65% of the people infected with the original SARS-CoV-2 (the COVID-19 virus) had symptoms. The asymptomatic 35% had no evidence of infection, but that did not mean they were not infected. It simply meant that the virus did not result in symptoms in those people.[6]

The same is true with computer viruses. If human-created tools do not detect malware, that does not mean there is not any malware present! It simply means that those tools did not find any. Tools generally detect only known threats. Someone or something has to identify the threat for it to be appropriately detected later by a signature-based tool. A new, novel threat could easily slip by a tool, and we would never know until it was too late.

Retrospective Security and Infinite Storage

When a new attack or malware is discovered, security researchers now commonly have historical data available, allowing them to look back in time to reveal where and how long that attack was being used before it was discovered.

For example, on January 6, 2021, Volexity first identified a new attack against Microsoft Exchange mail servers. Looking at historical data, it was later determined that attacks were occurring as early as January 3.[a] Such discoveries can identify new victims and better understand the origins of the attack.

This powerful capability is made possible by retaining logs and other historical data. It allows an investigator to reveal activity that seemed benign or unremarkable in the past but can now be seen as related to an incident.

There is a related misconception that everyone can indefinitely store unlimited data volumes, given the practically unlimited storage options today. A counterexample is Netflix. It has stated that “at some point in our business growth, storing device and server logs did not scale because the increasing volume of log data caused our storage cost to balloon and query times to increase.”[b] So now the company filters logs and selectively stores the ones it needs.

While we gain better insights about attacks and attackers from retrospective analysis, we should not rely on it to remedy the absence of evidence fallacy. If queries do not find any activity, it is more accurate to say that we *did not find any evidence in the dataset* than to say that *none exists*. There could have been activity that the sensors did not observe or did not record.

[a] www.volexity.com/blog/2021/03/02/active-exploitation-of-microsoft-exchange-zero-day-vulnerabilities/

[b] https://netflixtechblog.com/lessons-from-building-observability-tools-at-netflix-7cfafed6ab17

[6] https://epi.ufl.edu/articles/35-percent-of-all-covid-19-infections-never-show-symptoms.html

A related fallacy shows up in many real-world cases in another form. That is when there is a challenge to prove a negative. It is often quite simple to prove the existence of something, as in "Prove that gravity exists." If someone asks us that, we can simply drop a microphone in front of them and watch it fall. So, it would seem natural that if we can prove something exists, we should be able to prove the opposite. This is impossible outside some restricted domains because of the limits of experience and observation. If we do not have an example of something, is that because it does not exist or because we have not looked for it in the right places yet? (This is related to Black Swan events, further described in the section "Ignorance of Black Swans.") This has implications in many domains, from philosophy to physics. It is also used, sometimes unscrupulously, in politics ("Prove the election was not compromised—you cannot, so it was."), religion ("Prove Cthulu does not exist—you cannot, so it does"), relationships ("Prove you are not cheating on me"), and so on. Beware of anyone trying to win an argument with this approach!

Vendors have exploited this fallacy to set up "contests" around their products. They offer a reward to anyone who can defeat the product under some restricted circumstances.[7] Then, after no one succeeds (or even bothers to try), they advertise that their product is "unhackable by everyone on the Internet." It means no such thing; the absence of evidence to the contrary does not mean it is true.

Another example devised from an assignment one of us has presented to students: Prove to us there are no unicorns. There is no way to prove this negative, so it would be wrong for us to claim thereby that unicorns must exist because the students failed to show proof they did not! If that is not obvious to you, then we have a special offer: For $3 million, we will sell you a monitoring package guaranteed to keep your systems from ever being attacked by evil, rampaging unicorns. No one can prove it is not 100% effective, so send us the money!

The Straw Hacker Fallacy

Let's presume we know our network, our systems, and the worst thing a hacker could do to all of it. If we were the hacker going after our network, that high-value asset would be the first thing we would go after. In this thought exercise, the attacker would use the new and improved SomeWare rootkit. This is an interesting attack that is almost impossible to defeat once lodged in a system. Thus, that must be the attack that would be used.

In this scenario, we imagine that the attacker would only make the attack that we would make. We are reducing the list of possible attacks to only a few and assuming that's all we need to defend against or where we should apply the bulk of our defense resources.

Attackers, unfortunately, seldom think as we do. They will use the tools and knowledge they have at hand. They will attack the resources they believe are interesting or whatever is immediately available. The SomeWare malware might be fascinating to us, but if a social engineering phone call or simple phish gets attackers what they want, they will use that and not bother with the malware.

[7] This is not the same as setting up a bug bounty program.