Editors: Jens Grossklags, jens.grossklags@in.tum.de | Heather Richter Lipford, richter@uncc.edu | Jessica Staddon, jessica.staddon@gmail.com

Action Bias and the Two Most Dangerous Words in Cybersecurity Incident Response

An Argument for More Measured Incident Response

Josiah Dykstra, Jamie Met, and Nicole Backert | National Security Agency
Rebecca Mattie | Johns Hopkins University
Douglas Hough | Johns Hopkins Bloomberg School of Public Health

A Note From the Department Editors

The volume of cybersecurity incidents, start-ups and regulatory attention to incident reporting suggest this is an important time for cybersecurity response. To encourage more discussion around incident response strategy we are including a position paper in this issue. We welcome feedback on the position in this paper at sp-sociotech@computer.org.

An inherent aspect of cybersecurity is managing incidents where people expect to see an immediate response. Action bias is our tendency to favor action over inaction because it makes us feel better—even if we have made things worse.

People who work in cybersecurity share a common goal to prevent and mitigate harm. The opportunity to protect businesses and individuals as well as the resulting rewards and value of successfully preventing, detecting, and stopping cyber incidents draws people to the field. Most cybersecurity professionals acknowledge that these goals are challenging and that achieving perfect security is impossible. Nevertheless, they strive for perfection as the ultimate objective, and they feel loss, failure, and regret when incidents inevitably occur.

Human instinct, especially when faced with the crises that are a routine part of cybersecurity, is to react forcefully and immediately. The fight-or-flight response is a natural physiological reaction that occurs in response to a perceived harmful event, attack, or threat, whether physical or digital. Unfortunately, this sensitivity to the impacts of cyber incidents often leads to inefficient or inappropriate mitigations that come at a higher cost than necessary.

Digital Object Identifier 10.1109/MSEC.2022.3159471
Date of current version: 26 May 2022

Action Bias and Other Dangerous Mental Shortcuts

As far back as Aristotle, philosophers and psychologists have recognized that humans use at least two kinds of thinking: fast, automatic, "intuitive" thinking and slow, controlled, "analytic" thinking. Nobel Laureate Daniel Kahneman called these *system 1* (fast) and *system 2* (slow). People need and use both. For example, when we learn to drive a car, we must use effortful and controlled system 2 behavior, but, with experience and acquired expertise, we develop heuristics (or "rules of thumb") that allow us to drive with less mental time and energy (system 1). Heuristics allow us to ignore some information and make decisions quickly. Kahneman and his research partner, Amos Tversky, argued that "heuristics are highly economical and usually effective, but they lead to systematic and predictable errors."[10]

Cognitive biases occur when heuristics go wrong. Kahneman and many others have identified scores of

cognitive biases, each of which renders decision making using system 1 problematic. Most people will exhibit cognitive biases of some kind. For example, confirmation bias occurs when we search for and give more weight to evidence consistent with our hypothesis or prior beliefs. A cyber defender displays confirmation bias when assuming that an employee is an insider threat and concentrates on behavior consistent with that view, such as excessive printing, while ignoring other evidence, such as a looming project deadline that might explain the behavior.

Then there is action bias, which is our tendency to favor action over inaction because quick responses makes us look and feel better in the short term even if we have made things worse. Action, especially quick action, makes us feel decisive and in control.[9] Without preparation and practice for the ideal courses of action, decisions are noisy and error prone. New or unclear situations, such as cyber incidents, accentuate the action bias, which can occur when there is perceived pressure to show "leadership" or preempt second-guessing: "Don't just stand there, do something." One source of action bias is nature's flight-or-fight response. Though often beneficial, it can also derail rational decision making. Action bias causes an automatic knee-jerk reaction when the better action is a carefully calculated one, and it replaces a rational cost–benefit analysis with the impulse to do something.

One illustration of action bias is described in a 2007 study of elite soccer goalkeepers.[1] In the case of penalty kicks, when a striker and goalkeeper face each other one on one, the goalkeeper has a quarter second to react to a kick. The goalkeeper can move left, move right, or stay in the middle. On average, even elite goalkeepers block only about one in six attempts. The researchers found that, when the goalkeeper stayed in the center of the goal and did not jump to one side or the other, more goals were blocked (one in three); however, these elite goalkeepers moved left or right on 95% of penalty kicks. Goalkeepers—and their fans—believe that "doing something" is a better strategy than standing still, regardless of the data. The tendency to want to jump is action bias—when staying still can be the better, carefully calculated action.

Situational context can increase the likelihood of action bias. If others expect action, that pressure can influence the decision maker to act, just as it did for soccer goalkeepers. Individual and corporate characteristics, such as overconfidence, may make action bias more likely in those populations. Research also shows that people in a "hot" emotional state, such as during a crisis, make dramatically poorer decisions than when they are in a "cold," calm, and controlled state.[8] In the end, action bias causes people to rush decisions, which results in errors.

Action Bias in Cybersecurity

Action bias is seen in many aspects of cybersecurity. Phishing is an obvious example. Attackers lure victims into taking immediate action, such as opening a file or clicking a link, to get control over a seemingly urgent situation. The attacker hopes that the victim will not pause and carefully consider whether the situation may be malicious. The user's desire to quickly respond to the crisis is action bias. The knee-jerk reaction leads to worse actions with worse consequences than carefully calculated choices. Phishing training that involves teaching people to slow down and pay attention to red flags encourages system 2 thinking that can help avoid action bias.

Incident response highlights a common case of action bias, even though it has not been previously labeled as such. In 2014 and 2016, Uber suffered two similar breaches that demonstrate how action bias had negative consequences.

The first time, malicious cyber actors used stolen credentials and gained access to 100,000 Uber drivers' personal information. Uber investigated and disclosed the incident as legally required. The Federal Trade Commission cited Uber for "fail[ing] to provide reasonable security to prevent unauthorized access."[4] Uber settled the claim and agreed to implement a comprehensive privacy program.

Two years later, another attack was carried out in the same way, compromising driver and customer information for more than 57 million people stored in Amazon Simple Storage Service buckets, including names, email addresses, and driver's licenses. Instead of disclosing this, when Uber discovered the breach, its seemingly knee-jerk reaction was to pay the attackers US$100,000 to keep quiet about the event and delete (in theory) the personal data of the drivers and customers. Despite a legal obligation to report the breach of personal information, Uber kept quiet for almost a year and then incurred a lawsuit that cost it US$148 million in fines. Its action bias and lack of a thoughtful and well-rehearsed incident response plan resulted in significant negative consequences. Uber was unprepared to handle the data breach even though it had been faced with the same problem two years earlier.

Cybersecurity decisions frequently have inherent characteristics in common that encourage action bias, as illustrated by the response in the Uber case. These make nonrational behavior more likely and mitigations more critical. One characteristic is that the consequences of security and privacy choices are often uncertain, with long time lags and periods of latency until effects are felt. From a user's perspective, not knowing how bad the infection will be and hoping that speed will

remove the risk creates an urgency to fix the problem without seeking help. Hence, users will delete emails, reboot their systems, or even disconnect the network. Another characteristic of cybersecurity is its externalities. Cybersecurity outcomes are interconnected and often come from others' choices or the contributions of many parties. A user may not know how his or her immediate reactions interact with the preplanned responses of defenders or the disparate actions of other users on the network.

At first glance, the converse of action bias and preferable choice might appear to be "do nothing." Inaction or timidity is not what we are suggesting. There will always be circumstances that do require immediate and urgent action to gain control of a critical situation. The converse of an uninformed, rushed decision is following a thoughtfully created plan that is well rehearsed and carefully executed. High-affect situations, such as a data breach, malware spread, and ransomware, induce a reflexive response, but even the reflexive response can be premeditated and rehearsed. The key is to enhance the ability to distinguish materially between the benefits and risks of action across the spectrum of options.

"Never Again"

Two of the most dangerous words in cybersecurity consistent with action bias are *never again.* "That data breach we had? Never again!" Unsurprisingly, passion and desperation for protecting a valuable asset would lead someone to declare an ultimatum. Security professionals have known, long before the digital age, that security comes in layers and with risk/reward probability calculations—but never with guarantees of perfection. Senior leaders should know likewise.

The danger of *never again* is that it sets an impossible goal. It encourages the organization to try any and all actions to prevent crises that cannot be eliminated. It ignores the reality that hackers are incentivized to keep attacking and playing at an advantage to the defenders. A company clearly wants no data breaches, but it is irrational to think that the risk can ever go down to zero. There are no 100% solutions when turning off the computer is untenable.

Risk management, done appropriately, is about understanding the value of assets and risks to them as well as being appropriately prepared (not over- or underspending). This impossible goal of zero tolerance, or *never again*, leads to a ripe environment for hyper overreaction to events and opens the door wide for action bias. In other words, in the pressure cooker of a *never again* mentality and faced with constant threats, cyber defenders are much like the soccer goalkeepers—encouraged to act prematurely and attempt the heroic save. It's like telling the goalkeeper that he or she can never allow another goal.

One potential consequence of action bias is irrational resource allocation. Without strict adherence to risk management, a crisis can lead to over- or underspending on incident response. Another reason *never again* is such a dangerous phrase is that it sets a seemingly unlimited cap on the resources for security, which may be unchecked by the bounds of rationality. If a data breach is evaluated as a loss of US$10 million once per decade, the company would be irrational in spending US$10 million every year so that it "never" happens again. Request for Comments 4949 reminds us of Courtney's second law, "Never spend more money eliminating a security exposure than tolerating it will cost you," and the first corollary: "Perfect security has infinite cost." Societal costs, including externalities imposed on customers, may exacerbate the total burden even if the company does not bear them.

Not every leader makes an ultimatum. There is no evidence, for example, that leadership at Uber went so far as to declare that incidents must never happen again. Nevertheless, we see that Uber did not learn from the first incident and clearly suffered from action bias in the second breach. If its leaders were to take a *never again* stance, that would further encourage extreme action no matter the incident response plan.

Countermeasures to Action Bias

While there is no cure for action bias, measured approaches can help minimize the undesired effects. First and foremost among these is the preparation for and premeditation of threats, impacts, and an honest assessment of reasonable objectives. The following countermeasures are components to help mitigate the negative effects from action bias and improve your team's plans and overall cybersecurity. Whatever your role, from junior engineer to senior executive, seek to apply them where you have influence and advocate for them to your management.

Risk Management

Risk management is a proactive step to mediate action bias in a crisis. It allows an organization to make rational cybersecurity decisions by identifying assets, the value of those assets, risks, and probabilities. Various models for systematic cybersecurity can help avoid action bias. For example, the National Institute of Standards and Technology (NIST) framework for incident response identifies a logical process from preparation to detection and analysis; then containment, eradication, and recovery; and, finally, postincident activity. The framework emphasizes preparation and prevention "so that the appropriate actions are taken."[2]

In real-world situations, especially those where preparation has

not been done, it would be natural to jump to eradication without performing an analysis. In the Cyber Resilient Organization Report 2020 from IBM Security, 51% of respondents said their incident response plans were not applied consistently across the enterprise or that the plan was informal or ad hoc. Many frameworks exist, including the NIST Cybersecurity Framework, ISO/International Electrotechnical Commission 27001, and Factor Analysis of Information Risk. Other tools, such as the Gordon–Loeb model, can help calculate an optimal investment level for cybersecurity.[6]

Slow Down

Planning and preparation are vital to avoid action bias. Athletes sometimes describe the game as happening in slow motion. This phenomenon is a byproduct of expertise, experience, and deliberate practice long before the game. During an incident, the time between observation and action can remain short even when the mitigations are developed beforehand. Standard operating procedures and playbooks shift the decision making to a time before the crisis instead of during it, giving time for slow, careful, system 2 thinking.

Organizations must practice these plans, such as with tabletop exercises, to build experience and identify gaps in the plan. To prioritize tabletop exercises within organizations of high operations tempos, try showing leadership how preparation could have saved time or money in an incident they remember. You can also conduct a premortem by imagining that your response to a crisis has failed and anticipating what might go wrong. Chelsey "Sully" Sullenberger, who landed U.S. Airways Flight 1549 safely in the Hudson River in 2009 during an unprecedented emergency, kept a quote with him on every flight: "A delay is better than a disaster." He later explained this, saying, "This means having the integrity and courage to reject the merely expedient and the barely adequate as being—quite frankly—not good enough."[5] His preparation and careful thinking led to a safe outcome.

Culture Change

Everyone in the organization must be aware that action bias exists and can negatively impact decision making. Training and awareness are components that can help prepare for crisis and temper action bias behavior during a crisis. However, bias training alone is unlikely to have the most substantial impact on behavior change. Culture change will take time and require continual reinforcement to be effective. Organizations can measure these changes by documenting the occurrence of premortems and exercises, in after-action reports of incident management, and in cost–benefit analyses of the crisis response. Kahneman and others have also suggested that a decision observer with a bias checklist can help independently diagnose if biases, including action bias, are affecting decision making.

Senior Leadership Education

Senior leadership, boards of directors, and shareholders can be powerful forces for visible action after an incident. It is vital to assure stakeholders that the organization is appropriately prepared and that, sometimes, "wait and see" is the optimal choice. Stakeholders need constant reassurance that a crisis is being managed, even if it appears that you are not taking action. Lastly, help senior leaders refocus on their stated long-term goals and take a step back from their emotions during a crisis.

Mitigating action bias takes deliberate effort. Added effort is easy to dismiss when immediate action feels satisfying, but the consequences can be severe. A challenge, even to the countermeasures described, is giving the decision maker an incentive or motivation to think beyond his or her "normal" response of acting too quickly when trying to eliminate a problem. Each of us should have a healthy skepticism about decisions in cybersecurity and practice transparency in our own decisions. Organizations should reward preparation and strategic planning more than unrehearsed crisis responses.

Research is ongoing across the field of computing that has the potential to impact action bias, especially among end users. The study and development of mental models offer an approach to improving decision making. If users better understand what is happening inside their computers or online, they may be less prone to counterproductive actions. Another option is to simplify and normalize deep thought over reflexes. In 2020, researchers proposed "deliberate friction" as a way to help users slow down when making important security decisions.[3] Unfortunately, there is almost no analysis of the decision-making behavior of cybersecurity senior leaders or policymakers despite their enormous effect on users and security in practice.

Further study is also needed. Common metrics, such the mean time to recovery, might be contributing to action bias. Metrics can drive mindset and action, so realigning incentives may be required. Another area of potential exploration is the correlation between innovation and security and privacy incidents. In risk-tolerant environments, such as start-up companies and academia, positive behavior in innovation could increase security and privacy risks. Finally, there appears to be a noteworthy gap in the analysis of the decision-making assumptions and behavior of the roles that undermines both primary (e.g., business) and cybersecurity

outcomes. This gap may be leading to another source of error in cyber decision making: unwanted variability. Not only is there a tendency to do "something" after a cyberincident, such as ransomware or a data breach, but our hypothesis is that the judgment of what to do or how much to spend is wildly variable, even among professionals.

In his most recent book, Kahneman and his coauthors wrote that "Good decision making must be based on objective and accurate predictive judgments that are completely unaffected by hopes and fears, or by preferences and values."[7] Action bias is a systematic error that adversely affects cybersecurity and the outcomes of primary goals, from CEOs to home users. There is no evidence that bias is more or less prevalent in different settings, even in mature organizations; however, maturity may include better risk management, preparation, and rehearsal. Conscious awareness of this error is a necessary prerequisite for change, but more important is the need to carefully and thoughtfully prepare and rehearse before the next inevitable incident. ■

References

1. M. Bar-Eli, O. H. Azar, I. Ritov, Y. Keidar-Levin, and G. Schein, "Action bias among elite soccer goalkeepers: The case of penalty kicks," *J. Econ. Psychol.*, vol. 28, no. 5, pp. 606–621, 2007, doi: 10.1016/j.joep.2006.12.001.
2. P. Cichonski, T. Millar, T. Grance, and K. Scarfone, "Computer security incident handling guide," *NIST Special Publication*, vol. 800, no. 61, pp. 1–147, 2012.
3. V. Distler, G. Lenzini, C. Lallemand, and V. Koenig, "The framework of security-enhancing friction: How UX can help users behave more securely," in *Proc. 2020 New Security Paradigms Workshop*, pp. 45–58, doi: 10.1145/3442167.3442173.
4. "Uber technologies, Inc., In the matter of," Federal Trade Commission, Washington, DC, USA, Aug. 2017. [Online]. Available: https://www.ftc.gov/enforcement/cases-proceedings/152-3054/uber-technologies-inc
5. M. Gambino. "Q and A: Capt. Chesley 'Sully' Sullenberger." *Smithsonian Magazine*. Accessed: Mar. 23, 2022. [Online]. Available: https://www.smithsonianmag.com/science-nature/q-and-a-capt-chesley-sully-sullenberger-63542623/
6. L. A. Gordon and M. P. Loeb, "The economics of information security investment," *ACM Trans. Inf. Syst. Security*, vol. 5, no. 4, pp. 438–457, 2002, doi: 10.1145/581271.581274.
7. D. Kahneman, O. Sibony, and C. R. Sunstein, *Noise: A Flaw in Human Judgment*. New York, NY, USA: Little, Brown Spark, 2021.
8. G. Loewenstein, "Hot-cold empathy gaps and medical decision making," *Health Psychol.*, vol. 24, no. 4, pp. S49–S56, 2005, doi: 10.1037/0278-6133.24.4.S49.
9. A. Patt and R. Zeckhauser, "Action bias and environmental decisions," *J. Risk Uncertainty*, vol. 21, no. 1, pp. 45–72, 2000, doi: 10.1023/A:1026517309871.
10. A. Tversky and D. Kahneman, "Judgment under uncertainty: Heuristics and biases," *Science*, vol. 185, no. 4157, pp. 1124–1131, 1974, doi: 10.1126/science.185.4157.1124.

Josiah Dykstra is a technical fellow in the Cybersecurity Collaboration Center at the National Security Agency, Ft. Meade, Maryland, 20755, USA. His research interests include cybersecurity science, human factors and resilience in cyber, and the economics of cybersecurity. Dykstra received a Ph.D. in computer science from the University of Maryland, Baltimore County. Contact him at josiah.dykstra@cyber.nsa.gov.

Jamie Met is a researcher and innovation specialist at the National Security Agency, Ft. Meade, Maryland, 20755, USA with expertise in computer science, mathematics, and alternative analysis. His research interests include behavioral economics and human factors and their applications to security, operations, and leadership challenges. Met received a B.S. in mathematics from the Virginia Polytechnic Institute and State University. Contact him at jlmet@nsa.gov.

Nicole Backert is a network analyst at the National Security Agency, Ft. Meade, Maryland, 20755, USA with expertise in computer science and networking. Her research interests include theoretical foundations in computer science, data privacy, and secure software development. Backert received a B.S. in computer science from Towson University. Contact her at nmbacke@nsa.gov.

Rebecca Mattie is a master's student in Johns Hopkins University's Cybersecurity Program, Baltimore, Maryland, 21218, USA. Her research interests include cybersecurity and leadership. Mattie received a B.S. in computer science from Northeastern University. Contact her at rmattie2@jhu.edu.

Douglas Hough is a senior associate at the Department of Health Policy and Management, Johns Hopkins Bloomberg School of Public Health, Baltimore, Maryland, 21205, USA. His research interests include identifying the optimal size and structure of a physician practice and the application of the emerging field of behavioral economics to contemporary health care issues. Hough earned a Ph.D. in economics from the University of Wisconsin. Contact him at douglas.hough@jhu.edu.