Atlantic Council

CYBER STATECRAFT INITIATIVE

COUNTING THE COSTS:

A Cybersecurity Metrics Framework for Policy

by Stewart Scott

Atlantic Council

CYBER STATECRAFT INITIATIVE

The **Cyber Statecraft Initiative** works at the nexus of geopolitics and cybersecurity to craft strategies to help shape the conduct of statecraft and better inform and secure users of technology. This work extends through the competition of state and non-state actors, the security of the internet and computing systems, the safety of operational technology and physical systems, and the communities of cyberspace. The Initiative convenes a diverse network of passionate and knowledgeable contributors, bridging the gap among technical, policy, and user communities.

Authors
Stewart Scott

Editor
Samia Yakub

Atlantic Council
1400 L Street NW, 11th Floor
Washington, DC 20005

For more information, please visit
www.AtlanticCouncil.org.

May 2025

Design by: Donald Partyka and Amelie Chushko

Cover: Highsmith, Carol M, photographer. "One is Man Controlling Trade,"statue by Michael Lantz, at Federal Trade Commission, 600 Pennsylvania Ave., NW Washington, D.C. United States Washington D.C. District of Columbia Washington D.C, 2010. Photograph. https://www.loc.gov/item/2010641732/.

Atlantic Council

CYBER STATECRAFT
INITIATIVE

COUNTING THE COSTS:

A Cybersecurity Metrics Framework for Policy

by Stewart Scott

Table of Contents

Acknowledgements

The author would like to thank the many contributors to this piece, including peer reviewers Sara Ann Bracket, Alex Gantman, Stefan Savage, Emma Schroeder, Nikita Shah, and Adam Shostack. Thank you also to Nancy Messieh, Donald Partyka, Amelie Chushko, and Samia Yakub for their work editing, designing, and producing this report.

Executive Summary

US cybersecurity policy has a critical blind spot: the absence of reliable outcome metrics that can inform policymakers about whether the digital ecosystem is becoming more secure and which interventions are driving progress most effectively. Despite years of strategies, regulations, and best-practices campaigns, the field of cybersecurity metrics has room to grow, and policymakers still lack answers to fundamental questions. How much harm are cybersecurity incidents causing? Are things getting better or worse? Which policies deliver the greatest return on investment for reducing realized harm and the risk of future harm?

This report identifies two core problems holding back progress: first, the unknown state of the system, meaning policymakers cannot empirically describe how secure or insecure the digital landscape currently is; and second, unmeasured policy efficacy, which prevents policymakers from comparing which interventions are most effective at improving security and reducing harm. The result is a policymaking environment heavily reliant on intuition, anecdote, incomplete data, and proxy measures—all unsustainable for a domain with such systemic and escalating risks and so much security investment. To address these challenges, the report proposes a reframing of cybersecurity metrics along two dimensions:

1. Treating cybersecurity as a complex system—acknowledging that incident outcomes result from dynamic, probabilistic interactions between policies, technologies, adversaries, and users.

2. Focusing on harm as the key outcome metric—shifting emphasis from internal system attributes (e.g., the number of vulnerabilities discovered) to the real-world impacts of cyber incidents, such as financial losses, operational disruptions, and physical damage.

The report then explores the current limitations of available metrics, illustrating how wide-ranging estimates of incident costs and inconsistent data collection methods hamstring policymakers. It outlines the difficulty of measuring and interpreting harm data at scale due to factors such as silent failures, complex indirect costs, and underreporting, but it argues that such challenges are not insurmountable and that a desire for perfect metrics cannot impede progress toward better ones. Finally, the paper offers two actionable recommendations for near-term progress:

1. Strengthen existing reporting requirements (e.g., CIRCIA, SEC disclosures) to include consistent, updated measures of incident impact.

2. Centralize responsibility under a single federal entity to aggregate, analyze, interpret, and publish cybersecurity harm data across sectors.

While perfection in cybersecurity metrics may be impossible, measuring harms is the most direct way to track progress and guide investment and the most critical metric to bolster policymakers' toolkit. Without such measurement, the United States risks continuing to navigate a complex, evolving system with an incomplete map.

Introduction

A recurring theme in cybersecurity policy is the failure to quantitatively describe the end state toward which it aims, or even to enumerate what metrics should be measured to that end. How many incidents occur, how much damage do they cause, and to whom? If these are the metrics to consider, what is their desired level and by how much does cybersecurity need to improve to get there? And if not these metrics, then which?

In rare moments when policymakers clearly define cybersecurity outcomes, they tend toward absolutes of dubious achievability; for example, "prevent catastrophe" and "defeat ransomware."[1] Even complex legislation and national strategies,[2] while attempting to alter the incentives around building and using technology, rarely offer more than a glancing, qualitative description of what they strive for—a far cry from the clear, numerical state measurements and milestones in other spheres of public policy, such as inflation and unemployment rates for the Federal Reserve.

Even though more empirically developed policy fields such as economics still face routine crisis, US cybersecurity policymakers must adapt to the dizzying complexity, rate of change, and potential impact of failure in today's digital systems by taking exactly that step toward better measurement. It is critical to understand the current state of cybersecurity, set quantitative goals for its improvement, and assess the efficacy of government policies against those goals. "Intuition alone is insufficient to manage a complex system," as former National Cyber Director Chris Inglis put it.[3] Without specifying target outcomes, there is little incentive to establish critical baseline measures in the first place. Identifying the effectiveness of specific policies at improving security and the cost of their implementation is a step even farther, and the quantitative toolkit required for the US government to make that step has not yet been created. The novelty and dynamism of the digital domain mean that policy missteps will happen, but without that toolkit, identifying which remedies fall short and which succeed—let alone by how much—will remain extraordinarily difficult, if not impossible, all while the rapid integration of digital systems across all levels of society increases the impacts and risks of cyber incidents.

This paper aims to reboot and reorient a long-simmering debate around cybersecurity metrics for the policy community. It starts with context about the state of and need for better cybersecurity measurement by discussing two central and related problems created by the field's empirical immaturity:

- Insufficient cybersecurity metrics mean that government cannot empirically assess, across the digital domain, whether cybersecurity is good or bad, improving or deteriorating.

- Insufficient cybersecurity metrics also complicate the task of evaluating and prioritizing security practices and policies based on their efficacy.

After discussing these two problems, this paper offers two framings for cybersecurity metrics critical to improving their usefulness to policymakers: treating cybersecurity as a complex system and measuring harms. The guiding thesis of this paper is that the harms, in the broadest sense, caused by cyber insecurity are the most important outcome metrics for policymakers. Harms here refers to the bad things caused by cybersecurity incidents, from direct loss of money to intellectual property theft, from the compromise of national security information to the erosion of competitive economic advantage. Metrics for those harms at the macro level are an essential tool for policymakers seeking to manage and improve cybersecurity. After all, cybersecurity policymakers' driving mandate is to reduce realized harms and the risk of future harm as much as reasonably possible, whether through increasing economic competitiveness, securing critical infrastructure, imposing costs on adversary activities, managing strategic competition, or any number of methodological priorities.

This paper does not claim a lack of effort in policy or technical circles at quantifying security, and indeed elements

1 "National Cybersecurity Strategy," The White House, March 1, 2023, https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf.

2 "Secure by Design: Shifting the Balance of Cybersecurity Risk," Cybersecurity and Infrastructure Security Agency, October 25, 2023, https://www.cisa.gov/sites/default/files/2023-10/SecureByDesign_1025_508c.pdf; "Bicameral, Bipartisan Leaders Introduce Legislation to Strengthen Federal Cybersecurity," US Senate Committee on Homeland Security and Governmental Affairs, July 12, 2023, https://www.hsgac.senate.gov/media/dems/bicameral-bipartisan-leaders-introduce-legislation-to-strengthen-federal-cybersecurity/.

3 Katherine Golden, "National Cyber Director Chris Inglis: We Need to Become a 'Harder Target' for Our Adversaries," New Atlanticist, August 4, 2021, https://www.atlanticcouncil.org/blogs/new-atlanticist/national-cyber-director-chris-inglis-we-need-to-become-a-harder-target-for-our-adversaries/.

in both communities have been trying admirably for quite some time.[4] Moreover, even without a broad base of empirical data, policymakers make much use of threat intelligence, observed trends, risk assessments, and other sources of evidence. Instead, this paper suggests a starting point for identifying, measuring, and analyzing cybersecurity outcomes with the goal of reorienting and rebooting these debates rather than arriving at a final answer. After discussing cybersecurity as a complex system and outcomes in terms of harms, this paper analyzes different approaches to interpreting outcome data. Finally, this paper proposes initial policy steps toward improving cybersecurity outcome data.

Importantly, these recommendations do not aim at some final architecture for perfect cybersecurity statistics—such policy systems take time, trial, and error to create in any field. Instead, they combine practical changes and a broader policy reframing to move the needle of cybersecurity policy toward realistic empiricism, while recognizing the risks of both cynicism and perfectionism. Empirically characterizing cybersecurity at the macro level and the efficacy of specific security policies is difficult but not hopeless. And while no policy system for metrics is perfect—debates in more matured fields such as public health, law enforcement, and economics abound—that does not render them all useless.

Two problems

Unknown system state: What is "the problem"?

The first issue created by insufficient cybersecurity metrics is that they leave policymakers with no concrete way to describe the current degree of harm caused by insecurity. More than a decade ago, Dan Geer listed several fundamental cybersecurity questions offered in the context of a conversation with a firm's chief information security officer (CISO): "How secure am I? Am I better off than this time last year? Am I spending the right amount of [money]? How do I compare to my peers?"[5] These questions are as important for policymakers, and as difficult for them to answer, as when originally posed in 2003.[6] The primary US cyber policy coordinator, the Office of the National Cyber Director (ONCD) argued in 2024 that they were not answerable at all. A Government Accountability Office (GAO) report on the 2023 National Cybersecurity Strategy (NCS) criticized the NCS for its lack of "outcome-oriented performance measures," as well as ignoring "resources and estimated costs," to which the ONCD responded that "such measures do not currently exist in the cybersecurity field in general,"[7] and the claim rings true. Current cybersecurity metrics and the field's state have, after at least two decades, failed to provide policymakers with ways to answer the foundational question "how are we doing at cybersecurity?" at the highest level.

And yet, a general intuition that the current state of US cybersecurity is suboptimal animates industry, government, and the public alike. Headlines dominated by costly cybersecurity incidents and predictions that things will deteriorate without drastic change feed this perception. For example, former US Deputy National Security Advisor Anne Neuberger summarized data from the International Monetary Fund (IMF) and Federal Bureau of Investigation (FBI) data as suggesting that "the average annual cost of

4 Dan Geer, "Measuring Security," (Metricon 1.0, Vancouver, British Columbia, Canada, August 1, 2006), http://all.net/Metricon/measuringsecurity.tutorial.pdf; "Cost of a Cyber Incident: Systematic Review and Cross-Validation," Cybersecurity and Infrastructure Security Agency, October 26, 2020, https://www.cisa.gov/sites/default/files/publications/CISA-OCE_Cost_of_Cyber_Incidents_Study-FINAL_508.pdf; "Cross-Sector Cybersecurity Performance Goals (March 2023 Update)" Cybersecurity and Infrastructure Security Agency, March 2023, https://www.cisa.gov/sites/default/files/2023-03/CISA_CPG_REPORT_v1.0.1_FINAL.pdf.

5 Geer, "Measuring Security."

6 Dan Geer, Kevin Soo Hoo, and Andrew Jaquith, "Information Security: Why the Future Belongs to the Quants," IEEE Security & Privacy 1, no. 4 (July-August 2003): 24–32, https://doi.org/10.1109/MSECP.2003.1219053.

7 "Report to Congressional Addressees - Cybersecurity: National Cyber Director Needs to Take Additional Actions to Implement an Effective Strategy," US Government Accountability Office, February 1, 2024, https://www.gao.gov/assets/d24106916.pdf.

cybercrime worldwide is expected to soar from $8.4 trillion in 2022 to more than $23 trillion in 2027."[8] At appreciable fractions of global GDP, these are dire numbers that all but mandate extreme intervention. The hypothesis behind this metric is that the current amount of harm caused by cyber incidents could be reduced by interventions less costly than the consequences of their absence. But intervention against what, and how? Testing and refining that thesis with quantitative data is a critical first step too often overlooked—how much harm do cyber incidents cause? How much would it cost to implement recommended interventions? How much harm would they prevent? Is the cost of preventing security incidents actually lower than the costs that those incidents impose? And above all, if the current level of harms is deemed unacceptable, what would be considered acceptable? Current metrics are unable to provide answers at a scale useful to policymakers, leaving them with no baseline measures against which to judge policy efficacy.

In absence of this key outcome data, cyber policy conversations frame metrics as, at best, an after-action exercise for validating efficacy, rather than the first critical step in defining the problems they seek to solve. Even then, empirical impact assessments are rare. The NCS's "Assessing Effectiveness" section underlines this, providing just one paragraph on the strategy's final page, with a key progress report that failed to materialize before the change in administration.[9] The document's accompanying implementation plan (the National Cybersecurity Strategy Implementation Plan, or NCSIP) reduces assessment to determining whether proposed policies were enacted and whether a budget for them was created, and nothing more.[10] These are useful measures of output for policymakers, but do little if anything to track empirically how implementing the NCS changes the cybersecurity landscape; the strategy largely forgoes assessing its external impact, focusing instead on implementation—a familiar state for cyber policy, which more often concerns itself with adoption rates and completion progress than tangible effect on security outcomes.[11] If policymakers cannot, from the outset and at a high level, measure how they are doing at cybersecurity, all follow-on policy rests on a flawed foundation and it will be difficult to empirically demonstrate success.

Policymakers must use cybersecurity metrics as the foundation for characterizing the status quo, identifying specific problems with it, and shaping solutions. When the GAO asks what outcomes would demonstrate the success of the NCS, the ONCD should be able to respond by pointing to the very issues and data motivating the creation of NCS in the first place. The usefulness of measuring incident costs is relatively uncontroversial and has long frustrated policymakers—see for example a 2020 CISA study on just that problem and its associated challenges.[12] However, both cybersecurity policymaking writ large and efforts to imbue it with better metrics would benefit greatly from approaching metrics as a step toward problem definition first, then as solution assessment. Otherwise, the logical chain of cyber policymaking is broken, producing unbounded solutions with no clear, quantified statement of the problems they hope to solve, and thus no clear outcomes to strive for and measure success against.

Policymakers and practitioners are right to lament the dearth of cybersecurity statistics to inform their work, but they cannot afford to wait for the empirical field to mature on that same decades-long trajectory—they must proactively work to define, gather, and respond to cybersecurity metrics. It is unlikely that government can avoid a central role in gathering macroscale metrics and wait for the data they need to be developed for them. Monetary policy is guided by and assessed against the Consumer Price Index (CPI) and the unemployment rate, both of which are measured by the US Bureau of Labor Statistics. National crime statistics are collated and analyzed through the FBI's Uniform Crime Reporting Program. The Center for Disease Control's National Center for Health Statistics gathers a variety of public health metrics from across the country, as well as globally. Each of these programs is the result of decades of iterative policymaking and partnerships with experts in industry, academia, state and local governments, and civil society. The federal government has the clearest incentives and best means to gather metrics on a scale sufficient to describe the full ecosystem and assess policy efforts to shape it. Policymakers do require better cybersecurity metrics to guide them, but they have an active role to play in creating those tools.

8 "Digital Press Briefing with Anne Neuberger, Deputy National Security Advisor for Cyber and Emerging Technologies," US Department of State (transcript), October 18, 2023, https://2021-2025.state.gov/digital-press-briefing-with-anne-neuberger-deputy-national-security-advisor-for-cyber-and-emerging-technologies/.

9 "National Cybersecurity Strategy," The White House.

10 "National Cybersecurity Strategy Implementation Plan," The White House, July 13, 2023, https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/07/National-Cybersecurity-Strategy-Implementation-Plan-WH.gov_.pdf.

11 "Cybersecurity: National Cyber Director Needs to Take Additional Actions."

12 "Cost of a Cyber Incident."

For cybersecurity, some nascent policies might provide useful insight on data gathering and starting points for more matured, coordinated programs: for example, the FBI's Internet Crime Complaint Center (IC3) database,[13] the upcoming implementation of the Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA),[14] the Securities and Exchange Commission's (SEC) material cyber incident reporting requirements,[15] and so on. All either currently or soon will gather data on cybersecurity incidents, but there is little consensus about what to measure and how, and worryingly little progress toward data collection at the ecosystem scale.[16] For a young field—cybersecurity dates back to the 1970s as a defined field at the earliest, whereas econometrics began developing in the early 1930s—that status quo is understandable, but untenable.[17]

Unmeasured efficacy: What interventions address "the problem" best?

Second, insufficient cybersecurity metrics leave policymakers without measures of how effective specific policies are, meaning they can do little to prioritize or update policy interventions based on metrics. Policymakers are in the business of battling with long-perceived market inefficiencies that lead firms to under- and mis-invest in cybersecurity.[18] For now, they do so through recommendations and requirements about security practices and reporting for certain sectors, products, and entities. The past few years have seen a flurry of movement in cyber policy, from the National Cybersecurity Strategy and its dozens of implementation objectives to agency-led efforts such as CISA's Secure by Design (SBD) Initiative and the SEC's new cyber incident reporting requirements, several critical executive orders, and even an effort designed to harmonize the many existing and forthcoming regulations.[19]

Choosing the initiatives to pursue and those to reinvent or discard requires an understanding of their ultimate impact on cybersecurity outcomes. Determining which policies are effective—when measured against the cost of their implementation—requires quantifying the costs of incidents that they prevent or mitigate. A firm's ability to decide which SBD principles to prioritize necessitates understanding their cost and efficacy. And yet there are only early efforts at ranking these practices by their effectiveness, which challenges any attempt to identify the most urgent security practices or product security features to implement.[20] In short, no one knows what the best thing to do is, whether that be policymakers deciding what practices to require or industry deciding which to implement, only a great number of security practices that are probably good to try.

This is more than simply an optimization challenge. Seemingly potent security controls can lead to unexpectedly poor outcomes, especially in a complex system. For example, the National Institute for Standards and Technology (NIST) prescribes security practices for federal agencies and their contractors, and industry writ large often uses its guidance documents as a starting point for security policies even when a company's compliance is not required. One such publication, NIST SP-800-63B, offers recommendations on digital identity systems, including guidance about account credentials. Past versions of the document suggested the use of complex characters (a mix of numbers, capital and lowercase letters, and special symbols) and frequent password resets to prevent attackers from using dictionaries of common passwords to quickly guess their way into account access. The thinking was that complex characters would require attackers to brute force passwords (i.e. guess all possible combinations of characters in a password), and that the frequent rotation of credentials would limit the window of time in which attacks could guess a password successfully, since attackers would need

13 "Federal Bureau of Investigation Internet Crime Report 2023," Federal Bureau of Investigation Internet Crime Complaint Center, April 4, 2024, https://www.ic3.gov/AnnualReport/Reports/2023_IC3Report.pdf.

14 "Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA)," Cybersecurity and Infrastructure Security Agency, https://www.cisa.gov/topics/cyber-threats-and-advisories/information-sharing/cyber-incident-reporting-critical-infrastructure-act-2022-circia.

15 *Cybersecurity Disclosure*, US Securities and Exchange Commission (statement of Erik Gerding, Director of SEC's Division of Corporation Finance), December 14, 2023, https://www.sec.gov/newsroom/speeches-statements/gerding-cybersecurity-disclosure-20231214.

16 "Federal Bureau of Investigation Internet Crime Report 2023;" "Cybersecurity: National Cyber Director Needs to Take Additional Actions."

17 Olav Bjerkholt, "On the Founding of the Econometric Society," *Journal of the History of Economic Thought 39* (March 6, 2017): 175–98, https://doi.org/10.1017/S105383721600002X.

18 Ross Anderson, "Why Information Security Is Hard - An Economic Perspective," Keynote remarks, Seventeenth Annual Computer Security Applications Conference, New Orleans, LA, 2001, 358–65, https://doi.org/10.1109/ACSAC.2001.991552.

19 Jason Healey, "What the White House Should Do Next for Cyber Regulation," Dark Reading, October 7, 2024, https://www.darkreading.com/vulnerabilities-threats/what-white-house-next-cyber-regulation; "Request for Information on Cyber Regulatory Harmonization; Request for Information: Opportunities for and Obstacles To Harmonizing Cybersecurity Regulations," Office of the National Cyber Director, August 16, 2023, https://www.federalregister.gov/documents/2023/08/16/2023-17424/request-for-information-on-cyber-regulatory-harmonization-request-for-information-opportunities-for.

20 Daniel W. Woods and Sezaneh Seymour, "Evidence-Based Cybersecurity Policy? A Meta-Review of Security Control Effectiveness," *Journal of Cyber Policy* 8, no. 3 (April 7, 2024): 365–83, https://doi.org/10.1080/23738871.2024.2335461.

to start over after every rotation. The reality was different. Users rotated between similar passwords, often repeating old ones, and attackers developed dictionaries to quickly guess at common, easily remembered uses of complex characters like the suffix "123!" and substituting numbers for letters.[21] In other words, the intuition behind the practice was sound, but the ecosystem (users, here) reacted in a way that made the recommended practice insecure and costly.

Without metrics to provide an empirical understanding of the tradeoffs that recommended security practices create in practice, policymakers remain at risk for similar situations. For example, inconvenient authentication requests from multi-factor authentication (MFA) might lead users to share credentials in an insecure manner; rewriting software into notionally memory-safe programming languages might be effective at improving security but more costly than the incidents it prevents; or zero trust architectures might fail to meaningfully improve security across the digital ecosystem so long as they are not adopted past some unknown threshold. Without improving cybersecurity metrics, there is simply no way to know how new practices interact with the full ecosystem.

Reframing cybersecurity metrics

To address the connected problems cited above, policymakers must take two critical steps to reframe and develop their approach to empirical cybersecurity: to treat the digital domain as a complex system, and to measure incident harms as their key guiding outcome metric. These are closely related—understanding causality within a complex system and making predictions based on the arrangement of that system at any point in time are immensely difficult. Instead, focusing on the system's outcomes (here, incident harms) over the system's specific characteristics at a point in time (e.g., the adoption rate of memory-safe languages) will help policymakers avoid the trap of claiming progress in shaping behaviors without producing evidence that said behaviors have improved the cybersecurity status quo.

Treating cybersecurity as a complex system

Treating the cybersecurity landscape as a complex system-of-systems is key to assessing its status quo. This is the fundamental mandate for policymakers—to reduce bad cybersecurity outcomes across the board,[22] and not just for the handful of firms that can measure their own implementation and outcomes well. Accordingly, visibility into as much of the ecosystem as possible is critical. A systems approach also helps policymakers deal with the domain's complexity, which might lead to unforeseen interactions between policy interventions, technology design choices, and cybersecurity outcomes. The digital ecosystem has two key features that, unaccounted for, could mislead policymakers significantly as they approach improving its security: probabilistic incidents and extraordinary dynamism.

First, there is no deterministic formula to predict whether a cybersecurity incident will occur, when, or with what severity. An entity with extraordinary security practices might find themselves the target of an extremely sophisticated adversary or might remain critically vulnerable because of one simple oversight. Equally, a firm with poor security practices might avoid compromise by pure luck. While this probabilism is somewhat self-evident, it means that data with too small a sample size over too short a duration could significantly mislead policymakers. For example, observing fewer bad outcomes for a specific sector might indicate that changes to security practices in that field are stumping attackers who are now comprising fewer targets in general, or instead that attackers have simply moved on to another sector for any number of reasons without a net change in the ecosystem. There are hard limits on the usefulness and broad applicability of data provided on or by a handful of firms over a few years, and yet the majority of cybersecurity data available to the public today is often presented in the form of corporate annual reports.

Second, the ecosystem is constantly and rapidly changing and interacting with itself. Adversaries in the digital ecosystem are adaptive, the technologies they target change daily, the incentives of firms building technologies and those using them are in constant flux, and so

21 "The New NIST Guidelines: We Had It All Wrong Before," Risk Control Strategies, January 8, 2018, https://www.riskcontrolstrategies.com/2018/01/08/new-nist-guidelines-wrong/.

22 The precise meaning of "reduce" will be discussed later on.

on. Dynamism and unexpected interactions have consequences for measurement. By way of example, recall the NIST password guidance example cited earlier. All else remaining equal, passwords immune to dictionary attacks and changing too often to be brute forced would reduce account compromises, but all else does not stay equal in a complex system. The guidance changed user behavior in way that made accounts more vulnerable instead. Similarly, the relationships between security practices and outcomes are not immutable—techniques that stop would-be attackers one year might do little to slow them down the next as they refine their tactics and develop new tools. Capturing data on how outcomes in the entire digital domain shift over time is critical if policymakers hope to understand and manage it as a complex system. This should increase the urgency with which policymakers strive to better measure cybersecurity outcomes, as the relative lack of historical data means it will take time for newly gathered data to be of significant use.

To illustrate these dynamics in practice, consider the straightforward government-led disruption campaigns that the National Cybersecurity Strategy recommended,[23] in which law enforcement organizations or the military attack the infrastructure of malicious actors to prevent their campaigns from causing harm. Fewer attackers carrying out less malicious activity should be a boon to the ecosystem, and the US government (with international partners' assistance) accordingly increased the pace of its disruption operations through a combination of sanctions, prosecutions, and offensive cyber activities as part of its Counter Cybercrime, Defeat Ransomware strategic objective.[24] And yet, Microsoft measurements appeared to show that the volume of ransomware attacks nearly tripled in the final months of 2024.[25] Without vastly improved cybersecurity data, it is difficult what to make of these two facts. It might be that disruption campaigns mitigated some attacks, tempering cybercrime even as it continued to grow—if for example, without those disruption operations, ransomware attacks might have quadrupled. Alternatively, the expensive government interventions might have had little impact on the efforts of attackers who could easily buy or write new malware, procure new command and control servers, and move on to less well-defended targets. The disruption campaigns in this model might simply have prevented attacks against specific targets but shifted the attention of the attackers to undefended victims without a net effect. A third possibility is somewhere in between—disruption campaigns might work to reduce incidents at the ecosystem scale but with little return on investment. The thwarted incidents might've been drops in the ocean of cyber malfeasance not meriting the cost of disruption. Without macro-scale data or insight into specific adversary decision-making, there is no real way to know which of these models applies over a relatively narrow timeframe, let alone historical data.

The graphic illustrates a high-level mapping of the digital ecosystem as a complex system, sorting potential metrics into three categories: inputs, attributes, and outcomes. Inputs are forces, policies, and decisions that are largely external to the digital ecosystem, though no doubt shaped by it. These are the incentives that drive decision making within the ecosystem, its technological design and development, and so on. By far the two most dominant inputs are market incentives and policy choices, which drive investment, design, and decision making within the cyber ecosystem. Attributes are measures or descriptors of the ecosystem itself. Within the ecosystem, firms, attackers, defenders, IT infrastructure, and connected real-world systems all interact at a vast scale and rapid pace in a blend of technical, social, and economic subsystems. These attributes provide the vast majority of cybersecurity metrics available today—for example, vulnerability counts and severity, incident frequency, and the adoption rate of various security practices and products. Parsing the ecosystem and its specific components—its attributes—provides much utility, especially to specific entities within it, but that analysis must be taken with a grain of salt. The ecosystem is constantly changing, its various components interact with different degrees of coordination, and how those forces balance out in the long run is difficult to understand, let alone predict. This system-of-systems produces outcomes in the form of benefits (the efficiency, productivity, and innovation enabled by the digital ecosystem) and harms—the material damage caused by incidents.

23 "National Cybersecurity Strategy," The White House.

24 "US and UK Disrupt LockBit Ransomware Variant," US Department of Justice, February 20, 2024, https://www.justice.gov/archives/opa/pr/us-and-uk-disrupt-lockbit-ransomware-variant.

25 Matt Kapko, "Microsoft Reveals Ransomware Attacks against Its Customers Nearly Tripled Last Year," *Cybersecurity Dive*, October 16, 2024, https://www.cybersecuritydive.com/news/microsoft-customers-ransomware-attacks-triple/730011/.

Figure 1

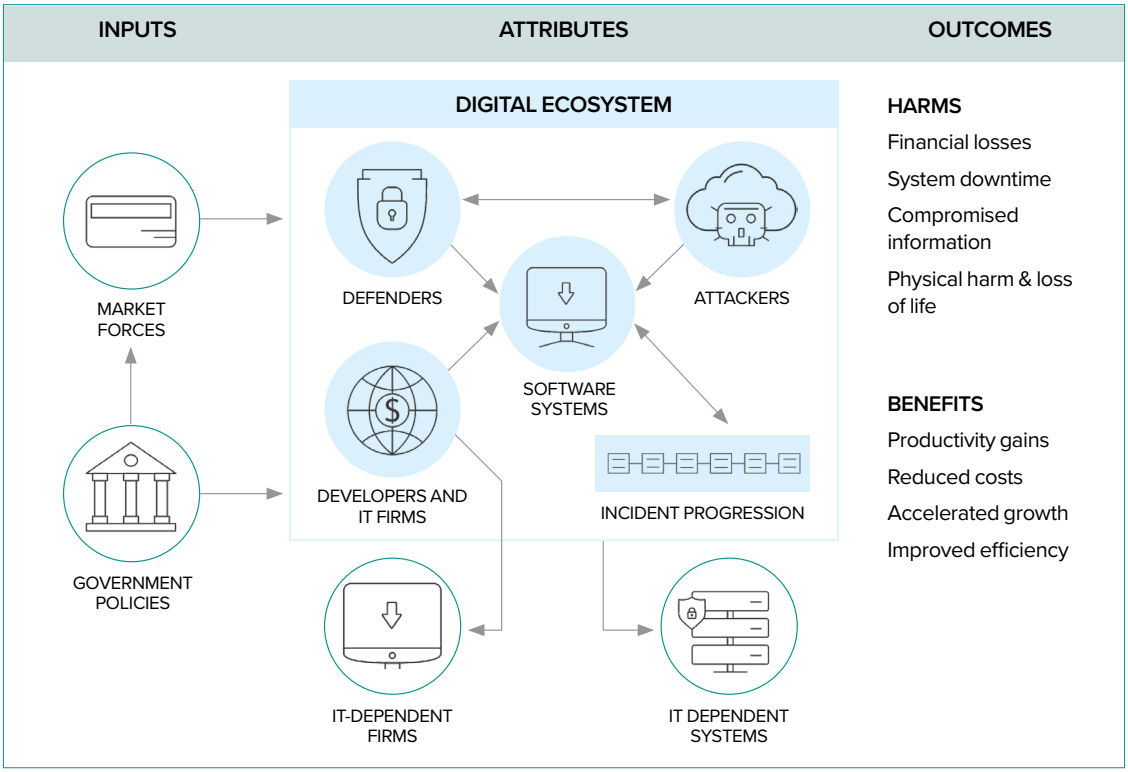

The goal of this mapping is to highlight how policymakers like those at the ONCD, CISA, or similar are interacting with the digital ecosystem at a different scale than firms and individuals. Many of the metrics useful to an individual firm are attributes, and they take on different meanings and behaviors for those concerned with system-of-systems security. For instance, vulnerability counts might tell a cloud provider what problems they have to fix, how often it creates those problems for itself, and how much effort to invest in patching. However, for policymakers, vulnerability counts indicate some vague blend of deficiency in technology design and success in vulnerability detection. Moreover, at the ecosystem scale, attributes interact with each other and outcomes in unpredictable or unknown ways—for example, it is unclear how attacker behavior adjusts to security practice changes at scale and with what effect on outcomes.

Importantly, this framing is not a call to anticipate all possible interactions or comprehensively measure all attributes. Such an approach to the management of a complex system is impractical. Rather, the complex system framing should highlight the importance of outcome measurements as a way for policymakers to navigate complexity or at least evaluate its consequences for the full set of stakeholders under their remit.

Measuring harms as outcomes

Taken together, the two abovementioned issues—an unknown system state and interventions with unmeasured efficacy—put policymakers in a difficult position. It is as if the Federal Reserve lacked data on unemployment rates and inflation while, at the same time, not knowing which policy tools most effectively influence those economic outcomes and how the rest of the economy reacts to their use. The task of assessing efficacy is difficult in the absence of data measuring realized harms. The Federal Reserve could not begin to know whether its interest rate hikes tempered inflation if the Bureau of Labor Statistics did not calculate the CPI. The cybersecurity arena resembles this, with policy more often being a response to singular incidents and anecdotes than to hard data, and with myriad vendors offering cybersecurity solutions in what could be charitably described as "a market for silver bullets" while at the same time producing much of the data currently available to inform policymaking.[26] Past incidents and subjective

26 Alex Gantman, "NDSS 2022 Keynote - Measuring Security Outcomes," April 27, 2022, by NDSS Symposium, YouTube, https://www.youtube.com/watch?v=qGD93mJ2ZAU.

anecdotes are helpful for policymakers, to a certain extent, and security products are not all ineffective. However, heuristics and hunches are only half a solution in managing the complexity of the cyber ecosystem. Metrics are the other critical and conspicuously absent component, and the first step to developing solid, ecosystem-wide metrics is figuring out what to measure and how.

The harms caused by cyber insecurity are the most important outcome metrics for policymakers, and measuring those harms at the macro level is essential if policymakers are to meaningfully manage and improve cybersecurity. Reducing bad cybersecurity outcomes in the form of harms, and mitigating the risk of future harm, is the implicit guiding principle of cybersecurity policy, and therefore measuring those harms broadly is the only path toward rigorous, empirical cybersecurity policymaking.[27] Nonetheless, key policymaking offices in the United States seem so far unable to agree on what a cybersecurity outcome even is. The GAO has suggested measuring tallies of CIRCIA reports—i.e., creating raw counts of inci-dents reported from specific sectors—and the frequency of government disruption campaigns; but both are attri-butes, not outcomes.[28] Few, if any, would disagree that reducing the harm caused by cyber incidents is progress, if not the entire point. Focusing on harms as outcomes in this complex system framing is critical to answering the core question about cybersecurity policy's progress for several reasons:

- Harms as outcomes do not depend upon untested hypotheses about the relationships between attributes or their impact on outcomes.
- Harms are distinct from the dynamic system-of-systems that produces them.
- Harms help reduce the breadth of units of measurement when compared to attribute metrics.
- Harms are more salient to the public than the specific security flaws that lead to them.

First, harms are independent of hypotheses about cybersecurity and key to evaluating them. While there is good reason to believe that many cybersecurity practices and policies improve security and thus reduce harms, the empirical evidence backing these beliefs—let alone describing the amount of harm reduction they are responsible for—is vanishingly thin, and sometimes proves those practices to be ineffective or even harmful.[29] It may be that currently identified best cybersecurity practices are indeed effective, but without knowing how the adoption of a practice interacts with the entire digital ecosystem, policymakers cannot make informed decisions about regulations or incentives. For example, MFA-secured accounts are almost certainly more secure than those backed by single passwords, all else remaining equal, but if the security offered by MFA requires a critical threshold of ecosystem adoption,[30] great effort would be wasted if policymakers were content with an adoption rate below this unknown threshold, and even more would be lost if the cost of pushing adoption past that threshold exceeded the losses prevented by the greater security such implementation might lead to. The fact that any given practice can make a given computer system more secure is necessary but insufficient to urge its broad adoption precisely because of both the possibility for unforeseen interactions within the cybersecurity ecosystem and the general lack of information about costs and benefits at the macro scale that single system adoption provides, especially when that system might be connected to a critical power plant or something far more innocuous.

Second, harms are distinct from the system that produces them, rather than descriptive of it. The complex cyber system, as discussed above, contains billions of machines and users interacting and changing at incredible speed across and above the entire planet. While understanding this ecosystem and its internal attributes at any point in time is useful, the fundamental question for policymakers is how much harm its insecurity enables (relative to the benefits it provides). Any description of the ecosystem—for example, the point-in-time adoption rate of security best practices—still requires outcome data to be meaningful, and as attackers find new routes to compromise, the relationship between best practices and the outcomes they influence are ever changing. In other words, ecosystem attributes alone are insufficient metrics. Attributes do not describe the cost of insecurity, but rather the probability of future harm, and even then unreliably until causal links between attributes and outcomes are better understood.

This reduces, over time and with no further context, the usefulness of measures of specific security practice adoption or of the reduction of the number of certain

27 Stewart Scott, "Counting the Costs in Cybersecurity," *Lawfare*, October 9, 2024, https://www.lawfaremedia.org/article/counting-the-costs-in-cybersecurity.

28 "Cybersecurity: National Cyber Director Needs to Take Additional Actions."

29 Woods and Seymour, "Evidence-Based Cybersecurity Policy?"

30 With enough unsecured accounts still accessible, attackers are able to avoid MFA protections entirely.

vulnerabilities.[31] For instance, in data about the types of memory safety vulnerabilities patched at Microsoft during an eight year period, use-after-free vulnerabilities dominated about 50 percent of vulnerabilities in 2015, compared to just 15 percent in 2022.[32] While this data represents discovered rather than exploited vulnerabilities, the corollary for either observation is the same—the digital system changes, so attacker practices change, and thus defensive measures that worked one year can fail to protect a target entirely the next. In this example, a naive analysis might argue that the reduction in use-after-free vulnerabilities over seven years is a sign of security improvement at Microsoft. This conclusion does not account for the concurrent increase in almost all other kinds of memory safety vulnerabilities, nor does it discriminate among which types or individual vulnerabilities led to the most harm. Microsoft's specific work to reduce use-after-free vulnerabilities succeeded, but what that meant for Microsoft's cybersecurity outcomes remains unclear from the data gathered. It might be that use-after-free vulnerabilities were critical to attackers, and their elimination required a costly pivot to other means. It might be that the discovery and exploit techniques used for use-after-free vulnerabilities were easily converted to other exploit paths. Or it might be that use-after-free vulnerabilities were never abused by attackers that much to begin with. Without outcome data, it is difficult to know (as with MFA) if the cost of reducing entire classes of vulnerabilities might exceed the value of reduced harms up to a certain threshold of coverage.

Third, many harms can be expressed in the common unit of dollars—from identity theft caused by data breaches to the value of stolen intellectual property and the costs imposed by system downtime for critical infrastructure providers. Such monetary losses are often measured or measurable by entities that fall victim to cyber incidents as they quantify incurred costs. Harms can be categorized relatively exhaustively:

1. Financial loss—such as ransomware payments, lost revenue, directly stolen funds, and the costs associated with an incident.

2. Physical harm—including loss of life and physical injury.[33]

3. System downtime or disruption—such as the time that a water treatment plant is taken offline, the time that a hospital operates at reduced capacity, or the inability to conduct government functions.[34]

4. Compromised information—including stolen intellectual property, compromised passwords, and emails stolen from government networks.

Harms can accumulate toward other effects too, often greater than the sum of their parts. These might include reputational damage to a firm or state that experience a sufficient number of harmful incidents, psychological harm to a population subject to repeated cyber incidents, the loss of strategic advantage when an adversary has compromised sufficient amounts of national security information, or similar. This last item, compromised information, highlights a critical nuance. While the act of stealing information might in itself be a harm—e.g., damaging the reputation or share price of a firm subject to a massive data breach or revealing to an adversary information about an upcoming operation—more often it creates the risk for future harm dependent on what the adversary does with that information. Stolen information might give an adversary insight into system flaws or offensive tooling they can later exploit, provide them with credentials or personally identifiable information (PII) that they can abuse later, expose intellectual property that can be leveraged for economic gain at the original owner's expense, or similar. Many other attributes of the complex cyber system contribute to the risk of future harm, from adversary prepositioning operations to the availability of data backups, or the average speed of patching critical vulnerabilities. Nonetheless, for policymakers, understanding how risks of future harm can manifest requires analysis of realized harms.

Overall, systematically measuring harms caused by cybersecurity failures can significantly contribute to

31 While these are not the only challenges that such measures face, they are the most definitional ones. For example, measures of known vulnerability struggle to account for unknown vulnerabilities or the potential for detected vulnerabilities to in reality be harmless given their context.

32 David Weston, "The Time Is Now - Practical Mem Safety," Slide presentation, Tectonics 2023, San Francisco, CA, November 2, 2023), https://github.com/dwizzzle/Presentations/blob/master/david_weston-isrg_tectonics_keynote.pdf.

33 There is often understandable distaste at lumping in physical harm with damages measured in dollars, but fortunately few deaths have ever resulted directly from cyberattacks. Moreover, a combined approach of tallying fatalities, financial damage, and injuries is how the impact of natural disasters is already measured. For more, see "How Can We Measure the Impact of Natural Disasters?," World Economic Forum, March 16, 2015, https://www.weforum.org/stories/2015/03/how-can-we-measure-the-impact-of-natural-disasters/.

34 Scott, "Counting the Costs in Cybersecurity."

understanding how much more or less secure the digital ecosystem is while helping to simplify the complexity and dynamism of the ecosystem, balancing and contextualizing the current focus on its attributes.[35]

The cyber metrics state of play

Policymakers today are not well equipped with the tools to help them describe the system state of cybersecurity over time, nor to measure and rank the efficacy of various interventions and practices in improving that state. Focusing cybersecurity metrics on harms as the key outcome metric for cybersecurity policy helps address these shortcomings while sufficiently navigating the ecosystem's complexity. However, cybersecurity metrics as of now are not up to the formidable task of outcome measurement. This section will detail the challenges of gathering and interpreting data on cybersecurity outcomes and the reality on the ground.

Despite the many industry reports and headlines discussing or predicting global and national costs of cybersecurity incidents,[36] no studies seek to examine differences between reported and forecasted losses, few estimates exhaustively describe their methodologies, cost estimates range significantly, and few predictions are adjusted for changes in the underlying ecosystem.[37] Critically, there is no single source that systematically tracks incident harms across a wide swathe of the ecosystem.

For example, the 2024 IMF Global Financial Stability Report estimated that reported 2022 cyber incident losses were around $5 billion,[38] while the FBI's IC3 report put 2022 losses for just the United States at $10.3 billion.[39] Statista, meanwhile, reports $7.08 trillion in losses for 2022 and projects $12.43 trillion in 2027, while then Deputy National Security Advisor Anne Neuberger's figures were $8.4 trillion and $23 trillion for the same years.[40] Two other reports, from Cybersecurity Ventures and Comparitech, esti-mate 2022 losses at $6.9 trillion and $42.8 billion respectively.[41] Importantly, only the FBI IC3 and IMF reports seem based entirely on confirmed incidents, though Comparitech's might aggregate similar such reporting. Rather than any specific estimate being wrong, the key issue is that few sources use the same methods or scoping, with differences in what is even considered a cyber incident. Additionally, many reports. or similar ones such as Verizon's Data Breach Investigations Report, originate in industry, presenting concerns about long-term availability in the event that a company removes old reports or decides to stop publishing new ones, as well as the potential for conflicting business incentives to shape methodology and reporting.

One 2019 study of the costs of cybercrime summarizes well how these estimates can be further misconstrued, writing "in our 2012 paper, we scaled UK estimates up to global ones…and presented them in a table. We warned that 'it is entirely misleading to provide totals lest they be quoted out of context…' Yet journalists happily ignored this and simply added up the columns, proclaiming large headline figures for global cybercrime—which were essentially twenty times our estimate of UK income tax evasion, as this was the largest figure in the table."[42]

35 Wasted in the sense that such efforts do not answer the macro question, "How secure are we?" These are useful measures in other respects, as enumerated below.

36 "Cybercrime To Cost The World $9.5 Trillion USD Annually In 2024," eSentire, https://www.esentire.com/web-native-pages/cybercrime-to-cost-the-world-9-5-trillion-usd-annually-in-2024; Steve Morgan, "Cybercrime To Cost The World $10.5 Trillion Annually By 2025," *Cybercrime Magazine,* November 13, 2020, https://cybersecurityventures.com/cybercrime-damages-6-trillion-by-2021/; "Unexpectedly, the Cost of Big Cyber-Attacks Is Falling," *The Economist*, May 17, 2024, https://www.economist.com/graphic-detail/2024/05/17/unexpectedly-the-cost-of-big-cyber-attacks-is-falling.

37 At the time of writing, the author was unable to find any source that revised predictive estimates up or down based on new policies, technologies, or geopolitical circumstance.

38 "The Last Mile: Financial Vulnerabilities and Risks," *International Monetary Fund,* April 2024, https://www.imf.org/en/Publications/GFSR/Issues/2024/04/16/global-financial-stability-report-april-2024.

39 "Federal Bureau of Investigation Internet Crime Report 2023."

40 "Estimated cost of cybercrime worldwide 2018-2029," Statista, https://www.statista.com/forecasts/1280009/cost-cybercrime-worldwide.

41 Morgan, "Cybercrime To Cost The World $10.5 Trillion;" Paul Bischoff, "Cybercrime Victims Lose an Estimated $714 Billion Annually," *Comparitech*, December 5, 2023, https://www.comparitech.com/blog/vpn-privacy/cybercrime-cost/.

42 Ross Anderson et al., "Measuring the Changing Cost of Cybercrime," The 18th Annual Workshop on the Economics of Information Security, Boston, MA, June 3, 2019, https://doi.org/10.17863/CAM.41598.

There are several systematic incident reporting processes in the United States that could usefully gather outcome data, but they are not fully realized. The SEC recently began requiring the reporting of material cyber incidents from publicly traded companies, which had already occasionally disclosed such incidents in their filings. However, of the nearly two hundred cyber incident reports (required or not) available at the time of this piece's writing, just seven contain cost estimates.[43] CIRCIA, which has yet to be fully implemented, seems intent on capturing incident impacts, though the tight timeframe within which to report an incident (seventy-two hours) likely means that accurate outcome measurement will have to rely on updates to initial reports.[44] While CIRCIA incident report updates are mandatory in its most recent proposal, whether they will capture outcome data remains to be seen, as full implementation will not begin until 2026.

Other useful incident reporting processes include (but are not limited to):

- FISMA, which requires federal civilian executive branch (FCEB) agencies to report incidents to CISA.[45]
- The US Department of Housing and Urban Development's (HUD) Significant Cybersecurity Incident Reporting Requirements, which covers mortgagees approved by the Federal Housing Administration.[46]
- The Gramm-Leach-Bliley Act requires a variety of financial institutions to report data breaches to the Federal Trade Commission.[47]
- The Federal Communications Commission's updated data breach notification rules, which cover telecommunications carriers.[48]
- The Department of Defense's (DOD) requirements for Defense Industrial Base contractors to report all cyber incidents involving "covered defense information."[49]
- The Department of Health and Human Services' Breach Notification Rule.[50]
- A tapestry of data breach reporting requirements across all fifty states and several US territories, as well as other sector-specific federal requirements both proposed and implemented.[51]

Together, these reporting requirements should notionally cover all publicly traded companies in the United States, critical infrastructure providers, FCEB agencies, and many smaller entities under state laws, with some entities facing multiple reporting requirements. Even more reporting requirements exist in the intelligence community, among defense contractors and recipients of federal grants, and others, while law enforcement captures at least an appreciable number of incidents targeting individuals through the FBI's IC3. Given this sample would represent a massive proportion of the US attack surface, it should provide a sufficient starting point for systematic cybersecurity outcome data, if properly arranged to gather such data and coordinated to arrive at central clearing agency for analysis. Even then, disincentives to accurate reporting have long plagued cybersecurity,[52] and the challenges in arriving at useful estimates of harms are significant.

Difficult numbers

Even with a robust reporting system tailored to capture incident costs from all the above sources while avoiding

43 "Cybersecurity Incident Tracker," Board Cybersecurity, last updated March 3, 2025., https://www.board-cybersecurity.com/incidents/tracker/.

44 "Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA) Reporting Requirements," Department of Homeland Security Cybersecurity and Infrastructure Security Agency,, April 4, 2024, https://www.federalregister.gov/documents/2024/04/04/2024-06526/cyber-incident-reporting-for-critical-infrastructure-act-circia-reporting-requirements.

45 "Federal Information Security Modernization Act," Cybersecurity and Infrastructure Security Agency, https://www.cisa.gov/topics/cyber-threats-and-advisories/federal-information-security-modernization-act.

46 Richard J. Andreano, Jr., "FHA Requiring Reporting of Significant Cybersecurity Incidents," *Consumer Finance Monitor*, May 24, 2024, https://www.consumerfinancemonitor.com/2024/05/24/fha-requiring-reporting-of-significant-cybersecurity-incidents/.

47 "FTC Safeguards Rule: What Your Business Needs to Know," Federal Trade Commission, last updated December 2024, https://www.ftc.gov/business-guidance/resources/ftc-safeguards-rule-what-your-business-needs-know.

48 "Data Breach Reporting Requirements," Federal Communications Commission, February 12, 2024, https://www.federalregister.gov/documents/2024/02/12/2024-01667/data-breach-reporting-requirements.

49 "Defense Industrial Base (DIB) Cybersecurity Portal - Cyber Incident Reporting," Defense Industrial Base (DIB) Cybersecurity Portal, https://dibnet.dod.mil/dibnet/#reporting-reporting-2.

50 "Submitting Notice of a Breach to the Secretary," US Department of Health and Human Services, last reviewed February 27, 2023, https://www.hhs.gov/hipaa/for-professionals/breach-notification/breach-reporting/index.html.

51 "State Data Breach Notification Chart," IAPP, March 2021, https://iapp.org/resources/article/state-data-breach-notification-chart/.

52 Seema Sangari, Eric Dallal, and Michael Whitman, "Modeling Under-Reporting in Cyber Incidents," *Risks* 10, no. 11 (October 22, 2022): 200, https://doi.org/10.3390/risks10110200.

disincentives that lead to underreporting—a far cry from the current status quo—the task of estimating incident outcomes is not easy, with two notable hurdles standing out: silent failures and complex costs.

Silent failures refers to the fact that in cybersecurity, when information is stolen, it often remains present on the victim's system, which makes noticing the compromise and its outcomes challenging.[53] Take for example the extraordinary lag time between the deployment of malicious SolarWinds Orion updates in late March of 2020, and the discovery of the intelligence gathering campaign in December 2020.[54] Attackers might have had access to target systems for at least nine months, with no "missing" data tipping off defenders. Such intelligence gathering is a fundamental feature of the cyber domain, and ensuring most of these compromises are discovered is ultimately a technical challenge, but it remains a key limiter on the value and feasibility of large-scale outcome data. Barring a complete technical solution, analysts will always need to assume that their data conveys an incomplete picture of ecosystem outcomes, especially when information theft is such a fundamental part of cybersecurity incidents.

Complex costs refer to the difficulties of quantifying many of the harms caused by cybersecurity incidents. Broadly, estimating the costs incurred by operational downtime, ransomware payments, and similar incidents is a tractable task for victim entities. However, attaching a dollar figure to harms resulting from stolen information is difficult, even when the extent of that compromise is definitively known, especially where that information might contribute to significant compromise but only when attached to other information (as in the case of linking phone numbers to email addresses to undermine MFA protections). Valuable information might include intellectual property, PII, information with national security value, account credentials, or similar. The quantity of information stolen by attackers and the sensitivity of that information can provide some insight into the risks of future harms, but precise measurement is difficult, especially when not all stolen data is abused successfully or when the abuse serves national security or intelligences ends, which are particularly hard (if not impossible) to quantify.

Complex costs also refer to other difficult-to-notice harms. For instance, the largest source of risk in the cyber ecosystem is its interconnection with effectively all layers of society: a cybersecurity incident can cause direct and immediate harms to any given sector with sufficient dependence on IT systems, affecting a huge number of entities even when only one entity was compromised. Even the most well-architected system for counting the costs of cyber incidents will struggle to accurately track total harms across sectors. These secondary costs can represent the bulk of harm caused by an incident but might remain buried in non-cyber reporting systems, if reported at all. Take, for example, the recent CrowdStrike outage, which led to flight cancellations globally as well as operational disruptions across many sectors. While one report from Parametrix Insurance estimated that the incident carried a net cost of $5.4 billion, tracking those costs all the way through different sector verticals is difficult.[55] The same Parametrix report assessed losses of $860 million for airlines, but the losses reported by just Delta Air Lines in an SEC filing amounted to at least $500 million.[56] This is not to criticize any particular estimate, but rather to highlight both the consequences of inconsistent methodologies and the challenges of tracking costs not funneled through established cyber incident reporting requirements. To the latter, Delta's disclosure came through Item 7.01 of a Form 8-K for reporting specific material events, effectively tagging it as a massive, unexpected cost. Generally, cyber incident disclosures through 8-K forms have been made through Item 8.01 for non-material incidents and the SEC's newly created Item 1.05 for material ones. In other words, accurately capturing all costs from cyber incidents is key to understanding their true impact, as cyber risk is generally a function of the critical role of systems connected to digital infrastructure. At the same time, such estimates are difficult to make and are difficult to capture by singular reporting mechanisms because of their appearance across all sectors.[57]

53 Dan Geer, "Prediction and The Future of Cybersecurity," Remarks, UNC Charlotte Cybersecurity Symposium Charlotte, NC, October 5, 2016, http://geer.tinho.net/geer.uncc.5x16.txt.

54 Trey Herr et al., Broken Trust: *Lessons from Sunburst*, Atlantic Council, March 29, 2021, https://www.atlanticcouncil.org/in-depth-research-reports/report/broken-trust-lessons-from-sunburst/.

55 "Crowdstrike's Impact on the Fortune 500: An Impact Analysis," Parametrix, 2024, https://www.parametrixinsurance.com/crowdstrike-outage-impact-on-the-fortune-500.

56 "Delta Airlines, Inc. Form 8-K Report on August 8, 2024," US Security and Exchange Commission, August 8, 2024, https://www.sec.gov/Archives/edgar/data/27904/000168316824005369/delta_8k.htm. It is alternatively possible that Delta systems were simply more severely impacted that other airlines.

57 "Cost of a Cyber Incident."

Reading the curves: Interpreting outcome data

If policymakers were able to measure with reasonable accuracy and precision the costs of cybersecurity incidents, they could use that data to begin addressing the two outstanding challenges with cybersecurity policy and metrics: assessing efficacy (or return on investment) and benchmarking system state. However, even with accurate measurement, interpretation of such data is not straightforward.

First, measuring return on investment requires the ability to answer two immediate, practical questions: How much harm does a specific practice reduce? How much do we spend where? While the latter is more tractable—expenditure is recorded somewhere, though general IT spend and cybersecurity spend can be difficult to separate in practice—at the micro level, robust outcome data would enable the study of return on investment for money spent implementing specific cybersecurity practices by revealing how much they reduced harms downstream. Heuristically, policymakers approach cybersecurity similarly, striving to maximize breadth and depth of impact against expenditure, but without a robust empirical body of evidence to back them. Such metrics would go a long way in helping prioritize the many different security controls recommended by both government and industry against their observed return on investment. There are some nascent efforts to carry out this analysis, including through CISA's revitalized Cyber Insurance and Data Analysis Working Group,[58] but they are primarily working with insurance claims data, which might not capture the full extent of costs given the above challenges in measurement and insurers' focus on policyholder claims versus net costs to claim holders (aside from the fact that they mainly have data on their customers rather than the ecosystem at large). Broadly, outcome data is the key to making attribute data about security practice implementation meaningful. It is the best way to point policymakers to both the best solutions and the right problems—for example, whether the harms of cybercrime results more from social engineering at scale or exploited vulnerabilities.

The second and more foundational application of complete outcome data is to give policymakers a macro-level picture of the size and nature of the cybersecurity challenge they face—and thus what scale of investment makes sense and what trends in success or failure at addressing cyber risk are worth pursuing. The first question that comes to mind when faced with net annual harms data is whether cybersecurity is improving or deteriorating. Interpreting outcome data is far from straightforward, and there are three broad approaches one might take, each with immediate policy consequences:

1. Uncontrolled metrics
2. Controlled metrics
3. Catastrophic risks

Uncontrolled metrics: More is worse

Uncontrolled metrics refers to simply using total harms figures without further context. Regardless of which existing source one uses, annual tallies of cyber incidents and their costs seem to be increasing, implying that, far from getting better, the state of cybersecurity is on the decline year by year at a more-than-linear rate. This framing of outcome data can be observed on the cover image of Verizon's 2023 Data Breach Investigations Report,[59] raw estimates of annual total incident costs such as Neuberger's figure referenced above, and the GAO's suggestion that ONCD use aggregated ransomware incident and loss data to assess the efficacy of the National Cybersecurity Strategy: incidents are more common year after year, as are best estimates of harms.[60] These are intuitive interpretations—more incidents causing more harm is bad—and, if the numbers are accurate, they do capture some objective truth about what occurs in the digital ecosystem. Such interpretations, however, are immature in comparison to other fields of empirical policymaking. Are harms growing per incident?

58 Nitin Natarajan, "Cybersecurity Insurance and Data Analysis Working Group Re-Envisioned to Help Drive Down Cyber Risk," Cybersecurity and Infrastructure Security Agency (blog), November 20, 2023, https://www.cisa.gov/news-events/news/cybersecurity-insurance-and-data-analysis-working-group-re-envisioned-help-drive-down-cyber-risk.

59 "2023 Data Breach Investigations Report," Verizon, June 2023, https://www.verizon.com/business/resources/T227/reports/2023-data-breach-investigations-report-dbir.pdf.

60 "Cybersecurity: National Cyber Director Needs to Take Additional Actions."

Are there simply more incidents? Or are we getting better at observing and counting more of the incidents that occur?

Controlled metrics: More is relative

A controlled metrics interpretation argues that meaningful cybersecurity metrics must account for the ecosystem's rapidly changing context, which uncontrolled metrics omit. Few other fields use uncontrolled metrics but instead account for changes in population or similar underlying variables. For example, public safety policy cares more about violent crime per capita than overall violent crime because a larger population in and of itself means more potential criminals and victims and therefore more crime in absolute terms. Similarly, the Federal Reserve cares more about the unemployment rate than raw unemployment counts. Parallel arguments could reasonably apply to cybersecurity—each passing day brings more potential cyber criminals, victims, and devices online as internet connectivity increases, and there are more dollars at stake in the digital ecosystem as more business grows intertwined with IT infrastructure. All else being equal, one could reasonably expect these trends to increase the overall number of cybersecurity incidents and losses year to year because, even if security remains constant, there are more people and dollars online. One 2015 study by Eric Jardine made such an argument and normalized cybercrime figures with data on the size of the internet and its userbase. In doing so, it found that most metrics improved year over year, or at least did not worsen.[61]

However, determining a reasonable denominator for cybersecurity is more challenging than in other fields where population is usually sufficient.[62] Financial harms can befall individuals, but also abstract entities like businesses or larger constructs like national economies. It is most likely that a rigorous approach to analyzing harms data will use different denominators for different harms. For instance, the cost of individually targeted cyber fraud works well per capita, while business ransomware payment costs would be more reasonably adjusted by gross domestic product or a similar dollar figure. Control metrics also highlight well the continued importance of attribute measures. This paper does not argue that attribute metrics are irrelevant, but that on their own, they can mislead policymaking in eliding a key part of the complex system—its external impacts.

Catastrophic risk: More to come

A third interpretation of outcome data borrows from the risk management experience of the financial sector by considering the role of catastrophic events. If there are a sufficient number of extremely costly cyber incidents, interpreting time-series outcome data into the future becomes difficult, especially given the relative novelty of the field, which leaves analysts with a limited historical record to study.[63] Similar to the economic growth preceding the Great Recession in 2008, years of improved outcomes might be interpreted as improved cybersecurity, but they might mean little if a significant catastrophe lies just around the corner. Unfortunately, without robust outcome data about past events, evaluating the possible severity, variance, and frequency of cyber catastrophes is challenging, particularly when potential harms might change suddenly with large shifts in geopolitical circumstance (e.g., the risks of cyber catastrophe might grow dramatically when two countries enter a formal war with each other).

One dataset sought to do just that by assembling a list of multi-firm cyber incidents estimated to have resulted in a loss of at least $800 million, inflation adjusted to 2023.[64] The dataset counted twenty-five total catastrophic events, with the worst costing $66 billion, and the average event reaching $14.8 billion. The author concluded that cyber catastrophes are not as significant a risk as often made out based on this data and the observation that these costs are only fractions of the costs that natural disasters can incur. However, things might not be so simple. The cost estimates used are subject to the same measurement challenges mentioned above, which the author notes well: "Unfortunately, many estimates come from popular media sites and corporate blogs."[65]

More specifically, the dataset omits the SolarWinds incident of 2019, for which one analysis estimates $100 billion in costs just for incident response across the thousands of victim organizations alone, not even accounting for the harms resulting from abuse of the information stolen during

61 Eric Jardine, "Global Cyberspace Is Safer than You Think: Real Trends in Cybercrime," Global Commission on Internet Governance, revised October 16, 2015, https://www.cigionline.org/publications/global-cyberspace-safer-you-think-real-trends-cybercrime/.

62 "Technical Report 22-02: Vital Statistics in Cyber Public Health," CyberGreen Institute, March 2022, https://cybergreen.net/wp-content/uploads/2022/04/Technical-report-22-02-Vital-Statistics-in-Cyber-Public-Health-FINAL.pdf.

63 Dan Geer, "For Good Measure: The Denominator," USENIX ;login: 40, no. 5 (October 2015), https://www.usenix.org/publications/login/oct15/geer.

64 Tom Johansmeyer, "Recent Cyber Catastrophes Show an Intensifying Trend – but They Are Manageable," *The Loop*, September 25, 2024, https://theloop.ecpr.eu/recent-cyber-catastrophes-show-an-intensifying-trend-but-they-are-manageable/.

65 Tom Johansmeyer, "Surprising Stats: The Worst Economic Losses from Cyber Catastrophes," *The Loop*, March 12, 2024, https://theloop.ecpr.eu/surprising-stats-the-worst-economic-losses-from-cyber-catastrophes/.

the intelligence gathering campaign, which for the reasons stated above is immensely difficult to quantify.[66] There are also reasonably costly single-firm incidents such as the Equifax breach, omitted by methodology—direct costs to the firm topped $1.7 billion, not to mention the costs of whatever identity theft and fraud may have resulted.[67]

Other data from the IMF about the distribution of cyber incidents by cost shows that, even if cyber catastrophes are less costly than natural disasters, they do present similar irregularity, with most incidents being mild while a handful reach disastrous extremes.[68]

Another method for assessing whether an ecosystem is prone to catastrophic events looks for near misses—almost-incidents that, fully realized, would have been catastrophic and were avoided by chance rather than systematic prevention. In an article about interpreting outcome data, Geer describes how relatively trivial changes to a 2001 malware could have allowed it to block 911 emergency services across the United States, which would certainly qualify as a catastrophic event, and one with difficult-to-quantify psychological harms on top of loss of life.[69] Moreover, given the rapid growth of the cyber ecosystem and its increasingly fundamental role in the functioning of all levels of society, Geer's warning in the paper should temper claims that cyber catastrophes are not that significant: "this proof (that we escaped such an attack by dumb luck) puts to bed any implication that every day without such an attack makes such an attack less likely."[70] In other words, he argues that cyber catastrophes might not have been comparatively as extreme as financial crises or natural disasters, but only so far, and the potential for extreme incident grows as more real-world services rely on relatively homogenous digital systems. This interpretation of cyber metrics holds two key lessons. First, attribute measures can be extremely useful in highlighting the potential for future catastrophe. Just as measures of debt ratios, leveraged capital, liquidity reserves, and more can help analyze financial catastrophes, measures of concentrated dependency, cloud systems resilience, vulnerability patch time, and more can describe the risk posture of the digital ecosystem. Second, while outcome metrics should not be used in an attempt to predict future harms, they are still key to establishing historical record of cyber incidents and catastrophes and understanding the true scale of cyber harms. Again, outcome metrics should not supplant attribute metrics, but instead, at the macro scale, are key for policymakers trying to understand and manage cybersecurity risks and harms.

Starting construction: Two changes

The result of the many measurement challenges and shortfalls in cybersecurity is a set of fundamental unknowns for cybersecurity policymakers. At the ecosystem scale, the cybersecurity status quo remains unmeasured, as does the efficacy of security practices at reducing harms, while a plan to address those quantitative lapses does not yet exist. These obstacles go well beyond making policy optimization difficult. Moreover, as the fundamental question of the size of the cybersecurity problem goes unanswered, the gap in historical outcome data increases and unproven policy and investments grow more entrenched. These challenges should not, however, prompt paralysis. More measurement, even if imperfect, can improve the empirical toolkit of policymakers, and there is good reason to believe that some policies and security interventions, even if not empirically shown, have improved cybersecurity.

With all this in mind, the US government should use the abundant reporting requirements already in existence to begin assembling a robust cybersecurity metrics system comparable to the already established thirteen federal

66 Gopal Ratnam, "Cleaning up SolarWinds Hack May Cost as Much as $100 Billion," *Roll Call*, January 11, 2021, https://rollcall.com/2021/01/11/cleaning-up-solarwinds-hack-may-cost-as-much-as-100-billion/.

67 Ben Lane, "Equifax Expects to Pay out Another $100 Million for Data Breach," *HousingWire*, February 14, 2020, https://www.housingwire.com/articles/equifax-expects-to-pay-out-another-100-million-for-data-breach/.

68 "The Last Mile: Financial Vulnerabilities and Risks," *International Monetary Fund*.

69 Geer, "For Good Measure: The Denominator."

70 Geer, "For Good Measure: The Denominator."

statistical agencies serving the fields of public economics, education, agriculture, public health, and more.[71] Building such infrastructure and pulling meaningful analysis from the data it assembles will take time, but waiting only delays a fundamentally necessary process. Additionally, developing a new policy lens is as important as creating new policy mechanisms, and questions about measurable efficacy and return on investment should become commonplace in policy conversations. Below are two small recommendations focused on existing reporting processes and offices

Counting harms

Given the importance of gathering outcome data both to understanding the cyber ecosystem and to making useful already-gathered attribute data, existing reporting requirements should incorporate impact estimates more rigorously. CISA's final implementation of CIRCIA should include explicit provisions requiring at least one update to incident reports that includes a revised estimate of incident impact and notes on the methodology used to reach that estimate. This information will help CISA weight incidents by their impact and provide a large inflow of outcome data from all critical infrastructure sectors. Similarly, the SEC should update its guidance on cyber incident reporting to include similar requirements—Item 1.05 reports in 8-K filings should be updated at least once with impact estimates from the reporting company in a similar format as above and updated when the reporting entity arrives at a final estimate. Given Item 1.05 reports only apply to material cyber incidents, they already require the information leading to the determination of materiality, which already should assess incident impact, although there is ongoing debate about the difference between a material event and an event with material impact.[72] Thus, not only should this data be generated already by the reporting company, but it is precisely the kind of information relevant to the shareholders that the item is designed to inform.

Like CIRCIA and SEC filings, all federal reporting requirements should include provisions mandating that outcome metrics and information be updated as an incident unfolds and an affected entity revises its estimates. Altogether, with tweaks to existing or forthcoming reporting requirements, the federal government can gather incident outcome data from publicly traded companies, critical infrastructure entities, DOD contractors, FCEB agencies, and others, creating a significant sample of high-quality outcome data without the need for new reporting regimes.

One office to count them all

Given the potential volume of outcome data from a wide variety of reporting sources and regulations, meaningful interpretation of that information requires that it flow to one entity, similar to how the Bureau of Labor Statistics collates price data from hundreds of goods and services in calculating the Consumer Price Index.[73] Fortunately, the US Department of Homeland Security's Office of Homeland Security Statistics (OHSS) is already on a course to assume this central role, with plans to report on cybersecurity incidents shared its way as a result of reforms to FISMA in 2025. This office should be enlarged and report annually on cybersecurity outcomes based not just on FISMA, but the myriad reporting systems through federal and state government. In collaboration with CISA's Office of the Chief Economist, OHSS should focus on:

- Developing a process for aggregating reports from disparate requirement systems with different timelines and data requirements
- Anonymized reporting on outcome data sourced from reporting systems that do not publicly reveal individual incidents, such as CIRCIA and FISMA for FCEB branches
- Researching and developing approaches to the gathering, analysis, and interpretation of cybersecurity harm data
- Recommending consistent scoping definitions for cybersecurity incidents, cyber-relevant harms, and similar components of the ecosystem

71 "Organization of the Federal Statistical System," in *Principles and Practices for a Federal Statistical Agency: Sixth Edition*, ed. Constance F. Citro (Washington, DC: National Academies Press, 2017), https://www.ncbi.nlm.nih.gov/books/NBK447392/.

72 Thomas Kim, letter to the Securities and Exchange Commission Division of Corporate Finance, "AT&T Inc. Form 8-K Filed July 12, 2024 File No. 001-08610," July 31, 2024, https://www.sec.gov/Archives/edgar/data/732717/000119312524190323/filename1.htm.

73 "Consumer Price Index Frequently Asked Questions," US Bureau of Labor Statistics, December 18, 2024, https://www.bls.gov/cpi/questions-and-answers.htm.

Conclusion

Cybersecurity policy has matured significantly in recent years, but as steady as the flow of executive orders, legis-lation, strategy, and guidance documents has been, cyberattacks have continued with shocking consistency and significant impact. With the previous administration witness to the aftermath of the SolarWinds campaign, Colonial Pipeline, the United Healthcare hack, two Microsoft Exchange compromises, Volt Typhoon, and now Salt Typhoon—to name only a few—the question, "Are we getting better at cybersecurity?" is far from an academic exercise in empiricism.

The state of metrics for cybersecurity policy is insufficient to meet two core functions today: to assess the status quo of the cybersecurity ecosystem at the macro level, and to provide insight into the relative efficacy of different security controls, practices, and requirements at the micro level. Without these dual capacities, cybersecurity policymakers are left with intuition and risk assessments to guide them. These are necessary but insufficient tools for approaching the monumental task of improving cybersecurity, which will require measuring the harms caused by cyber insecurity as key outcome metrics, and understanding those harms as the product of a complex, dynamic system is critical to meaningfully interpreting them.

Unsolved challenges to interpreting outcome data, assuming its successful measurement, remain. Knowing how much harm cybersecurity incidents have caused over a given timeframe is a start toward understanding trends in improvement, but nuanced questions about what "better" and "worse" look like, and what the data can and cannot reveal about the future still persist. In the near term, the need for this data to be systematically gathered at all and for continued progress toward interpreting it demand consistently reported outcome measures and some degree of centralization within the federal government of that information. Those embarked on improving cybersecurity can no longer afford to guess as to the best remedies for insecurity and hope that they work once implemented—policymakers will benefit immensely from measuring the harms caused by cyber incidents to see how well their remedies have worked, too.

About the author

Stewart Scott is a deputy director with the Cyber Statecraft Initiative, part of the Atlantic Council Tech Programs. He works on the Initiative's Cybersecurity Strategy and Policy portfolio, with focuses on software supply chain and open source software security policy. He earned his BA from Princeton University at the School of Public and International Affairs along with a minor in computer science.

Atlantic Council

Board of Directors

CHAIRMAN
*John F.W. Rogers

EXECUTIVE CHAIRMAN EMERITUS
*James L. Jones

PRESIDENT AND CEO
*Frederick Kempe

EXECUTIVE VICE CHAIRS
*Adrienne Arsht
*Stephen J. Hadley

VICE CHAIRS
*Robert J. Abernethy
*Alexander V. Mirtchev

TREASURER
*George Lund

DIRECTORS
Stephen Achilles
Elliot Ackerman
*Gina F. Adams
Timothy D. Adams
*Michael Andersson
Alain Bejjani
Colleen Bell
Sarah E. Beshar
*Karan Bhatia
Stephen Biegun
Linden P. Blue
Brad Bondi
John Bonsell
Philip M. Breedlove
David L. Caplan
Samantha A. Carl-Yoder
*Teresa Carlson
*James E. Cartwright
John E. Chapoton
Ahmed Charai
Melanie Chen
Michael Chertoff
George Chopivsky
Wesley K. Clark
*Helima Croft
Ankit N. Desai
*Lawrence Di Rita
*Paula J. Dobriansky
Joseph F. Dunford, Jr.
Richard Edelman
Stuart E. Eizenstat
Tara Engel
Mark T. Esper
Christopher W.K. Fetzer
*Michael Fisch
Alan H. Fleischmann
Jendayi E. Frazer
*Meg Gentle
Thomas H. Glocer
John B. Goodman
Sherri W. Goodman
Marcel Grisnigt
Jarosław Grzesiak
Murathan Günal
Michael V. Hayden
Robin Hayes
Tim Holt
*Karl V. Hopkins
Kay Bailey Hutchison
Ian Ihnatowycz
Deborah Lee James
*Joia M. Johnson
*Safi Kalo
Karen Karniol-Tambour
*Andre Kelleners
John E. Klein
Ratko Knežević
C. Jeffrey Knittel
Joseph Konzelmann
Keith J. Krach
Franklin D. Kramer
Laura Lane
Almar Latour
Yann Le Pallec
Diane Leopold
Jan M. Lodal
Douglas Lute
Jane Holl Lute
William J. Lynn
Mark Machin
Marco Margheri
Michael Margolis
Chris Marlin
William Marron
Roger R. Martella Jr.
Judith A. Miller
Dariusz Mioduski
*Richard Morningstar
Georgette Mosbacher
Majida Mourad
Mary Claire Murphy
Julia Nesheiwat
Edward J. Newberry
Franco Nuschese
Joseph S. Nye
*Ahmet M. Ören
Ana I. Palacio
*Kostas Pantazopoulos
David H. Petraeus
Elizabeth Frost Pierson
*Lisa Pollina
Daniel B. Poneman
Robert Portman
*Dina H. Powell McCormick
Michael Punke
Ashraf Qazi
Laura J. Richardson
Thomas J. Ridge
Gary Rieschel
Charles O. Rossotti
Harry Sachinis
C. Michael Scaparrotti
Ivan A. Schlager
Rajiv Shah
Wendy R. Sherman
Gregg Sherrill
Jeff Shockey
Kris Singh
Varun Sivaram
Walter Slocombe
Christopher Smith
Clifford M. Sobel
Michael S. Steele
Richard J.A. Steele
Mary Streett
Nader Tavakoli
*Gil Tenzer
*Frances F. Townsend
Melanne Verveer
Tyson Voelkel
Kemba Walden
Michael F. Walsh
*Peter Weinberg
Ronald Weiser
*Al Williams
Ben Wilson
Maciej Witucki
Neal S. Wolin
Tod D. Wolters
*Jenny Wood
Alan Yang
Guang Yang
Mary C. Yates
Dov S. Zakheim

HONORARY DIRECTORS
James A. Baker, III
Robert M. Gates
James N. Mattis
Michael G. Mullen
Leon E. Panetta
William J. Perry
Condoleezza Rice
Horst Teltschik
William H. Webster

**Executive Committee Members*

List as of March 24, 2025

1400 L Street NW, 11th Floor,
Washington, DC 20005
(202) 778-4952
www.AtlanticCouncil.org