

Identifying Conflicting Incentives in United States Federal Cybersecurity Policy:

A Sociotechnical Systems Approach

by

Jamie Winterton

A Dissertation Presented in Partial Fulfillment  
of the Requirements for the Degree  
Doctor of Philosophy

Approved August 2023 by the  
Graduate Supervisory Committee:

Andrew Maynard, Chair  
Diana Bowman  
Katina Michael

ARIZONA STATE UNIVERSITY

December 2023

## ABSTRACT

Despite increased attention and funding from companies and governments worldwide over the past several years, cybersecurity incidents (such as data breaches or exploited vulnerabilities) remain frequent, widespread, and severe. Policymakers in the United States have generally addressed these problems discretely, treating them as individual events rather than identifying commonalities between them and forming a more effective broad-scale solution. In other words: the standard approaches to cybersecurity issues at the U.S. federal level do not provide sufficient insight into fundamental system behavior to meaningfully solve these problems.

To that end, this dissertation develops a sociotechnical analogy of a classical mechanics technique, a framework named the Socio-Technical Lagrangian (STL). First, existing socio/technical/political cybersecurity systems in the United States are analyzed, and a new taxonomy is created which can be used to identify impacts of cybersecurity events at different scales. This taxonomy was created by analyzing a vetted corpus of key cybersecurity incidents, each of which was noted for its importance by multiple respected sources, with federal-level policy implications in the U.S..

The new taxonomy is leveraged to create STL, an abstraction-level framework. The original Lagrangian process, from the physical sciences, generates a new coordinate system that is customized for a specific complex mechanical system. This method replaces a conventional reference frame –one that is ill-suited for the desired analysis – with one that provides clearer insights into fundamental system behaviors. Similarly, STL

replaces conventional cybersecurity analysis with a more salient lens, providing insight into the incentive structures within cybersecurity systems, revealing often hidden conflicts and their effects. The result is not a single solution, but a new framework that allows several questions to be asked and answered more effectively.

Synthesizing the findings from the taxonomy and STL framework, the third contribution involves formulating reasonable and effective recommendations for enhancing the cybersecurity system's state for multiple stakeholder groups. Leveraging the contextually appropriate taxonomy and unique STL framework, these suggestions address the reform of U.S. federal cybersecurity policy, drawing insights from various governmental sources, case law, and discussions with policy experts, culminating in analysis and recommendations around the 2023 White House Cybersecurity Strategy.

## DEDICATION

To Ezra – thank you for your patience and support. I love you.

## ACKNOWLEDGEMENTS

I'd like to gratefully acknowledge the guidance, support, and encouragement of my Chair, Dr. Andrew Maynard. Andrew, it was unbelievably great to have another physicist to ideate with – I would not have been able to pursue this particular course of research without your mentorship. I am also deeply appreciative of my committee members, Dr. Diana Bowman and Dr. Katina Michael, who provided unique insights and made this a truly interdisciplinary work. I have truly enjoyed learning from you all and creating new ideas together! Dr. Heather Ross and Andra Williams were also phenomenally helpful in bringing this effort to fruition. (Without you, Andra, I'd still be trying to figure out iPOS.)

I am also grateful for the support and encouragement of the Global Security Initiative, in particular Dr. Nadya Bliss, GSI's Executive Director. Nadya, thank you for inspiring me and empathizing with me through this journey. It's an honor to work with you and to be your friend.

When I started this program, I bought a refrigerator magnet that said, "My decisions are not always right, but they are always interesting." Pursuing a PhD as a full-time executive and a parent turned out to be both. I'm so grateful for the team of amazing people who walked with me on a similar path – Dr. Barbara Whye, Dr. Elizabeth Garbee, Lucy Tournas, J.D. (and almost-PhD), Jessica Barnett (M.S.) and Corinne Dixon (almost-DNP). Your camaraderie in this effort means more to me than you know!

Finally, a heartfelt “thank you” to my parents for teaching me the value of education and hard work, and to my son Ezra, who makes it all worthwhile. I love you, kiddo.

TABLE OF CONTENTS

	Page
LIST OF TABLES .....	ix
LIST OF FIGURES .....	x
CHAPTER	
1. THE ROOTS OF "CYBER" .....	1
Modern Definitions And Their Drawbacks .....	2
Cybersecurity Impacts At Multiple Levels.....	4
Three Unique Impacts Of Research .....	6
Bounding The Research Scope.....	10
Structure Of The Dissertation.....	11
2. FOUNDATIONS OF CYBERSECURITY .....	13
Origins Of The Internet .....	13
A Note On Complexity.....	18
Updates To U.S. Federal Computer Laws.....	43
Van Buren V U.S. – U.S. Supreme Court Decision And Dissent .....	47
A Post- ‘Van Buren’ Internet – Expectations And Remaining Questions .....	52
Addressing Cybersecurity Complexities, Not Just Problems.....	56

CHAPTER	Page
3. BUILDING A RELEVANT TAXONOMY .....	58
Creating A Taxonomy .....	61
Research Methods .....	64
The Data Set .....	73
Categorizing Events.....	74
Assessing Levels Of Impact .....	82
From Taxonomy To Framework .....	89
4. A NEW FRAMEWORK .....	91
Revisiting The Motivation.....	91
Complex Mechanical Systems .....	92
Creating A Sociotechnical Analogy .....	101
“Conventional” Cyber Coordinates .....	102
New Cyber Coordinates .....	108
The Socio-Technical Lagrangian .....	111
Validating The Socio-Technical Lagrangian.....	132
Contributions To Knowledge .....	146



CHAPTER	Page
5. POLICY RECOMMENDATIONS .....	149
The White House National Cybersecurity Strategy .....	150
Implementation.....	167
General Recommendations.....	168
6. CONCLUSION .....	172
Knowledge Contributions Of The Work .....	174
Stakeholders .....	175
Future Research Directions .....	177
Conclusion.....	181
REFERENCES.....	181
APPENDIX	
A DESCRIPTION OF EVENTS IN DATA SET .....	209
BIOGRAPHICAL SKETCH .....	246

## LIST OF TABLES

Table	Page
1. Events Used To Create The Taxonomy, The Year In Which They Occurred, And Their Categorization .....	80
2. General Schema Of Impact Rankings .....	83
3. Personal Impact Scores.....	84
4. Infrastructure Impact Scores.....	86
5. Political Impact Scores .....	88
6. Terms And Definitions Used In The Lagrangian Mechanics Example Of A Double Pendulum With Spring (Figure 12). .....	101
7. Terms And Definitions Used In Both The Double Pendulum, And The Socio-Technical Analogue Of Cybersecurity Systems.....	112
8. Events From The Data Set With Assigned Impacts In The Three Primary Categories, As Well As Average And Variance Of Impact. ....	113
9. Validation Events Have Been Chosen To Match Features Of Key Events From The Primary Data Set .....	134

## LIST OF FIGURES

Figure	Page
1. A Timeline Of Select Events In The Development Of The Internet. ....	13
2. Arpanet On A Napkin. (This Napkin Drawing Is A Reproduction (Computer History Museum, Accessed August 1, 2020).).....	16
3. Packets Pass Through Several Different Machines When A User Visits A Website .....	25
4. The Screen That Greeted Many Sony Employees After The North Korean Intrusion. (Zetter, 2014) .....	36
5. Percentage Of Households With Computer And Internet Use: 1984 To 2015. (Ryan & Lewis, 2017) .....	39
6. Mobile Phone And Smartphone Adoption, 2000 – 2020 (Pew Research Center, 2021).....	40
7. Stakeholders In The Cloud Computing Value Chain (Kolevski Et. Al, 2020). 63	
8. The Number Of Breached Records (In Millions) For Four Select Events In The Data Set .....	75
9. Personal Impacts Assigned To Four Selected Events .....	85
10. Infrastructure Impacts Assigned To Four Selected Events. ....	86
11. Political Impacts Assigned To Four Selected Events.....	88
12. A Double Pendulum, One Pendulum Is Attached To The Base By A Spring (Jankowski, 2011).....	94

Figure	Page
13. An Object (Green) Traversing A Linear Path (Red), Shown On A Cartesian Grid.....	97
14. An Object (Green) Traversing A Circular Path (Red), Shown On A Cartesian Grid.....	98
15. The Same Object And Path As Figure 14, Shown On A Polar Grid.....	99
16. Occurrences Of The Term “Ransomware” In The U.S. Congressional Record, 2015-2022.....	130

## CHAPTER 1

### THE ROOTS OF "CYBER"

The prefix “cyber-” is ubiquitous. Prepend to everything from punk and sex to attacks and even tanks, nearly everything has a cyber-analogue. Given the current prevalence of internet connectivity, the broad inclusion of “cyber” in the English language is understandable.

Before all these terms were invented, even before the internet was born, there existed the word *cybernetics*. The term dates to Plato – the Greek Κυβερνήτης (*kybernetes*), meaning steersman or governor, referring to the “science of effective government” (Heylighen & Joslyn, 2001). The word then became popular in the 1940’s to describe the study of systems and their feedback loops (Novikov, 2016). In 1948, Norbert Wiener, professor of mathematics at Massachusetts Institute of Technology, defined *cybernetics* as “the scientific study of control and communication in the animal and the machine” (Weiner, 1948) – an elegant definition that evokes the complexity, nuance, and interdisciplinary perspective needed to approach these systems (Adamson, Kline, Michael, & Michael, 2015). Wiener’s work set the foundation for a series of cybernetics conferences, sponsored by the Macy Foundation between 1946 and 1953, bringing researchers together from a wide variety of disciplines, with the goal of critiquing strict disciplinary approaches and creating systems-level understanding that transcended the disciplines on their own (Visvanathan, 2002). Over the years, Wiener’s conceptualization of cybernetics has been used to study a wide variety of systems – mathematical (Weiner, 1948), biomechanical (Magin, 2017), literary (Hayles, 1999), and

others. In Wiener's definition and the work that followed, the tug of forces between human and machine are made explicit; the unstable equilibrium that results is anticipated and interrogated.

Wiener's work influenced the development of early technologies, from machine learning algorithms to prosthetics (Adamson, Kline, Michael, & Michael, 2015), but one of his most notable impacts was on the development of the internet. J.C.R. Licklider, a research psychologist by training, attended Wiener's cybernetic conferences (named "the Macy conferences" for their sponsor) as well as a weekly working group at Norbert Wiener's home in the mid-1950's (Lee & Rosin, 1992). Inspired by these sessions, Licklider began thinking about the relationship between humans and machines, developing ways in which they could work collaboratively (Licklider, 1960). Licklider managed the Information Processing Techniques Office in the Defense Advanced Research Projects Agency (DARPA) from 1962-1964, sponsoring several programs that led to the predecessor of the modern internet (ARPANet) in 1966 (Markoff, 1999).

### **Modern definitions and their drawbacks**

Today's internet resembles the blend of "animal and machine" from Wiener's *cybernetics*, a complex interplay between human designers, human users, and the technology that connects them. Of all the cyber-terms in our current discourse, "cybersecurity" may be the most widely recognized, and certainly one that is relevant to society and humans. Despite the internet's inherently sociotechnical nature, efforts to make it more secure have generally focused on technology alone, unfortunately losing the nuance and interdisciplinarity of the original "cybernetics". The word "cybersecurity" is

frequently defined by U.S. government organizations and technology companies as some variant of ‘technologies to keep computers safe on the internet’ (e.g, (National Institute of Standards and Technology (NIST), 2018), (Cybersecurity and Infrastructure Security Agency (CISA), 2019). The most recent version of the Committee on National Security Systems Glossary, created by the U.S. National Security Agency (NSA), defines “cybersecurity” as:

*Prevention of damage to, protection of, and restoration of computers, electronic communications systems, electronic communications services, wire communication, and electronic communication, including information contained therein, to ensure its availability, integrity, authentication, confidentiality, and nonrepudiation. (Committee on National Security Systems , 2015) pg. 40*

These definitions focus exclusively on the technology aspects of cybersecurity – the wires and electrons that push data packets from one place to another. Humans – the inventors, developers, and beneficiaries of internet connectivity – have no representation in the previous description. Human, legal, and social dynamics are critical to consider, in addition to technology, despite the prevalence of one-dimensional definitions of cybersecurity. A definition of “cybersecurity” that lacks heterogeneous elements is incomplete – and dangerously so, when definitions, like those cited above, are used to create standards that form the basis of government policy.

There are a few definitions of “cybersecurity” that do explicitly include humans. The Freedom Online Coalition (FOC), a group of 35 countries<sup>1</sup> focused on human rights

---

<sup>1</sup> As of 2022, the FOC consisted of the following countries: Argentina, Australia, Austria, Canada, Chile, Costa Rica, the Czech Republic, Denmark, Estonia, Finland, France, Georgia, Germany, Ghana, Ireland, Italy, Japan, Kenya, Latvia, Lithuania, Luxembourg, the Republic of Maldives, Mexico, Moldova, Mongolia, the Netherlands, New Zealand, Norway, Poland, Spain, Sweden, Switzerland, Tunisia, the United Kingdom and the U.S.

and the internet (Freedom Online Coalition , 2022), was founded in 2011 and will be chaired by the U.S. in 2023. The group promotes democratic ideals globally, with a focus on the internet, and has been cited in US federal documents such as the 2023 White House Cybersecurity Strategy. FOC provides this definition of “cybersecurity”:

*“Cybersecurity is the preservation – through policy, technology, and education – of the availability, confidentiality and integrity of information and its underlying infrastructure so as to preserve the security of persons both online and offline.”*  
(Internet Free and Secure Initiative, 2021)

Here, the policy aspect is at the forefront, and the technical components (availability, confidentiality, and integrity (Neumann, Statland, & Webb, 1977)<sup>2</sup>) are in support of “the security of persons both online and offline”. But this coalition definition lacks the veracity of government agency definitions, and even this forward-thinking definition neglects the security of social groups, which have a different threat surface than individuals alone. The Coalition’s definition also assumes an existing state of security for individuals, through its use of the word “preserve”. Unfortunately, before security can be preserved, significant work must be done in repairing and rebuilding it.

### **Cybersecurity Impacts at Multiple Levels**

Neglecting human, social, and legal elements of cybersecurity, and the interplays between them, leads to deep and persistent issues with real-world effects (Spafford &

---

<sup>2</sup> The “CIA model” – confidentiality, integrity, and availability (no relation to the Central Intelligence Agency) – is pervasive in the field of information security but is not without its critics. A fuller exploration of the failures in definitions and standards will require a deeper critique of the CIA model and some analysis of alternatives. For example, some proposed models explicitly include the concepts of privacy and non-repudiation as critical elements of security (Samonas & Coss, 2014).



Antón, 2008), (Malatji, Von Solms, & Marnewick, 2019). Given a very general definition of cybersecurity as “keeping people safe on the internet”, then one can define failures in cybersecurity as “incidents wherein the safety of individuals or groups has been compromised as a result of an internet-enabled action”. The internet has provided incredible benefits to humanity, but cybersecurity failures have also been numerous, and span the space from individual, to organizational, to global levels.

A brief example of cybersecurity impacts at multiple levels is the 2014 data breach of the United States (U.S.) Office of Personnel Management (OPM). OPM is the U.S. government entity that manages civilian employees, and as such, it holds extensive data on personnel, including those with security clearances. In 2014, an adversarial actor breached the OPM infrastructure, and exfiltrated over 21 million security clearance records. This is not the largest data breach as measured by number of records violated, but the OPM breach is one of the most concerning from a U.S. perspective, given the sensitive personal data lost, the extensive failures in protecting those data, and the potential impacts of that loss on national security (Winterton, 2015). A fractured maze of government oversight and years of technical debt (outdated systems and procedures) were no match for a sophisticated nation-state attacker that spent several weeks undetected in the system, collecting sensitive data on cleared personnel and their families (Gallagher, 2015). These records include not only social security numbers and addresses, but the deeply personal life stories that the U.S. government uses to assess a person’s clearance eligibility. These data include interviews with the clearance holder and their family or friends on sensitive topics such as family issues (divorce or abuse), alcohol and drug use, or financial insecurity. In the hands of a foreign adversary, some of this

information is likely being used to infer which technologies are advancing in classified spaces, but the personal information can also be used to identify which U.S. clearance holders may be susceptible to bribery or blackmail. The OPM breach was a failure of technology, but also a failure to understand the human, legal and social aspects of cybersecurity, which led to irreparable harm – a near-peer adversary’s ability to map the technical and social landscape of U.S. defense research and development. The full extent of the damage from this breach is still not fully understood, in part due to the slow pace of the legal system (The Cyberwire, 2022)<sup>3</sup> and in part due to the assumed motives of espionage behind the attack.

The failures across multiple dimensions reflected in the OPM data breach – technological, political, human, and social – are not unique to this event. Without a multidisciplinary approach to the problem - a “scientific study of control and communication in the animal and the machine”, to reprise Wiener’s definition – needed improvements on the state of internet security are unlikely for individuals, organizations, or nations.

### **Three unique impacts of research**

Responding to the need for effective interdisciplinary interventions, the research herein explores new methods by which the multidimensional problems of cybersecurity can be more clearly understood. This dissertation draws on deep domain-specific

---

<sup>3</sup> A class action lawsuit was resolved in October of 2022 in the U.S. District Court of the District of Columbia, awarding U.S. \$63M to a group of plaintiffs on the basis of increased risk of future harm due to their sensitive data being breached (n re U.S. Office of Personnel Management Data Security Breach Litigation, 2022).

knowledge and a perspective informed by socio-technical studies to reveal new ways of understanding and addressing the impacts of cybersecurity failures in the larger system.

This dissertation includes three unique knowledge contributions:

1. A novel taxonomy to describe and evaluate harms incurred by a spectrum of cybersecurity failures.

Describing a complex system is challenging. There is no unique definition of a “complex system”, but definitions across the literature generally include the notion of a large number of components interacting with one another in a multitude of ways (Ladyman, Lambert, & Wiesner, 2011), (Philips & Austad, 1996). Failure to adequately describe these components will result in a flawed description and analysis of emergent behavior in the larger system. No taxonomy for cybersecurity incidents presently exists that is adequate for an anticipatory and broad-scale understanding. Globally, current efforts are primarily focused on narrower remediation and response efforts. The Cybersecurity and Information Systems Information Analysis Center (CSIAC), a U. S. Department of Defense research center, has created a detailed list of cybersecurity taxonomies, focused on threat detection and incident response (Launius, 2020). This work is important in creating a shared language between organizations and facilitating information sharing programs between companies and government entities. The European Cybersecurity Taxonomy takes a different perspective, categorizing existing institutions and their areas of expertise, for the purpose of identifying gaps and overlaps in regulation (European Commission , 2021). These efforts are important in their contexts, but these existing

categorizations do not scale up to support the creation of the abstraction-level framework pursued in this research.

This dissertation introduces a new taxonomy for cybersecurity events, focused on impacts on three primary groups. This taxonomy was created by analyzing a vetted corpus of key cybersecurity incidents, each of which was noted for its importance by multiple respected sources, with federal-level policy implications in the U.S.. Due to the challenges of attributing the source of an attack and the global nature of the impacts, it is difficult to create a rigorous jurisdictional guide for these events, but defining a rough geographic scope is necessary in order to make reasonable recommendations. The taxonomy is based on a broad spectrum of evidence, from a diversity of English-language sources, and it includes new ways to evaluate and more clearly communicate about harms of security failures. The disciplines of law, public policy, systems thinking, and socio-technical fundamentals each provide relevant insights, which are uniquely powerful in concert with one another. A diverse literature review underpins this work, ranging from academic writings in socio-technical theory to news media, legal opinions from U.S. circuit courts and the U.S. Supreme Court, U.S. federal law, and associated regulatory and policy documents. In addition to literature, this effort has been informed deeply by interactions with the security community, including policy experts, technologists, and hackers - who are often a blend of both. Chapter 3 describes this process and its outcomes in depth.

The taxonomy has intrinsic value but also provides a necessary step in creating the second knowledge product: a new framework that can be applied to achieve a more

complete understanding of the various incentives and impacts in the cyber socio-technical realm.

2. Develop new way of analyzing the incentive structures present in these systems to discover oft-hidden conflicts and their effects

There is no lack of discussion on insecure technology, ineffective security policy, or a lack of “cyber hygiene” (an unpleasant term often used to describe the behavior of individual users). But the current policy conversation focuses almost exclusively on the *effects* of these failures, often ignoring the *incentives* within these systems that drive the failures. And there is very little discussion on how the incentives of one category conflict with those of another, affecting behaviors throughout the system. One explanation is the relative invisibility of incentives in cybersecurity. Generally, incentives in a complex system are often obscured, due to the absence of central control and variety of elements comprising multiple levels (Ladyman, Lambert, & Wiesner, 2011). While it may not be easy to view the direct cause-and-effect in complex systems, methods of abstraction have been created in some domains (e.g. physics), which allow for a clearer understanding of the system’s central characteristics. While not directly applicable in all cases, these approaches can be studied and used as a model to inspire approaches in new domains. This dissertation includes such an approach: an abstraction-level framework, built on the previously described taxonomy, to reframe cybersecurity issues in a more accessible and insightful way. This method describes incentives within the system, how they interact, and where efforts could be best allocated for positive change. Chapter 4 introduces an example of reapproaching complex dynamic systems in the physical domain, Lagrangian

mechanics (Fowles & Cassiday, 1999). The Lagrangian mechanics approach allows easier analysis of the forces and effects within a complex mechanical structure through the creation of a new coordinate system specifically for that complex structure. From here, the dissertation proceeds through development of a new analogue of Lagrangian dynamics specifically for cybersecurity issues in the socio-technical domain. Several examples of new insights from this approach are developed, and a separate set of events is used for validation of the approach. This work is supported by the Appendix of incidents, as well as the corpus of case law and associated literature cited throughout.

3. Synthesize findings and formulate recommendations for reasonable and effective “next steps”, given the current state of the system.

By leveraging a contextually appropriate taxonomy and a powerfully unique framework for understanding misaligned incentives in cybersecurity, a set of recommendations for system improvements can be generated. These suggestions address current and prior efforts to reform U.S. legislation, incorporating insights from the U.S. congressional record, White House guidance such as the National Security Strategy, case law, and discussions with policy experts. The ultimate goal of this dissertation work is to improve the state of cybersecurity at the federal level in the U.S., so this effort must include implementable knowledge and a transition path outside the academy. These recommendations are described at length in Chapter 5.

### **Bounding the research scope**

This project does not directly apply principles of physics to complex socio-technical systems. The inputs to Lagrangian mechanics are purely quantitative; reducing

the broader cybersecurity system to numerical values would continue to ignore the social and political challenges that are at the root of many issues (Clarke & Primo, 2012). Nor does this work employ a singular, established social science method – this would result in an incomplete approach given the practical and integrated nature of the challenge.

However, this research has been informed by a variety of appropriate methodologies from the domain of science and technology studies, which provide a scholarly socio-technical foundation for the physics-inspired approach. This dissertation is not a prescription for resolving conflicts in these systems – rather, the intention is to find a way to better identify the conflicts. It may be that the techniques developed for this dissertation can be applied to additional complex socio-technical systems, but the analysis here will be scoped specifically to cybersecurity at the U.S. federal level.

### **Structure of the dissertation**

Chapter 2 builds a foundational knowledge of the current cybersecurity landscape. A short history of the internet is provided, from its origin in U.S. federal research and development to its current usage. The evolution of legal structures (primarily U.S. federal, but international soft law approaches are also discussed when they provide relevant insights, such as regulation of website domain names) alongside internet technology is also described, particularly through case law as unique complexities in geography, authority, time, and society are examined and oft-assumed boundaries are challenged.

Creating a new approach often requires new building materials. Chapter 3 calls out the need for a taxonomy of cyber incidents, defines a rubric for relevant data to create

it, identifies those relevant data, and justifies the use of news media as a sufficiently accurate and necessarily timely set of sources. The taxonomy is then built in detail. An appendix of case studies accompanies this chapter.

Chapter 4 uses the taxonomy to build the cyber-specific “socio-Lagrangian” – a new framework that illuminates the gaps and conflicts in cybersecurity incentives. This is the “crux” of the research. The framework is used to analyze the primary data set, identifying insightful comparisons and groups of events, with new opportunities for progress. Chapter 4 concludes with validation of the new framework by applying it to a new data set and comparing the outcomes with the primary analysis.

The case studies of Chapter 3 and framework of Chapter 4 are then used in Chapter 5 to create a set of policy recommendations. This section applies the new framework to the cybersecurity issues and policy recommendations described in the most recent U.S. National Cybersecurity Strategy released by the White House in March 2023.

Chapter 6 summarizes the achievements of the dissertation and identifies potential areas for future scholarly contributions.

Appendix A describes the case studies used in Chapter 3.



## CHAPTER 2

### FOUNDATIONS OF CYBERSECURITY

#### Origins of the internet

In understanding current issues in cybersecurity, it helps to know how the internet began. Many people associate Silicon Valley with the internet, but its story began elsewhere. It's true that the initial concept and structure of the internet were expanded considerably by private industry, and then leveraged for a dazzling array of applications such as the World Wide Web, but much of the foundational development came (as it often does) from the U.S. Federal Government (Mazzucato, 2013). There is another popular conception that the internet was invented by the Advanced Research Projects Agency (ARPA - predecessor to DARPA). This is a bit too succinct, but it is more accurate, and this position allows one to reflect on design decisions made in the ARPANet era that are at the core of many current internet security issues today<sup>4</sup>. A timeline of select events in the internet's development is shown below in Figure 1.



Figure 1. A timeline of select events in the development of the internet.

---

<sup>4</sup> For additional information on the internet's evolution, Janet Abbate's book 'Inventing the internet', MIT Press, 1999, and Claire L. Evans' book, 'Broad Band: The Untold Story of the Women who Made the internet', Portfolio/Penguin, 2018, are two excellent in-depth resources.

When ARPANet was first designed in 1969, it wasn't envisioned to be the internet as experienced today. Its original intention was much more limited – an auxiliary system to connect military research sites into the Department of Defense's existing network, the Automatic Digital Network System, or AUTODIN (Fidler, 2017), (Ryals, 2017). If one considers ARPANet to be the mother of our modern internet, AUTODIN may be considered its grandmother. The AUTODIN system had been in place for many years prior, primarily managing intercontinental ballistic missile logistics for the Air Force (McKenzie & Walden, 1991). However, AUTODIN required specially trained personnel to operate it, which meant the system was costly as well as slow (Office of Inspector General, 1996) and its messaging and resource allocation methods were inflexible. Instead of sending a message in smaller pieces (as the current internet does) that could be prioritized or duplicated if dropped, AUTODIN could only send entire messages in the order they were requested. Any errors in the messages had to be remediated manually – meaning, a person had to find, correct, and re-queue the message to be sent - and one recollection indicates that up to 50% of these messages required such intervention (Elkins, Retrieved 2022). The AUTODIN network was revolutionary when it launched in 1962 - as the Air Force's first automated, electronic network, AUTODIN connected over 300 users, replacing a slow and disjoint manual punch-card system (Sibit, 2018). Demands on the network grew quickly, with eleven new switching stations across the globe added to the network's original five, all of which supported 24/7 operations (ibid). AUTODIN couldn't always match the operational pace of the military. The attack on the USS Liberty in 1968 is one such example. Mistaken by Israeli forces for an adversarial Egyptian vessel, the USS Liberty was directed several times via the

AUTODIN system to withdraw from its position. Those instructions were unfortunately stuck in the queue for several critical hours. During that time, Israeli forces attacked the ship, and as a result, 34 servicemembers were killed, 75 were wounded, and the Liberty was damaged beyond repair (Gorzoch, 1979). Israel paid a total of \$12.9M USD to the United States between 1968 and 1980, as compensation for those killed and wounded, as well as for the material damage to the ship. This incident was unfortunate, but it's worth remembering that AUTODIN represented the state of the art for the time, and it would be many more years before the Department of Defense considered potential alternatives.

ARPANet was not initially designed as a replacement for AUTODIN. The first ARPANet was an experimental network designed for military research, not operations. It was a prototype that radically re-thought the way messages were sent and received. This experiment was done in an academic environment instead of in a mission context, with graduate students instead of generals, and absent any connection to the AUTODIN or telephone network communities who had different ideas about the fundamentals of communications and how to secure them (Fidler, 2020). The initial network is shown in Figure 2, comprised of the University of Utah, University of California locations in Santa Barbara and Los Angeles, and the Stanford Research Institute. There was also a different level of authority attached to each system. A message sent through the official AUTODIN system was a military directive that could be acted upon, but a message sent through the experimental ARPANet research system was not. AUTODIN had been certified for military use, while ARPANet existed initially as a proof-of-concept.

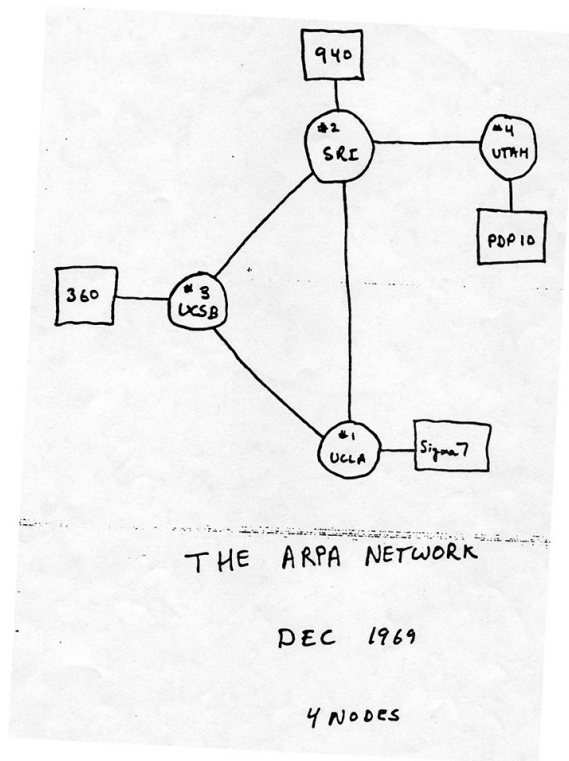


Figure 2. ARPANet on a napkin. (This napkin drawing is a reproduction (Computer History Museum, Accessed August 1, 2020).)

For a while, it was intended that the two networks co-exist. The researchers on ARPANet wanted more flexibility than AUTODIN could provide; AUTODIN’s military users wanted to keep potentially disruptive academic research off an operational military network (Abbate, 1999). The ARPANet added more university nodes and proceeded with research, and work on the AUTODIN II network began in the late 1970’s. But by 1981, the armed services had pressing concerns about AUTODIN II’s high cost and performance problems. A 1974 report from the U.S. General Accounting Office claimed that the AUTODIN system was costing the Department of Defense \$190M USD annually - the equivalent of \$1.18B USD in 2023 currency (U.S. General Accounting Office,

1974). These concerns led to a three-day competition between AUTODIN II and ARPANet. As a result, AUTODIN II was cancelled in 1982, and ARPANet then took over as the backbone of the new Defense Data Network (Fidler, 2017).

The data transfer and security functions of these early networks, researched in detail by historian Dr. Brad Fidler (Fidler, *Cybersecurity governance: a prehistory and its implications*, 2017), were separated in the early design to support a typical military “god’s eye view.” In this design, the network as a whole is protected from the outside, but all internal nodes are open to one another as well as open to the system’s administrators. This design reflects the general military perspective wherein a single observer has the final authority and responsibility for a given system, a person with a well-defined scope who operates within a structured organizational hierarchy. In the initial concept, data transfer between nodes of the network was simplified by putting the onus of security only on the devices at the network boundary, not in between. If each endpoint was guaranteed to be secure and observable by the proper authority, as was the case in the original military ARPANet system, data could flow freely and securely through the network without the burden of additional interstitial security measures. Thus, Fidler observes, “the absence of cybersecurity technologies on the early internet was not an oversight, but a necessary compromise” between the disparate purposes of the system (Fidler, 2017) pg 450.

The modern internet evolved from a handful of connections owned and managed by the Air Force, through the more flexible architecture (complicated, not yet complex) of ARPANet. The benefits of scaling the network beyond the Department of Defense to the civilian world were obvious to ARPANet’s designers and were pursued in tandem

with military development (Fidler, 2017). At this point, it was too late to reverse or even truly understand the impact of security decisions made in the early design phases. These choices altered the course of cybersecurity, enshrining certain capabilities while extinguishing others (DuPont & Fidler, 2016). And as the potential of the network exploded through novel uses – or to quote Trevor Pinch and Ronald Klein, its “users of technology acted as agents of technological change” (Kline & Pinch, 1996) – the ARPANet grew into the internet, the complex system that is familiar to today’s users.

### **A Note on Complexity**

It’s important to differentiate the terms “complex” and “complicated”. Complicated problems are hard, but can be solved with a sufficiently detailed set of rules. Filing one’s taxes in the U.S. is a good example. Depending on the context of the person filing, different parts of the tax code come into play, but even if that person’s situation is unusual, there are rules that govern what should be paid and how. It may not be easy, but it is possible to arrive at a single unique solution.

Complex problems, on the other hand, “involve too many unknowns and too many interrelated factors to reduce to rules and processes” (Nason, 2017). The dense interconnectivity of cybersecurity – the network of individual human lives and the corresponding societal properties, the legal aspects, the technology itself – and the rapidity with which this heterogeneous network changes, means that cybersecurity problems are *complex* (Abbas, Michael, Michael, & MG, 2014). These couplings between society, human behavior, and technology then result in emergent behavior (“resonances”) that can be constructive or destructive. Much of this behavior is

unpredictable. And as eloquently noted by Miller and Page, these resonances in the overall system can't be understood simply by understanding their component pieces. A complex system must be approached holistically, to include emergent behaviors and the system's relationships with its environment (Miller & Page, 2007). This is certainly the case in cybersecurity. Increasing the security knowledge of the individual users does not make the internet more secure, nor does the development of purely technologic or pure policy-based security measures. In fact, many of these additions to the system's components cause friction between these components, generating vulnerabilities and at times decreasing the internet's security.

Given its history, securing the internet may seem like a technical problem – one that is difficult, but essentially manageable through improvements in computer science and technology design. In the early days, that was almost true! Unfortunately, given the internet's current intersections with individual lives and broader society, tied in with legislation and policy, improvements in cybersecurity will need to extend beyond the technical. Cybersecurity problems have both roots and reverberations in areas like local, state and federal policy, social influences, human behavior, and even human emotion. These problems are evolving and emergent, and they don't stay neatly within boundaries. Assumptions about boundaries in geography, authority, time, and social structure should be specifically challenged, as misconceptions often lead to flawed technology and/or policy solutions. The next four sections will address these complexities.

### 1. Complexity in Society: Hacking People

All the routers, firewalls, packets and passwords of the internet exist for a very human purpose: to connect people to one another. Whether it's grandparents and their grandkids,

or a physical therapist with an insurance company, this infrastructure would not exist without people. That sounds obvious to the point of absurdity, but this aspect is often neglected. The human aspects of the cybernetic network are hard to predict and challenging to measure. An individual may inhabit multiple “worlds” simultaneously, having one persona online and another in the physical world, and the actions they perform in these worlds may reflect different values.

This human-centric perspective adds a new dimension to hacking. It’s not just routers and firewalls that have vulnerabilities and can be exploited – the humans can, as well. Hacking humans is often referred to as “social engineering,” and 20% of data breaches are estimated to have some element of social engineering (per the Verizon Data Breach Incident Report, an industry standard which included analysis of over 5,000 global incidents in its 2022 report) (Verizon, 2022). Software tools have been created to facilitate social engineering (e.g., Social-Engineer Toolkit (Pavković & Perkov, 2011)<sup>5</sup>), but these tools require a human target and enough human insight to construct a believable pretext for the attack. These are important components of cyber-offense that, for the foreseeable future, can’t be scripted or automated. Solutions to cybersecurity problems that don’t acknowledge the humans (individuals and groups) are very common, but their scope is extremely limited.

How should these human vulnerabilities be approached? Spam filters and corporate education programs are numerous, but social engineering has become more sophisticated. Malicious actions are more difficult to discern from legitimate purposes (Cohodas, 2015), which makes training even more difficult. An annual social engineering competition at

---

<sup>5</sup> The Social-Engineer toolkit is available online at <https://www.trustedsec.com/social-engineer-toolkit/>



DEF CON, the world's largest hacker convention, highlights the severity of the social engineering attack vector and show how poorly companies are defending against it (Hadagny & Fincher, 2014). In the competition, approximately 15 contestants attempt to obtain potentially sensitive information from previously selected target companies. The competitors may only use openly available online data and telephone elicitation. The data gathering is performed in the two weeks leading up to the conference, and phone calls are performed live in front of an audience at the event<sup>6</sup>.

The “flags” obtained by competitors include information about a company's IT infrastructure, services such as catering or trash pickup, and for the largest point value, convincing an employee to navigate to the Social-Engineer website<sup>7</sup>. While seemingly innocuous, answers to these questions could be used to infiltrate a company in a variety of ways. One group of flags focuses on physical intrusion – by knowing about a company's catering services, for example, an intruder could pose successfully as a delivery person and allowed access to the facility. Other flags could be used to electronically compromise the company, such as which anti-virus software is used and whether or not software patches are up to date. One of the final flags of the competition is getting a human target from the company to type in a specific web address (created for the event) and describe what's on the webpage. The competitor who gathers the most sensitive information from the largest number of companies on the list wins the competition. Analysis of each contestant is done after the call by the Social-Engineer team and the audience has the opportunity to ask questions.

---

<sup>6</sup> The Electronic Frontier Foundation oversees the event to ensure that no laws are broken, and a core rule of the competition is that no one – competitors, employees, or companies – are victimized or bullied.

<sup>7</sup> <https://www.social-engineer.org/>

The Social Engineering challenge highlights the fact that not all hacking relies solely on direct machine-to-machine exploitation. After the competition, the organizers provide target companies (and the public) with a report on how to better educate their employees to mitigate these attacks. The DEF CON challenge is one of the only ways to observe successful – and unsuccessful – social engineering attacks and responses to them, as both sides of the phone call are broadcast to the audience.

The competition is a unique window into the human side of cybersecurity— not just the constructed pretext of the attacker and how human vulnerabilities are used to exploit target companies, but how these human features simultaneously play an important role in keeping a company and society functioning. Most of these features are related to trust, and the desire to help another person. These qualities are essential aspects of a society, particularly a democracy. The reason social engineering attacks work is because they are posed as legitimate human needs. Successful attackers often pose as fellow employees or customers in duress. Faced with the opportunity to help someone to whom they feel a responsibility, employees of targeted companies will often let down their guard, which compromises the network. Can technology assist humans in detecting and repelling these attacks, without undermining humanity itself?

Social engineering is a method of manipulating an individual to gain access to a larger network. Oftentimes, getting a single person to click on a link or divulge a password is the first step in a larger campaign. Ransomware is one example. Once an unsuspecting user provides access to a network, ransomware code either blocks a user login or encrypts the file system, rendering the network unusable until the victim pays a

ransom, at which point a key is provided to log in or un-encrypt the data and restore the system to normal.

The first known ransomware attack, back in 1989, traveled to its first victim via floppy disk, and demanded that the computer's owner pay \$189 USD to recover access to their files (Waddell, 2016). In the decades since, ransomware attackers have gradually shifted focus from individual computer users to larger and more lucrative targets in the infrastructure sector, such as hospitals, cities, and schools, leveraging the global connectivity of the internet and the anonymity of cryptocurrency to scale up their enterprises.

Ransomware is effective specifically because it interrupts key functions of individual lives and of society. It is a highly scalable method of "hacking people" – whether locking an individual out of their home computer so they are unable to access personal items or freezing a larger scale enterprise like a hospital or a city, ransomware is a lucrative business model because of its immediate personal impacts across a community.

The ransomware attack on the city of Atlanta in 2018 is one of many examples (Whittacker, 2018). For over a week, municipal functions were paralyzed – online services like utility billing, parking payments, and court proceedings were halted as city officials scrambled to restore backup systems or implement manual methods (Hutcherson, 2018). An estimated 6 million people were affected, and three months after the initial attack, 30% of the city's systems were still offline (Reuters, 2018). Atlanta's mayor likened the attack to a hostage situation.

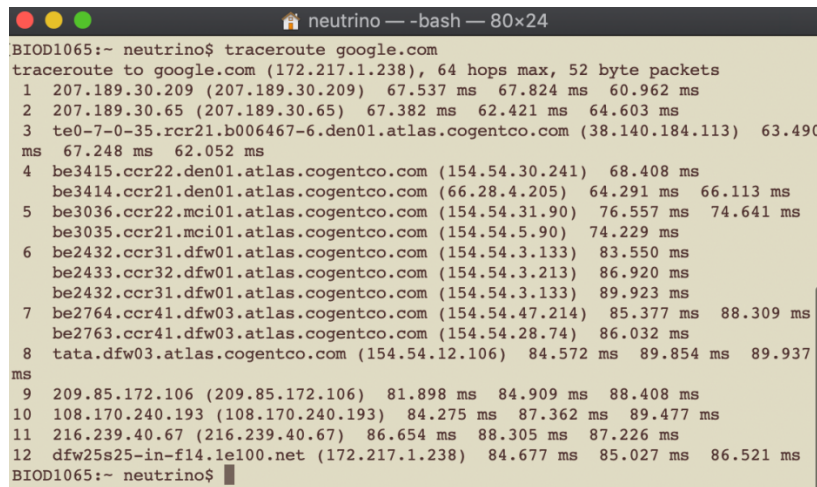
As disruptive and pervasive as ransomware attacks have been, they didn't garner much interest from the Federal Government until the 2021 attack on Colonial Pipeline. Colonial transports 3 million barrels of fuel per day, or 45% of the East Coast's fuel from the Gulf Coast (Sanger, Krauss, & Perloth, 2021), and with the pipeline out of commission for 5 days, fuel reserves dried up and gas prices rose across the country by 18-21 cents per gallon (Turton & Mehrotra, 2021). It was at this point that ransomware suddenly became a topic of unprecedented popularity at the federal level. The White House convened an international working group and the Congressional discourse focused on the ransomware problem to a greater extent than ever before. (The U.S. federal response will be analyzed in greater detail in Chapter 4.) Constituents got angry, businesses were negatively impacted, and two class action lawsuits were filed – one by consumers, alleging that Colonial's failure to “*maintain reasonable security measures, procedures, and practices*” resulted in higher gas prices (Dickerson v CDPQ Colonial Partners, LP, 2021), another by a group of over 10,000 gas stations with the same assertion (EZ Mart 1, LLC v. Colonial Pipeline Company, 2021). This language is often used in data breach cases, but in this instance, the harm is not exposed personal information but the higher cost of gasoline (Bryan & Lonergan, 2021). (Both lawsuits were dismissed in 2022 for failure to state a claim.)

Ransomware is effective because it targets many individuals at once, seizing critical services on which they depend. These populations are often co-located, when the entity under attack is a city, a hospital, or a school district. The Colonial Pipeline incident showed that a single attack has the potential for a widespread geographic effect. But there

are other ways in which geography can cause complications that are far more interesting than geographic extent.

## 2. Complexity in Geography: Russian Hackers and Chinese News

Physical location and cyberspace location have an interesting relationship that contributes to complexity. One frequent misconception among users is how traffic travels through the internet. When a person types in a web address – google.com, for example – the packets of information (an ARPANet-era invention) sent by the user and received from the target website don't go straight from one to the other. Instead, the packets travel through several different intermediaries. Figure 3 shows, using the Unix command “traceroute,” an example of a machine connecting to google.com.



```
neutrino -- -bash -- 80x24
BIOD1065:~ neutrino$ traceroute google.com
traceroute to google.com (172.217.1.238), 64 hops max, 52 byte packets
 1  207.189.30.209 (207.189.30.209)  67.537 ms  67.824 ms  60.962 ms
 2  207.189.30.65 (207.189.30.65)  67.382 ms  62.421 ms  64.603 ms
 3  te0-7-0-35.rcr21.b006467-6.den01.atlas.cogentco.com (38.140.184.113)  63.490
ms  67.248 ms  62.052 ms
 4  be3415.ccr22.den01.atlas.cogentco.com (154.54.30.241)  68.408 ms
   be3414.ccr21.den01.atlas.cogentco.com (66.28.4.205)  64.291 ms  66.113 ms
 5  be3036.ccr22.mci01.atlas.cogentco.com (154.54.31.90)  76.557 ms  74.641 ms
   be3035.ccr21.mci01.atlas.cogentco.com (154.54.5.90)  74.229 ms
 6  be2432.ccr31.dfw01.atlas.cogentco.com (154.54.3.133)  83.550 ms
   be2433.ccr32.dfw01.atlas.cogentco.com (154.54.3.213)  86.920 ms
   be2432.ccr31.dfw01.atlas.cogentco.com (154.54.3.133)  89.923 ms
 7  be2764.ccr41.dfw03.atlas.cogentco.com (154.54.47.214)  85.377 ms  88.309 ms
   be2763.ccr41.dfw03.atlas.cogentco.com (154.54.28.74)  86.032 ms
 8  tata.dfw03.atlas.cogentco.com (154.54.12.106)  84.572 ms  89.854 ms  89.937
ms
 9  209.85.172.106 (209.85.172.106)  81.898 ms  84.909 ms  88.408 ms
10  108.170.240.193 (108.170.240.193)  84.275 ms  87.362 ms  89.477 ms
11  216.239.40.67 (216.239.40.67)  86.654 ms  88.305 ms  87.226 ms
12  dfw25s25-in-f14.1e100.net (172.217.1.238)  84.677 ms  85.027 ms  86.521 ms
BIOD1065:~ neutrino$
```

Figure 3. Packets pass through several different machines when a user visits a website

The large number of machines used to execute this straightforward command increases the potential threat surface. Instead of two secure machines – the one sending the request and the one receiving it – twelve different machines (in this case) must be

secure. Considering the number of commands and other operations that the internet handles on a daily basis – this network is much more dense than most people realize.

In the prior example, all the machines transmitting information were located in the U.S., and there the user did not attempt to modify the route the packets took. But geographic boundaries are incredibly permeable when it comes to internet traffic. An internet user may be physically in a coffee shop in Phoenix, Arizona, but routing their traffic through Denver, Colorado, by a virtual private network (VPN)<sup>8</sup> company headquartered in Panama. With a few clicks, the author could send their packets through Ukraine, Bulgaria, or (and!) Malaysia. Were the user in question a malicious actor, they could perhaps be using an online service<sup>9</sup> to identify unprotected webcams in yet another country and infect them with spyware, instead of more innocuous commonplace activities. The global nature of the internet is undisputed, but the ease with which one can cross geographic boundaries, and the legal consequences of doing so, are much less clear.

Consider the case of *U.S. v. Aleksey Vladimirovich Ivanov* (United States v. Ivanov, 2001), a Russian hacker who infiltrated the networks of several businesses in the U.S.. After his intrusions, Ivanov then posed as a “white hat” hacker (one with good intentions), offering security assistance to a financial company he had breached in exchange for \$10,000 USD and persistent access to the network. An arrest warrant was issued, charging Ivanov with computer fraud (18 U.S. Code §1030), conspiracy (18 U.S. Code §371), and extortion (18 U.S. Code §1051). The FBI needed Ivanov on U.S. soil to arrest him, so they set up a counterfeit security company in Seattle, Washington, and

---

<sup>8</sup> Using a VPN protects a machine from people on a network that might have bad intentions. A coffee shop network is often unprotected – meaning no password is required – so using VPN is a good defensive measure.

<sup>9</sup> <https://www.shodan.io/> is a search engine that indexes internet-connected devices instead of websites.

invited Ivanov to apply for a job in person. Before the visit, Ivanov was requested to provide proof of his hacking skills – which he did, unknowingly, on a target machine (“honeypot”) set up by the FBI to capture all of Ivanov’s keystrokes and network traffic. The captured credentials were used to access Ivanov’s machines in Russia, confirming his complicity in the prior hacks. Within hours of arriving in Seattle, Ivanov was arrested (Koerner, 2002).

Ivanov was indicted on 8 counts – one count of conspiracy, three counts of computer fraud (CFAA violations), one count of extortion (the Hobbs Act), one count of disrupting commerce and one count of threatening to cause damage. Ivanov filed a motion to dismiss all of these, based on the fact that he was physically in Russia when the machines were compromised and therefore immune to U.S. law. Ivanov’s attorney, C. Thomas Furniss, stated “The U.S. doesn’t have the power to decide this case... The allegation is that he did a bunch of stuff from Russia using the internet. No country, including the U.S., owns the internet.” The court denied Ivanov’s motion to dismiss, based on the fact that his intent was to breach computers in the U.S., and his actions had an effect on assets within U.S. borders. In addition, the U.S. District Court for the District of Connecticut also claimed that “each of the statutes under which Ivanov was charged with a substantive offense was *intended by congress to apply extraterritorially*”. Having his motion denied, Ivanov pled guilty to several charges, and was sentenced to 48 months in federal prison. He was not only prosecuted in Connecticut for computer fraud, but also in five other district courts for similar crimes.

As a side note, is important to recognize that what the U.S considers “computer fraud and abuse” is radically different from how many other countries define these

activities, particularly in nations where economic circumstances can suppress legitimate career trajectories for technologists. "There is a very large group of educated individuals in Eastern Europe, people that have degrees in computer science, in mathematics," says Arif Alikhan of the computer crimes section at the U.S. Attorney's office in Los Angeles. "And I think the economic circumstances sometimes make it very, very attractive to commit crimes." (Koerner, 2002) These cultural differences – and assumptions about them – potentially undermine international cooperation and the establishment of global norms that could deter harmful activity like Ivanov's.

The interesting piece of Ivanov's story is not that he was convicted. It is clear that his actions harmed U.S. interests. However, the extent to which the FBI went to get Ivanov onto U.S. soil for the arrest is remarkable – an effort Russia claims is "illegal and criminal". (While Russia declared U.S. actions to be illegal, they did not fight for Ivanov's return.) The court's statement that the laws under which Ivanov was convicted were "intended by Congress to apply extraterritorially" is interesting as well, given the lack of international cooperation on cybersecurity and internet governance. The statement calls Ivanov's lawyer's statement into question – does the U.S., in some form, "own the internet"? To investigate this question, it is valuable to look at the ways in which a variety of international issues can be addressed.

Not every international conflict on the internet is related to criminal intrusion. Consider the case of Cable News Network (CNN) v Cnews.com (Cable News Network LP, LLLP v. CNNews. Com, 2001), a domain name registered by Maya HK, a Chinese news company. (Note that the domain name breaks down to *cn news dot com* – 'cn' being China's top-level internet domain code – not *cnn news dot com*.)



In the lawsuit, CNN claimed that the domain name ‘cnnews.com’ was an infringement of CNN’s trademark in both the U.S. and China, that it diluted their brand recognition, and was a violation of the Anti-cybersquatting Consumer Protection Act (ACPA), a U.S. law implemented in 1999 to deter the creation of misleading website names<sup>10</sup>. Maya HK moved to dismiss the case, claiming (in part) that the court lacked jurisdiction and that a court in China would be a more appropriate venue, given that all entities and evidence (except of the plaintiff) were located in China.

The court disagreed, siding with CNN. The judge’s ruling states that “the relevant public and private interests favor this forum [the U.S. District Court for the Eastern District of Virginia] because it is the [location] both of the domain name in issue and of the pertinent registry.” The website’s name was registered in the U.S., and the registration managed by a company called Verisign, which is the exclusive registry for all websites ending in a .com suffix. (In addition to the .com websites, Verisign also is the authoritative registry for all .net addresses, and manages the back-end systems for all .gov and .edu addresses, which gives them a considerable amount of global power.) The court admitted that it had no jurisdiction over the company or the individual who registered the website – but that it was not necessary, as the legal action was against the domain name itself (in legal terms, an in rem instead of an in personam proceeding).

When a complaint about a domain registration is filed, the ACPA requires that the domain name certificate be transferred to the court while the case is underway. Having established that the Eastern District Court of Virginia was the appropriate venue, based

---

<sup>10</sup> At that time, many website names were being created using trademarked names or phrases “in bad faith,” for the purpose of selling the website name to the trademark holder for significantly more money than the original name cost.

on the registration location, the certificate for cnnews.com was transferred accordingly, providing the court with complete authority over the domain name during the trial.

Maya HK made several arguments to dismiss, claiming that the ‘cn’ in the website name was shorthand for Chinese News, and that the intended readership of the website wouldn’t be familiar with CNN’s brand, so no confusion was likely. However, the Court decided that there was sufficient evidence that the Chinese company acted in “bad faith,” that they were intentionally leveraging CNN’s brand to advance their own economic interests, which put them outside the safe harbor provision of the ACPA. The matter was settled in favor of CNN and the domain registration transferred accordingly.

One could imagine different outcomes to this case, had the registry been located outside the U.S. There are no .com domain registries in China; Maya HK would have had to register cnnews.cn instead of cnnews.com if the company wanted to register inside their home country. In the decision, Judge Ellis noted the challenges: “the... issues this case addresses underscore the difficulties that inhere in the effort to use [U.S.] national laws to protect intellectual property rights and to police commercial activity on the global internet.” Much of the internet was developed in the U.S., so several of its key pieces are still within U.S. legal jurisdiction, even though they now perform international functions. Even conflicts that aren’t resolved in court may be decided by groups that are ostensibly international but have deep U.S. history and connections. But geography is only one element in a murky and volatile governance environment – even when parties in conflict are from the same country, authority can still be a complex issue.

### 3. Complexity in Authority: Blurring the Lines between Private and Public Sectors

Not all domain name conflicts go straight to court; many are arbitrated via a group called ICANN – the internet Corporation for Assigned Names and Numbers. ICANN is a nonprofit multi-stakeholder company intended to ensure internet stability and global participation, although its ability to arbitrate in a globally just fashion has been questioned, since ICANN reported to the U.S. Department of Commerce from its instantiation in 1998 until 2016 (Yu, 2009), (ICANN, n.d.). ICANN provides many organizational services that keep the internet functional, and one of its key contributions is the Uniform Domain-Name Dispute-Resolution Policy (UDRP). The UDRP provides a way for trademark disputes over domain names to be resolved outside the judicial system (faster and more cost-effective). Complaints are arbitrated by one of several international committees approved by ICANN to follow the UDRP guidelines. These committees use five “burdens of proof” to evaluate complaints. Per Diane Cabell, at Harvard’s Berkman Center for internet and Society (Cabell, 2000), the necessary factors are:

- (1) that the Complainant has trademark or service mark rights in the name,
- (2) the domain name is identical or confusingly similar to the mark,
- (3) that there was both registration and use of the domain name,
- (4) such registration and use was in bad faith, and
- (5) the registration and use was without rights or legitimate interest.

A person or company who feels they have a right to a particular domain name that has been registered “in bad faith” by another party can submit a complaint to ICANN, with supporting evidence and their choice of arbitration venue. The domain name holder is notified and must respond to the complaint within 20 days, or they automatically forfeit the name. A panel is appointed to review the complaint and response, and the decision is rendered within 45 days of the initial complaint (Babbitt, n.d.). This method was used in 2000 by American recording artist Bruce Springsteen in an attempt to claim the domain name *brucespringsteen.com*, which had been registered in 1996 by an unrelated party, Jeff Burgar. Springsteen’s complaint addressed each of the five criteria above and requested that the domain be transferred (WIPO, 2001). Springsteen had not trademarked his own name, but his fame was sufficient, his team maintained, to establish a common law trademark. The “bad faith” argument focused on the fact that Burgar had not just claimed *brucespringsteen.com*, but had registered more than a thousand domain names linked to celebrities, companies, and other organizations, which then redirected to Burgar’s own website, *Celebrity1000.com*.

Several of Burgar’s domain name registrations had been contested at this point, most recognizably by Mariah Carey and Hewlett Packard (Smith, 2001). Countering the assertion that his website was misleading to the public, Burgar claimed that using a celebrity’s name in a URL is no more confusing to the public than printing their name on a magazine cover. For example, upon seeing a Rolling Stone cover with Bruce Springsteen on it, no reasonable person would conclude that Bruce himself owned the magazine, and Burgar equated that magazine cover with the website’s URL. Burgar also pointed out that Springsteen’s record company had registered *brucespringsteen.net*

several years prior, instead of *brucesteen.com*, which he felt indicated that the domain name was not of official interest and therefore should be uncontested.

The dispute between Bruce Springsteen and Jeff Burgar is an example of interconnected complexities and the messiness that can result – it is a conflict between two U.S. parties, arbitrated by an international committee (the World Intellectual Property Organization), which ostensibly follows policy set by a U.S. corporation, which reports to a U.S. federal entity via a Memorandum of Understanding. This particular WIPO committee decided 2-1 in favor of Jeff Burgar, saying that in spite of Burgar’s history of registering domain names of celebrities and companies, and using them to redirect to his own website, that generally, “[the internet] is a valuable source of information in many fields, and any attempt to curtail its use should be strongly discouraged”. The domain name was therefore left in Jeff Burgar’s possession. As of this writing (August 2022), Bruce Springsteen’s official website remains *brucesteen.net*, rather than *.com*. (Typing the *.com* address into a web browser resulted in several security warnings, and the author elected not to proceed to the site, so it is unclear who currently owns it.)

This messiness in authority leads to considerable variance in how the different ICANN committees interpret the UDRP, and similar cases will end up with different results. Two weeks after WIPO’s decision on Springsteen, Burgar’s registration of *celinedion.com* was evaluated by a different WIPO panel. The circumstances behind the complaint were nearly identical, but this time, the panel decided the name had been registered “in bad faith” and Burgar was forced to reassign the domain name to Celine Dion’s record company (WIPO, 2001). Like the U.S. judicial system, different panels (or

courts) can interpret policy in radically different ways. But unlike the U.S. judicial system, ICANN has no hierarchy of evaluation committees to consider appeals, provide guidance, or remand cases for review. Those who go through the UDRP process and are unsatisfied with the outcome may bring their complaint to court (which is considerably more time consuming and expensive), but there are no pathways to appeal the decision through ICANN (Babbitt, n.d.).

Many key internet functions are utilized globally, but still effectively under U.S. control. The U.S. enjoys significant advantage from its role in developing the internet, as seen in the CNN case, and has thus far resisted most attempts to move these functions to a neutral governing body. Even though ICANN is not technically a U.S. government entity, it did report to the U.S. Department of Commerce for nearly 20 years, and the corporation is currently registered in California. For decades, there have been discussions in various fora on how to effectively provide international governance for the internet – an arguably international resource – but very little progress has been made. In 2015, President Barack Obama hoped to collaborate with Chinese President Xi on an “international framework” that would set basic standards for internet use and bring other countries to the table (Jackson, 2015), but little came of the intention. There are still no effective governing bodies that cross geographic borders, and in many ways, security suffers as a result, without commonly accepted standards of accountability and response.

Another complex element of cybersecurity is the ways public and private sector responsibilities resonate with one another. While they operate in different ways, the two are inextricably intertwined. The Sony Pictures breach is one example. In 2014, Sony

produced a movie, 'The Interview', that mocked North Korean dictator Kim Jong Un. The movie was perceived by the North Korean government as a grave insult against their leader, and that the offense deserved retaliation. North Korean military hackers broke into Sony's network, took over many of the machines and displayed a threatening message (see Figure 4), and exfiltrated sensitive and embarrassing email from Sony leadership (Zetter, 2014). Unlike many other attacks, North Korea openly took credit for the incident. These exfiltrated emails, along with password lists, unreleased footage from multiple movies, and other documents, were released publicly in an attempt to force Sony to abandon the movie. Sony initially conceded to these demands and cancelled the movie's release, until the Obama administration convinced them to reconsider, on the grounds that acquiescing to a foreign dictator would be a security issue for the entire U.S. – not just Sony. "The cyber attack against Sony Pictures Entertainment was not just an attack against a company and its employees. It was also an attack on our freedom of expression and way of life," explained U.S. Secretary of Homeland Security Jeh Johnson in a December 2014 statement. The movie was released to the public, and Secretary Johnson continued to encourage the private sector to work closely with the Department of Homeland Security to share information on cyber threats and protect U.S. interests.

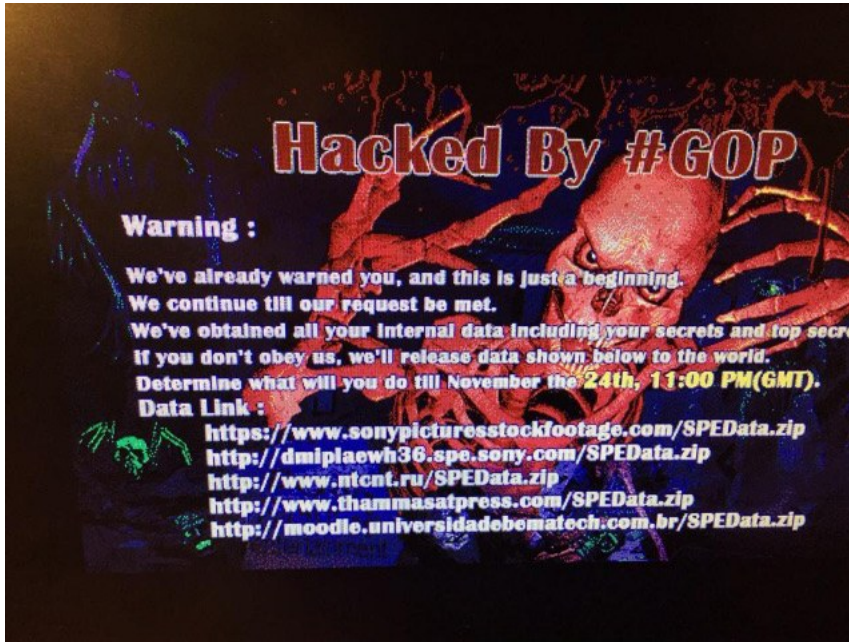


Figure 4. The screen that greeted many Sony employees after the North Korean intrusion. (Zetter, 2014)

The focus on public/private partnerships did not end with the Obama White House. The subsequent national security strategy of the U.S. under President Donald Trump (U.S. Office of the President, 2017) also specifically addresses cybersecurity partnerships with the private sector. “Since threats transit globally, passing through communications backbones without challenge, the U.S. Government will work with the private sector to remediate known bad activities at the network level to improve the security of all customers”, states the document in a section called ‘Deploy Layered Defenses’. Government and industry are encouraged to collaborate “in a way that respects free markets, private competition, and the limited but important role of government in enforcing the rule of law.” An executive order by President Joe Biden in 2021 echoes the need for the Federal Government and industry to collaborate closely (The White House, 2021). This encouragement is not a mandate, and many companies



don't participate in federal data sharing programs, due to a lack of trust or because of the additional burden (and therefore cost) that these programs invoke for the companies (Schwartz E. , 2018). The potential security benefits are negated, because the aims of the government and the responsibility of private companies don't align.

Cross-sector complexities can have enormous effects. When private companies are hacked, it's not just the individual and the company that are jeopardized; but there may be national security implications as well. Consider the value to a foreign adversary of 134.5 million credit records – the number of records lost in the Equifax data breach in 2017 (Winterton, 2017). An individual's credit history tells a lot about them – their financial security, their patterns of life, what they're willing to go into debt over. Scale that up by 134.5 million, and you get the same picture of a nation – the financial security and patterns of life of a nation's entire population, with their strengths and vulnerabilities. This is concerning in and of itself, but combined with other data breaches, is dire. Foreign hackers exfiltrated 21.5 million security clearance forms from the U.S. Office of Personnel Management – not just social security numbers, but exhaustive family information, financial details, and areas of technical expertise (Winterton, 2015). Combine these two data breaches, and a foreign adversary could have a very detailed picture of the U.S. national security enterprise and how it might be exploited (Winterton, 2017). As Senator Jeff Flake said in a November 2015 hearing on the Experian data breach, “What may not be sensitive as one item may become sensitive in the aggregate” (U.S. Senate Committee on the Judiciary, 2015). Given the history of data breaches, across multiple sectors, individuals, collectives, and countries are facing the potential

consequences of an aggregate-of-aggregates, a detailed multi-dimensional data set that affects individuals to nation states and all the levels in between.


#### 4. Complexity in Time: Technology Governance and the “Pacing Problem”

The term “pacing problem” refers to the rapidity with which technology moves forward and the much slower response of the mechanisms that govern it. Gary Marchant observes that, rather than closing, this gap is accelerating – that traditional legal tools are struggling to address issues in yesterday’s technologies as well as today’s (Marchant, 2011). internet-related technologies are no exception – their development and adoption has spread rapidly through society, but the laws that govern them have generally stagnated, resulting in a legal landscape that is ill equipped to govern the technologic reality.

The Computer Fraud and Abuse Act, the U.S. federal law addressing crimes that specifically include computer misuse and intrusion, was enacted in 1986. (Prior to the CFAA, computer-related crimes were often prosecuted using mail or wire fraud laws.) Concerned about the impact hackers might have on national security, inspired by the movie Wargames (Kaplan, 2019), President Ronald Reagan requested legislation that specifically addressed this threat. Since its instantiation in 1986, the CFAA has been amended to expand some of its definitions, but the core of the legislation remains the same. Therefore, U.S. federal cybersecurity governance in many ways resembles that of the year 1986, even though computers now play a radically different role in government, society, and individual lives.

The plot in Figure 5 shows the dramatic increase of the percentage of households with a computer, and the percentage with internet connectivity, since 1984 (Martin, 2021). From less than 10% in 1984, to over 80% (depending on the survey) in 2018, the connection between computers to households and individuals has grown considerably. One can also see that by 2013, nearly all households with computers also had internet subscriptions according to the Current Population Survey.

**Percentage of Households With Computer and Internet Use: 1984 to 2018**




Note: More information can be found at <[www.census.gov/cps](http://www.census.gov/cps)> and <[www.census.gov/acs](http://www.census.gov/acs)>. Sources: U.S. Census Bureau, 1984–2017 Current Population Survey (CPS) Computer and Internet Supplement, 1993 CPS Education Supplement, 1994 CPS Voting and Computer Use Supplement, and 2013–2018 American Community Survey (ACS), 1-year estimates.

Figure 5. Percentage of households with computer and internet use: 1984 to 2015 (Martin, 2021).

The fundamental definition of what a “computer” is has also changed since 1986, when the term generally referred to a beige box on a desk that allowed a person to create documents, perform mathematical calculations, and play basic games (Computer History Museum , 2021). A 2021 study by the Pew Research Center found that 75% of American adults own a desktop or laptop computer, 50% own a tablet, and 85% own a smartphone (Pew Research Center, 2021), as shown in Figure 6.

### Mobile phone ownership

% of U.S. adults who say they own a ...



Note: Respondents who did not give an answer are not shown.  
Source: Surveys of U.S. adults conducted 2002-2021.

Figure 6. Mobile phone and smartphone adoption, 2000 – 2020 (Pew Research Center, 2021).

In addition to computers and smartphones, computational power – and internet connectivity – are now built into Christmas tree lights (Peters, 2018) and coffee mugs

(Peters, 2017), and while one may not think of these things immediately as “computers”, it would be irresponsible to leave them out of the cybersecurity space, as they contribute to the overall threat surface. According to a Deloitte survey, the average household in 2020 had 11 internet-connected devices. Only 7 of these devices include screens for user interaction (Westcott, et al., 2021), but any internet-connected device can have vulnerabilities with serious security impacts.

Humanity’s relationship with computers has grown from the box-on-the-desk model of 1986 to include omnipresent, wearable devices, as well as those that enable critical societal functions. Families and friendships remain strong through internet-enabled communication, even when people are separated geographically (Boase, Horrigan, Wellman, & Rainie, 2006). Broadband internet adoption has positively impacted jobs and income in rural areas (Whitacre, Gallardo, & Strover, 2014). Financial transactions, like paying bills or acquiring credit, used to take several days but now can be done within minutes, if not seconds. Educational resources, like libraries and classes, are now open to a wider range of people, both socioeconomically and geographically (Selwyn, 2014). Remote monitoring of implanted devices, such as pacemakers, has been more effective than traditional methods for routine assessments (Roberts & ElRefai, 2019).

In addition to these intentional and visible examples, there are many positive interactions between humans and computers that are less observable. Even those who shun internet connectivity still have their credit history stored online, their medical records transferred via internet, their utilities managed and billed online, and so on. Much of this happens without an individual’s express involvement, although they still benefit

from this connectivity. For example, a person walking into a bank in person can apply and be approved for a loan in minutes, because their credit can be checked online. A home's utilities can be monitored remotely, regardless of how a homeowner pays their utility bill. A patient's medical imaging and history can be shared quickly with their medical care team, expediting treatment. Since 1986, the internet has grown into one of the major arteries of society, enhancing and expediting the necessary functions of our civilization. Given this density of connectivity between humans and computers, whether visible or not, the space to secure is much larger and more complex than it was in 1986 when the CFAA was enacted.

This growth has also come with considerable downsides. Multiple large-scale data breaches, each exposing more than a hundred million records, have happened in the past few years. (A "record" is a piece of personal information. It could be a username and password, it could be a credit card or social security number, or it could be a person's medical file.) Yahoo! has the dubious honor of the largest global breach by number of records – 3 billion email addresses were exposed in 2013 (Newman, 2017). As of April 2020, Wikipedia listed<sup>11</sup> 29 separate breaches where 100 million or more records were exposed, either because the breached party was intentionally hacked, or because they accidentally exposed their own sensitive data<sup>12</sup>. In July 2022, the list contained 33 such events. The problem is so extensive that the term "breach fatigue" emerged to describe consumer's waning concern over these incidents (Zorabedian, 2019).

---

<sup>11</sup> This dynamic list of breaches over 30,000 records is composed of press releases, news articles, and government reports, and is updated frequently.

<sup>12</sup> [https://en.wikipedia.org/wiki/List\\_of\\_data\\_breaches](https://en.wikipedia.org/wiki/List_of_data_breaches)

Critical functions of several cities and hospitals have been severely impacted when their computer systems were encrypted and held hostage by hackers, who demanded a ransom before de-encryption keys were provided. According to Josephine Wolff, assistant professor of cybersecurity policy at Tufts, more than 1,000 health-related organizations have been hit with ransomware since 2016, costing the U.S. health care system more than \$157 million USD (Wolff, Slate, 2020). Vulnerabilities have been found in medical devices like pacemakers (Newman, 2018) and insulin pumps (US Food and Drug Administration, 2019), and while it is not yet known whether or not these vulnerabilities have been exploited, hundreds of patients and doctors wonder when – not if – they’ll have to face an *in vivo* cyberattack that could be a matter of life and death (Frank, 2017). Companies see their own intellectual property in knockoff products that suddenly emerge overseas: a loss of millions of dollars of research and development money annually (Rogin, 2012).

### **Updates to U.S. federal computer laws**

Given these monumental changes in internet-related technologies and the space they occupy in society, it is surprising that the CFAA has repeatedly been amended to be more general, rather than more specific. Most updates to the law have been expansions to its scope, broadening definitions that were already vague and adding new prohibitions and penalties (Sarian, 2021). The most recognizable attempt to reshape the CFAA’s scope to fit modern internet usage is Rep. Zoë Lofgren’s proposed amendment, H.R. 2454, introduced in 2013 and referred to as “Aaron’s Law” after the late internet activist Aaron Swartz.

Aaron Swartz was 26 years old when he died by suicide, having been pursued by the U.S. Department of Justice for two years over illegally downloading 4.8 million journal articles from JSTOR, an online database (United States v Aaron Swartz, 2011). JSTOR provides free access to articles for those with an organizational subscription, as many universities have, but these organizational licenses are expensive, and they prohibit scripted downloads and dissemination. Much of Aaron's activism focused on making information freely available; he was appalled by gatekeeping of knowledge. Swartz accessed JSTOR from a closet on the MIT campus, using a guest account, and his script at times requested up to 200,000 journal articles per hour. Swartz was identified after a video camera was placed in the closet, and he was arrested by MIT police and a Secret Service agent. JSTOR declined to pursue a civil lawsuit once Swartz turned over the downloaded articles, but federal prosecutors pursued a total of 13 felony counts under the CFAA, with maximum penalty of 35 years in prison and \$1M in fines. During negotiations, prosecutors offered a plea deal – 6 months in federal prison, if Swartz pled guilty to all 13 counts. Swartz and his team rejected the deal. They had strong counter-arguments from key experts, and Swartz's lawyer felt increasingly confident they could win in court. Four months before the trial, though, Swartz hanged himself, setting off a flurry of criticism from tech and legal experts, and inspiring Lofgren's proposed changes to the CFAA. Federal prosecutors dropped the charges after Swartz's death. Nearly ten years later, Swartz's ambitions for opening knowledge would be realized – in August of 2022, the Biden administration released an executive order mandating that taxpayer-



funded research be freely and immediately accessible to the public (White House, 2022).<sup>13</sup>

Lofgren's bill proposed a more specific definition of "unauthorized access," based on the security measures of the technology rather than policies such as user agreements and terms of use. Aaron's Law would have also removed redundant sections of the CFAA that had been used by prosecutors to attach multiple penalties to a single offense, significantly decreasing the sentences that could be pursued, especially for a first-time offender (Lofgren, 2013). Aaron's Law stalled in subcommittee in 2013, and was reintroduced in 2014 only to suffer the same fate. Swartz's biographer, Brian Knappenberger, claimed that the software company Oracle was to blame for lobbying against the bill (Dekel, 2014). Big technology companies had long since benefited from the CFAA's extensive reach. Even before Lofgren's bill was introduced, the Software and Information Industry Association – which represents companies like Google, IBM, and Oracle – had been lobbying Congress to retain the CFAA's broad scope, stating its *"opposition to proposals that would limit the definition of "exceeds authorized access" in the CFAA in any way that would prevent its application to violation of contractual obligations or agreements"* (Maas, 2013). Aaron's Law, defining unauthorized access in a purely technologic way, ran counter to the SIAA's desire to retain terms of service and other policy agreements as effective methods of corporate prosecution under the CFAA.

As computing and internet connectivity became more pervasive, interpretations of copyright law evolved to address issues around reverse-engineering computer code (Sega

---

<sup>13</sup> Justin Peters' 2016 book 'The Idealist: Aaron Swartz and the Rise of Free Culture on the internet' is an excellent resource, detailing Swartz's pursuits, and the larger context of information control in the internet age.

Enterprises, Ltd. v. Accolade, Inc, 1992), and anonymous online file sharing (Well Go USA, Inc. v. Unknown Participants in Filesharing Swarm Identified by Hash B7FEC872874D0CC9B1372ECE5ED07AD7420A3BBB, 2012). Questions around the CFAA’s scope – such as “Does breaking a company’s terms of service constitute ‘unauthorized access’?” (United States v. David Nosal, 2011), and “Does collecting publicly accessible information constitute ‘unauthorized access’?” (United States v Andrew Auerheimer, 2011) – have been answered inconsistently across the Circuit courts. To resolve some of these differences, the Supreme Court recently considered Nathan Van Buren v U.S. (Van Buren v United States, 2021), another case involving questionable computer access.

As a police officer in Georgia, Nathan Van Buren had the ability to access specific law enforcement information via the Georgia Crime Information Network in his police vehicle. Van Buren accepted a \$5000USD bribe from an outside party to use his police credentials to research a local exotic dancer and find out if she was an undercover police officer. She was – but the person offering the bribe was an FBI informant, and Van Buren was convicted under the Computer Fraud and Abuse Act for unauthorized access of the network (Van Buren v United States, 2021).

Prior to the Van Buren decision, U.S. circuit courts were split as to what constitutes “authorized access.” The 11<sup>th</sup> circuit took a policy-driven approach – what information and actions are allowed, based on terms-of-use and other situationally dependent documents? The 4<sup>th</sup>, 9<sup>th</sup> and 2<sup>nd</sup> have followed a tech-based approach – what information and actions are allowed, based on the system’s configuration and access controls (O'Connor, 2021)? The U.S. Supreme Court has ruled with the 4<sup>th</sup>, 9<sup>th</sup> and 2<sup>nd</sup>

circuits, 6 to 3, adopting a “code-based” definition of “authorized access” (Van Buren v United States, 2021), deciding that while Van Buren’s behavior violated department policy, he *did* have authorized access to the computer systems in question as defined by the CFAA because the technology did not prohibit him from doing so. In this decision the Court emphasizes the technical requirements of the system – the information the user can potentially access using their proper credentials – when determining “authorized access,” rather than terms of use or corporate policies. (Note that this interpretation aligns with Rep. Lofgren’s proposed amendment to the CFAA from 2013, which sought to define unauthorized access as “*knowingly circumventing one or more technological or physical measures that are designed to exclude or prevent unauthorized individuals from obtaining that information*”.)

### **Van Buren v U.S. – U.S. Supreme Court decision and dissent**

What effects might this decision have, going forward? The Van Buren decision clarifies that responsibility for authorized access rests with those who design and manage a system’s technology, rather than the individual users of that system. This “code-based approach,” long advocated for by legal scholar Orin Kerr (Kerr O. , 2021), does provide a clearer assessment of “authorization” – what is and what is not allowed for a given user’s credentials.

Generally, the SCOTUS decision redefines the responsibility for authorized access from individual users to a more central management function: system administration. A technical interpretation (rather than policy-centric) of “authorization to access” defines boundaries based on the user’s credentials and system configuration

rather than by corporate guidelines or terms-of-use. Instead of each user relying on their own understanding of a network's policy, the job of protecting data now rests clearly on system administrators, individuals with specific authority and technical knowledge, to define and manage access. Per the majority opinion on the case, written by Justice Barrett: "*An interpretation that stakes so much on a fine distinction controlled by the drafting practices of private parties is hard to sell as the most plausible.*" In other words, boundaries defined by technology are more clear, and therefore more defensible, than those set forth in terms of service or other guidance documents. Nathan Van Buren did not have to break through any technical boundaries to access the police records database; he had been given authorized access through his credentials as a police officer. His intentions were subversive, but not a violation of the CFAA.

It's worth noting that most of the highly visible CFAA cases involve behavior that would be hard to justify morally. Nathan Van Buren accepted a bribe to disclose the identity of an undercover officer; David Nosal used another employee's password to undermine his former employer (United States v. David Nosal, 2011); Gilberto Valle used a federal police database to find addresses of women he had fantasized about kidnapping (United States v. Gilberto Valle, 2015); Andrew Aurenheimer ("weev") has been labeled an extremist by the Southern Poverty Law Center for being a notorious antisemitic troll who continually spews violent pro-Nazi hate around the internet (United States v Andrew Aurenheimer, 2011), (Southern Poverty Law Center, Accessed 11 September 2022). The question is not whether these individuals are "good people", innocent of any crime, but whether or not their actions specifically constitute

unauthorized access under the CFAA. Aurenheimer's personal world view, while horrendous, does not change the fact that the data he scraped from AT&T's website was openly available, and no "hacking" methods were necessary to obtain it. The CFAA, per the SCOTUS decision, should be used to punish computer intrusions (or "hacking") from a technical standpoint, not to make decisions based on the alleged intruder's motivations or purpose.

Per the court's decision, this perspective removes the "*arbitrariness into the assessment of criminal liability.*" With the authorization question now answered by the technology, there are some known solutions that can be applied. Organizations can strengthen role-based access controls, dividing users into groups based on their need for information, and defining permissions based on those roles. Network segmentation is another technology that could be employed to secure a system. If data are sequestered in various locations across the network, rather than all in one place, it becomes more difficult for a bad actor to steal all of them.

In its decision, the court also acknowledges that, due to the expansion of human/computer connections, "*the prohibition now applies – at a minimum – to all information from all computers that connect to the internet.*" This is a very broad "minimum" – almost to the point of being amusing – but it is an explicit and necessary realignment of the CFAA with modern internet usage. This direct acknowledgement differentiates the 2021 SCOTUS interpretation of the CFAA from its original assumptions in 1986, when only certain computers were connected to the internet and even fewer were used for financial transactions.

Shortly after the SCOTUS decision on the Van Buren case, the U.S. Department of Justice (DoJ) issued a new policy on how the Federal Government would approach the CFAA (U.S. Department of Justice, 2022). The policy provides guidance to DoJ attorneys on how to apply the SCOTUS ruling, with additional detail on how and when the CFAA should be invoked, and what evidence should be collected to support the decision. The policy also creates a “safe harbor” for security researchers, instructing federal prosecutors to “*decline prosecution if available evidence shows the defendant’s conduct consisted of, and the defendant intended, good-faith security research.*” The policy does not restrict the definition of “research” to an academic sphere – any researcher who can show a good faith intent is covered. The policy goes on to define “good faith” research as accessing of a system to identify and correct flaws, without causing harm, and without the intention of exploiting the vulnerabilities that have been found.

The concept of a “safe harbor” in the CFAA has been advocated for by security researchers for years (Ellis, 2018), such as Rianna Pfefferkorn from the Stanford Internet Observatory (Pfefferkorn, America’s anti-hacking laws pose a risk to national security, 2021), and Jeff Moss, the originator of the DEF CON hacking conference (Detch, 2016). The DoJ policy has been cited by the Electronic Frontier Foundation as a positive step for security researchers (Crocker, 2022). It should be noted, however, that the policy only applies to Federal Government attorneys, and would be much easier to repeal than a formal amendment to the law. The DoJ policy also explicitly calls out “cease and desist” letters as an acceptable barrier for authorization, which seems to run counter to the “code-

based” definition of authorized access, and which companies often use to shut down competition or inquiries they don’t like (Facebook v Power Ventures, Inc, 2019), (hiQ Labs, Inc. v. LinkedIn Corp., 2022). Given these weaknesses, it's likely that security researchers will continue to pursue true CFAA reform in the years ahead.

Van Buren v U.S. was not unanimously decided by the Supreme Court justices, so it is both fair as well as instructive to analyze the dissenting opinion, written by Justice Clarence Thomas. The majority opinion, written by Justice Amy Coney Barrett, claims that a policy-based definition of “authorization” would criminalize many commonplace activities (such as checking a personal email account from a work computer). Justice Thomas makes the counterargument that the CFAA’s intent can be understood by “an ordinary reader of the English language”, and that the “plain meaning” of the statute clearly does not include such activities. However, terminology is rarely so deterministic, particularly around technology (as discussed in Chapter 1), and to rely on a person’s reading level in English as being “ordinary” or otherwise seems discriminatory when large fines and jail time may be potential outcomes.

The dissent also claims that the majority’s interpretation would provide bad actors with immunity, invoking the analogy of a scientist providing an enemy nation with information on nuclear weapons. If the scientist has technologically sanctioned access to the information, Justice Thomas claims, they could not be prosecuted for malicious use of the material. It is true that no technical barrier can adequately determine a person’s intent. But there is little reason to believe that this hypothetical situation would be treated solely as a CFAA violation, as several other laws would apply, such as export control,

classification authority, and federal acquisition. The CFAA does not need to cover every conceivable access control issue, particularly when the ramifications of doing so would criminalize standard, non-malicious user behavior, such as using a work laptop to pay bills. Very few people segment their work and personal lives so carefully, and this entanglement has gotten even more dense due to the pandemic of the early 2020's, which pushed many people to work from home to ensure their physical safety.

### **A post- 'Van Buren' internet – expectations and remaining questions**

It is valuable to reconsider key CFAA events of the past in the new light of Van Buren, such as Justin Shafer and Patterson Dental. Justin Shafer, a security researcher and software technician, discovered that a dental group was storing unencrypted patient information online (Doe, 2016). The server was not password-protected and easily accessed via a browser, and personal information on 22,000 patients had been openly available for years. Shafer appropriately alerted the dental company of his findings and worked through a responsible disclosure process, only to be the subject of a 6am raid by the FBI at his home, being arrested in front of his children and taken to jail in his boxer shorts. The FBI's justification for the raid was that Shafer had "exceeded authorized access" by finding the information. To be clear, Shafer did not sell the data or otherwise exploit his finding. Instead, he worked with DataBreaches.net, a group that aids researchers in responsible disclosure, to inform the company of their problem. Patterson Dental didn't respond to Shafer or DataBreaches (Doe, 22,000 dental patients' info exposed on unsecured Eaglesoft FTP server , 2016), but instead claimed that Shafer should be criminally charged under the CFAA and called the FBI.



The personal health data that Shafer found were publicly accessible on the web; no technical intrusion method (or “hacking”) was needed to obtain them. In the absence of any gates to keep people out, Shafer (and anyone else on the internet) would have had “authorized access” to those data, per the Van Buren ruling, so his disclosure would not today be considered a violation of the CFAA. This is a considerable step forward for security researchers and other tech-savvy internet users who may find themselves in a similar position in the future.

While the Van Buren decision provides some needed clarifications on “authorized access”, it doesn’t address the question of how to assess loss or damage, nor does it rescop the onerous punishments in section (c), all of which are key considerations in the U.S. v Aaron Swartz incident. It would be challenging, even given the Van Buren decision, to reasonably claim that Aaron Swartz had authorized access to the MIT JSTOR archive, having allegedly circumvented several protections that MIT and JSTOR put into place once his downloads were discovered (Kerr, 2021). Swartz downloaded 4.8 million journal articles, and JSTOR’s response to his aggressive downloading cut the MIT campus off from the JSTOR archive for a handful of (non-subsequent) days. However, the articles weren’t deleted from the repository, nor was JSTOR access impeded for MIT for more than a day at a time. JSTOR declined to prosecute, but the Federal Government persisted, charging Swartz with 13 felony counts (Peters, The Idealist, 2013). How was the damage of the Swartz incident assessed, and how were the proposed punishments (\$1M in fines and 35 years in federal prison) justified? The problematic pieces of the CFAA that enabled the prosecution’s onerous response still remain.

It is too early to tell how the new onus of responsibility for “authorized access” will be implemented by industry and government. Per the Van Buren decision, access to a computer must now be governed by a technical solution, rather than by policy. Companies or government groups that don’t have such safeguards in place will have to create them and ensure that they are implemented correctly. These technologies do exist – as mentioned previously, role-based access control (RBAC) is one of these approaches. RBAC assigns roles to each user of the system, allowing them to view relevant data for their position and hiding data that is sensitive or irrelevant. While this seems simple, some roles overlap, and some change over time. An individual’s mother-in-law shouldn’t have access to their health records, but their dentist should, and if their mother-in-law is also their dentist, she will need some of that information to provide the patient with the proper care. But should she also have access to other health information, like mental health records? And if this patient changes providers – perhaps due to a divorce – at what point is the dentist’s access revoked?

The archaic posture and vague definitions in the CFAA, as well as the inconsistent interpretation between Circuit courts, has caused a “chilling effect” on independent security researchers, or even just security-aware users that may notice vulnerabilities in the systems they use on a day-to-day basis. The ruling in Van Buren v U.S. provides some welcome clarification on the definition of “authorized access” of data (as were found by Justin Shafer), but leaves untouched other problematic aspects of the law. The section on punishments for unauthorized access remains archaic and heavy-handed. There is no broadly accepted definition of other key terms, like “loss” or “intentional.” Now that “unauthorized access” is defined as crossing a technological

barrier, does it follow that any action subsequent to the intrusion is considered “intentional”? With the question of “authorized access” resolved (to a degree), questions like these may emerge in the future.

As of this writing, this decision is relatively new – its effect in the lower courts has yet to be explored (although some cases have been remanded for further consideration, e.g., *LinkedIn v HiQ Labs*). These issues with the CFAA have resulted in the arrest of well-intentioned researchers, even after attempts to follow a responsible disclosure procedure (Doe, 2020). In the weeks after the arrest of Justin Shafer, the privacy researcher who attempted to alert a dental office about personal health data they were inadvertently storing online, one anonymous privacy advocate received hundreds of undisclosed vulnerabilities from other researchers, who refused to go through a responsible disclosure process due to the potential for prosecution under the CFAA (Doe, 2016). This kind of social precedent among researchers cannot be instantaneously reversed, even though the SCOTUS decision has been praised by some in the hacker community (Barth, 2021). Experience has shown that companies often have extensive legal resources at their disposal, and those that feel threatened are likely to find an adequate prosecutorial hammer other than the CFAA if necessary. It will also be interesting to see if the *Van Buren* decision applies to flawed code, in addition to accessible information. Do the same “gates” apply for discovering problems in the software itself, such as potentially exploitable buffer overflows, or are those discoveries somehow different from accessible information?

While the SCOTUS ruling on *Van Buren v U.S.* does resolve an important conflict in the system, it is not yet known whether or not this change will result in

meaningful cybersecurity improvements, from a legislative perspective. It could allow for more disclosures by the security community, unearthing previously unreported problems that can now be corrected. However, may also remove an important barrier to detrimental behavior that may no longer be prosecutable under the CFAA. Recall that Nathan Van Buren used his account on the police database system in an attempt to expose an undercover officer. The ruling means that web scraping – mass collection of openly available data – is no longer a CFAA violation if that information wasn't put behind adequate technical gates. Will the Van Buren ruling embolden malicious actors to search for and exploit those data, now that one of the primary penalties for doing so is gone?

When the CFAA was introduced, the ways in which a computer could be accessed were limited and deliberate, and “unauthorized access” could be more clearly defined. The technological landscape has changed, but the position of the U.S. Federal Government has failed to keep pace, allegedly resulting in fewer security issues being reported and corrected.

### **Addressing Cybersecurity Complexities, Not Just Problems**

The world is not hurting for new cybersecurity products. Per a report by Gartner, a regarded tech consulting firm, the global cybersecurity market was estimated at \$150B U.S. dollars in 2021 (Gartner, 2021). This value doesn't include government efforts to secure systems, either through legislation or information sharing programs that connect industry. But the world is hurting for effective cybersecurity *solutions*. Investments by the public and private sector continue to grow, but there has been no abatement in the rate of data breaches or high-profile attacks (Wolff, 2018). The current perception of

cybersecurity as a purely technical problem drastically limits the space of potential solutions, specifically excluding some of the most critical elements. Without a more balanced approach, international cooperation will continue to be stunted, and the dangers of cyberwar will accelerate. Gaps in public/private oversight will remain – some segments will suffer from conflicting and competing authorities, and others will suffer from a lack of oversight. Human trust will remain a vulnerability in any online system, and will continue to be exploited.

An approach that leverages multiple disciplines is important, but more specifically, effective solutions will be scarce without a better understanding of conflicts between the technical, legal, and social domains. Of specific importance are the incentive structures of the three aspects– what risks and rewards motivate them, and how those elements contribute to clashes between the three domains. Research into these misalignments and the resulting systemic vulnerabilities will provide a fresh lens through which potential solutions can be conceptualized and evaluated, and the three elements brought into better harmony.

## CHAPTER 3

### **BUILDING A RELEVANT TAXONOMY**

Despite increased attention and funding from companies and governments worldwide over the past several years, cybersecurity incidents remain both frequent and severe. The current approach to solving these problems is to address them one by one, treating them as individual events, rather than looking for the commonalities between them and forming a more effective general solution.

To borrow a common analogy, when it comes to addressing cybersecurity problems, “we can’t see the forest for the trees” – and unfortunately, the forest is currently on fire. It is possible in theory (although not always in practice) to know quite a bit about an individual cybersecurity incident – the number of personal records exposed, the dollar amount paid in a ransomware attack, the range of IP addresses used in a coordinated campaign, and possibly the source of the attack. A substantial amount of time and effort goes into analyzing a specific attack, and for good reason. The breached organization could be called out as negligent by the Federal Trade Commission, and therefore subject to considerable fines and onerous federal oversight processes, if they did not mitigate the known factors that led to an incident. Various incident response frameworks have been created, such as those published through the Department of Homeland Security’s Cybersecurity and Infrastructure Security Agency (CISA) (National Institute of Standards and Technology, 2018). They are valuable in their context, but returning to the analogy, these approaches are akin to determining how a specific tree

caught fire and the process by which it burned, while the rest of the forest remains aflame.

In addition to being absorbed by the details of an individual event, the current approach to cybersecurity incidents often ignores the commonalities between different events. Incidents are often referred to by the type of attack, such as “ransomware” or “SQL injection”, but these labels are of limited utility when it comes to understanding the impact on individuals, communities, and society at large. Attack-based descriptions are valuable for tactical approaches, such as incident response, but on their own they lack the power to describe who or what was affected, and the extent of the harm inflicted.

Grouping incidents only by the type of attack also fails to show the evolution in the threat landscape. For example, the Colonial Pipeline attack was a ransomware incident, and ransomware has been a common attack vector for many years. The way ransomware has been deployed has been evolving, though, and unlike any previous event, the Colonial Pipeline incident in 2021 caused a significant shift in federal policy (more on this in Chapter 4). The example of Colonial Pipeline indicates that there is room to create more impactful and insightful ways to parameterize the space of cybersecurity incidents, in addition to the labels that are needed for short-term tactical mitigations. In terms of the forest fire analogy: how should the trees in a (cybersecurity) forest be most effectively categorized to slow the spread of the current fire? It may be that the number of affected trees is less insightful than their species, the mix of tree types, how closely they are spaced, whether or not they are on a certain type of terrain, or the fact that they are actually a single organism with a shared root structure like the Pando aspen grove in Utah

(DeWoody et. al., 2008). Standard language, such as labels or metrics, is often insufficient to create new descriptions of existing problems. Using pre-existing categories can impose boundaries on our ability to describe things, and therefore hinder our understanding. A general, broad-scale approach often requires new verbiage, which can then be used to design and build new analytic methods.

Current U.S. federal (and often state) approaches to cybersecurity are tuned to the immediate mitigation of a single event, and in the rush to understand and contain the damage, the context (or landscape, to continue the forest fire analogy) is often neglected. Russian cybercrime syndicates, such as REvil and DarkSide, have been responsible for many visible attacks on U.S. infrastructure (and probably many more attacks which aren't as visible). Attribution of an attack – defining the source of the attack and possible motivations behind it - is a very challenging technical problem, but the real challenges begin once attribution has been determined (Singer & Friedman, 2014). The political, economic, and social landscapes of attacker and victim are often the dry vegetation around the trees that make the difference between a small fire and one that continues to spread out of control. This detritus that propels and perpetuates the fire is separate from the tree (or trees), but it remains a key element of understanding the true nature of the problem and mitigating negative effects.

Finally, much as the aggressive wildfire policies of the past have inadvertently increased the spread and lethality of today's forest fires, existing cybersecurity policy in the U.S. is also contributing to a dangerously flammable digital ecosystem. These well-intentioned actions of the past contribute to and exacerbate the problems experienced



today. It is difficult to see these connections, however, when there is too much focus on the incidents themselves, rather than the relationships and connections that comprise the system and have shaped it over decades. Without a sufficient temporal understanding of the problem, whether it be the history of wildfire management or the policy decisions made in the early days of the internet, the system's emergent behaviors will remain a mystery, and proposed solutions will be incomplete.

Creating a generalized description of any complex system is a challenge, whether it be a forest fire or the increasing trend of cyber attacks on U.S. critical infrastructure. But failure to adequately describe and then address the general system will result in a continuation of the current trend, which has already wrought considerable negative impacts on the security of individuals, society, and nations. Fortunately, methods of abstraction do exist for certain complex systems, from a diversity of disciplines, and these can be studied to inspire approaches in novel domains.

### **Creating a Taxonomy**

Deciding on the relevant building blocks – and defining what is “relevant” – is a valuable first step in abstraction. Often this means creating a unique set of categories, based on what is common between relevant parts of the system. A taxonomy is one important, initial step in creating a generalization, to understand the fundamental behaviors of a system, to approach a wider variety of problems in an efficient way, and to make comparisons that help build new knowledge. A taxonomy for approaching

cybersecurity incidents proactively does not yet exist<sup>14</sup>, and it is necessary to create one in order to make sense of the landscape. For one, the number of known cybersecurity incidents is so large that, even given the specific scope of this specific inquiry, it would be untenable to investigate each of them independently and expect to come to any generalizable conclusions<sup>15</sup>. Second, the number of existing incidents is quite large, but the speed at which they are occurring is also considerable. Third, the process of creating a taxonomy allows the opportunity to specifically assess what should be considered an “incident” worthy of study, and what features make it so. The process of creating a defensible set of categories requires that one think carefully about which potential data speak to the research question and which do not. A taxonomy of relevant incident types allows a purposeful approach, building with intentionally curated groups of diverse incidents, and permitting new events to be assigned to existing categories.

Choosing the features by which incidents will be categorized is of the utmost importance. What one considers “significant” in creating a taxonomy is driven by the goal. One example of a purpose-driven socio-technical schema is the Cloud Computing Value Chain (Kolevski et. al., 2020), which developed a novel view of stakeholders,

---

<sup>14</sup> Some cybersecurity incident classification guides have been created, such as the ‘Cybersecurity Incident Taxonomy’ by the European Union Agency for Network and Information Security (European Union Agency for Network and Information Security, 2018), or the SWIFT Institute’s ‘Cyber Security Ecosystem’ (Ferdinand, 2017), but these catalogues have been created specifically to aid organizations with incident response. As such, they are reactive instead of proactive, and therefore ill-suited to the specific purpose of this work. Similarly, there are guides that detail the types of security flaws that may lead to an incident (e.g. (Aslam, Krsul, & Spafford, 1996)), but the need here is to classify the incidents themselves.

<sup>15</sup> The taxonomy is necessarily limited to cybersecurity incidents which are known as of this writing, but it is important to acknowledge that there are many unknown incidents as well, and the distribution of these unknown incidents may be sufficiently different from those that are known. Significant future changes in network monitoring technologies and/or disclosure policies may impact the findings in this dissertation, but as of this writing, this taxonomy represents the body of currently known incidents.

operational and non-operational, and the interactions between them. A diagram of these stakeholders, from the paper, is shown in Figure 7:

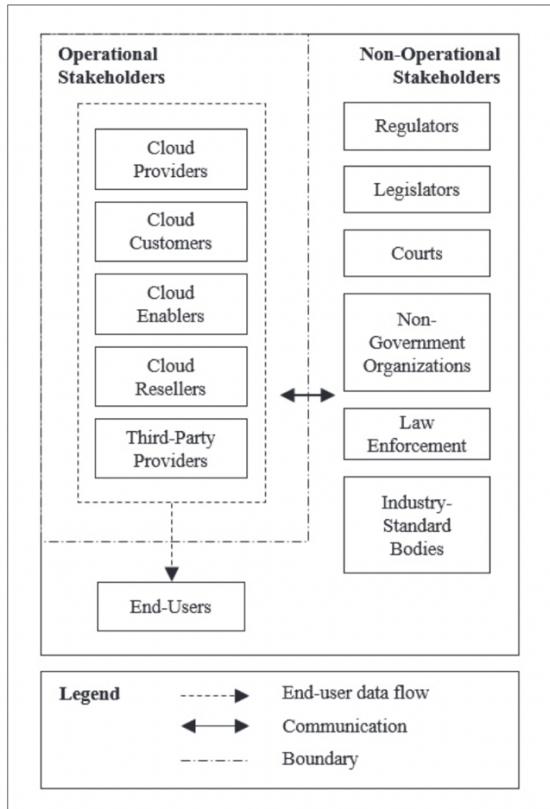


Figure 7. Stakeholders in the cloud computing value chain (Kolevski et. al, 2020).

From this view, one can easily identify how the traditional view of cloud computing, which does not include users as stakeholders, obfuscates the harms individuals experience during data breaches. The Cloud Computing Value Chain crisply illuminates “that cloud services encapsulate more than just technology services” (Kolevski D. , Michael, Abbas, & Freeman, 2020), and from this perspective, a more responsible socio-technical approach to regulation and remediation can be developed.

In the same spirit as Kolevski et. al., the purpose of this research is to uniquely illuminate the incentive conflicts and gaps in federal level cybersecurity, as well as their consequences, and the taxonomy has been designed accordingly. There will certainly be other goals for which this taxonomy is valuable, and those for which it will be inappropriate, as no taxonomy can be sufficient for all conceivable purposes (Harcup & O'Neill, 2016).

### **Research Methods**

There are several examples of academic research pursuits that primarily collect data through news media sources (The Ohio State University, 2018), (Kolevski D. , Michael, Abbas, & Freeman, Forthcoming). Many recent examples primarily use online news sources, as they are widely accessible, and easily translated if the primary language differs from that of the researcher. This approach allows for aggregation and scholarly comparison with traditional academic sources, such as journal articles, books, and book chapters (e.g. (Cross, Parker, & Sansom, 2019), (Kolevski D. , Michael, Abbas, & Freeman, 2021) (Kolevski D. , Michael, Abbas, & Freeman, 2020)). Leveraging news media as a rich data source allows scholars to analyze events in quickly evolving fields, such as cybersecurity, and build them into longitudinal analyses of the socio-technical consequences beyond the individual events themselves (Kolevski D. , Michael, Abbas, & Freeman, Forthcoming).

Government documents, such as U.S. White House strategy documents, U.S. Congressional records, emergency directives from the Department of Homeland Security, and Federal Trade Commission consent decrees, have also been used to supplement

traditional academic sources. By referencing these original documents, one can interpret them immediately within the broader context, rather than waiting for academic interpretations.

Information security is no longer an issue that is confined neatly to the technology space (Bunker, 2012). As security problems impact nearly all individuals (regardless of their technical expertise) as well as larger groups such as communities, societies, and nations, public interest is increasing. In a personal conversation via e-mail, Joseph Marks of the Washington Post stated, *“Public interest in cybersecurity issues has definitely surged in recent years... Watching people wait in line for gas telegraphs to people that cyberattacks are like wars or economic shocks -- things that can happen suddenly, that they have no control over and that can profoundly affect their lives and finances”* (Marks, 2022). Given the pervasiveness of security coverage in the news media, a substantial amount of data for this analysis can be found through these sources. In addition to the volume of data, information on these incidents through news media sources will be more timely than traditional academic journal articles. Data breach incidents are generally not represented in the academic literature until more than 5 years after they are reported in the media (Kolevski D. , Michael, Abbas, & Freeman, 2021). One purpose of news media outlets is to bring urgent issues to the public in a timely manner, raising awareness and increasing the public’s connection with the topic and events (Meijer & Bijleveld, 2016), (Monahan, 2010). These are important qualities in socio-technical research generally, but since creating clear and valuable policy recommendations is part of this dissertation, it is also critical that the data used herein are broadly accessible and relevant. Leveraging the

news media as a source of knowledge allows bridges to be built between academic researchers and practitioners, specifically resulting in more actionable outcomes of this dissertation.

The potential bias of a news outlet, and the implications of that bias on research, are concerns that must be addressed when building a data set from mass media publications. Several academic studies indicate that political bias in media outlets is fairly common (Spinde, et al., 2021) (Groseclose & Milyo, 2005), however, it has also been shown this bias may not translate directly to the predisposition of a reader (Wolton, 2019). Certain types of publications, such as blog posts, editorials and op-ed pieces, and letters to the editor may provide insight into social perspectives on a cybersecurity event but should not be used as primary sources when building a data corpus. Note that the term “bias” here refers to a news organization’s position on a political spectrum and assumes that the bias is manifest in how truthful information is framed. Disinformation, or “fake news”, is considered as separate from political bias. The veracity of a source’s information on an event is assessed through corroboration with other sources, as well as the outlet’s reputation for journalistic integrity, particularly in the technology space.

The cybersecurity incidents used to create the classification can be adequately researched using news media – however, the sociotechnical frameworks and theories that provide perspective on these events are represented by a more traditional academic corpus (i.e., journal articles, books, and chapters). Building scholarly research using a variety of sources opens the doors for academics to apply their expertise to emerging issues, and to connect research and methods of knowledge production to an application

space (Piccoli & Wagner, 2003). New insights into these theoretical foundations can then be achieved by applying them to recent events, and vice-versa –current events can be more thoroughly understood when they are situated within an established socio-technical paradigm.

As mentioned previously, the number of cybersecurity incidents is large and is growing rapidly. This provides a challenge to anyone wishing to study this space – namely, how can a set of objects be defined to study that sufficiently spans the space, yet still be achievable within a reasonable time period? One must necessarily select incidents to investigate deeply, which means one must also choose which incidents (both in the past and future) will be considered “out of scope”. This decision must be done with great care and intention, lest the research findings be jeopardized by mischaracterization. The selection criteria for key cyber incidents for this research are as follows:

**1. The incident must be reasonably recent – within the past 10 years.**

This criterion bounds the population of study in a reasonable and actionable way. The cybersecurity threat surface of today is significantly different than it has been in prior decades (which is one of the fundamental motivators of this research agenda), and one may derive more actionable knowledge by building models from temporally relevant events. The broad historical timeline is important to understand the modern context, but in constructing a socio-technical model it is less insightful to consider the incentive forces around the spread of the Morris

Worm in 1988 (FBI, 2018), for example, than it is the Equifax data breach of 2017 (Fruhlinger, 2020).

To be clear, the historical significance of the Morris Worm incident should be acknowledged. Notably, it led to the first felony conviction under section 1030(a)(5)(A) of the CFAA (FBI, 2021). The young Robert Morris attempted to write a virus just to see if he could – and he was successful! – but he underestimated the virality of his creation, which brought networks across the country to a halt. The Morris Worm became the first Distributed Denial of Service (DDoS) attack, leading to the decision of the U.S Court of Appeals, Second Circuit (United States of America, Appellee, v. Robert Tappan Morris, Defendant-Appellant, 1991). Morris escaped jail time, sentenced instead to probation and community service, and eventually became a computer science professor at MIT. Robert Morris’ story provides a striking comparison point to modern era CFAA incidents, such as the heavy-handed prosecution of Aaron Swartz or Marcus Hutchins (the young British hacker who stopped the internationally rampant WannaCry virus and then was arrested by the FBI a few months later at DEF CON (Greenberg A. , The Confessions of Marcus Hutchins, the Hacker Who Saved the Internet, 2020)). But it doesn’t make sense to include Morris specifically in the *taxonomy* because both the technology and the policy have evolved such that his story would be unlikely to unfold in the same manner today. It would therefore be problematic to incorporate it into a model alongside modern events. Therefore, a temporal boundary is set on which objects are used to



construct the model (see Table N), but the longitudinal context in which this model is placed is not limited.

**2. There must be enough publicly accessible data around the event to be insightful.**

One problem with cybersecurity incidents is that, even after their initial disclosure, their scope and impact tend to grow. Very recent incidents, therefore, may not have enough reliable detail to include in this research. The Yahoo! breach of 2017 was initially disclosed at 1 billion users (Vindu, 2016) and grew to a total of 3 billion affected users over ten months - for the record, 3 billion was the total number of Yahoo! accounts in existence (McMillan & Knutson, Yahoo Triples Estimate of Breached Accounts to 3 Billion , 2017). In the days after the 2016 Arizona voter registration database hack, the initial belief was that the intruders gained access through a specific method, called a SQL injection attack, that had also been used in the Illinois voter registration database breach. That belief was reported by many in the media – myself included (ASU News, 2016). As the investigation proceeded, this method of intrusion was eventually ruled out (and we corrected the ASU News piece cited above). Giving an incident some time to mature allows one to be more confident that the relevant statistics and facts of the event will remain relatively static.

However, some breaches have very little public data, regardless of the time elapsed since their discovery. Facebook, for example, has been diligent about

keeping many of their intrusions shrouded from public scrutiny, obeying the letter of data breach notification laws but staying quiet about the ways and means of the breach. Generally, there will be some missing information around a cybersecurity incident, but to be a valuable object of study, there must be a reasonably detailed description of the event and its impact, beyond the number of records. Focusing solely on quantitative data will be misleading – much like airport advertisements for cyber response companies promising “100% protection” against “billions of attacks per day”, the numbers are disingenuous without any surrounding context.

To ensure relevance, this research focuses on incidents which have enough detail to be insightful. To fit within the relevant date range, data collection must also span a wide variety of sources, as most peer reviewed papers on data breaches are published at least five years after the event (Kolevski D. , Michael, Abbas, & Freeman, 2021). Court records, interviews, U.S. federal documents (including White House and Congressional records), and news articles from reputable outlets supplement the traditional academic journal articles for this project.

### **3. There must be some consensus that the incident is important**

Creating an adequately representative corpus of events requires one to intentionally consider the broader landscape, not only the incidents with which one is immediately familiar. Using a singular, personal vantage point alone to select events for a taxonomy would not result in a complete set, so there must be an unambiguous way to include other viewpoints.

Crowdsourcing is one method of incorporating a diverse set of views, and one way to find a crowd is via social media. Recently, Joe Uchill (a cybersecurity reporter in Washington, D.C.) asked via Twitter: “What are the “wake-up” moments for cybersecurity policy since 2015?” (Uchill, 2021). Joe’s larger point was that there hasn’t truly been a “wake up moment” - that U.S. government entities (and many companies) have been hitting the snooze button for several years - but the thread serves as an interesting list of significant events. I used this thread as an initial grouping of notable incidents, weighting more heavily the responses by practitioners who have been working in the field for a considerable amount of time, and who have garnered the respect of the security community. From this weighted list, the duplicates were removed, and combined different variants of the same source incident (for example, the NotPetya malware impacted multiple sectors across the globe; it makes sense to consider the malware as the root problem instead of the separate victims if the goal is to uncover fundamental mechanics of incentives in cybersecurity). This approach leverages my expertise as a scientist and cybersecurity expert, as well as the crowdsourced wisdom of cybersecurity professionals from technology, policy, and academia.

#### **4. The incident has policy implications at the U.S. federal level**

An important part of solving a problem in physics is defining the scope of the analysis, and that remains true for this research as well. Without creating proper bounds, the pursuit of a solution will be either stunted or endless. This project

centers on the state of cybersecurity in the U.S. at the federal level, but it is important to note that cybersecurity problems often do not respect jurisdictional boundaries, and the perception that they do is problematic. The fractured (and fractious) nature of global cybersecurity policy may benefit from the approach established herein, but that approach must first be designed, developed, and validated before advancing to an increased level of complexity.

Cybersecurity issues may not respect geographic boundaries, but a global-scale policy assessment is not within the scope of the current project. For clarity, and establishment of the theoretical approach, this work centers on how incentive conflicts can be viewed from and addressed by U.S. federal policy, both codified law and regulatory action. It may be possible to apply the techniques developed during this research to policy around computer access and security in other countries, but these perspectives have great variance at a global scale. Replicating this research outside the U.S. would require an independent taxonomy of relevant data and interpretation of those data through another nation's lens.

While the Computer Fraud and Abuse Act is the primary U.S. federal-level law addressing unauthorized computer access, there are cybersecurity issues that are not within the purview of the CFAA. Privacy in the online age is one key example (Michael & Clarke, 2013). A 2019 survey by Pew Research of randomly selected U.S. adults found that 81% of those surveyed believe that the risks of personal data collection outweigh the benefits (Auxier, et al., 2019). Certain facets of privacy are covered by different constitutional amendments (Warren & Brandeis,

1890), and broader efforts towards privacy have been made at the state level in the last several years (Stauss, 2021). However, no single federal-level law addresses an individual's right to privacy (Nissim & Wood, 2018). Given that the policy landscape is herein defined as both the federal legal code and the actions of federal organizations, the capabilities and responses of organizations such as the Federal Trade Commission (FTC), the Cybersecurity and Infrastructure Security Agency (CISA) within the Department of Homeland Security, and to some extent the Securities and Exchange Commission (SEC), will be considered, in addition to documents from the U.S. executive, legislative, and judicial branches.

### **The data set**

To generate the taxonomy for this research, 15 separate events (listed in Table 1) were selected that fit the criteria above. Each event is sufficiently recent, and each has an adequate amount of detailed information available about the means of intrusion and the results. Each event was noted for its importance by multiple respected sources, and each has federal-level policy implications in the U.S.. Events that were excluded from this corpus include those that fall outside of the chronological boundaries – for example, the Morris worm was rejected for being too early, and the log4j vulnerability for being too recent. Also excluded are those with primary impacts outside the U.S. (e.g., the Stuxnet attack on Iranian nuclear reactors), or events that primarily served as a demonstration (car hacking demonstrations at conferences such as DEF CON are often cited as important events, but at this time the real-world impacts remain mostly theoretical). Even as this data set was being built, information on new cybersecurity incidents was being made

available on almost a weekly basis – several would be fascinating additions, but again, the analysis must be chronologically bounded for the subsequent analysis to be feasible. The incidents are listed in Table 1, and details on each of these incidents is included in Appendix A.

### **Categorizing Events**

A list of events is not in itself a classification scheme. To create a true taxonomy, these events must be analyzed, and from them key features discerned to create new categories. Once primary categories have been determined, these can be used to describe events. There may be several valid ways to categorize the events in this data set, but to specifically illuminate the incentive structures that led to and in many cases aggravated these occurrences, it is most sensible to focus on *who or what bears the primary impact* in each event.

Often, media coverage focuses on quantitative aspects of breaches (such as number of records or the financial impact of the event) but these categories fail to encapsulate important aspects of these events for policymaking. Qualitative methods can be used to provide unique insights which are inaccessible through the more traditional quantitative means. As an example, consider Figure 8, a plot of four different events and the number of records breached in each one.

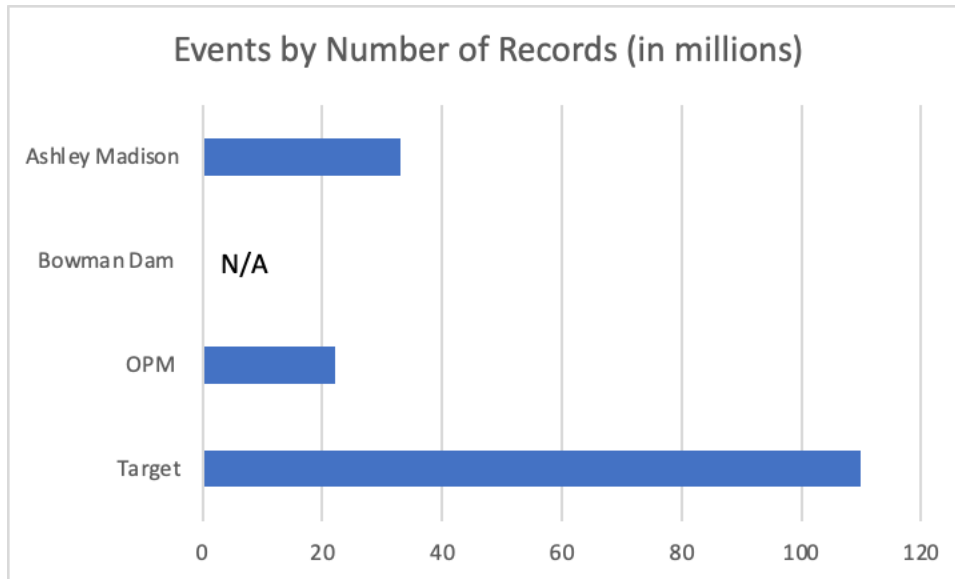


Figure 8. The number of breached records (in millions) for four select events in the data set

These events were selected specifically to show the limitations of traditional quantitative measures, such as the number of records of an incident. By looking at Figure 8, one might easily conclude that the Target breach is the most impactful of the four, since it includes the most records (110 million). However, the figure does not – and cannot – indicate the *importance* of these records. The Target data breach consisted primarily of credit card numbers, which can be cancelled and reissued with a minor amount of effort. The OPM breach, however, was comprised of security clearance forms that included very sensitive personal data – far more valuable to a foreign adversary than credit card information. While OPM’s breached records are fewer, the content of these records is much more sensitive than the data from the Target breach.

The Ashley Madison breach also looks, from Figure 8, to be less impactful than the Target breach. This breach also included data which, on the surface, may look similar

to those from the Target incident – name, credit card number, email address. However, considering that Ashley Madison’s business model exists solely to facilitate infidelity, these names and credit card numbers represent much greater potential for personal harm. Having one’s information leaked from a neutral vendor like Target is less damaging than if it came from a company that exists solely to facilitate infidelity.

There are also cybersecurity incidents for which the “number of records” metric doesn’t make sense. A reasonable taxonomy must be able to describe these events as well. No records were lost in the Bowman Dam hack, but the dam was one of the first pieces of infrastructure that was successfully and demonstrably hacked by a foreign adversary. No names, credit card numbers, or personal information were available to be stolen, but Iranian hackers were able to obtain real-time operational status of the dam, and would have been able to control it if the main controller had not serendipitously been taken offline for maintenance (Cohen, *Throwback Attack: How the modest Bowman Avenue Dam became the target of Iranian hackers*, 2021). The number of records involved in a cybersecurity event is a valuable metric for incident response, but is a myopic approach to take when evaluating incentives and harms, and designing policy improvements across the larger-scale system.

After analyzing the full corpus of 15 events and experimenting with a variety of schemas, three categories of primary impact emerged: personal, infrastructure, and political, defined below.

#### Category 1: Personal Impacts



Events in the ‘**personal impacts**’ category have effects that are felt primarily by individuals. The key events in this category have impacted millions – or in the case of the Yahoo! breach, billions – of people. The generalization of these events is as follows: a company or government organization is breached, often due to insufficient security practices, and the personal data of large groups of people is siphoned off by malicious actors. Despite the large number of individuals involved, the onus of responsibility remains with the affected individuals, who had no insight into or responsibility over the events that caused their exposure. These individuals are asked to monitor their credit, enact credit freezes, change their passwords, and remain vigilant of misuse of their own data. Occasionally they are provided with a year of credit monitoring service, but these services can only alert someone to the misuse of their personal data – it is not possible to provide any active prevention of misuse once data has been exposed. (The identity theft company LifeLock was fined \$12M by the Federal Trade Commission in 2010 in part due to “misrepresenting that its services offer absolute prevention against identity theft because there is unfortunately no foolproof way to avoid ID theft,” to quote Illinois Attorney General Lisa Madigan in the FTC press release (Federal Trade Commission , 2010).)

It has been suggested as recently as May 2021 (Isadore, 2021) that the free market provides an incentive for companies to protect individuals’ data, but this notion has proven false. In general, the stock prices of the companies in this data set experienced a momentary dip in stock price, before recovering and returning to an upwards trajectory since their breaches. Research has shown that negative impacts of data breaches on stock prices generally subside after three days, after which the stock price returns to its prior

trajectory (Rosati, Deeney, Cummins, van der Werff, & Lynn, 2019). The question of market effects also doesn't apply to groups which are not listed on the stock exchange, such as government functions (U.S. Office of Personnel Management, Democratic and Republican National Committees), and privately-held companies (Colonial Pipeline).

## Category 2: Infrastructure Impacts

Events in the “**infrastructure impacts**” category reflect incidents that primarily impact larger-scale systems, such as companies, infrastructure, hospital networks, or municipal services. These services may be in the public or private domain, and may include a combination of entities (e.g. the U.S. power grid). While individuals are certainly impacted by many of these – the effect of hospital ransomware or hijacked implanted devices can be a literal life-or-death issue – the response to these events is taken on at an organizational level. For example, cities, hospitals, and educational systems afflicted with ransomware must decide whether to pay the ransom (thereby feeding the ransomware economy) or go through an arduous restoration process on their own. Even when organizations have created system-wide backups, it can be more disruptive to restore a network from scratch than it is to pay hijackers to decrypt the existing network. Ransomware has become such a widespread issue that the White House created a global Counter-Ransomware Initiative, with initial participation from 30 countries<sup>16</sup>, in October 2021 (White House, 2021). However, in a follow-up interview

---

<sup>16</sup> Countries participating in the first meeting were, per the White House press release: Australia, Brazil, Bulgaria, Canada, Czech Republic, Dominican Republic, Estonia, the EU, France, Germany, India, Ireland, Israel, Italy, Japan, Kenya, Lithuania, Mexico, the Netherlands, New Zealand, Nigeria, Poland, the Republic of Korea, Romania, Singapore, South Africa, Sweden, Switzerland, Ukraine, the UAE, and the UK.

with PBS, Anne Neuberger, U.S. Deputy National Security Adviser for Cyber and Emerging Technology, estimated that the U.S. federal government only has insight into approximately a quarter of ransomware attacks<sup>17</sup>, which is problematic when working towards a coherent federal-level solution (Neuberger A. , 2021). In the near future, ransomware generally may elevate to the next category – “political impacts” – but at the moment, the ramifications reside primarily at the organizational rather than the federal level.

### Category 3: Political Impacts

Events in the “**political impacts**” category are these incidents or combinations of incidents that induce a response at the governmental level – in this case, the U.S. federal government - particularly due to national security implications. The intent behind these attacks is often not for individual gain, but for foreign intelligence gathering, manipulation of democratic processes (such as elections) or promotion of malicious political ideals. Many of the victims are government organizations themselves (OPM, DNC/RNC), but some are companies (SolarWinds, Sony). Sometimes events will start in the infrastructure category but rise to the political – for example, when the entertainment company Sony was breached by North Korean hackers with a political agenda, Sony’s initial response was to bend to the attackers’ demands. The U.S. State Department and White House intervened, with the justification that acquiescing would incur further

---

<sup>17</sup> In 2020, the FBI seized the foreign servers from which a particular ransomware attack (NetWalker) was launched. This exercise allowed them to contrast the number of attacks reported with the total number of attacks from this group (Grieg, 2022). The 25% estimate from Neuberger’s PBS interview is an application of this statistic across all ransomware attacks. Given that this statistic is derived from a single incident, it may not be highly accurate, but it is currently the best possible estimate.

politically motivated attacks and ultimately restrict free speech (Zetter, 2014). The key differentiator here is whether or not the U.S. federal government is compelled to respond to the event, not whether or not the attacked entity is itself a government organization. The 2017 national security strategy of the U.S. specifically addresses cybersecurity partnerships with the private sector, claiming that the global nature of communications requires new kinds of partnerships to “remediate known bad activities” (United States Office of the President, 2017). Government and industry are encouraged to collaborate “in a way that respects free markets, private competition, and the limited but important role of government in enforcing the rule of law.” The definition of this category therefore reflects how responsibility and authority are deeply intertwined between public and private sectors, and the complexity that results.

*Table 1. Events used to create the taxonomy, the year in which they occurred, and their categorization*

<b>Event</b>	<b>Year</b>	<b>Category</b>
Target data breach	2013	Personal
Bowman Dam hack	2013	Infrastructure
Sony Pictures Entertainment	2014	Political
Ashley Madison	2015	Personal
Office of Personnel Management (U.S. government)	2015	Political
Anthem	2015	Infrastructure
Election system hacks (Democratic National Committee, Republican National Committee, voter registration databases)	2016	Political
NotPetya malware	2016	Political
Yahoo data breach	2017	Personal
Equifax data breach	2017	Personal
Facebook data breach	2019	Personal
SolarWinds	2020	Political

Colonial Pipeline	2021	Infrastructure
Ransomware attacks in hospitals, municipal services, education	2017-2022	Infrastructure
Microsoft PowerApps	2021	Infrastructure

Appendix A describes each incident in detail, with key statistics and facts up front, a narrative description of the event, and reasoning for the event’s primary categorization.

There are those who make the case that events in the personal category should have an industry or government response (Jain & Ropple, 2018) (Eggers, 2016), and some that claim personal and industry cases should never be addressed through federal intervention (Thierer & Szoka, 2009). These debates are rarely conclusive. But the taxonomy developed here is based on what *is*, not what *should be*, because that is the lens through which one can discover the source of problems in existing systems. Once these new insights have been discovered, then an improved future state can be intentionally designed.

It is relevant to revisit the cloud computing value chain from Kolevski et. al. (2020). Cloud computing services are extensively utilized in both government and industry, due to storage capacity and professional support functions that greatly exceed what most of these groups could hope to achieve locally. The joke “the cloud is just someone else’s computer” is inaccurate for many reasons (Branscome, 2016), but what many security professionals allude to with this one-liner is that security concerns do not diminish or disappear when processing moves “to the cloud” (King & Raja, 2012). Comparing the categories of this taxonomy with the value chain description from Kolevski et. al., one can see that impacts don’t fit neatly within operational/non-

operational stakeholder boxes. Events with primary personal impacts may stem from intrusions on the operational side, and often include other non-operational stakeholders such as courts (e.g. Ashley Madison) and some element of law enforcement (e.g. ransomware, OPM). The two schema, particularly in concert with one another, show the power of abstracting a system to key elements when searching for clarity in a complex system.

### **Assessing Levels of Impact**

The taxonomy consists of a set of categories, identified from examining a diverse data set for commonalities and differences, with an eye towards who or what was primarily harmed by each. These categories – personal, infrastructure, and political – can now act as a “basis set”; meaning that they can define a space wherein cybersecurity issues can be described in a new way. The impacts of events in these three categories can be used instead of traditional ineffective metrics such as the number of records.

One hurdle remains, however. Three primary types of harm have been identified, but not each incident of a particular type is of an identical severity. The Target and Ashley Madison breaches of Figure 8 - both categorized as being of primarily personal harm – provide a useful comparison. As discussed, the number of records from the Target breach exceeds that of the Ashley Madison breach by a factor of 3, but the Ashley Madison data are far more sensitive, and the loss was arguably more harmful to those affected. Through analyzing a diverse set of events, three primary categories have been determined, but there must also be a way to indicate the magnitude of severity within each category. There is no reasonable quantitative way to compare the loss of one’s

credit card with the exposure of one’s infidelity, so a qualitative assessment is a more appropriate approach. I have chosen a scale from -3 to +3, where -3 is considerable negative impact, 0 is neutral, and +3 is considerable positive impact. (It is rare, but not impossible that events may contain positive impacts in any category.) A general description of impact rankings is shown in Table 2:

Table 2. General schema of impact rankings

General Impact Rankings	
3	The event resulted in unique benefits or long-term advantage to the impacted entity or related group (i.e., not the attacker)
2	The event uniquely contributed to or accelerated positive outcomes
1	Positive outcomes outweighed negatives, but both were present
0	No impact
-1	Negative outcomes outweighed positives, but both were present
-2	The event uniquely contributed to or accelerated negative outcomes
-3	The event resulted in unique harms or long-term damage

The numeric score is a comparative assessment of impact in each area. The purpose of assigning a number is two-fold – first, a simple numeric scale allows for accessible and intuitive comparisons of events of different types. Second, adopting a consistent numeric approach allows for the creation of larger models wherein the impacts of multiple events can be combined. The choice of 3 categories for this analysis was intentionally aimed at designing more intuitive models. Humans generally interpret their surroundings in 3 dimensions – a skyscraper is taller than a single home, a warehouse is wider and deeper than a house – even without quantitative measurements of these building dimensions, one can visualize them and make quick relative assessments of their volume and therefore their purpose. Defining three categories – analogous to the dimensions of our spatial existence – and assigning relative values to events in each

category unlocks the potential for us to leverage the way in which humans understand and act upon their world.

A few examples of assigning scores and comparing events in each of the categories illustrates the utility and power of this approach.

### Personal Impact Scores

Different groups are harmed in different ways, so the general rankings must be translated to apply to each specific group. Personal impacts are considered from an individual perspective. What burden do they bear, if any, as a result of the event?

Examples are given in Table 3.

*Table 3. Personal impact scores*

Personal			
0	No impact		
-1	Passive measures needed (e.g. credit monitoring) or non-critical data exposure	1	Positive outcomes outweighed negatives
-2	Active measures needed or moderate data exposure	2	Positive contributions to individual security
-3	Highly personal data exposure that leads to or provides unique opportunities for personal harm/exploitation	3	Considerable and unique positive impacts to individual security



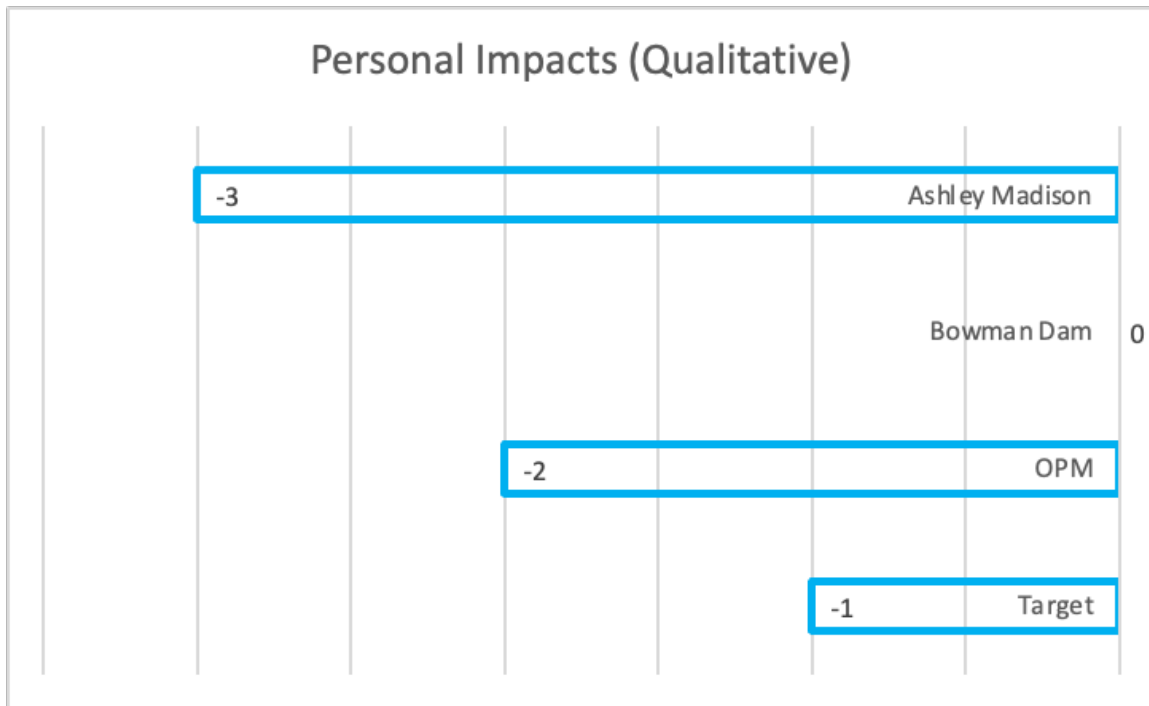


Figure 9. Personal impacts assigned to four selected events

The four events from the prior section remain an insightful comparison. As shown in Figure 9, the Ashley Madison data breach is defined as the most personally harmful of the four events, with a score of -3. The OPM breach has been given a score of -2, due to some of the sensitive information in a security clearance application, which is of concern but is less likely to have the same kinds of direct familial or workplace impacts. The Target breach is ranked a -1, since credit cards are replaceable and credit monitoring is a reasonable option for remediation. Finally, the Bowman Dam hack has a personal impact score of 0, which is a more insightful metric than the ‘N/A’ number of records.

#### Infrastructure Impact Scores

Figure 10 shows an assessment of the infrastructure impacts of the four events, with the general rubric of Table 2 adapted to this category in Table 4.

Table 4. Infrastructure impact scores

Infrastructure		
0	No impact	
-1	Temporary negative impacts, remediation requires minor changes to business processes	1 Positive outcomes outweighed negatives
-2	Remediation requires considerable changes to status quo	2 Positive contributions to operations
-3	Remediation unknown; long-term damage to operations ; losses in the tens of millions of dollars (IP, lost productivity, not FTC fines)	3 Considerable and unique positive impact to operations

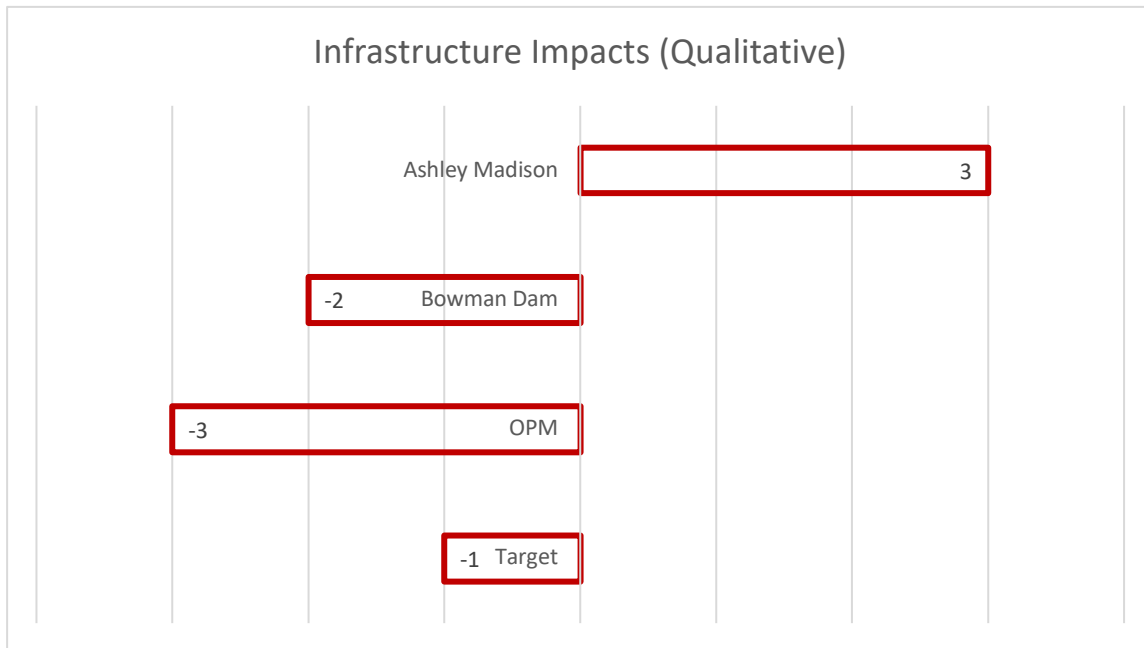


Figure 10. Infrastructure impacts assigned to four selected events.

As mentioned, it is rare for positive impacts to occur in cybersecurity events, but in the case of the Ashley Madison data breach, the media coverage of the breach brought considerable attention to the website and its services, and the number of people signing up increased, from 34 million to over 43 million in the 5 months after the incident (Moscaritolo, 2015). The company profited - quite literally - from the attention stemming from this event, even though substantial harms were incurred by some of their customers.

Infrastructure impacts of the other three are negative, but in different degrees. The Bowman Dam hack is scored a -2 for infrastructure impacts. This event had the potential for considerable harm to the dam and the network of people who depend on its proper function for both water and safety. Even though the dam's primary controller was taken offline shortly before the intrusion, this event was a shocking reminder that anything connected to the internet can be hacked. The OPM data breach is scored a -3 on infrastructure impacts – this event will have lasting negative effects on the U.S.' national security infrastructure, due to the sensitive information that can be derived from these lost data. The 2013 Target hack is scored with a -1 in this category. The harm was relegated to Target and a few of its third-party vendors, and did not spread through the entire retail ecosystem. For the company itself, the damage was also limited. Target experienced temporary bad press, and their stock price dipped briefly, but the effects were short-lived. The company's image recovered quickly, and even after paying \$18.5M in a consolidated class action settlement in 2017, Target's stock price has continued back on its positive trajectory.

### Political Impact Scores

Figure 11 shows an assessment of the political impacts of the four sample events, with the general rubric adapted to this category shown in Table 5. Federal government interventions are often justified on the basis of national security, which can have a broad interpretation. For this analysis, "national security" includes the threat of espionage and intelligence gathering by foreign adversaries, the impact of an event on the U.S. position on the global stage, and generally the degree to which federal level intervention is required to remediate or address a particular incident.

Table 5. Political impact scores

Political	
0	No impact
-1	Potential damage to national security when combined with other events
-2	Potential damage to national security without other events
-3	Assumed long-term damage to national security
1	Positive outcomes outweighed negatives
2	Positive contributions to national security or global standing
3	Considerable and unique benefits to national security or global standing




Figure 11. Political impacts assigned to four selected events

Much like the prior 3 categories, a -3 indicates severe negative political outcomes of the incident, 0 is a politically neutral event, and a positive value would indicate that the U.S. government benefited from the incident at the federal level. In this context, the OPM breach and the Bowman Dam attack were both overtly negative. The data siphoned off during the OPM breach describes the U.S. national security apparatus in unprecedented detail. The Bowman Dam attack was the first true realization that U.S. critical

infrastructure was not beyond the reach of hackers, and at that time there was no mitigation strategy in place. Neither the national security infrastructure issues represented by OPM nor the critical infrastructure issues represented by the Bowman Dam could be quickly patched or mitigated, which earns them both the maximum negative score for a single event. None of the four events used in this exposition were scored as positive or neutral. Even if certain events, such as Target or Ashley Madison, don't have obvious impacts at the U.S. federal level, all data breaches include some loss of personal information which can be used by adversaries to complement or supplement information from prior breaches. Given the long history of breaches from a variety of sources, the ways in which data aggregation by an adversary could be used to undermine U.S. national security must be considered during the scoring process.

### **From Taxonomy To Framework**

A new taxonomy has been created from data that are relevant, timely, and acknowledged as significant by the cybersecurity community. This taxonomy consists of three primary categories – individual, infrastructure, and political - which can be used to identify impacts of cybersecurity events at different scales. Numerical values have been mapped to the impacts, ranging from -3 (negative impact) to +3 (positive impact) via a specific rubric. These fundamental levels of impact can now be used to help create new and insightful methods of describing cybersecurity events. This taxonomy has been constructed to ensure that each event from the data set fits within one of the categories – given the diversity of the data set, one can be confident that the taxonomy can sufficiently define the overall space of cybersecurity events. At the next level of granularity, each event in reality has multiple impacts; this next step will be explored in Chapter 4.

In Chapter 4, the taxonomy will be used to build a new framework helpful for pursuing unique insights into cybersecurity. This framework is inspired by a method from classical Physics, which provides deep insights into complex mechanical systems by creating new coordinate systems, called Lagrangian Mechanics. The next chapter will provide a high-level description of the Lagrangian approach, including a specific example of a complex mechanical application of the method. A cyber-specific analogue of the Lagrangian framework is then constructed, concluding with the validation and verification of this new approach.

## CHAPTER 4

### A NEW FRAMEWORK

#### Revisiting the Motivation

At this point, it's worth revisiting the fundamental motivation for this work: namely, cybersecurity problems (such as data breaches, ransomware, and intrusions into critical infrastructure) have been around for a long time, and despite high levels of attention and funding, many of these problems not only persist, but are getting worse. Phrased in the most general way: *commonly used approaches do not provide enough insight into fundamental system behavior to help solve certain problems.*

Given the entrenchment of the status quo, how can the system be re-examined to provide the necessary insights for progress? Defining what is meant by “the system” was a critical first step. Scoping the definition properly prevents the analysis from being unnecessarily large (and therefore too vague to be applicable) or from being overly narrow (thus brittle and unactionable). For this analysis, cybersecurity at the US federal government level is a well-scoped space – large enough to be robust, but narrow enough to be functional. Once that scope was defined, a reasonable number of examples were chosen that represent the system (or “span the space”, in the parlance of Physics). Again, setting relevant boundaries and key features for these examples is critical, ensuring that the data set properly represents the problem space and the identified concerns. Finally, those data were organized into a schema or “taxonomy” that facilitate the work described

in this chapter: creating a novel framework that facilitates a clearer understanding of the fundamental system's behavior.

Generally, this is how science progresses (Kuhn, 2012) – questions are raised that can't be answered by existing means, so new methods are developed to ask these specific questions. These new methods are tested to establish whether they are correct, and if so, to identify under which circumstances they should be applied. This approach is employed across scientific disciplines, but some of the most concrete examples come from the field of physics – specifically the theory of classical mechanics, which describes macroscopic objects and their motions. Outside the field of physics, these examples can serve as valuable analogies, because they describe common objects that people can intuitively understand. Like springs, weights, and pendula. These literal objects can be related carefully to complex systems constructed of less tangible elements.

A physical and accessible example of a complex mechanical system (specifically, the double pendulum) will be used to show how a conventional classical mechanics process doesn't answer certain kinds of questions. A new process, Lagrangian mechanics, will be examined to show how it was developed to specifically ask and address these questions. Once a foundational understanding of this process has been established, it can be abstracted (where appropriate) to the socio-technical case of cybersecurity issues.

### **Complex Mechanical Systems**

In the classroom, physics students usually start learning from specific quantitative examples and through practice, and from here build to general abstract concepts that can



be broadly applied. This approach is often the reverse of physics in the research laboratory (or the notebook, for the theorists), where physicists make qualitative observations before they approach specific quantitative problems. More than 100 years prior to the development of Lagrangian mechanics, Isaac Newton pondered a falling apple in his garden, observing an object and an action, which inspired him to create a mathematical relationship between the two (Conduitt, 1727). Without a solid qualitative understanding of objects and their behavior, one risks the improper application of mathematical tools and techniques, and mischaracterization of the objects of study. While the field of physics has earned its reputation for being mathematically intensive, a period of careful and informed qualitative observation precedes the development and application of specific numerical approaches.

Lagrangian mechanics (Fowles & Cassiday, 1999)<sup>18</sup> is a process by which one can describe the motion of a complex mechanical object by creating a new framework<sup>19</sup>, based on the forces which have the greatest impacts on the object's behavior, and discarding those with little to no impact. It is a valuable technique for certain complex systems where a conventional framework (Newtonian mechanics) results in equations of motion that are intractable without computational means. It's valuable to walk through a specific example of Lagrangian mechanics, to establish terminology and justify

---

<sup>18</sup> See also: [https://www.reddit.com/r/physicsmemes/comments/d1cig6/lagrange\\_vs\\_newton/](https://www.reddit.com/r/physicsmemes/comments/d1cig6/lagrange_vs_newton/)

<sup>19</sup> I choose the word "framework" here intentionally; per the Cambridge Dictionary, the term means "*a supporting structure around which something can be built; a system of rules, ideas, or beliefs that is used to plan or decide something*". This definition encapsulates the description of the knowledge product as well as the intent for how it will be used (for planning or deciding new policy directions in U.S. federal cybersecurity).

development of the socio-technical analogue. Consider the specific system shown in Figure 12 – two pendula, one of which is attached to the other by a spring<sup>20</sup>.




Figure 12. A double pendulum, one pendulum is attached to the base by a spring (Jankowski, 2011)

The motions of this double pendulum are challenging to describe precisely. Gravity is pulling on both masses ( $m_1$  and  $m_2$ ) as well as the spring, which is oscillating – and the overall resulting motion depends on where the masses and the spring started out relative to one another. Writing the equations of motion in a conventional Cartesian coordinate system (using  $x$  and  $y$  positions as coordinates, labeled in the figure) results in something that is hard to understand intuitively. The  $x$  and  $y$  motions depend on one another, and that coupling is also time dependent. From the Cartesian perspective, it is cumbersome to calculate the system’s motion, and even harder to build intuition of how to change it. What happens if the spring is pulled in a certain direction or another? What

---

<sup>20</sup> The double pendulum is a very interesting problem in physics, and there is considerable research into the numerical methods used to fully describe it. I am using the double pendulum here as an example of a complex mechanical system, giving short shrift to the actual physics involved, to avoid a lengthy mathematical tangent that does not support the aims of this dissertation.

if one wishes to affect the system to achieve a specific kind of behavior? Neither question can be answered (or even estimated) using  $x$  and  $y$  coordinates without handing the task over to a computer. Using this framework to develop the equations of motion provides no intuition for how this system behaves now or in the future.

However, one can get a much clearer understanding of the system by changing perspective. Instead of using the Cartesian framework, in which the equations of motion don't lend themselves easily to understanding, Lagrangian dynamics gives us a process by which one can create a new set of coordinates to describe the object more elegantly. In this specific case, the new coordinates are the angles at which the pendula are offset from their pivot points, labeled in the figure as  $\psi_1$  and  $\psi_2$ . The new framework is called a generalized coordinate system. It is specific to this object, and it enables a more intuitive exploration of the system. The equations of motion are more streamlined from a mathematical perspective, and they also allow for greater intuition about the system and how its behavior can be modified.

The numerical results of these two coordinate systems is the same, just as the results of a Newtonian and a Lagrangian calculation will also be identical. This is expected, as they are different ways to describe the same object! But the effort involved and the insight available from the equations, prior to calculation, is considerably different. Equations are often thought of as a tool for calculating a specific number, but they also have the power to provide a quick qualitative sense of how the system behaves before plugging any numbers in. That's what "intuition" means in this case: a conceptual foundation that informs further exploration and allows for intellectual progress (Wilder,

1967). Some scholars of mathematics pedagogy claim that building intuition is essential for active understanding – without it, a student is limited to rote memorization of facts (Fischbein, 1982). Similarly, in physics, it's easier to build intuition about a system from a simple equation than one with interdependent variables. (Human brains don't handle that level of complexity very well.) The Lagrangian process that leads to this intuitive development requires some initial understanding of the object before diving into calculation, and it can be tricky to think about it in this unconventional way. But once physics students learn how to do so in their advanced classical mechanics courses, they achieve a deeper and more intuitive understanding of complex mechanical systems.

The equations of motion look rather intimidating to the uninitiated, even after applying the Lagrangian formalism, but an accessible example of the power of intuition is warranted. Fortunately, one can be found by comparing two pre-existing coordinate systems: *Cartesian* ( $x$  and  $y$ ) and *polar* (radius and angle). Imagine an object that is moving in a straight line – as shown in Figure 13. The object is represented by a green dot, and its path by the red line. The simplest way to describe this path is in  $x$  and  $y$  coordinates. The path can be simply described as the number of steps in the  $x$  direction, and number of steps in the  $y$  direction. (Source: <https://www.desmos.com/calculator>).

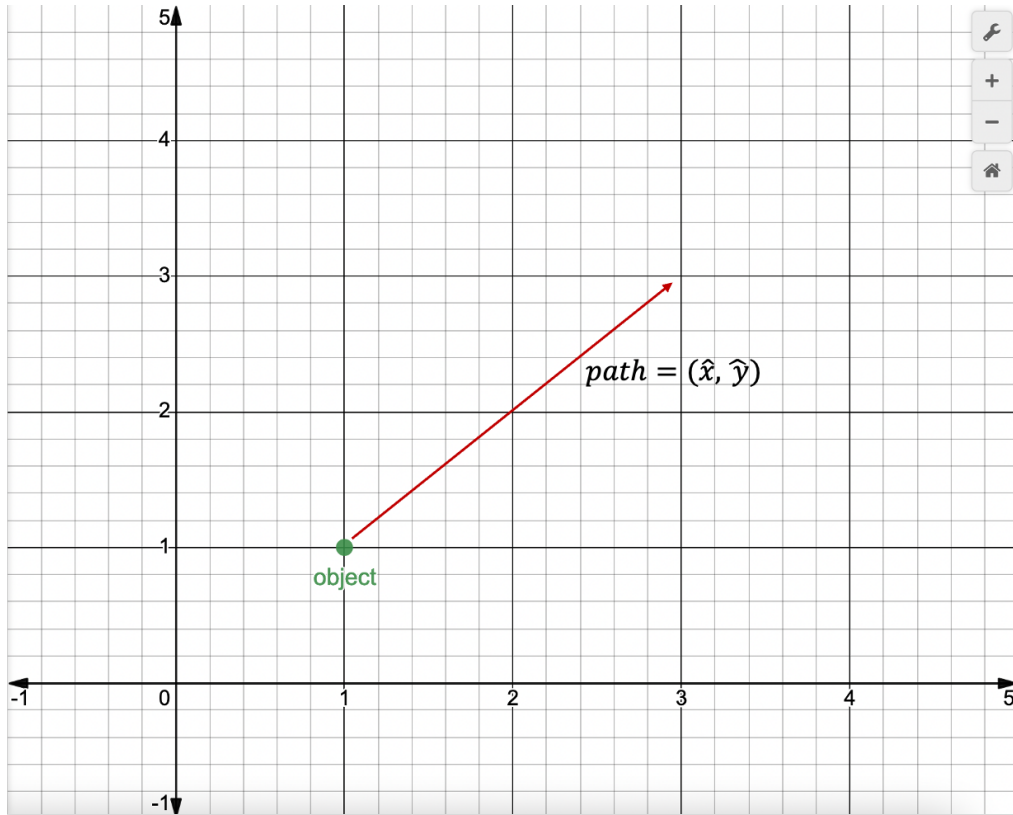


Figure 13. An object (green) traversing a linear path (red), shown on a Cartesian grid.

But if this particle is instead moving in a circle, as shown in Figure 14, a description in  $x$  and  $y$  is more complicated. Describing a circular path in Cartesian coordinates requires square roots and trigonometric functions, and the variables are not independent.

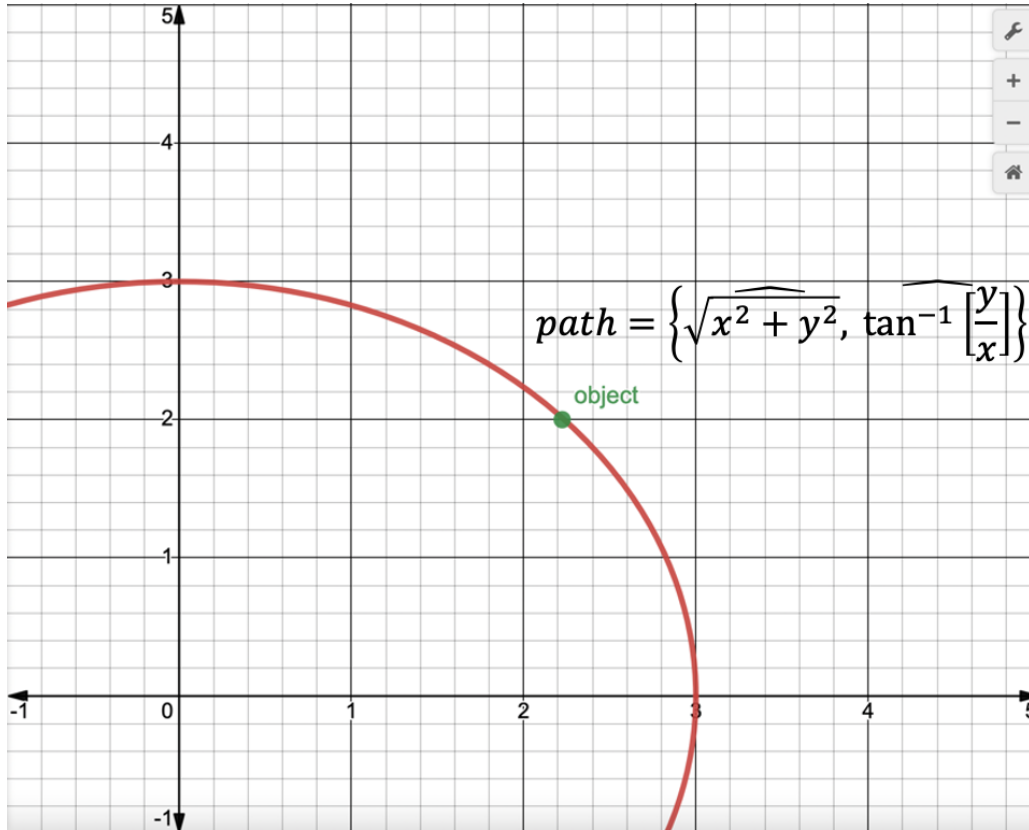


Figure 14. An object (green) traversing a circular path (red), shown on a Cartesian grid.

Instead, using the distance from the center ( $r$ ) and the angle from the starting point ( $\theta$ ) instead of  $x$  and  $y$  provides an equally accurate, yet much more accessible description of this object's behavior. The Cartesian grid of Figures 13 and 14 has been replaced by a polar grid, shown in Figure 15, and the path re-written in polar coordinates. Now the path can be defined simply by the radius of the path (the circles of the grid) and the angle the object has traversed (the spokes). (Source: <https://www.desmos.com/calculator>)

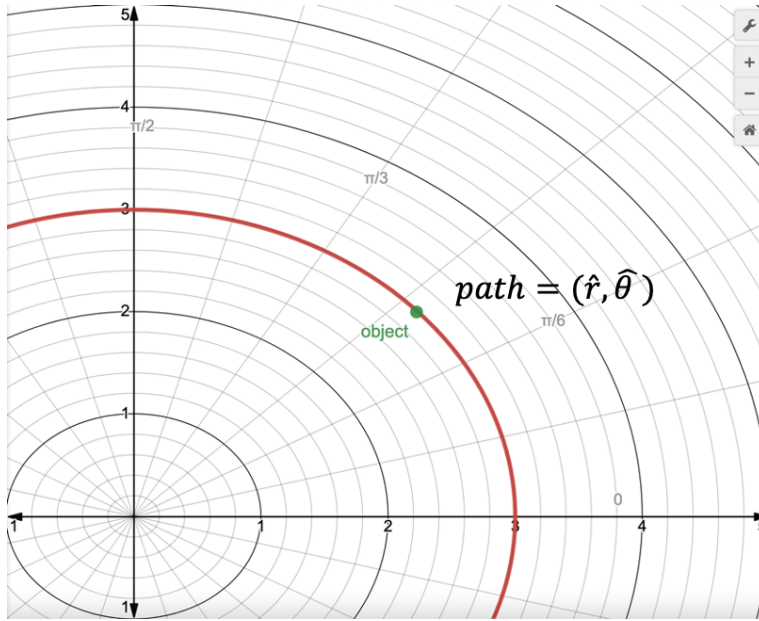


Figure 15. The same object and path as Figure 14, shown on a polar grid.

With this example in mind, we return to the Lagrangian method, which goes beyond a coordinate transform between two known systems to develop a new coordinate system that is specific to the object of study. One key benefit of the Lagrangian approach is the ability to identify small effects in the system and neglect them when appropriate. In a Cartesian coordinate system, one may spend a considerable amount of time calculating a behavior that ends up being physically insignificant. Consider a pendulum made from a heavy bowling ball at the end of a rope. If a person is in the path of this pendulum, the ball's velocity in their direction is of much greater concern than whether or not it is also spinning as it speeds towards them! Certain frames of reference, or ways of looking at the problem, can sometimes obscure the difference between what's insignificant and what is meaningful. Effective frameworks allow one to spend time on the important parts of a problem and discard minor perturbations in an accurate and intentional way.

How did Lagrange (and others) know that his process resulted in a framework that is correct? One can test the process on simple examples (a single pendulum, for example) and verify that this new process leads us to the same answer as the conventional method. One should also check that the results are commensurate with the known laws of physics. The Lagrangian process can provide new insights on complex mechanical systems, but the behavior of these systems are still governed by the fundamental laws of the universe.

The language used in the prior description was chosen intentionally, to facilitate an analogy that can be used for socio-technical complex systems. For reference, the key terms and how they are defined in the specific physics example are summarized below in [Table 6](#). Throughout Chapter 4, we'll create corresponding definitions for each term as the socio-technical analogy develops, which are summarized alongside their physical counterparts in [Table 6](#).

<b>The Type of System</b>
complex mechanical systems
<b>The Specific System</b>
a double pendulum
<b>Conventional Coordinates</b>
x, y, z
<b>Conventional Coordinate System</b>
Cartesian
<b>The Problem</b>
Conventional approach does not provide insight into fundamental system behavior and how to change it
<b>The Process (to build a more appropriate framework)</b>
Lagrangian Dynamics



<b>New coordinates</b>
$\psi_1, \psi_2$
<b>New coordinate system/framework</b>
Generalized coordinate system
<b>New abilities</b>
<p>The new framework enables a more intuitive exploration of the system: how do individual pieces come together to form the collective system? Which interesting behaviors are possible, and under what conditions?</p> <p>The result is not a single solution, but a framework that allows several questions to be asked and answered more effectively.</p>
<b>Validation</b>
<p>Can use simple cases (single pendulum) to check results w/ conventional method.</p> <p>Results are confirmed to obey the laws of physics.</p>

Table 6. Terms and definitions used in the Lagrangian mechanics example of a double pendulum with spring (Figure 12).

### Creating a Sociotechnical Analogy

While the “forces” that drive behavior in sociotechnical systems are not literal, components within the system do affect each other – often in ways that are not easily understood from a conventional reference frame (Zadeh, 1973). These components are not pendula and springs, but the government entities that set policy, the courts that interpret laws and policies, the organizations and groups that respond to these policies (in a variety of ways), and the individuals and communities that are connected to them all. The forces in this system can be considered as the incentives that drive their behaviors. And the resulting motions that emerge from the interactions between components include a myriad persistent security concerns that span the space between individual and global

effects, (e.g. large-scale data breaches (Wolff, You'll See This Message When It's Too Late: The Legal and Economic Cost of Cybersecurity Breaches, 2018), identity theft (Owaida, 2021), stolen intellectual property (Rogin J. , 2012)).

Despite the efforts of industry and government, the interplay between the incentives at different levels isn't understood in a meaningful way, and stakeholders have not been able to meaningfully change the behaviors of the overall system. The internet is a large-scale sociotechnical system in which seemingly small actions result in large systemic impacts, and large efforts sometimes provide no impact at all. For example, a weak password or misapplied software patch can lead to national-scale security incidents - a small action with large impact. Conversely, billions of dollars are spent to secure systems, but the number of incidents has not decreased - a large effort with small impact. Many of the deeply human elements of the system are obscured by standard methods of approaching cybersecurity policy, which may jeopardize the security of individuals, communities, and nations. Given that conventional approaches haven't worked, perhaps we can follow a similar path as Lagrange, and create a new process that provides a clearer, more intuitive view of the complex system of cybersecurity.

### **“Conventional” cyber coordinates**

One challenge in creating a sociotechnical analogy of a physics concept is that physical systems include well-defined coordinates and coordinate systems, but complex socio-technical systems do not. There is some work to do in developing appropriate socio-technical parallels of the key concepts in Table 6, but the data set and taxonomy

from Chapter 3 can be leveraged, just as data and prior definitions from the physical sciences informed Lagrange’s process.

The Lagrangian mechanics process begins when one recognizes that the conventional coordinate system for analyzing a given object is insufficient to answer their questions. If one claims that a socio-technical analogue will similarly make cybersecurity incidents easier to understand, can the insufficient “coordinate system” currently used to describe cybersecurity be identified? Applying the notion of a coordinate system to cybersecurity events is a unique approach, so these coordinates must be deduced. This can be done by analyzing which features and metrics are common in how the events of interest have conventionally been reported and discussed<sup>21</sup>.

#### Conventional metric 1: number of records

One commonly reported metric around cybersecurity events is the number of records impacted by the event. A “record” is not a precisely-defined unit – it can be a name, a credit card number, or a larger group of data belonging to a single person. Although it is a commonly cited metric for cybersecurity events, the number of records is often a poor indicator of the event’s impact. It is an easy metric to understand, but the number of records involved in an event does not indicate its severity (as shown in Chapter 3). One can imagine, given the loose definition of a “record”, a data breach

---

<sup>21</sup> I noted the prevalence of these metrics in the materials used for this research, as well as through my general experience in cybersecurity and media. It is worth noting that, when asked about these incidents, ChatGPT also reported number of records, financial impact, and attribution for nearly all cases as well. ChatGPT is a statistical representation of a much larger data corpus (approximately 45 TB as of March 2023) than any individual could ever process, and came to the same general conclusions, which adds credibility to this assessment.

which is very large but effectively irrelevant - not all personal data are harmful if disclosed, or of value to an adversary. “Number of records” is also a poor metric because it doesn't address the number of companies or range of sectors impacted in business-to-business incidents, such as the SolarWinds or PowerApps events.

The number of records in a breach can be impressively large (recall the 3 *billion* records exposed in the 2016 Yahoo! breach) but is not always insightful. A big number could indicate two things: first, the breached entity is large, but this can generally be determined without a data breach. Second, the breached entity likely did not properly segment their data storage, which would allow that large number of records to be accessed more easily in a single event. The second point is a rather narrow one – it may be actionable for incident response and small-scale technical system improvements, but of little utility in the broader socio-technical context. And as previously discussed, some significant events do not have any number of records associated with them at all, such as Colonial Pipeline and Bowman Dam.

#### Conventional metric 2: attribution of the incident

Another commonly reported “metric” is the source of the incident, or the “attacker” (although some incidents, such as the Equifax data breach of 2017, could be better attributed to a lack of defense than a committed offense). Identifying the source of an attack is known as “attribution”. It a difficult process, a blend of technical tools, organizational culture and political will, one that is often more of an art than a science (Rid & Buchanan, 2015). As shown in Chapter 1, an attacker is likely to hide their machine’s address (unless they are explicitly trying to take credit for their work), and can

easily obfuscate their geographic location by using a variety of proxy machines. Some offensive cyber groups use specialized tools and techniques or pursue specific targets that may provide clues to attribution (Mandiant, n.d.), but these often fall short of a definitive “smoking gun”.

As Peter W. Singer wrote in his book ‘Cyberwarfare: what everyone needs to know’, attribution is a challenge, but deciding what to do once an attacker has been identified (or assumed) is much more difficult (Singer & Friedman, 2014). For some companies, claiming “a sophisticated nation-state attacker” or “an advanced persistent threat” is to blame may be to the company’s benefit, whether or not the evidence support that assertion. From the perspective of the public (or the perspective of regulators) being attacked by a heavily resourced foreign intelligence service feels more excusable than simply having a poor password policy that a less sophisticated hacker could exploit. Some security experts have called into question the attribution claims of certain companies for this reason (Mimoso, 2016). Given that attribution is rarely a certainty, political or financial motivations may be more influential in the decision to point a finger than technical data-driven justifications.

Aside from the uncertainty around attribution, there are circumstances in which knowing the attacker’s identity is simply not insightful. The Center for Strategic and International Studies (CSIS), a prominent American think tank, maintains a timeline of “significant cyber incidents<sup>22</sup>” that have occurred since 2006 (Center for Strategic and

---

<sup>22</sup> In this case, “significant” is defined as attacks on government agencies (U.S. federal, state, local, and international), large companies, or economic crimes with impacts over \$1M USD. Implied in this definition

International Studies, 2023). As of March 7, 2023, this list included more than 870 entries, and nearly all include the national affiliation of the attackers. (There are a few incidents on this list for which no reasonable speculation can even be made.) While attribution information - or informed assumptions - can be valuable in certain contexts, they can also be distracting or misleading. For example, some variants of ransomware code can travel between systems indiscriminately, encrypting any network they encounter, rather than intentionally selecting victims. ‘Ransomware as a Service’ further complicates the idea of an “attacker”, as those performing the attack are doing so at the request of a third party (for a fee, of course). For example, the CSIS timeline states that the Colonial Pipeline hack was “attributed to DarkSide, a Russian speaking hacking group”, but the for-profit group released a statement shortly afterwards saying that they are “*apolitical*”, and that their “*goal is to make money, and not creating problems for society*”. The Russian government may be turning a blind eye to DarkSide’s activities, but there is no evidence the Russian government is funding or directing the group’s activities (Rivero, 2021). Even when the attack can be clearly attributed to a nation-state group (e.g. the Sony Pictures Entertainment event, for which North Korea directly claimed responsibility), this information is of little utility outside the realms of national security or international relations. Attribution – or attempts to attribute - may make for an interesting story but cannot reliably indicate the impact of an attack at a personal or organizational level.

### Conventional metric 3: financial impact

---

is a sufficiency of data around the incident – an inestimable number of events would not qualify for this list.

Another large number that is often discussed alongside cybersecurity events is the amount of financial damage incurred as a result of an event. Financial impacts can occur in a wide variety of ways – a (generally brief) drop in stock price (Rosati, Deeney, Cummins, van der Werff, & Lynn, 2019), loss of intellectual property value, lost revenue due to interrupted operations, fines from government entities (SEC or FTC), cost of remediation and/or compliance with a regulatory consent decree, class action lawsuits or other settlements, or outright theft of funds. These costs are often associated with a company or government organization or rolled into a vague global aggregate value, but the financial burden on individuals is rarely explored.

For example, Colonial Pipeline paid 75 bitcoin to recover its operations – at the time an equivalent of \$5M USD, but valued at \$2.1M USD at the time of this writing (March 17, 2023)<sup>23</sup>. But while the pipeline’s operations were offline, gas prices for individuals were reported to spike by 18-21 cents per gallon (Tsvetanov & Slaria, 2021), and in some areas of the country gas stations ran out of supply completely (Turton & Mehrotra, Hackers Breached Colonial Pipeline Using Compromised Password , 2021). Neither the cost of the ransom nor the cost of the operational interruption reflects the expense at an individual level.

Ransomware attacks can be costly even without paying the ransom. When hit with ransomware, Sinclair Broadcast Group was able to restore their operations independently and did not pay the attackers. However, the disruption still resulted in a loss of \$63M USD in advertising revenue and an additional \$11M in remediation costs.

---

<sup>23</sup> Half the ransom was then recovered by the FBI, but it is very rare to recover any ransomed funds from a ransomware event. (Romo, 2021))

Cyber insurance covered \$50M USD of the loss but left Sinclair with a deficit of \$24M USD uncovered (Brumfield, 2022).

In short, a dollar amount can't sum up the entire financial impact of an event. It may indicate one facet of the incident's cost, such as the size of a particular penalty or an estimate of loss in revenue. But there are too many different types of costs at different levels of the system, incurred over extended time periods, to make a comprehensive, accurate, and therefore insightful estimate of the total dollar amount.

In forming the classical Lagrangian, variables with little to no effect on the system can be discarded, which reduces computational effort and improves understanding of the fundamental (or first-order) motions of the physical system. It is with a similar motivation that we disregard the number of records, the source of an attack, and financial impact of an event as key features in constructing a new framework. Too much focus on these factors can complicate our understanding of the system, obfuscating and distorting the fundamental behavior we strive to discover. They are valuable in some contexts, but do not contribute meaningfully to this research trajectory.

### **New cyber coordinates**

The conventional coordinates of number of records, attacker, and financial impact have been evaluated and deemed insufficient for our needs. To develop a more appropriate set of coordinates for cybersecurity incidents, let's continue to borrow inspiration from the physical sciences, returning to the double pendulum example from Figure 12.



Upon examining the object, two pendula and a spring have been identified. The specific identity of the spring is deemed unnecessary since it has been established through previous observations that all springs function similarly, characterized by a "spring constant" which denotes the amount of energy stored by the spring when compressed or extended. For the current scenario depicted in Figure 12, the spring constant is the only pertinent piece of information that needs to be considered, with other attributes such as its color, material composition, or origin being deemed irrelevant. It should be noted that the term "spring" does not exclusively refer to coiled metal pieces but encompasses any object that conforms to the characteristics of a spring, such as rubber bands. This insight, which has been derived from meticulous observations and precise measurements, has enabled physicists to circumvent the cumbersome and superfluous process of generating new classifications for objects that share similar behaviors. By adopting a collective nomenclature, "springs," to describe objects with analogous behaviors, researchers can concentrate on the fundamental characteristics of the system, which represents an indispensable phase towards constructing generalized models.

Groups are chosen based on the kind of problems that are being solved. In a classical dynamic system, treating a rubber band like a spring is appropriate; treating a metal rod like a spring is not. These groups are chosen based on outcomes – when displaced from equilibrium, the spring and rubber band both produce oscillatory motion. Replacing a rubber band or a metal spring with a metal rod would fundamentally change

the dynamics of the system. Based on the interest in kinematics, one can defensibly group rubber bands and springs into the same category, and exclude metal rods<sup>24</sup>.

Similarly, can some useful groups be defined, based on a robust data set and centered on outcomes, that can be used to define a new set of “cyber coordinates”? Fortunately, this work has been done: the different levels of impact described in Chapter 3. The categories from this taxonomy were chosen to facilitate this kind of parameterization; important properties of the system were defined based on this specific research interest. For both the physical and socio-technical systems, the appropriate care has been taken to understand which objects share important properties of interest and create categories accordingly.

What one considers “important” is driven by the goal – namely, to uniquely illuminate the incentive conflicts and gaps in federal level cybersecurity. The new “cyber coordinates” – personal impact, infrastructure impact, and political impact – can now be utilized to achieve this goal. The double pendulum was described in the new coordinates  $\psi_1$  and  $\psi_2$ ; new cyber coordinates can now be assigned to the cybersecurity events in the initial data set. Recall that in Chapter 3, a range from -3 to +3 was designed to indicate the severity of the impact (see Table 2 for the general rubric, and Tables 3, 4, and 5 for category-specific rubrics). The values for each event are shown in Table 8. A framework now exists that can be used to examine ways to rebalance the system, and

---

<sup>24</sup> A different research inquiry, perhaps into the thermodynamic properties of the system, would not find rubber bands and springs to be equivalent in their important properties. A physicist working in this area would argue that the spring and metal rod should be grouped together, given their reactions to heat and cold, so the rubber band is the dissimilar object, as it has fundamentally different properties of interest.

achieve a more stable equilibrium. In the physical example, once the equations of motion have been put into new coordinates using Lagrangian mechanics, they can be used to investigate specific system behavior or predict future behavior. Similarly, the new socio-technical coordinates for cybersecurity incidents create a useful framework which can be used to investigate specific system behavior or predict future behavior. In recognition of its origins and the inspiration behind it, the name Socio-Technical Lagrangian (or STL) seems appropriate.

This chapter proceeds with analysis of the initial data set using STL, investigating the interesting features that have emerged and what these features imply about systemic incentives. Then the new framework is applied to a separate set of validation data to test the framework to ensure it is robust.

### **The Socio-Technical Lagrangian**

A novel framework, the Socio-Technical Lagrangian (STL), has now been built to enable a more intuitive exploration of cybersecurity issues in the U.S. The result is not a single solution, but a framework that allows new questions to be asked and answered more effectively: how does the system behave as a collective? Where do interesting or problematic behaviors occur and what changes could be made to affect them? The pieces of this new framework are analogous to the steps leading to Lagrangian Dynamics formalism – the steps of both the physical example and the sociotechnical analogue are aligned in Table 7 below:

Table 7. Terms and definitions used in both the double pendulum, and the socio-technical analogue of cybersecurity systems.

	<b>The Type of System</b>	
complex mechanical systems		complex socio-technical systems
	<b>The Specific System</b>	
a double pendulum		cybersecurity issues in the U.S.
	<b>Conventional Coordinates</b>	
x, y, z		Number of records, Attribution, Financial impact
	<b>Conventional Coordinate System</b>	
Cartesian		Number of records, Attribution, Financial impact
	<b>The Problem</b>	
	Conventional approach does not provide insight into fundamental system behavior and how to change it	
	<b>The Process (to build a more appropriate framework)</b>	
Lagrangian Dynamics		Sociotechnical Lagrangian
	<b>New coordinates</b>	
$\theta_1, \theta_2$		harms at individual, infrastructure, and political levels
	<b>New coordinate system/framework</b>	
Generalized coordinate system		Primary Impacts
	<b>New abilities</b>	
	New framework enables a more intuitive exploration of the system. The result is not a single solution, but a framework that allows several questions to be asked and answered more effectively.	

Each event has been assigned a score in all three categories. The mean impact across all three categories for each event is then calculated, as well as the variance<sup>25</sup> of the impacts across the three categories. A large variance indicates that there are significant differences in the degree to which each level was harmed; a small variance indicates that the groups were impacted at approximately the same degree. Since the power of this approach is in comparing similar events or looking at outliers, rather than in the numerical values themselves, color scales have been added to Table 8 to facilitate, where green indicates a lower mean harm or a lower variance between harms for a particular incident, and red indicates higher values.

*Table 8. Events from the data set with assigned impacts in the three primary categories, as well as average and variance of impact.*

---

<sup>25</sup> In statistics, “variance” is formally defined as the standard deviation *squared*, whereas the value calculated here is the standard deviation itself. “Range” may be a more appropriate term here, as “standard deviation” may imply a larger set of values.

Event	Year	Overall Category	Number of Records (M)	personal impact	infrastructure impact	political impact	average	variance
Bowman Dam hack	2013	Infrastructure	N/A	0	-2	-3	-1.7	1.5
Anthem medical data breach	2015	Infrastructure	80	-2	-3	-1	-2.0	1.0
Colonial Pipeline ransomware	2021	Infrastructure	N/A	-2	-2	-3	-2.3	0.6
Ransomware attacks in hospitals, municipal services, education	2017-2022	Infrastructure	N/A	-3	-3	-1	-2.3	1.2
Microsoft PowerApps data exposure	2021	Infrastructure	38	-2	-3	-1	-2.0	1.0
Target data breach	2013	Personal	110	-1	-1	-1	-1.0	0.0
Ashley Madison breach	2015	Personal	33	-3	3	-1	-0.3	3.1
Yahoo data breach	2017	Personal	3000	-3	-3	-1	-2.3	1.2
Equifax data breach	2017	Personal	143.5	-1	0	-2	-1.0	1.0
Facebook data breach	2019	Personal	533	-3	0	-1	-1.3	1.5
Sony Pictures Entertainment attack	2014	Political	minimal	0	-2	1	-0.3	1.5
SolarWinds malware	2020	Political	N/A	0	-3	-3	-2.0	1.7
Office of Personnel Management (U.S. government)	2015	Political	21	-2	-2	-3	-2.3	0.6
Election attacks (Democratic/Republican National Committees; voter registration databases in multiple states)	2016	Political	N/A	-3	1	2	0.0	2.6
NotPetya malware	2017	Political	N/A	0	-3	1	-0.7	2.1

It's important to walk through a simple example to show the methodology by which the impacts are assigned: the Target data breach of 2013. The Target impact has been scored -1 for personal, infrastructure, and political impact. Individual customers may have replaced their credit cards, or signed up for credit monitoring services, but these are relatively low burdens compared with some of the other incidents. This event also scores a -1 for infrastructure – Target did respond to the breach and made some operational changes based on the event (Kassner, 2015), but there were no insurmountable impacts at this level. Finally, the event was assigned a -1 in the political impacts category as well. While a foreign adversary may have been able to extract some intelligence value from the data when combined with other efforts, the potential harm from this type of data is minimal. Because all three impacts are the same, the variation is zero for this event (the smallest variation in the data set).

Now that the baseline procedure for assigning impacts is understood, the focus can now turn to groups of events, and the interesting contrasts and commonalities

between them. The new parameterization immediately highlights areas that are worthy of further exploration in a way that the traditional discourse does not enable. The power of this method comes from these comparisons, and not simply from the impact values themselves. Another expert in the field may not concur with the exact numerical assignments in the three categories but would still be able to use the STL methodology to create useful comparisons. In other words, it isn't the numbers themselves that matter, but the method by which they are assigned and subsequently used in the analysis.

#### Highest Variation: Ashley Madison

Compare the values of the Target breach with another event from the Personal Impacts category, the Ashley Madison breach. This event scores the maximally negative value of -3 for personal impact. Revelations of infidelity, actual or intended, caused unique and irreparable harm to the individuals affected, including harassment, hate crimes, and in at least one confirmed case, suicide (Segall, Pastor outed on Ashley Madison commits suicide, 2015). Experts noted that the psychological trauma was not limited to those named in the breach, but spouses and children suffered as well (Gregoire, 2015). Contrast this with the infrastructure impact score - an extraordinary +3, since the breach effectively acted as free advertising for the company, which saw a 30% increase in new accounts in the five months after the breach. Ashley Madison's brand recognition grew, more traffic was driven to their website due to the news coverage of the breach, and the company hit the 60 million member mark in 2019 (Dodgson, 2019). A class

action lawsuit<sup>26</sup> (In Re: Ashley Madison Customer Data Security Breach Litigation, 2017) was settled in 2017 for \$11.2M - down from the original filing of \$576M – a small percentage of the company’s profits that year<sup>27</sup>. This huge disparity in impacts highlights a concerning imbalance in systemic incentives. The Ashley Madison company lied to their customers about their security posture, and once they were breached (because the promised security measures were nonexistent), made significant profit directly from an event that was devastating to individuals.

Averaging the impacts, the Ashley Madison breach scores a -0.3 overall, but a variance of 3.1 – the largest in this data set. A large negative average impact indicates that an event was harmful at all levels, but a large variance reveals that harms are significantly unbalanced in the system. STL allows us to first discover these imbalances, then to consider what policy changes could be made to decrease them. Let’s imagine a policy environment in which company growth could be inhibited in the wake of a serious security incident. In this notional world, one can envision a consent decree from the FTC that prohibits the creation of new accounts on insecure infrastructure, or sends profits from new business into some kind of escrow account – specifically how this might be accomplished is a question for Chapter 5, but for now it’s enough to imagine that some kind of policy exists. How would that impact the scoring of the Ashley Madison breach

---

<sup>26</sup> The number of potential claimants was deemed “sufficiently numerous” to fit the requirement for a class action lawsuit. However, a motion to proceed under pseudonyms was denied by the court, requiring plaintiffs be identified by their real names, and disbursement of awards may have also been complicated by the need for self-identification.

<sup>27</sup> Two types of losses (and associated damages) were available to claimants as a result of the breach. Those who purchased “Full Delete” but had their data exposed could claim between \$19USD and \$500USD, depending on the number of accounts. Those who experienced identity theft or other unreimbursed losses as a result of the breach could claim up to \$2000USD.



in this alternate universe? Assuming the breach still happened, the personal harm ranking would be unaffected at -3. However, in this notional world, the positive infrastructure impacts are impeded by the new policy – so this category’s score moves from a +3 to at most a 0. Political impacts remain the same at -1, since data breaches can always be mined by a foreign adversary for intelligence value. In this speculative reality, the Ashley Madison breach now scores an average impact of -1.3 instead of -0.3, and 1.5 variance instead of 3.1. The average impact score is more negative, which seems to more accurately reflect the situation. The variance has been cut in half, indicating that the wildly disparate impacts between people and infrastructure have been brought into better alignment through our imagined policy.

In the previous scenario, a hypothetical policy was considered that would trigger federal intervention after a data breach, either removing the potential for profit or delaying it until security issues have been remediated. Would such a policy impact the probability of the data breach happening in the first place? FTC intervention is not a new concept, but the ways in which it has been applied have been vague and not always proportional to the harms done.

In the real world, the FTC fined ALM \$1.6M USD and enacted a consent decree that required the company implement a more rigorous security program. (\$1.6M may seem significant, but as a comparison, Ashley Madison offered \$10M to the City of Phoenix in 2010 to change the name of Sky Harbor International Airport to ‘The Ashley Madison International Airport’ for a period of five years (Fisher, 2010).) The FTC’s engagement may seem like a mere slap on the wrist, but worse, it could be a disincentive

to promote security. The FTC is a user advocacy group which engages when companies are dishonest about their practices (50 CFR part 216 Subpart H), and the FTC did not get involved in the Ashley Madison breach because of a lack of security; they engaged because the company promised robust security features that they did not deliver. Had Ashley Madison not overtly promised these security features to their customers, the FTC would have been in a much weaker position and may not have had the legal authority to intervene. Instead of an incentive to create meaningful security for users, the threat of FTC action may currently be acting as a disincentive to promise and enact specific security measures for customers.

At best, the current mode of FTC engagement is an insufficient incentive for companies to secure their systems. It's unlikely that any company thinks about the potential positive outcomes of a data breach – the Ashley Madison event is unusual in this sense – but a policy that invokes regulatory action in a more direct and predictable way overall may act as a better motive for companies to spend the time and money on meaningful security measures.

#### Another High Variance: Election Infrastructure Attacks

Let's look at another event with a relatively low average impact and a high variance – the 2016 attacks on the electoral system. This is a case where infrastructure and political impacts are both positive (1 and 2 respectively). Infrastructure in this case is defined as the local and state level election systems, such as voter registration databases. These systems were breached but there is no evidence that any data were changed, nor were the vote tallies impacted by this event (Greenberg J. , 2016). What did change is the

important realization of electoral systems as critical infrastructure with unique cybersecurity needs – not just the vote tabulation machines but supporting technologies such as state and local databases of voter information.

Since the attacks in 2016, all 50 states have improved their election security posture, many through cooperation with federal-level entities such as DHS or the National Guard (Root, Kennedy, & Souzan, 2018). While it would be preferable to not have had *any* intrusions into the U.S. electoral system, the response from state and local governments resulted in some meaningful improvements to election infrastructure, which aggregate to a stronger national security position. There are still many improvements that should be made to election technologies, but the infrastructure and national security posture have both been improved as a result of these incidents, hence the positive rankings at these levels.

The personal impacts of this event have been scored maximally negative at -3. It may not be possible to know the full intentions behind the electoral hacks in 2016, but it is quite possible that one of the primary motivations was not just to exploit the vulnerabilities in machines, but to exploit human vulnerabilities as well. Specifically, human trust is required to make democracy functional, and trust in election systems has been damaged significantly in the past several years (Rinehart, 2020). The 2016 election system hacks are not the only avenue by which voter trust has been eroded; the problem is broader than cybersecurity and continues to evolve. But from the individual perspective, these attacks were powerful and effective, undermining many peoples' trust in democracy. The negative impacts persist to this day, which is why this category received a score of -3.

In this case, it would not be useful to balance the system by negating positive impacts at the infrastructure or political levels! Mitigating issues at the personal level is the only reasonable option. However, it is more difficult to envision a policy solution to this particular problem, especially given the irretrievable confluence of election security with partisanship and disinformation since the event in 2016. This is an example of how the impact of cybersecurity events can extend beyond technology to the human and social components of the system, and how difficult it is to remediate once that happens.

Hopefully it is clear at this point that the Socio-Technical Lagrangian's value is not in the rankings themselves, nor in the subsequent numerical calculations, but in framing events from a different perspective that allows one to pose more useful questions. A large variance for an event does not prescribe a certain policy response, but it spurs one to ask, "what imbalances exist in the system to cause this effect?". A mean impact of zero doesn't mean that the event truly had *no impact*, but that the negatives in the system are offset by positives elsewhere. This may indicate a pathology in policy, like the Ashley Madison case, or it may indicate that the event instigated positive change in some ways, which are also worth recognizing. The analogy with complex mechanical systems holds – describing the object of interest in a new (and purposefully designed) way allows for new insights into its behavior.

### **High Average Impacts: Comparing Infrastructure Events (SolarWinds and NotPetya)**

The category of “infrastructure” impacts is more loosely defined than personal or political, which reflects the myriad of ways in which public companies, private companies, government organizations, and others interact to form and operate much of the internet. To quote cybersecurity expert and Mandiant CEO Kevin Mandia, “*Any conflict in cyberspace, whether motivated by a criminal element or motivated by geopolitical conditions, it's going to involve both the government and the private sector.*” (Temple-Raston, 2021) Recapping from Chapter 2, events in the “**infrastructure impacts**” category reflect incidents that primarily affect groups such as companies, infrastructure, hospital networks, or municipal services. These services may be in the public or private domain, or a blend of both, and they may include a combination of entities (e.g. the U.S. power grid). Some of the most far-reaching cybersecurity events occur when infrastructure entities are impacted – particularly when one entity provides services to a wide and diverse web of sectors.

All but one of the infrastructure events in this data set ranked in the most negative 50% of average impacts, which indicates that impacts to infrastructure strongly affect individuals and national security as well. The significantly negative average impacts of this category suggest that policy efforts focused on infrastructure may be able to affect positive impact across the system as a whole. This is not to say that all infrastructure-centric events follow the same pattern, however. Comparing three events with maximally negative infrastructure scores – the Microsoft PowerApps data exposure, the SolarWinds breach, and the spread of NotPetya malware – provides a more granular level of insight into the incentives that drive behaviors at this level.

## Data Exposure By Design: Microsoft PowerApps

Microsoft PowerApps is an excellent example of a single entity that serves a large and disparate population of other entities, who then serve individuals (or even other intermediaries). PowerApps is a tool that allows software developers to quickly spin up new applications – both the front-facing site for users, and the back-end database to record and store user data, hosted in the cloud. PowerApps is the engine behind a huge number of applications that manage a diverse set of data including flight records (American Airlines), medical information (COVID vaccine information from U.S. government entities), drug screening information including social security number (J.B. Hunt), information of minor students (NYC schools), and more (UpGuard, 2021). Some of the data should be public, like the addresses of vaccination clinics, but some should be protected, like the names, addresses, and vaccination status of their patients. PowerApps' default setting was to make data publicly accessible, and although the software documentation explained how to change that setting to private, security researchers from UpGuard found thousands of databases that had been inadvertently left open to public sharing (Newman L. H., 38M Records Were Exposed Online—Including Contact-Tracing Info, 2021). UpGuard submitted a vulnerability report to Microsoft, but the response was that this configuration was “considered to be by design” – putting the responsibility back on the software users - and the issue was closed (UpGuard, 2021). UpGuard continued to reach out to Microsoft as well as the secondary impacted companies, and Microsoft eventually changed the default setting to private. although nearly 50 separate corporate and government entities had been impacted by that point (Vincent, 2021).

The PowerApps event scores a -2 in personal impacts, in part due to the medical information and personal data of minors that were exposed. As of March 2023, there was no evidence of the data being misused, which prevents the event from scoring a -3, but there is also no way to ensure that it hasn't been misused in some way. The large cross-sector impact is the motivation behind a -3 scoring in the infrastructure category. For one software setting to expose such a wide variety of data, from travel to medical to education, indicates a considerable imbalance in the system, and is reflected as such in STL.

What policy intervention might create a more balanced system of incentives? The data exposure from PowerApps was due to a software setting, not due to an attacker who intentionally breached the perimeter. Public accessibility was likely chosen as a default setting because it made building applications easier, but the PowerApps developers likely didn't consider the impact of that choice during the design phase. Overall, misconfigured cloud storage is cited as a dominant trend in data breaches by the 2022 Verizon Data Breach Report (Verizon, 2022), an annual analysis that is held in high regard in the information security industry. The report goes on to say that "*the fallibility of employees should not be discounted*", but given the extensive nature of the problem, it seems more appropriate to connect this fallibility with software design, rather than the users.

The problem of misconfigured cloud storage, inadvertently making data public when it should be private, could be solved by a standard that privacy advocates have been discussing for years. Whether individual users on the web or software developers using tools like PowerApps, an opt-in standard would require that data sharing be turned off by default, and a reasonable explanation of the risks provided before sharing can be turned

on. There are many reasons to share data, depending on who it's being shared with and for what purpose, but these choices should be made in an informed way, on a case-by-case basis.

## SolarWinds

It is valuable to compare two large business-to-business events in the data set which have a different scoring of impacts. SolarWinds is not as recognizable a name as Microsoft, but the company provides a similar service – they are a software vendor that specializes in IT/network management tools for customers in both industry and government sectors. Like the PowerApps misconfiguration, the SolarWinds event affected a centralized software platform that few people outside of IT had heard of, impacting a large web of industry and government entities.

Unlike the exposed data from PowerApps, the SolarWinds incident was an intentional and sophisticated piece of malware created by a foreign adversary for espionage purposes, and as such impacted a different group of targets than the PowerApps exposure. “Fewer than 18,000 customers” were impacted, but these customers included companies like Microsoft and FireEye, as well as federal agencies including the Department of Homeland Security, the Treasury Department, and the Commerce Department. The attackers were able to monitor and steal data from these networks for several months without detection (Cimpanu, 2020).

It's not clear when the company became aware internally of the breach, but it was first reported by the Washington Post on December 13, 2020, and for several days after, SolarWinds continued to distribute compromised code (Varghese, 2020). Additional



security issues came to light, including the use of hardcoded passwords (“solarwinds123”) and official documentation recommending that customers disable their antivirus software to make the SolarWinds products more effective (Seals, 2020). The company’s new CEO, addressing Congress in February 2021, took a page from the Equifax playbook and blamed an individual (an intern, no less) for the poor password choice (Johnson, SolarWinds CEO expresses regret for ‘blame the intern’ defense during Orion hack investigation, 2021), although neither a simple password nor a default of any complexity should have been allowed in a company with reasonable security standards.

SolarWinds is one of the few companies to have a longer-term impacted stock price due to a cyber event, likely due to continuing high-profile legal action. The Securities and Exchange Commission has filed a lawsuit accusing SolarWinds executives of not properly disclosing security issues, and trading \$280M USD in stock just days before the breach was made public (Harwell & MacMillan, Investors in breached software firm SolarWinds traded \$280 million in stock days before hack was revealed, 2020). A separate class action lawsuit has been filed on behalf of shareholders (Bremer v SolarWinds, 2021), neither of which has been resolved as of this writing.

### Policy Design for Diverse Events: SolarWinds and PowerApps

The SolarWinds and PowerApps incidents share a maximally negative score of -3 in the infrastructure category, but they are very different in the individual and political categories. PowerApps scores more negative for individuals, due to the personal data that were exposed, but less negative for political, as the exposure was unintentional, not due to a malicious actor breaching the system, and there has been to date no evidence that the

exposed data were abused. SolarWinds scores a zero for individual impact and a maximally negative score for political impacts, due to the adversarial purpose and specific targets that were selected. Infrastructure-related events don't have the same outcomes – therefore, an effective policy approach shouldn't focus solely on user privacy or national security but must be capable of addressing both.

Designing such a policy is a challenge, particularly when it comes to implementation. But addressing harm reduction for individuals and national security in a coherent policy will make sure that neither is sacrificed for the benefit of the other. Considering both goals during policy design will hopefully also result in guidance that is easier for companies and government organizations to follow, rather than having a myriad of individual policies in place that cause confusion or interfere with one another.

#### Collateral Damage at a Global Level: The Spread of NotPetya

The NotPetya malware incident stands out in Table 8 for its large variance in impacts. Like SolarWinds and PowerApps, this event scores a -3 in infrastructure for its broad range, but it scores a 0 in personal impact and an unusual +1 in political impact. The uncontrolled spread and wide-ranging consequences of the NotPetya malware is a good example of why cybersecurity issues can't accurately be described by geographic boundaries. Like ransomware, this malicious code travels through a victim's computer network, encrypting each machine it can access. Unlike ransomware, however, there is no way to pay a ransom and recover the machines – they are permanently and irretrievably locked. The NotPetya malware was designed by Russian military hackers specifically to target Ukrainian accounting software, one part of a continuing Russian campaign to

damage Ukrainian infrastructure (Greenberg A. , The Untold Story of NotPetya, the Most Devastating Cyberattack in History, 2018). (This attack was part of the extended pre-kinetic conflict before Russia's invasion of Ukrainian geographic territory on February 24, 2022.)

As discussed extensively in Chapter 2, however, computer networks aren't segmented along geographic boundaries. Ukrainian computers connect to machines all over the world, and many of them share the same software configuration that the NotPetya malware was targeting. NotPetya moved extremely quickly, knocking out significant infrastructure in Ukraine (banks, hospitals) and moving across the globe in a matter of hours. The malware seized machines in hospitals in Pennsylvania, factories in Africa, halted global operations of the shipping company Maersk, and even reached back to infect the Russian state-owned oil company Rosneft. Overall, the extent of NotPetya's damage, as well as the speed at which it inflicted the damage, were both astonishing. For this malware's speed, extent, and damage, it scores a -3 in infrastructure.

It is unusual for specific individuals to be charged in global cyber events, but in 2020 the U.S. Department of Justice charged six Russian Intelligence officers with a spate of cybercrimes, "wantonly causing unprecedented damage to pursue small tactical advantages and to satisfy fits of spite" (The United States Department of Justice, 2020), including the global NotPetya malware infection. It's unlikely that the named Russian officers will be arrested, but the detailed attribution at the individual level makes a powerful statement about U.S. intelligence gathering abilities in cyberspace. (The Kremlin officially responded that claims that U.S. attribution of NotPetya as a Russian offense are part of a "Russophobic campaign that is not based on any evidence." (Wood,

2022)) It is also interesting to note that the U.S. DOJ cites several companies as being integral to their investigation, including Google, Twitter, Facebook, and “some private sector companies”. While the NotPetya malware was incredibly damaging, the U.S. was able to leverage it to show political strength against an adversary, which is why it scores a +1 for political impacts.

That being said, it would be unreasonable to suggest that the NotPetya malware is “good”, or that additional malware attacks should be encouraged so that countries have the opportunity to flex their intelligence capabilities. The decision to publicly release intelligence data is delicate and nuanced, with significant potential to impact global relationships, and may not always be the right choice. It is far more sensible to address the potential for widespread damages through intelligent policy decisions for organizations in the infrastructure category, including guidelines for cyber insurance coverage that more accurately reflect the operational environment.

### The Evolution of Ransomware: Colonial Pipeline’s Impact

Ransomware attacks are generally represented in the data set as a single category, due to their prevalence and the similarity of their impacts. This “event” includes attacks on hospitals, municipal services, and educational systems (from elementary schools to universities). But there was a moment when the response to ransomware changed course within the U.S. federal environment – the Colonial Pipeline attack. The response to Colonial Pipeline was significantly different than the decades of ransomware events that preceded it. What about this event is different, and what policy implications might this have?

The first ransomware attack, back in 1989, traveled via floppy disk, and demanded that the computer's owner pay \$189 USD to recover access to their files (Waddell, 2016). Since then, ransomware attackers have shifted focus from individual computer users to larger and more lucrative targets in the infrastructure sector, such as hospitals, cities (Atlanta, Baltimore, Washington D.C.), and schools, leveraging the global connectivity of the internet and the anonymity of cryptocurrency to scale up their enterprises. As disruptive and pervasive as these attacks were, they didn't garner much interest from the federal government until the ransomware attack on Colonial Pipeline. As mentioned in Chapter 2, the Colonial Pipeline is responsible for transporting 45% of the East Coast's fuel from the Gulf Coast (Sanger, Krauss, & Perlroth, 2021). The disruption caused a spike in gas prices nationwide (Turton & Mehrotra, Hackers Breached Colonial Pipeline Using Compromised Password, 2021), raising the ire of constituents across the country.

For the first time, the federal government responded. Whether it was the geographic extent that made a difference, the disgruntled populace, or the hot-button issue of fuel prices is unclear (Turner, 2021) (Wolff, 2021), but ransomware suddenly became a topic of unprecedented popularity at the federal level. The White House convened an international Counter-Ransomware Initiative, a virtual meeting including 30 countries to discuss "everything from efforts to improve national resilience, to experiences addressing the misuse of virtual currency to launder ransom payments, our respective efforts to disrupt and prosecute ransomware criminals, and diplomacy as a tool to counter ransomware" (White House, 2021) - although notably neither China nor Russia were included in the talks. Occurrences of the word "ransomware" in the U.S.

congressional record (the official record of congressional proceedings) quadrupled between 2020 and 2021. “Colonial Pipeline” is specifically mentioned 48 times in the 2021 congressional record, with only four mentions prior (in 2018, 2017, 2015, and 2012) – as shown in Figure 16.

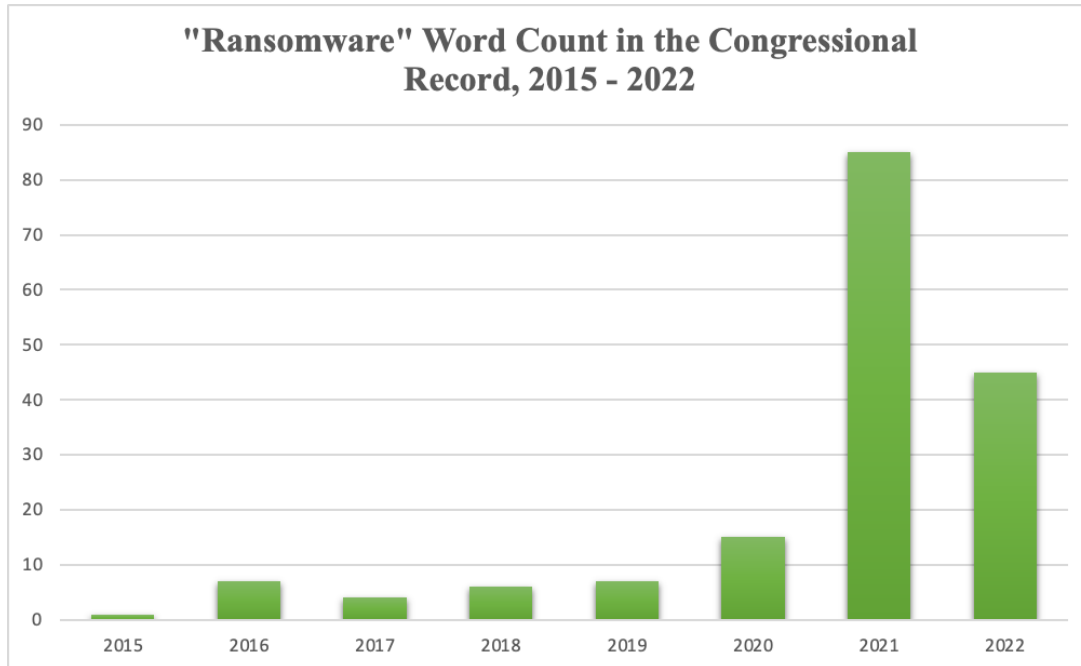


Figure 16. Occurrences of the term “ransomware” in the U.S. congressional record, 2015-2022

In July, 2021, the U.S. leveraged the ‘Rewards for Justice’ program, established in 1984 to fight international terrorism, to offer a \$10M bounty for the Colonial Pipeline hackers (or anyone engaging in “malicious cyber activities against US critical infrastructure”) – a significant step farther than the federal government had engaged in any prior ransomware engagement (Vincent, U.S. government offers \$10 million bounty for information on Colonial Pipeline hack, 2021).

Interestingly, the Colonial Pipeline attackers offered up an apology note, claiming ***“We are apolitical... Our goal is to make money and not creating problems for society.”*** (Clark, 2021) The group also promised to vet their customers’ targets more carefully in the future (DarkSide runs a “Ransomware-as-a-Service” model, providing ransomware packages to affiliates which then perform the intrusion and share the payment with the provider (U.S. Department of State, 2021).)

The ransomware attack on Colonial Pipeline doesn’t look different from prior ransomware attacks, from a technology perspective. However, the U.S. federal government’s response deviated significantly from prior behavior, and as a result, ransomware is now starting to be viewed as a national security threat. This pivot in perspective is likely to lower the personal and infrastructure impacts of ransomware, as companies and municipal governments now have more information and support available to counter the threat. Evolutions in ransomware have accelerated since the Colonial Pipeline event, from both the attacker and the policy perspective.

#### General Features of Interest: Positive Impacts and Zero Impacts

Two general features in this analysis deserve some specific attention – positive numbers, and zeros. My initial perspective was that positive values indicated a particularly concerning pathology in the system. If the impact at one level was positive, did that necessarily come at a cost to another level? The positive impact the Ashley Madison company got from their data breach did come at extraordinary expense of their individual users. Analyzing additional data indicated, however, that there are some

instances of positive impact that don't require incurring damage in other parts of the system. Specifically, there are some events where the U.S. federal government has been able to achieve some positive value by either using the event to create positive change, or for global-level strategic positioning, neither of which causes harm to the infrastructure or personal categories. The three categories, therefore, do not define a closed system.

While rare, there are a few zeros assigned in the table of events and their impacts. The Bowman Dam hack scores a zero for personal impact; while an adversary was able to penetrate the dam's infrastructure, they were not able to actually operate the dam and affect individual lives of the local population. The Sony Pictures Entertainment event also scores a zero for personal impact – while a few Sony executives were embarrassed by the leaked data, it was in context of the operations of the company, and the personal aspect didn't merit an assignment of -1 to this category. (Perhaps a -0.1 would be appropriate, were fractional values to be used.) Equifax and Facebook both scored zeros in the infrastructure category since the companies suffered minimal to no impact to operations or financial position (Owens, 2018) . But due to the intelligence value of aggregated data to a foreign adversary, no event scores a zero for political impact. Either the overall political impacts were positive enough to offset the negative to some extent, or they were simply negative.

### **Validating the Socio-Technical Lagrangian**

What is the goal of creating a new framework? Ultimately, to develop policy that results in not just fewer incidents, but fewer harms from the incidents that do occur. When incidents happen – and all should accept that they will – it would be optimal to



have policy in place such that negative impacts are minimized to the extent possible across the entire system. A framework based on relative harms, at multiple scales, is a more actionable way to identify how and where policy solutions could achieve this goal. This view helps us understand what incentives are at work, how they manifest and how to change them as to address root causes and not cause additional harm elsewhere in the system.

Thus far, examples from the original data set have been used to identify incentive conflicts and consider notional policy avenues to resolve these conflicts. STL should also be tested on an independent data set to ensure it is both robust and extensible. To be valuable, a framework must be insightful for events beyond those used in its construction. To that end, a set of additional events have been chosen for validation purposes (see Table 9). Each of the validation events will be scored without referencing the original data. Once their scores have been assigned in all three categories, they will be compared with the original data to ensure coherence of the method and outcomes.

In the validation step, two questions will be addressed. First, can the initial categories and scoring rubric be used to coherently describe events outside the initial data set? In other words, do similar events in a similar context score the same way? Second, if there are notable differences in scoring or analytics in this validation data, can those differences be explained by changes in context since the initial events were selected? These changes could refer to updated policy or technological developments, but the reason for a difference in scores should be identifiable. This research is not performed in a vacuum –there are plenty of new events to choose from when selecting a validation set,

but there have also been policy changes at work in the system as well. Therefore, the events themselves and the sociotechnical context in which they occurred will be compared during validation.

### Validation Data and Process

Five events were chosen to test the applicability and extensibility of STL, listed below in Table 9:

*Table 9. Validation events*

Event	Year	Overall Category	personal impact	infrastructure impact	political impact	Analogous Event
Planned Parenthood hack	2021	Personal	-3	-1	-1	Ashley Madison
Halfnium attack on MS Exchange	2021	Infrastructure	0	-1	-1	SolarWinds, NotPetya
Twitch data breach	2021	Infrastructure	-1	-2	-1	Sony, Ashley Madison
Medi-Cal data breach/Russian ransomware	2023	Personal	-2	-1	-1	Anthem
Meltdown and Spectre	2018	Political	-1	-3	-3	OPM, SolarWinds

The process for choosing validation data is the same as was used to select primary data, detailed in Chapter 3. First, the event must be recent – within the past 10 years. This criterion was originally designed to establish a relatively stable policy landscape for the analysis, but neither technology nor policy are truly stagnant for long. Some key policy changes have occurred in the U.S. federal government since this effort was undertaken, and the impacts of these changes is evident in some of the scoring. These impacts have been called out specifically in the following analysis.

Second, there must be enough publicly accessible data about the event for a robust analysis. Multiple credible sources for each event are necessary to build a defensible description and subsequent scoring. Third, the event must be in some way important to

the cybersecurity community. I revisited Joe Uchill's crowdsourced list (referenced in Chapter 3) for three of the validation events that fit the other criteria (Halfnium, Meltdown/Spectre, and Planned Parenthood). The other two events (the MediCal and Twitch data breaches) were selected for their prominence in community discussions and some key similarities to events from the primary data set. Finally, the event must have policy impacts at the U.S. federal level. While this list includes a diverse set of targets, from software to hardware, from personal health records to the finances of social media influencers, each is relevant to aspects of U.S. federal policy.

As in the prior section, the validation events will be briefly described, then scored according to the rubric of Chapter 3. These scores have been assigned without referencing scores of the primary data, however, so that coherence in scoring may be properly assessed. The primary data scores were assigned several months prior to the validation effort, and the time between these efforts helps ensure that the rubric is being accurately applied, rather than the similar scores repeated from memory. (A poor memory is as close as one can get to a blind comparison of a single person's assessments.)

After the validation cases have been scored, these scores are compared with the analogous events from the primary data set. Most scores are similar, indicating that STL is robustly applicable outside the initial list of events. In one case (the Halfnium attack on Microsoft Exchange), there is a notable difference in scores, but these differences are easily explained by changes in federal policy (CISA) that have been enacted since the initial event.

## Validation Case 1: Planned Parenthood

In October 2021, the Los Angeles Planned Parenthood was hit by a cyber attack that impacted 400,000 individuals. The attack included the same variant of ransomware that was used in the Colonial Pipeline attack, but data were exfiltrated from the network in advance of the ransomware's deployment. The attack compromised sensitive personal and medical data, including names, procedures, prescriptions, and other medical data. Given the ransomware aspect, this event may have been financially motivated rather than politically motivated, although to date, neither the attacker nor their motivations have been discovered. A prior data breach in 2015 was executed by critics of the organization, who posted names and email addresses of Planned Parenthood employees online. (No patient data was exposed in the 2015 event.) (Schaffer, Marks, & Knowles, 2021)

The incident scores a -3 for personal impacts. This scoring is coherent with the previously scored data. Planned Parenthood hack shares some similarities with the Ashley Madison and Anthem breaches. Like these incidents, the attack compromised personal data beyond credit card numbers or social security numbers, potentially putting individuals at risk of targeted harassment in addition to identity theft. Additionally, the attack may have been politically motivated, similar to the ALM breach, which was carried out by activists targeting the site's employees.<sup>28</sup>

This incident scores a -1 for infrastructure. Unlike the Ashley Madison event, Planned Parenthood did not benefit from the breach, but to date there has been no evidence that the data have been used for malicious purposes. While this breach is

---

<sup>28</sup> Descriptions of validation events were written with the assistance of ChatGPT (March 23, 2023 version), with significant guidance, as the events in question happened after ChatGPT's primary training period.

concerning, a score of -2 or -3 can't be reasonably justified without evidence of specific harm to the organization. It is, however, important to acknowledge that organizations with smaller budgets (like Planned Parenthood) often don't have the requisite funding to pursue incident response to the same degree as a larger organization (like Ashley Madison).

This incident scores a -1 in the political category. Again, without evidence of specific harms to national security, a more negative score is not defensible. Per the rubric, however, any data breach may contribute to the aggregate knowledge available to a foreign adversary.

The independently assigned scores are consistent and explainable in each category, which adds confidence to STL's utility.

#### Validation Case 2: Halfnium Attack on Microsoft Exchange

In early 2021, Halfnium (assumed by Microsoft security to be a Chinese state-sponsored hacking group) took advantage of multiple existing Microsoft Exchange vulnerabilities to install malware on hundreds of thousands of machines across the globe. While larger groups were able to address the issue relatively quickly, smaller-scale or rural organizations faced a greater burden of patching. The attack highlighted the importance of timely patching and the challenges faced by smaller organizations with limited resources. (Duffy, 2021)

In response to the attack, the White House formed a task force to address the issue and coordinate a broad communication effort involving the federal government and the Cybersecurity and Infrastructure Security Agency (CISA). These concise and technical

recommendations were considered an unusually public response from the federal government. The CISA emergency directive was issued promptly and provided specific, actionable steps for federal agencies to take to mitigate the threat (CISA, 2021).

The Halfnium attack scores a 0 for personal impacts, as no individual harms have been recorded to date from this event. This scoring is coherent with two similar events from the primary data set, NotPetya and SolarWinds.

This event scores a -1 for infrastructure impacts, due to the speed at which the issue was resolved. Similar to the NotPetya malware, the attack was widespread, and not relegated to a single industry. This incident also resembles the SolarWinds attack, since it was spread broadly from a centralized software provider. Unlike SolarWinds, however, Microsoft addressed the problem quickly. As a result of these efforts, Microsoft announced on March 22, 2021, that the exploit had been patched or mitigated on 92% of Exchange servers. Minimizing the time available for attackers to exfiltrate data or pivot within a system is a key part of incident response, and explains why the Halfnium incident scores lower on infrastructure impacts than either NotPetya or SolarWinds.

This event scores a -1 for political impacts. The U.S. did not overtly pursue the individuals responsible for the attack, which was the reasoning behind the +1 scoring for NotPetya in this category. The government did create a clear, coherent, and rapid response to the event, in partnership with the affected company, which resulted in a much more effective remediation than the Solar Winds event. This is an example of how changes in the policy landscape – namely, an effective and empowered CISA in partnership with an impacted industry partner – can mitigate harms. While CISA was stood up in 2018 (elevated from a lower-level directorate within DHS), the agency was

tangled in a bitter feud with then-President Donald Trump on the topic of election interference and disinformation. Under the subsequent presidential administration, CISA has been able to pursue its goals of securing critical infrastructure more effectively.

### Validation Case 3: Cl0p Breach of Santa Clara Family Health Plan

Defining an “event” for the purposes of this analysis is not always straightforward. For example, in 2023 a ransomware gang known as Cl0p claimed to take advantage of a single software vulnerability and used it to allegedly breach more than 130 organizations (Gatlan, 2023). A diverse set of victims were impacted, including (among others) U.S. consumer goods company Procter and Gamble, a U.K. pension fund (Starks, The latest mass ransomware attack has been unfolding for nearly two months, 2023), and a California-based medical provider that offers services primarily to low-income residents (Greschler, 2023). The “event” could be defined as the initial vulnerability - a flaw in a file transfer program used by organizations from different sectors across the globe. In this sense, the 2023 Cl0p mass ransomware attack looks like Solar Winds or NotPetya, in that it is a single vulnerability that was exploited at scale, enabled by widespread usage of the software. Each individual victim, with its own impacts and outcomes, could also be defined as an “event”. These approaches are complementary, and both are valid. Since the prior validation case was comparable to Solar Winds and NotPetya, we will complement this case by instead examining one of the organizations impacted by this attack – the medical provider in Southern California.

The U.S. Department of Health and Human Services (HHS) released an alert on the topic of the ClOp ransomware activity, noting that *“healthcare is particularly vulnerable to cyberattacks, owing to their high propensity to pay a ransom, the value of patient records, and often inadequate security”* (U.S. Health and Human Services, 2023). The Santa Clara Family Health Plan is part of the California state-level version of Medicaid, providing medical insurance to low-income residents. The data of nearly 280,000 patients were assumed to be compromised in the attack, including names, social security numbers, contact information, and insurance credentials. It isn’t yet clear whether or not this information has been used for identity theft, insurance fraud, or other activity, but patients have been instructed to sign up for credit monitoring and to change their passwords on Medi-Cal systems.

This event scores a -2 in personal impact, as health care data cannot easily be tracked through credit monitoring, nor can it be changed, like a credit card number. The Mercury News, a local Southern California news outlet, interviewed affected individuals, who did not learn about the breach until a month after their data were exposed (Greschler, 2023). Many were frustrated by the vague and inadequate guidance from the company, and anxious about how their family data were being used. They reported spending significant time attempting to remediate the issue, which is an additional burden for lower income families.

This event also scores a -2 in infrastructure impact. As mentioned by the HHS alert, patient records are financially valuable, as they can be used for a variety of fraudulent activities and are not easily replaced. While not an existential threat to the



organization, the provider will have to upgrade infrastructure and processes, as well as continually monitor for adversarial usage of the stolen data.

While there are no direct national security impacts of the Santa Clara Family Health Plan breach, any data breach can be used by a foreign intelligence service to gain insight, so this event scores a -1 on political impacts per the rubric in Chapter 3. These scores align well with the Anthem data breach, which is the closest analogue in the primary data set. The Anthem event scored a -3 in infrastructure impact, primarily due to the large number of medical records exposed (80M, or 285 times more than the records exposed from the Santa Clara Family Health Plan). Otherwise, the scores are identical. The matching scores for similar events gives credibility to the STL framework, as the scoring for validation was assigned without reference to the primary data set, and only compared after the assignment.

#### Validation Case 4: Twitch data breach

The beauty of the internet is that it allows people to connect with one another over shared interests, regardless of their geographic location. Social media platforms must collect and manage personal data to foster these connections, which means they will always be a target for breaches. Add the ability for users to monetize their personalities and achieve a level of fame based on secretive algorithms, and the incentives to steal data become even more appealing. Such was the case for Twitch, a video streaming service that primarily caters to the video gaming community. Like many social media platforms, Twitch has had issues with content moderation and harassment, specifically “hate raids”

where users (many of whom are actually bots) flood a streamer's channel with incessant vitriolic messaging to drive them off the platform (Parrish, 2021).

In October 2021, an anonymous user posted a large data file to the website 4Chan, claiming that it was the first part of a Twitch data breach. A note accompanying the data dump referred to the Twitch community as "a disgusting toxic cesspool" and claimed the leak would "foster more disruption and competition" in video game streaming (Tidy & Molloy, 2021). The file included more than user names and login information; it also contained the site's source code, security information, product development plans, and the amount of money that every streamer had earned since 2019 (Browning, 2021). The data posted was confirmed by several sources to be real.

A number of "employees" (in Twitch's case, the streamers) had their privacy violated, but the site's users (over 30M in the U.S.) were not impacted (Clement, 2022). For this reason, the event scores a -1 for personal impact. The Sony Pictures Entertainment breach is a similar event, in that a company's secrets were published openly in order to cause harm to the organization, rather than for financial gain, but the goal was to cause embarrassment to a few, rather than exploit the many.

There is a reasonable comparison to the Ashley Madison breach, in that the attacker is calling out objectionable and harmful behavior by engaging in objectionable and harmful behavior themselves. A significant amount of information was spilled on the company's internal operations. The exposure of future product plans, security information, the site's source code, and payment information for streamers all impact the

company's posture negatively. Any competitive advantage for upcoming products is minimized, adversaries now know how the company protects its infrastructure, and top streamers can easily move to other platforms. While not an existential threat to Twitch, the data exposed were significant, so the event scores a -2 via the established rubric.

There is no overt impact to national security of these specific data, but per the STL framework, any breach has the potential to be exploited by a foreign adversary who can combine the information with existing data to extract additional insights. The event scores a -1 in political impacts for this reason. These scores are coherent with those from the primary data set.

#### Validation Case 5: Meltdown and Spectre

Software vulnerabilities have been the focus of the analysis to this point, but computing hardware can also be susceptible to attack, and these issues can be as widespread as those found in software.

The Meltdown and Spectre vulnerabilities affect central processing units (CPUs), which allow an unauthorized third party to read information - such as passwords or encryption keys - that normally would never leave the machine's hardware (Hautala, 2018). These vulnerabilities are particularly concerning, as they are not relegated to a specific type or brand of CPU, and therefore an astonishing array of devices and machines are exploitable. In fact, Meltdown was given its name because it "melts" expected hardware security measures, and Spectre because these kinds of attacks are impossible to detect and "will haunt us for quite some time" (Graz University of

Technology, 2018). The security researchers who created the first proof of concept initially believed that their findings were a mistake, because the implications were both so wide-ranging and serious (King, Kahn, Webb, & Turner, 2018).

These issues were initially discovered by academic researchers and Google's 'Project Zero', a group dedicated to finding zero-day (or previously undiscovered) vulnerabilities. The vulnerabilities were privately disclosed to the chip manufacturers, as well as software developers and cloud computing companies, who worked together in secret to develop a fix. This kind of process is common, as it allows the companies to develop and deploy software patches before the public disclosure, mitigating the potential for exploitation of unpatched vulnerabilities. However, one of these patches contained such specific notes that other independent researchers were able to identify the root problem, which unfortunately accelerated the public release before many of the patches had been deployed (Ars Technica, 2018).

Disclosure is a complicated issue, even more so when the vulnerable technology is connected to a wide variety of organizations (Manion, 2018). The impacted CPUs were not just built into devices in the U.S., but also into products manufactured by foreign companies. Intel's pre-public outreach included some Chinese tech companies, but did not include the U.S. government, who were not aware of the issue until the public disclosure, seven months later (McMillan, 2018), (Ng, 2018).

Given the widespread nature of the issue, and the covert nature by which sensitive data could be exfiltrated, one can easily imagine the value of these vulnerabilities to U.S.

near-peer adversaries for intelligence collection and espionage purposes. Several months of advance notice would give a well-resourced, technically sophisticated adversary sufficient time to exploit these vulnerabilities. While there has been no direct confirmation that Meltdown or Spectre were exploited, threat researchers do not always expect to find concrete evidence of foreign intelligence collection, as opposed to other types of adversarial activity such as hacktivism, ransomware, or identity theft (Schukai, 2023). For this reason, the Meltdown and Spectre issue scores a -3 in political impacts.

Meltdown and Spectre disproportionately impact organizations that are running multiple CPUs, particularly those that are using them to create “virtual machines” that in theory are kept separate. These attacks would allow a malicious user to pivot from a single virtual machine in a cloud computing service to the system that runs all the virtual machines in the environment, providing a much broader surface for potential attack and data exfiltration. Additionally, the Meltdown and Spectre vulnerabilities have several variants, each of which requires its own remediation. Many of these patches have a negative impact on processing speed, or have caused system instabilities such as unintended rebooting. Others require significant adaptation of the software (Ars Technica, 2018). For these reasons, the organizational impacts score is -3.

While individual machines (laptops, smart devices, and phones) are also vulnerable to Meltdown and Spectre, they are a much less complex environment than a cloud computing provider. A single patch to address the specific hardware is a sufficient solution for individuals. Beyond the standard security practice of installing recommended

software patches, there is no additional burden on the individual user. The individual impacts, therefore, score a -1.

Comparing with the primary data set, this event matches closely with the Solar Winds event – a vulnerability in commonly used technology that transcends a particular industry or sector, the impacts of which are significant in the organizational and national security realms but relatively minor to individuals. Solar Winds scored a 0 in personal impacts and Meltdown/Spectre scores a -1, since individuals do not tend to use IT management software but they do have several CPU-powered devices. Otherwise, the rankings are the same.

One might also compare the Bowman Dam incident with Meltdown/Spectre, since both events include impacts in the physical world. Personal data are nebulous, they exist in multiple places at the same time, but the Bowman Dam and the population of CPUs are physical objects, vulnerable to attackers at a distance due to the interconnected nature of the internet. The Bowman Dam scored a 0 for personal impacts, as the attack was unsuccessful. That score would have been different if the attackers had been able to control the dam's operations. Instead, like Meltdown/Spectre, the burdens of the attack are substantial for organizations and national security, requiring significant action at both levels to prevent future harms.

### **Contributions to Knowledge**

In this chapter, a new framework – the Socio-Technical Lagrangian - was created to enable innovative perspectives on cybersecurity events and the sociotechnical context

in which they occur. This framework was inspired by the Lagrangian Mechanics concept from physics, which identifies the most important features of a physical system and creates a coordinate system based on these features. The socio-technical analogue is similar in that it is built on important features from a policy perspective – harms at the personal, infrastructure, and political levels – rather than conventional measures, which are appropriate for incident response but have been shown to be misleading in a policy context.

A diverse and robust set of cybersecurity events were analyzed using STL, with scores assigned to describe harms at the various levels. Groups of events were then created, based on high impacts in certain categories, or high variance in scores. The insights of this method come from these comparisons - not in the numerical assignments, but in framing events from a different perspective that allows one to gain broader intuition and pose more useful questions, such as “what imbalances exist in the system to cause this effect?” For example, the high variance cases identified by STL showed that some policy approaches (such as FTC orders) may unintentionally act as security disincentives. Analysis of these cases also showed that policy to address personal harms will also need to address infrastructure impacts if it is to be broadly successful. The high average impact cases highlight some of the challenges in creating effective policy at the infrastructure level. Infrastructure-based events often have disparate outcomes in personal and national security impacts - an effective policy approach aimed at the infrastructure level must be capable of addressing personal and national security effects as well.

The cybersecurity problems described in this chapter are not new, but the STL framework developed in this chapter provides a new way of analyzing the true impacts of

these problems and understanding how they might be mitigated. Designing effective policy is a challenge, particularly when it comes to implementation. But addressing harm reduction coherently, at multiple levels, will ensure that no group is sacrificed for the benefit of the other. This balanced approach will hopefully also result in guidance that is easier for individuals, companies, and government organizations to follow, “lifting all boats” and improving the cybersecurity posture for all.

The next chapter uses the new STL framework to address current policy issues, as set forth in the 2023 White House National Cybersecurity Strategy. This document from the U.S. executive branch will be analyzed through the lens of STL, and suggestions provided for effective implementation.



## CHAPTER 5

### **POLICY RECOMMENDATIONS**

In the previous chapter, the Socio-Technical Lagrangian framework was built using a data set of historical cybersecurity events. This framework was also validated using a separate data set of existing cases within the predefined scope. The purpose of building STL was not simply to analyze historical events, but to generate abstraction-level insights that will help proactively identify areas for potential policy improvement, and meaningfully evaluate courses of action. The key insights included:

- A. Harms at the personal, infrastructure, and national security levels are insightful “features” to assess, as opposed to number of records, financial loss, and attribution. These new features provide the basis for a more insightful inquiry.
- B. Some intended incentives in the current system (such as certain actions by the Federal Trade Commission) may in some cases incentivize against meaningful security measures.
- C. For all three levels of harm, emphasizing security in the design phase is important. Attempting to secure a system after it has been created is much more difficult and can create unintended negative consequences.
- D. Vulnerabilities in centralized infrastructure services, such as cloud storage, file sharing, or business-to-business software, can have negative global impacts at all three levels. Security improvements to these sectors will positively impact all three sectors.

- E. Policy approaches must address all three levels of harm in some manner in order to be truly effective.
- F. Lessons can be learned from patterns of small to medium-scale threats. To minimize harms, implement these lessons before these patterns grow to national or global scale.
- G. Without intentional and proactive policy measures, history will repeat itself.

Now that the STL framework has been built and tested, and insights derived regarding U.S. federal cybersecurity policy, a forward-looking analysis can be undertaken. The primary object of this analysis will be the 2023 National Cybersecurity Strategy. The aims of this chapter are to use the knowledge generated via STL to evaluate the suggested courses of action and to provide recommendations for effective implementation.

### **The White House National Cybersecurity Strategy**

The purpose of White House strategy documents generally is to encapsulate the executive branch's intent for policy development and assign responsibility where appropriate. These documents do not appropriate funding or create legislation, but they inform relevant stakeholders (e.g. federal offices and agencies, the legislative branch, and industry) of the executive branch's priorities and intentions. The concept of a National Cybersecurity Strategy is not new – the first was produced by the George W. Bush administration in 2003 (Lemos, 2003). Neither Bill Clinton's nor Barack Obama's administration produced a separate National Cybersecurity Strategy, but both addressed cybersecurity issues in other strategic documents. The Trump White House created a

separate cybersecurity strategy document at the behest of Congress, which incorporated many efforts of prior administrations (Starks, 2018).

As discussed previously, the U.S. federal cybersecurity policy landscape is rife with gaps and conflicts, a result of the “pacing problem” wherein technology is developed much more quickly than the regulatory structures that govern or influence it (Marchant, 2011). The most recent National Cybersecurity Strategy as of this writing, released by the White House in March 2023, is an effort to close some of those gaps and remove or ameliorate misaligned incentives (National Cybersecurity Strategy , 2023). The Biden cybersecurity strategy was built in concert with the 2022 National Defense and National Security Strategy documents; the three are coherent in their recommendations and build on existing policies and initiatives (some from prior administrations). The Congressional Research Service review of this Strategy identifies common objectives across a wide range of additional federal-level cybersecurity efforts as well (Congressional Research Service, 2023). The 2023 document assigns the overall responsibility for implementing this strategy to the Office of the National Cybersecurity Director, a position established by Congress in 2021 (116th Congress of the United States, 2020). However, new funding to accomplish these objectives would require congressional approval (Jaikaran, 2023) and some recommendations fall outside the authority of the executive branch and would require Congressional action (such as the creation of new federal-level legislation).

Many elements of White House cybersecurity strategy remain remarkably similar across the years and the various administrations. For example, emphasizing the

importance of public/private partnerships, or recommending additional information sharing programs, have been consistently featured in the past and are part of the most recent document as well. The Biden-Harris plan of 2023 does have some key differences from its predecessors – what it calls “*two fundamental shifts... to change those underlying dynamics that currently contravene our interests.*” The first major change is to realign defensive responsibility from individual users to organizations. As this analysis has also noted, there is currently an overreliance on end users to protect large-scale cyber systems. Per the White House document, the inadvertent actions of an individual “*should not have national security consequences*”. The second shift is to realign incentives for long-term security outcomes. In a framing reminiscent of the “objects and forces” analogy of Chapter 3, the strategy aims to “*focus on points of leverage*” to maximize positive outcomes with minimal interventions. These are much needed realignments of the overall policy perspective. But the strategy’s most specific policy departure is the explicit call for regulation of critical infrastructure – until 2023, White House recommendations had only recommended voluntary measures for compliance (Nakashima & Starks, U.S. national cyber strategy to stress Biden push on regulation, 2023).

Overall, the 2023 National Cybersecurity Strategy is a clear and straightforward document that specifically addresses current cybersecurity concerns. The premise of the Strategy aligns well with the objective of this research agenda – to “*realign incentives to favor long-term investments in security, resilience, and promising new technologies*”. The document is organized into five pillars around which the White House seeks to build

collaboration and consensus, and a group of Strategic Objectives supports the aims of each pillar. At this more granular level, however, some of the recommendations are vague (particularly when it comes to implementation) and a few seem misguided. The next five sections will address each pillar, using STL to analyze proposed approaches and desired outcomes, discuss the aforementioned points of leverage, and predict the outcomes of applying pressure to these points.

### **Pillar One: Defend Critical Infrastructure**

Of all three categories of primary harm – individual, infrastructure, and political – the STL analysis shows that infrastructure has the widest impact. All but one of the infrastructure events in this data set ranked in the most negative 50% of average impacts. In other words, impacts to infrastructure strongly affect individuals and national security as well. The significantly negative average impacts of this category suggest that policy efforts focused on infrastructure may be the most effective way to positively impact the system as a whole. The name of the first Strategic Objective - SO1.1, ‘Establish Cybersecurity Requirements to Support National Security And Public Safety’ – reinforces this idea.

The first Strategic Objective also notes that market forces currently disadvantage proactive cybersecurity measures. Companies are penalized either by the costs of security, or by being beat to the market by less secure options. In the Chapter 4 discussion on the Ashley Madison breach, a new policy was envisioned that would induce federal intervention after a data breach, which would assign a cost to actual security deficits - either removing the potential for profit or delaying it until security

issues have been remediated. This is one way to partially rebalance some of the market forces. But it does not address all misaligned economic incentives. Policy like FTC action, which penalizes companies for not achieving stated levels of security (e.g. Ashley Madison) also may be undermining meaningful and proactive security measures. Economic forces are important to consider, and they are manifest in many different ways in the system. For optimal results, the implementation of this policy should address these holistically.

This first Strategic Objective goes on to say that the Administration will “encourage and support” companies that exceed cybersecurity standards. This is an interesting idea, but there is little information on how it would be accomplished. How would these companies be evaluated and identified? It’s unclear what “exceeding” standards means – although STL could be valuable in defining these metrics. If one considers the baseline for an infrastructure organization as a ‘0’ in the framework, one way to exceed the standard may be to improve scores in the individual and national security categories as well. The intention is clear and achievable at the infrastructure level, and incentives are also put in place to improve outcomes for individuals and national security as well. To implement this part of the strategy, however, one would need a repeatable and reliable way of ranking organizational security posture, rather than events, which tend to be security failures.

In the spirit of shifting incentives for long-term security outcomes, this policy must also address the fact that smaller organizations often don’t have the same funding available to pursue security measures as larger ones - yet these smaller organizations should have the same opportunity for incentives. Recall that the Target breach occurred

when attackers pivoted from a compromised third-party refrigeration contractor – a much smaller organization. Attackers often look for these smaller targets (no pun intended) that are easier to infiltrate and can be used as a pivot point for a larger-scale operation. This adds an additional challenge to the construction of metrics and evaluation – they must be relevant to a highly diverse set of organizations.

As one might expect, the first pillar strongly recommends that industry and government work more closely together. The nature of infrastructure in the U.S. is a blend of government entities, public and private companies. To make meaningful progress on infrastructure security, policies will have to be created with the incentives of this heterogeneous group at the core. One positive example is the mitigation of the Halfnium attack on Microsoft Exchange, wherein CISA provided fast, actionable, and direct guidance to impacted groups in industry and elsewhere in government. This relationship would be even more powerful if this approach could be scaled to include a wider range of vulnerabilities, attack vectors, and organizations.

Strategic Objective 1.2, ‘Scale Public-Private Collaboration’, addresses one hurdle of public/private partnerships by recommending a machine-to-machine data sharing program by which threats can be shared between companies. The “machine-to-machine” part is unique and important – as discussed previously, companies are often reluctant to share information directly about their security posture with other companies, as it may reveal competition-sensitive information. But building an extensive library of data, including “indicators of compromise” and specific behaviors from threat actors from multiple companies, would be beneficial to all. In this case, the potentially sensitive information could be aggregated and disseminated via “bots”, which could also protect

systems based on this new information more quickly than humans. This may be one way in which a technology solution to a policy problem could be the best choice, and might create new and powerful incentives for cooperation. As discussed in Chapter 4, security issues will never disappear completely, but they can lead to broad positive change if the proper incentives exist.

The prior Strategic Objective addressed business-to-business incentives. Strategic Objective 1.5 – ‘Modernize Federal Defenses’ – starts addressing government-to-business incentives. This section proposes that federal cybersecurity advance to the level where it can serve as a model for industry. The proposed execution for this goal – removing all legacy systems that are incapable of sufficient modernization within ten years - inadvertently highlights how far the federal government needs to advance before they can set an example for industry. But rather than simply dismissing this goal as unrealistic, it is more productive to reflect on the significant positive potential it represents. The analysis of Chapter 4 showed that impacts are frequently transmissible between the three levels of harm. Upgrading the security standards of U.S. federal systems will positively influence any industry that wants to do business with the federal government. (This incentive is also promoted in Strategic Objective 3.5, ‘Leverage Federal Procurement to Improve Accountability’.) From this perspective, the top-level goal is admirable and should be pursued. To realize the lofty impact of this goal, the various sectors of government/industry collaboration should be assessed via STL, and modernization should be prioritized by those sectors that currently exhibit the greatest vulnerability to cross-sector harms.



## **Pillar Two: Disrupt and Dismantle Threat Actors**

The strategic objectives in Pillar Two focus on disrupting adversarial activity and removing incentives to engage in cybercrime. Public/private coordination also features heavily in this pillar, for information sharing between the private sector and government and disrupting active adversarial campaigns. The language of this section is strong, calling out the need to mobilize “all elements of national power to counter the threat”, including law enforcement, financial infrastructure, and diplomatic efforts, as well as both cyber and kinetic military action.

The goal of Strategic Objective 2.1 is to add friction to “disruption campaigns” such that they are not as profitable as they are currently. This part of the strategy is correct in that threats cannot be stopped, but they can be disincentivized. However, this objective focuses solely on financial incentives, and it is clear from the prior analysis that threat actors are not always motivated by profit. STL can be used to illuminate these additional motives and help balance the system against them as well.

As mentioned, information sharing programs promoted by the federal government have not always been welcomed by the private sector. Strategic Objective 2.2 encourages private sector partners to work with nonprofit third-party intelligence sharing organizations, which will then connect with federal partners. In addition to the machine-to-machine option set forth in Pillar 1, industry may find this option more appealing than providing information directly to the government.

The private sector has important insights on the threat surface, but the federal intelligence community also has uniquely powerful methods to understand threats and the groups that perpetrate them. Strategic Objective 2.3 acknowledges that the speed at which this information is shared with industry is insufficient and commits to improving the process. One recommended improvement in this section is creating methods by which classified intelligence can be shared appropriately, by expanding clearances to certain programs for critical infrastructure operators. This option should be very carefully considered from an STL perspective. Information is often classified because of the ways and means by which it was collected, and methods could be derived from the intelligence information itself. Information regarding these methods should be restricted to a minimal audience, based on “need to know”, lest they be leaked and rendered ineffective, and the intelligence posture of the U.S. compromised as a result.

Strategic Objective 2.5 focuses solely on cybercrime and ransomware. As discussed previously in this dissertation, ransomware has evolved from an individual nuisance to a national security issue. This objective recommends that victims not pay the ransom, in an effort to make ransomware less profitable, but it doesn’t actually prohibit the payments. While paying a ransom does unfortunately incentivize ransomware (particularly ransomware-as-a-service operations), there are often more pressing incentives for the victim to pay the ransom and recover operations quickly. Ransomware actors have generally honored these “agreements”, releasing an organization’s data or system once the ransom has been paid. (They want their methods to remain profitable!) However, one emerging trend in ransomware is for malicious actors to now request

multiple payments of their victims – one for decrypting the information, and another for keeping the stolen data private (Chapman, 2023). This is likely to drive the cost of the ransomware epidemic up even further, adding extra incentives for organizations to pay. Multiple ransoms for a single incident could also skew the cyber insurance market, limiting coverage to a single payment per incident, leaving targeted companies open to significant financial risk.

Ransomware is a significant problem, no longer relegated to one or two levels of harm, and it will likely be so for the foreseeable future. If approached carefully, though, significant improvements to the system could be made which aggregate to a stronger security position for all levels. Perhaps security improvements to election infrastructure could be a model for these kinds of changes. Election systems are more robust due to investment and cohesive action at local, state, and federal levels. There are still issues with election security systems that need attention, but overall, the security of U.S. elections has been significantly improved. Could resilience to ransomware be established through a similar trajectory?

Finally, it is notable that both election security and ransomware issues existed for years at a smaller scale before growing to the national security level. Defeating (or at least mitigating) ransomware impacts is important to the national and global economy, and to security at federal, infrastructure, and personal levels. There may be other threats which have not yet hit critical levels but are on a trajectory to do so. These problems may be identifiable through analysis of small to medium-scale threats, and mitigated before they reach the critical broad-scale impacts of ransomware.

### **Pillar Three: Shape Market Forces To Drive Security And Resilience**

The National Strategy notes in the ‘Emerging Trends’ section that the *“quantity and intimacy of personal data is growing rapidly... malicious activities are disproportionately impacting those without the resources necessary to protect themselves, recover, or seek recourse”* (pg. 2). The last part of the statement - “seek recourse” – is particularly interesting. This phrasing elevates the position of individuals and seeks to empower them beyond the usual “patch your systems and make good passwords” recommendations that are commonplace but incomplete.

To that end, Strategic Objective 3.1 is “Hold the stewards of our data accountable”, and security by design is a critical part of that accountability. Consider the example of Microsoft PowerApps from Chapter 4. The data exposure from this event was due to a software setting that made the data public by default, not an attacker who had to exploit a vulnerability to breach the perimeter. Public accessibility was likely chosen as a default setting because it made building applications easier, which is an important business consideration. However, the impact of that choice on individuals whose data would reside in these applications was not considered during the software’s design phase. The first step of accountability for data stewards is creating an “opt-in” vs. “opt-out” setting, providing individuals with some control over how their data are used by companies, for what purposes, and for how long they will be stored. From a technical perspective, this is not particularly challenging, but from a policy perspective, there are no current incentives to do so. Control over data is perceived as a competitive advantage;

implementing this strategic objective would likely require top-down legislation from the federal government.

The National Strategy acknowledges this reality, stating that “the administration supports legislative efforts” in this area. Whether or not Congress will be able to come together and create meaningful legislation is another question. After all, Congress has held hearings for years on this topic, bringing tech CEOs to contentious and highly visible hearings on the Hill for a few days before allowing them to go back to their business, meanwhile refusing to implement meaningful federal legislation to actually protect user privacy. Data brokers (such as Equifax and Experian) and social media companies (most notably Facebook, but also TikTok) have been called to hearings to account for their lax security and privacy standards, but at the end of the day, no harms have been incurred by the companies other than a few unpleasant hours of questioning by legislators with a debatable understanding of technology, and no new federal-level legislation to protect privacy has emerged.

Strategic Objective 3.3 aims to ‘Shift Liability for Insecure Software Products and Services’ to companies rather than end-users. A formal legislative approach to establish liability for insecure software is encouraged in this section. The concept of liability implies the existence of harm, which can manifest in different ways. Per STL findings in Chapter 4, infrastructure-related events (such as SolarWinds or PowerApps) have a wide range of outcomes – therefore, an effective policy approach must be capable of addressing individual, corporate, and national security simultaneously, as well as

usability of the software. But without a well-balanced approach that considers these multiple angles, the implementation of this objective may backfire.

One potential unintended consequence: to avoid liability, companies may subject users to an untenable web of Terms of Service and End-User License Agreements that only allow the software to be used if individuals waive the rights that this objective seeks to establish. According to one study, the average software user would need 40 minutes per day, every day, to read all the Terms of Service and end-user agreements that they are presented with in a year (Vedantam, 2016) This is already a significant burden that most people don't have the time or expertise to engage in. One experimental survey (n = 543) showed that 98% of participants agreed to Terms of Service that asked for their first-born child as payment for access to a job-sharing website (presumably because they were not reading the Terms, not because they want to divest from their children) (Obar & Oeldorf-Hirsch, 2018). Not all Terms of Service contracts have been found to be legally enforceable, but that doesn't mean they can be safely ignored, as there are also several cases of company's terms being upheld by the courts, even though individual plaintiffs reasonably claimed that the terms were vague, confusing, or hidden (Neuberger J. D., 2020). For this Objective to be realized, policy must be created to intentionally protect individual interests, rather than relying on the courts.

As mentioned in Pillar 1, Strategic Objective 3.5 seeks to leverage the federal government's position in the marketplace to drive accountability. This Objective is complementary to Strategic Objective 1.5 - the "carrot" is the financial benefit of working with the federal government, and the "stick" is the Civil Cyber-Fraud Initiative

(CCFI), announced by the Department of Justice in October 2021 (Department of Justice, 2021). The CCFI is empowered to leverage the False Claims Act to pursue civil action against companies who receive federal funds and who fail to adhere to cybersecurity requirements. Individuals are allowed to bring suit on behalf of the government and can receive the damages if they are awarded (André, Ferber, & Alexander, 2021). The specific cybersecurity requirements have not yet been defined, and Deputy Attorney General Lisa Monaco has asked Congress to create a national standard for reporting incidents, focusing on the damage that ransomware has inflicted on global infrastructure (Monaco, 2021). This new organization may provide incentives which are complementary to the FTC's. The FTC can pursue action against a company for making false claims about security to customers, but the CCFI can now bring civil action against companies that fail to actually meet certain standards. If the requirements are reasonably defined, this action may help balance out the currently misaligned incentive of the FTC in that the CCFI can address the implementation of security, not just how it is advertised.

Strategic Objective 3.6 assigns the responsibility for investigating a federal cyber insurance backstop to the executive branch. This section is one of the least detailed in the document, saying simply that the Administration will “assess the need for and possible structures of a Federal insurance response” to large-scale cybersecurity events. Unlike a flood or a fire, a cyber incident has the potential for global-scale cross-sector impact, which means cyber insurers are either left open to potentially devastating risks, or must protect themselves by creating exclusions that move exposure back to the policyholder (Shokrai, 2023). As discussed previously, several insurance companies attempted to

claim an “act of war” exclusion for damages related to the NotPetya malware. These claims were taken to court and ultimately refuted, but the point remains: if a large portion of policyholders in an insurance company’s portfolio are impacted by a single event (like NotPetya), how can a company remain viable and continue to underwrite its other policies? To maintain critical infrastructure services in a blended public/private system (such as gasoline distribution), the operational environment and risk surface need to be considered, which includes insurance carriers. This Strategic Objective is vague, but a better understanding of the federal government’s role in cyber insurance could pave the way for additional balance in the current system.

#### **Pillar Four: Invest in a Resilient Future**

Pillar Four lists several areas in which the White House recommends specific investments in order to improve internet security. The first area, in Strategic Objective 4.1, is the foundation of the internet itself. This section starts with a sentence that is simply too good to paraphrase: “*The internet is critical to our future but retains the fundamental structure of its past.*” As seen in Chapter 2, many of the initial decisions in how the internet was built and governed are present today, in the form of recurring weaknesses or vulnerabilities in its structure (Fidler, 2017), (DuPont & Fidler, 2016). As more services and capabilities are built on top of the current internet, these issues not only persist but can manifest in new ways. Investments to improve existing functions are not as popular or exciting as those to create brand-new technologies, and maintenance projects are rarely incentivized by the market, but they will be necessary to improve the



state of cybersecurity and allow for innovation. Without these interventions, history will continue to repeat itself, as seen in the primary and validation data sets.

Strategic Objective 4.2 specifically calls for increased investments in fundamental research. The goal is to produce an innovation strategy that goes beyond industry, providing multiple viewpoints and approaches that are independent of the commercial market, resulting in a more balanced approach. Post-quantum security and secure clean energy are two specific topic areas recommended for federal investment. One hurdle for these recommended investments is that the executive branch does not have the power to directly fund them. Congressional action is required to allocate and appropriate funds, and in the current environment, Congress may not be able to reach consensus on these topics. For example, the National Strategy claims that “strong encryption is foundational” and should receive increased support, but for years Congress has been attempting to undermine encryption through a variety of means (Pfefferkorn, 2022). Those efforts have not yet been successful, but any attempts to establish federal funding for encryption would be subject to heavy criticism and most likely to fail.

Strategic Objective 4.5 recommends investment in a less contentious topic than encryption or clean energy: the “development of a trusted digital identity ecosystem”. The “ecosystem” isn’t clearly defined in this section, but the high-level goal is to stem the tide of data breaches and identity theft and, from the perspective of this research, reduce harm for individuals, businesses, and public benefit programs that rely on identity validation.

Without a definition or scope of this “ecosystem”, it is hard to evaluate whether or not this Objective could be successful. STL is one method by which the government could envision multiple definitions and evaluate which would likely result in the lowest overall harms. The right set of stakeholders could be convened to build an identity verification system that would work for a variety of use cases – a single digital identity that authenticates a person’s access to social security, medical records, and education (to name a few). Such a system would need to accommodate existing regulations regarding health data and educational data but must also have exceptionally rigorous security requirements from a technology perspective, to protect the various facets of identity combined in this system. Interoperability is a well-founded concern mentioned in this section – for a digital identity ecosystem to function, it will need to interface securely with a wide variety of programs in public, private, and nonprofit sectors – some of which are likely to be based on legacy software or hardware. The potential outcomes of such a system would be positive, but the technical, legal, and social challenges are all considerable, and the risks are high. If such a system were breached, it could cause immense harms at all three levels of the framework. Such an endeavor should be approached with great care.

### **Pillar Five: Forge International Partnerships to Pursue Shared Goals**

Pillar 5 contains a set of Strategic Objectives to scale collaboration methods between U.S. stakeholders to a global level. The research in this work specifically focuses on U.S. federal level policy, and while it is clear that an international cooperative approach is necessary, the data and framework of this dissertation were not crafted with

international engagement in mind. Future research on this topic would better inform analysis of this section and is encouraged in Chapter 6.

## **Implementation**

The National Cybersecurity Strategy of 2023 contains five primary Pillars and a total of 28 Strategic Objectives. Given the volume of recommendations, it is disappointing that the implementation section of the document is slightly more than a single page in length. This section assures the reader that a data-driven approach will be taken and the progress and outcomes measured, but does not address the metrics or data that would lead to meaningful assessments. A fundamental claim of this dissertation is that the wrong metrics are often used to discuss cybersecurity events and create policy, resulting in the obfuscation of harms and perpetuation of conflicting incentives. Without careful and intentional design of metrics, including the language used to evaluate them, the laudable objectives of this Strategy may fall to the wayside, as so many good intentions to reshape this space have done before.

The Office of the National Cyber Director in the White House is a relatively new position, established November 2021 (The White House, n.d.). This Office has been directed to work with government partners and publish an implementation plan for this Strategy. As of June 2023, the plan has not yet been released to the public and given the extent of the Strategy it must cover, one can reasonably expect it should take several more months to complete. A full implementation plan will require buy-in from a wide range of federal stakeholders, all of whom have areas of responsibility and incentive structures that drive their behavior. These structures may push towards or pull away from

the fundamental aims of the Strategy, and those “motions” may depend on circumstance. From the perspective developed in this work, one can see a new complex system evolving from within the existing complex system, one that will surely have influence in all three levels of impact defined previously. To ensure that these motions are positive, and no inadvertent harms occur in the system (for example, those posited in the section on software liability) – or even to improve the overall security and harmony of the larger system! - a framework such as STL, focused on impacts and harms, should be leveraged throughout, and revisited during implementation.

### **General Recommendations**

This chapter leveraged the 2023 U.S. National Cybersecurity Strategy as a specific example of how STL can be applied to a policy strategy. Key themes can be applied to cybersecurity policy more generally, such as:

- 1. Consider the system holistically when designing and implementing policy, to avoid unintended negative consequences.**

To circumvent negative consequences, one must be able to anticipate which actions will result in negative outcomes, and for whom. Anticipation does not guarantee positive outcomes, but is a key principle of responsible innovation (Maynard & Garbee, 2019). Projecting policy impacts at the three levels defined in this dissertation – individual, organizational, and national security – helps uncover and remediate potentially hidden risks.

- 2. No single policy will achieve all objectives – prioritizing is necessary.**

Broad-scale policy documents, such as the 2023 U.S. National Cybersecurity Strategy, are often ambitious and widely-scoped. This is an expected and appropriate feature of a federal-level document that will be used to design a variety of implementations. However, not all the objectives in such a document can be achieved simultaneously, due to the realities of time, budget, and competing priorities. Whether enacting federal-level legislation or setting internal company policy, it is necessary to set priorities. STL can help assess which policies will have the greatest positive impact in the short term, and avoid spreading efforts (and budgets) too thinly.

**3. Identify (and address) small or mid-scale trends that have potential to accelerate quickly.**

Recall that ransomware started as a relatively small-scale attack (from a financial perspective) against individuals, but evolved to impact companies, government organizations, and nations themselves. The advice to “think like a hacker” is often given to system defenders in order to help them imagine how a threat might evolve or a vulnerability might be exploited (Vigna, 2019). The hacker perspective need not be relegated to defending computer networks, but can be applied to less tangible concepts such as democracy, security, or privacy (Winterton, 2018). Part of this mindset is being able to envision how a threat might grow or be used in new ways. STL can help shed light on how threats could potentially impact different scales, and how policy could ameliorate those effects.

#### **4. Without intentional and forward-looking efforts, history will repeat itself.**

The unfortunate cycle of data breaches and resulting fatigue from impacted individuals was a primary motivator for this work. During the course of this research, dozens if not hundreds of additional breaches, attacks, and inadvertent data exposures occurred. It isn't enough to "patch and pray" (to borrow another information security phrase), particularly when human lives, societal cohesion, and democratic ideals hang in the balance – some fundamental course corrections must be made. This is part of the motivation behind Pillar 4 of the U.S. National Cybersecurity Strategy, the emphatic recommendation to fund research and development on topics pertaining to current and future cybersecurity. The objectives around secure technology by design in Pillar 3 also support a fundamental change in the status quo. Generally speaking, disrupting the system can have long-lasting benefits, as long as these changes are done with a wide variety of stakeholders in mind, and potential ramifications identified across the system as a whole.

The 2023 U.S. National Cybersecurity Strategy provided a collection of specific recommendations to which STL could be applied, but a well-structured singular document is not a necessity for STL to be useful. STL provides new methods of gaining insight into cybersecurity issues more broadly, prompting new and more insightful questions around how to mitigate problems and reduce harm at multiple levels. While there are many additional avenues of valuable inquiry in cybersecurity policy (which will

be discussed in the next chapter), it is my hope that STL will eventually lead to a safer internet for all who use it.

## CHAPTER 6

### CONCLUSION

This dissertation has explored the multifaceted landscape of cybersecurity, highlighting its intricate nature as a complex dynamic system comprising humans, social groups, technology, and policy at various levels. To recap, this research perspective was inspired in part by Norbert Wiener's definition of "cybernetics" - "the scientific study of control and communication in the animal and the machine" (Weiner, 1948) - and the essential role that human, social, and policy aspects play in this system. U.S. federal-level cybersecurity policy approaches have to date been overly focused on the technological aspects, patterned primarily after incident response techniques rather than proactive measures. These approaches often overlook the substantial harms inflicted on individuals, organizations, and national security by cybersecurity issues, such as data breaches, exploited vulnerabilities, and exposed data resulting from insecure design. These "blind spots" have constrained the impact of policy responses.

To identify and address these limitations, I developed a novel taxonomy that categorizes cybersecurity incidents based on how they impact individuals, organizations, and national security. This new taxonomy provides a new understanding of the diverse range of harms that may be incurred by cybersecurity failures. I then used this taxonomy as a unique "coordinate system" as the next step in developing a framework to interrogate and evaluate the dimensions of harm and how they combine and evolve within the cybersecurity landscape.



The Socio-Technical Lagrangian framework, or STL, was built from this new coordinate system, a novel approach that enables a deeper inquiry into the sociotechnical nature of cybersecurity events. By applying the concepts inspired by Lagrangian mechanics (and acknowledging the limits of the physical analogue), one may gain more profound insights into the interplay of incentives within the system, revealing the often-hidden conflicts and their effects in various domains.

STL was then put to practical use in evaluating the relevant pillars and strategic objectives of the 2023 U.S. National Cybersecurity Strategy. Through this analysis, areas of improvement were identified in an effort to ensure the strategy's positive implementation, minimizing harm at each of the identified levels of the “coordinate system” - individual, organizational, and national security. Some of these objectives were well aligned with the premise of this research – namely, that individuals currently bear too much responsibility for large-scale cyber events, and that liability for insecure software needs to be more realistically and equitably aligned. Other objectives were found to be well-intentioned but misguided, such as the concept of a cross-platform digital identity that would work across all government functions. STL was also useful in assessing how the recommended policy measures might be implemented, since the implementation plan for the strategy has yet to be designed or released to the public.

In essence, this dissertation has made significant strides in enhancing the understanding of cybersecurity complexities and has contributed innovative tools to address the shortcomings of current U.S. federal-level cybersecurity policy. By emphasizing a proactive and anticipatory approach, considering the multifaceted harms, and applying the STL framework, it is possible to pave the way for a more resilient,

secure, and inclusive cybersecurity ecosystem. By embracing a holistic perspective and fostering collaboration between diverse stakeholders, we can collectively fortify our cybersecurity defenses and create a safer digital environment for generations to come.

### **Knowledge Contributions of the Work**

This dissertation presents three unique knowledge contributions aimed at enhancing our understanding of cybersecurity failures and improving the state of cybersecurity at the federal level in the U.S. The first contribution is a novel taxonomy developed to describe and assess harms resulting from various cybersecurity incidents. Existing taxonomies are limited in scope, focusing on remediation and response efforts rather than providing a comprehensive understanding. The new taxonomy categorizes impacts on three primary groups, drawing insights from diverse sources, including academic writings, legal opinions, policy documents, and interactions with the security community.

The second contribution involves the creation of a new framework, STL, which can be used to analyze the incentive structures within cybersecurity systems, revealing often hidden conflicts and their effects. While existing discussions concentrate on the consequences of security failures, they overlook the underlying incentives driving such failures and their interplay across different categories. The dissertation introduces STL as an abstraction-level framework, inspired by Lagrangian mechanics in physics, to reframe cybersecurity issues. This approach allows a clearer understanding of incentives within the system, their interactions, and identifies areas where efforts could lead to positive changes.

Synthesizing the findings from the taxonomy and framework, the third contribution involves formulating reasonable and effective recommendations for enhancing the cybersecurity system's state for multiple stakeholder groups. Leveraging the contextually appropriate taxonomy and unique framework, these suggestions address the reform of U.S. federal-level policy, drawing insights from various governmental sources, case law, and discussions with policy experts, culminating in analysis and recommendations around the 2023 White House Cybersecurity Strategy. One primary goal of this work was to drive improvements in cybersecurity at the federal level, distributing the newly generated knowledge outside academia to positively impact relevant stakeholder groups.

### **Stakeholders**

Which stakeholder groups could potentially be impacted by this research? Federal-level policymakers in the U.S. are the first group that comes to mind, as this dissertation includes specific recommendations on how to implement the most recent cybersecurity strategy of the Executive branch. The 2023 White House Cybersecurity Strategy includes a wide range of objectives, many of which require significant investments from the federal government (and therefore Congressional approval) and from industry. Knowing which objectives to prioritize and how to gain consensus between public and private entities will be essential in the Strategy's success, and whether or not it endures.

Policymakers are not the only group who could benefit from the STL framework. As shown in this research, harms incurred at the organizational level are very likely to

have impacts on individuals, communities, national security, and may even result in global effects. The spectrum of potential harms, as well as how they can evolve and spread, has been clarified through the STL framework and its application to the test and validation data sets. U.S. federal policy is important, but the private sector has a uniquely important and necessary role to play in protecting individuals and promoting national security. A company's security posture may have highly scalable (and hard to anticipate) impacts. If thoughtful, comprehensive, and well-designed, an organizational-level security policy can have globally positive influence, whether the company is a national-scale gas pipeline, a social media platform, or a local health care provider. STL can help translate federal-level policy into coherent and reasonable guidance for companies, which in turn can then build appropriate safeguards for their business partners and customers.

Individuals are stakeholders as well, but in a different way. They certainly have a vested interest in staying safe online, having their personal data protected, and experiencing fewer data breach notifications. Individuals deserve to partake in the benefits of the internet without having to worry about being neglected or exploited by the platforms they use (particularly when they have no way to opt out, as in the case of Equifax). Overall, people are exhausted by the continual notifications of data breaches, as well as the confusing and manipulative terms of service and privacy policies. Too much emphasis has been placed on their personal responsibilities for internet security, when in fact the responsibility should reside at much higher levels. Very few individuals have the influence or ability to cause large-scale cyber events like the ones analyzed in this dissertation, but literally billions of people have been affected by them. Rebalancing

incentives for internet security would be positive at all levels, but would provide the greatest relief to individuals, who have been saddled with this responsibility for too long.

### **Future Research Directions**

As the ever-evolving digital landscape continues to present new challenges and opportunities, the insights and recommendations presented in this dissertation may be able to serve as a foundation for future exploration. During the course of this research, some aspects of cybersecurity have changed quickly from technical, policy, and social viewpoints – and some have stayed the same. Some changes have been overtly negative, like the broader proliferation of ransomware. Some have been positive, such as the emergence of CISA as an positive and (moderately) impactful federal-level agency. And in some ways, the situation feels remarkably stagnant. Large-scale data breaches continue to happen, large-scale vulnerabilities are still being discovered and exploited. By and large, industry continues to abuse the privacy of individuals for monetary gain, Congress has responded with great sound and fury, but unfortunately, signifying nothing. Given this intriguing combination of momenta and inertia in the cybersecurity system writ large, it is somewhat unfortunate to conclude this research without addressing the most recent developments. But unpredictable and continual change is a feature of complex systems. It would be impractical to think that, at some point in the foreseeable future, the complex socio-technical landscape of cybersecurity will settle into a predictable pattern of behavior and the work will be “done”. Hopefully, the research in this dissertation provides a foundation for new and impactful explorations that go beyond the scope of this

document (as initially described in Chapter 1). This concluding chapter outlines some of those potential areas.

### International Dynamics

It's been stated many times in this dissertation that cybersecurity is a global problem that does not respect the lines drawn on a map. This is true from technical, social, and policy aspects. The framework created in this dissertation is focused on cybersecurity policy at the U.S. federal level, which is only one important lens. International cooperation in cybersecurity is a difficult challenge, but one with tremendous potential benefits if it is successfully addressed. Analysis of harms and incentives at the global scale could inform positive realignment of incentives internationally, supporting strategy and implementation of the fifth pillar of the Biden-Harris National Cybersecurity Strategy.

### State and local connections to Federal policy

The interplay between local, state, and federal cybersecurity policy is another system that would be an interesting object of study. For example, can proactive state-level laws impact incentives within the large-scale system? Do they ameliorate harm at multiple levels, or is their impact constrained? There are many important questions that could be asked of a framework that illuminates the connections between levels of government, and the answers may provide guidance for a more cohesive and manageable regulatory system.

## Soft law approaches

As discussed previously, federal-level laws are difficult to change, and they often do not reflect the current socio-technical reality. Top-down federal legislation may not be the most appropriate or effective solution to many socio-technical cybersecurity issues. Are there specific “soft law” approaches, either in this domain or others, that have had positive impact? Can these approaches be relied upon to re-shape market forces, disrupt threat actors, and decrease harms generally within the system? There is a considerable body of work on soft law that could be leveraged in concert with a corpus of cybersecurity events to create a framework analogous to STL that may help answer these questions.

## Additional perspectives

One limitation of the approach outlined in this dissertation is the reliance on a single perspective to evaluate the three types of harms for each event. Now that the STL framework has been established and validated, a valuable next step would be to solicit additional expert evaluations for the same examples in the STL framework (in the primary and validation data sets) to both gain additional validation and provide a more robust basis for applying the framework. As there are many kinds of expertise relevant to this field, one potential area of future research could be to investigate whether or not experts from certain subfields (governance and policy, information security, military intelligence) would code these events similarly, or if other unexpected groupings might occur. This would shed light on the perception of harms within key stakeholder groups, and possibly lead to more coherent policymaking.

## Future evolutions

The socio-technical cybersecurity landscape continues to evolve. Complex systems exhibit new emergent behaviors which cannot be predicted, but which can have large impacts. The framework created in this dissertation was designed specifically to understand the complexity of the socio-technical cybersecurity landscape, but not all of the future “motions” of the system can be understood or designed for in advance. As the landscape changes, there may be a point at which STL can no longer provide the necessary insights to design and influence policy positively at the U.S. federal level. These changes could come from technological disruptions, such as quantum computing, which is anticipated to radically change computing in general but also impact current security measures such as cryptography. The policy landscape changes as well – several measures proposed in the most recent National Cybersecurity Strategy could be considered “stretch goals”, but would have significant impact if they can be successfully implemented. And as always, people change, as does their relationship to technology. There is no linear or predictable path for how humans will use future technologies, nor how they might use today’s technology in the future. The social impacts of computing are less about the shiny new developments in the tech itself, but about how people themselves determine what will be successful, what will fail, what will be used for positive or for negative outcomes. Given the unpredictability of the future, reassessing the validity of the STL approach would be both warranted and welcomed.



## **Conclusion**

Cybersecurity problems are far from being solved, much less fully understood. The research in this dissertation provides a first step towards clarity – if nothing else, showing that a new vantage point, intentionally designed, can shed light on previously unseen conflicts and point towards ways to resolve them. I hope this dissertation will serve as a foundation for additional research, as well as a method to establish more realistic and responsible policies in cybersecurity. No one can predict the future of the internet or humanity’s relationship with it, but I am confident that through creative research and effective communication, that relationship can evolve in positive and productive ways. Continuing to incorporate insights from diverse disciplines and engaging with real-world cybersecurity incidents (and practitioners), it is possible to pave the way for a more anticipatory and effective approach to address the complexities of cybersecurity in today's interconnected world.

## REFERENCES

- Merck & Co. Inc. vs. Ace American Insurance Co. et al, L-002682-18 (New Jersey Superior Court January 13, 2022).
- 116th Congress of the United States. (2020, June 25). *H.R.7331 - National Cyber Director Act*. Retrieved from <https://www.congress.gov/bill/116th-congress/house-bill/7331>
- Abbas, R., Michael, K., Michael, & MG. (2014). The regulatory considerations and ethical dilemmas of location-based services (LBS) : A literature review. *Information Technology & People, Vol. 27 No. 1, 2-20*.
- Abbate, J. (1999). *Inventing the Internet*. Cambridge: MIT Press.
- Adamson, G., Kline, R., Michael, K., & Michael, M. (2015). Weiner's Cybernetics Legacy and the Growing Need for the Interdisciplinary Approach. *Proceedings of the IEEE, Vol. 103, No. 11, 2208-2214*.
- André, J. L., Ferber, S., & Alexander, T. (2021, October 12). *DOJ Announces New Civil Cyber-Fraud Initiative*. Retrieved from National Law Review: <https://www.natlawreview.com/article/doj-announces-new-civil-cyber-fraud-initiative>
- Ars Technica. (2018, January 3). *"Meltdown" and "Spectre:" Every modern processor has unfixable security flaws*. Retrieved from Ars Technica: <https://arstechnica.com/gadgets/2018/01/meltdown-and-spectre-every-modern-processor-has-unfixable-security-flaws/>
- Ars Technica. (2018, January 5). *Meltdown and Spectre: Here's what Intel, Apple, Microsoft, others are doing about it*. Retrieved from Ars Technica: <https://arstechnica.com/gadgets/2018/01/meltdown-and-spectre-heres-what-intel-apple-microsoft-others-are-doing-about-it/>
- ASU News. (2016, August 29). *ASU cybersecurity expert says hacked database controls who is allowed to vote*. Retrieved from ASU News: <https://news.asu.edu/20160829-arizona-impact-asu-risk-hack-arizona-voter-database-trust>
- Auxier, B., Rainie, L., Anderson, M., Perrin, A., Kumar, M., & Turner, E. (2019, November 15). *Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information*. Retrieved from Pew Research Center: <https://www.pewresearch.org/internet/2019/11/15/americans-and->

privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/

- Babbitt, C. (n.d.). *The UDRP Process*. Retrieved from Harvard Berkman Center for Internet and Society: <https://cyber.harvard.edu/udrp/process.html#remedies>
- Barth, B. (2021, June 4). *Supreme Court narrows interpretation of CFAA, to the relief of ethical hackers*. Retrieved from SC Magazine: <https://www.scmagazine.com/news/security-news/supreme-court-narrows-interpretation-of-cfaa-to-the-relief-of-ethical-hackers>
- Bell, L. (2023, March 12). *Meta is deleting the Facebook and Instagram accounts of hacking victims – and with them, years of irreplaceable memories*. Retrieved from Metro: [https://metro.co.uk/2023/03/12/meta-is-deleting-the-facebook-and-instagram-accounts-of-hack-victims-18373835/?ico=zone-post-strip\\_item\\_2\\_news](https://metro.co.uk/2023/03/12/meta-is-deleting-the-facebook-and-instagram-accounts-of-hack-victims-18373835/?ico=zone-post-strip_item_2_news)
- Berger, J. (2016, March 25). *A Dam, Small and Unsung, Is Caught Up in an Iranian Hacking Case*. Retrieved from The New York Times: <https://www.nytimes.com/2016/03/26/nyregion/rye-brook-dam-caught-in-computer-hacking-case.html>
- Boase, J., Horrigan, J., Wellman, B., & Rainie, L. (2006). *The Strength of Internet Ties*. Pew Internet and American Life Project Report. Pew Research.
- Bowman, A. (2021, April 9). *After Data Breach Exposes 530 Million, Facebook Says It Will Not Notify Users*. Retrieved from National Public Radio: <https://www.npr.org/2021/04/09/986005820/after-data-breach-exposes-530-million-facebook-says-it-will-not-notify-users>
- Branscome, M. (2016, July 12). *Stop saying the cloud is just someone else's computer - because it's not*. Retrieved from ZDNet: <https://www.zdnet.com/article/stop-saying-the-cloud-is-just-someone-elses-computer-because-its-not/>
- Bremer v SolarWinds, 1:2021cv00002 (US District Court for the Western District of Texas January 4, 2021).
- Brennan Center for Justice. (2019, October 23). *Recommendations to Defend America's Election Infrastructure*. Retrieved from The Brennan Center for Justice: <https://www.brennancenter.org/our-work/research-reports/recommendations-defend-americas-election-infrastructure>
- Browning, K. (2021, October 6). *A 'potentially disastrous' data breach hits Twitch, the livestreaming site*. Retrieved from The New York Times : <https://www.nytimes.com/2021/10/06/technology/twitch-data-breach.html>

- Brumfield, C. (2022, March 17). *SEC filings show hidden ransomware costs and losses*. Retrieved from CSO Online: <https://www.csoonline.com/article/3654293/sec-filings-show-hidden-ransomware-costs-and-losses.html>
- Bryan, K., & Lonergan, D. (2021, July 14). *Second Colonial Pipeline Data Incident Litigation Filed on Behalf of . . . Over Ten Thousand Gas Stations?* Retrieved from National Law Review: <https://www.natlawreview.com/article/second-colonial-pipeline-data-incident-litigation-filed-behalf-over-ten-thousand-gas>
- Bunker, G. (2012). Technology is not enough: Taking a holistic view for information assurance. *Information Security Technical Report*, vol. 17, no 1-2, pp. 19-25.
- Cabell, D. (2000, April 20). *Overview of Domain Name Policy Development*. Retrieved from Berkman Klein Center for Internet and Society at Harvard University : <https://cyber.harvard.edu/udrp/overview.html>
- Cable News Network LP, LLLP v. CNNNews. Com, 177 F. Supp. 2d 506 (E.D. Va. 2001 2001).
- CBS News. (2016, March 24). *U.S. charges Iran in cyberattacks against banks, New York dam*. Retrieved from CBS News: <https://www.cbsnews.com/news/us-charges-iran-in-cyberattacks-against-banks-bowman-avenue-dam-in-new-york/>
- Center for Strategic and International Studies. (2023, March 7). *Significant Cyber Incidents*. Retrieved from Strategic Technologies Program : <https://www.csis.org/programs/strategic-technologies-program/significant-cyber-incidents>
- Chaffetz, J., Meadows, M., & Hurd, W. (2016). *The OPM Data Breach: How the Government Jeopardized Our National Security for More than a Generation*. House Committee on Oversight and Government Reform. Washington D.C.: 114th Congress.
- Chapman, R. (2023, May 3). (J. Winterton, Interviewer)
- Cimpanu, C. (2020, December 13). *Microsoft, FireEye confirm SolarWinds supply chain attack* . Retrieved from ZDnet: <https://www.zdnet.com/article/microsoft-fireeye-confirm-solarwinds-supply-chain-attack/>
- CISA. (2021, March 3). *Emergency Directive 21-02*. Retrieved from Cybersecurity Directives: <https://www.cisa.gov/news-events/directives/emergency-directive-21-02>
- Clark, M. (2021, May 10). *Colonial Pipeline hackers apologize, promise to ransom less controversial targets in the future*. Retrieved from The Verge: <https://www.theverge.com/2021/5/10/22428996/colonial-pipeline-ransomware-attack-apology-investigation>

- Clarke, K. A., & Primo, D. M. (2012, March 30). Physics Envy. *The New York Times*.
- Clement, J. (2022, October 18). *Twitch*. Retrieved from Statista:  
<https://www.statista.com/topics/7946/twitch/#topicOverview>
- Cohen, G. (2021, August 12). *Throwback Attack: How the modest Bowman Avenue Dam became the target of Iranian hackers*. Retrieved from Cybersecurity Pulse:  
<https://www.industrialcybersecuritypulse.com/throwback-attack-how-the-modest-bowman-avenue-dam-became-the-target-of-iranian-hackers/>
- Cohen, G. (2021, August 12). *Throwback Attack: How the modest Bowman Avenue Dam became the target of Iranian hackers*. Retrieved from Industrial Cybersecurity Pulse: <https://www.industrialcybersecuritypulse.com/facilities/throwback-attack-how-the-modest-bowman-avenue-dam-became-the-target-of-iranian-hackers/>
- Cohodas, M. (2015, October 2). *Poll: Employees Clueless About Social Engineering*. Retrieved from Dark Reading: <https://www.darkreading.com/perimeter/poll-employees-clueless-about-social-engineering>
- Committee on National Security Systems. (2015). *Committee on National Security Systems Glossary (CNSSI:4009)*. Ft. Meade: National Security Agency.
- Computer History Museum . (2021). *Timeline of Computer History*. Retrieved from Computer History Museum: <https://www.computerhistory.org/timeline/1984/>
- Computer History Museum. (Accessed August 1, 2020). *Blavak and white image of pencil drawing of the four-node ARPAnet map*. Retrieved from Computer History Museum: <https://www.computerhistory.org/collections/catalog/102658020>
- Conduitt, J. (1727). *Draft account of Newton's life at Cambridge*. Retrieved from The Newton Project :  
<http://www.newtonproject.ox.ac.uk/view/texts/diplomatic/THEM00167>
- Congressional Research Service. (2023, March 8). *The National Cybersecurity Strategy—Going Where No Strategy Has Gone Before*. Retrieved from Congressional Research Service: <https://crsreports.congress.gov/product/pdf/IN/IN12123/2>
- Crocker, A. (2022, May 19). *DOJ's New CFAA Policy is a Good Start But Does Not Go Far Enough to Protect Security Researchers*. Retrieved from Electronic Frontier Foundation: <https://www.eff.org/deeplinks/2022/05/dojs-new-cfaa-policy-good-start-does-not-go-far-enough-protect-security>
- Cross, C., Parker, M., & Sansom, D. (2019). Media discourses surrounding 'non-ideal' victims: The case of the Ashley Madison data breach. *International Review of Victimology*, 53-69.

- Cybersecurity and Infrastructure Security Agency (CISA). (2019, November 14). *Security Tip (ST04-001): What is Cybersecurity?* . Retrieved from National Cyber Awareness System: <https://us-cert.cisa.gov/ncas/tips/ST04-001>
- Dekel, J. (2014, May 1). *Swartz doc director: Oracle and Larry Ellison killed Aaron's Law*. Retrieved from Canada.com: <https://web.archive.org/web/20140502003344/http://o.canada.com/technology/swartz-doc-director-oracle-and-larry-ellison-killed-aarons-law>
- Department of Justice. (2021, October 6). *Deputy Attorney General Lisa O. Monaco Announces New Civil Cyber-Fraud Initiative*. Retrieved from The Department of Justice: <https://www.justice.gov/opa/pr/deputy-attorney-general-lisa-o-monaco-announces-new-civil-cyber-fraud-initiative>
- Detch, J. (2016, July 14). *Influencers: Antihacking law obstructs security research*. Retrieved from The Christian Science Monitor: <https://www.csmonitor.com/World/Passcode/Passcode-Influencers/2016/0714/Influencers-Antihacking-law-obstructs-security-research>
- Dickerson v CDPQ Colonial Partners, LP, 1:21-CV-02098 (U.S. District Court, Northern District of Georgia (Atlanta) May 18, 2021).
- Dockray, H. (2016, September 26). *Donald Trump identifies new cybersecurity threat: 400-pound guy sitting on bed*. Retrieved from Mashable: <https://mashable.com/article/trump-400-lb-guy-dnc-hack>
- Dodgson, L. (2019, April 22). *Ashley Madison now has 60 million users. Two men told us why they use it*. Retrieved from Insider: <https://www.insider.com/why-men-use-ashley-madison-online-dating-2019-4>
- Doe, D. (2016, February 15). *22,000 dental patients' info exposed on unsecured Eaglesoft FTP server* . Retrieved from DataBreaches.net: <https://www.databreaches.net/22000-dental-patients-info-exposed-on-unsecured-eaglesoft-ftp-server/>
- Doe, D. (2016, May 27). *FBI raids dental software researcher who discovered private patient data on public server*. Retrieved from The Daily Dot: <https://www.dailydot.com/debug/justin-shafer-fbi-raid/>
- Doe, D. (2016, June 1). *Security researchers stop disclosing vulnerabilities after FBI raid on fellow researcher*. Retrieved from The Daily Dot: <https://www.dailydot.com/debug/justin-shafer-security-researcher-chilled-speech/>
- Duffy, C. (2021, March 10). *Here's what we know so far about the massive Microsoft Exchange Hack*. Retrieved from CNN: <https://www.cnn.com/2021/03/10/tech/microsoft-exchange-hafnium-hack-explainer/index.html>

- DuPont, Q., & Fidler, B. (2016). Edge Cryptography and the Codevelopment of Computer Networks and Cybersecurity. *IEEE Annals of the History of Computing*, 55-74.
- Eggers, W. (2016, July 25). *Government's cyber challenge: Protecting sensitive data for the public good*. Retrieved from Deloitte Insights: <https://www2.deloitte.com/us/en/insights/deloitte-review/issue-19/protecting-sensitive-data-government-cybersecurity.html>
- Elkins, W. (Retrieved 2022, June 26). *Automatic Digital Network*. Retrieved from U.S. Army in Germany: [http://usarmygermany.com/Sont.htm?http&&usarmygermany.com/Units/Signal/USAREUR\\_SignalCorps%20AUTODIN.htm](http://usarmygermany.com/Sont.htm?http&&usarmygermany.com/Units/Signal/USAREUR_SignalCorps%20AUTODIN.htm)
- Ellis, C. (2018, October 11). *The Future Of Safe Harbor Is Now*. Retrieved from Forbes: <https://www.forbes.com/sites/forbestechcouncil/2018/10/11/the-future-of-safe-harbor-is-now/?sh=c3fb977eb61c>
- Esser, M. (2015). *Statement before the Committee on Oversight and Government Reform on "OPM: data breach"*. Washington, D.C.: United States House of Representatives.
- European Commission . (2021). *JRC Cybersecurity Taxonomy*. Retrieved from Joint Research Centre: <https://cybersecurity-atlas.ec.europa.eu/cybersecurity-taxonomy>
- EZ Mart 1, LLC v. Colonial Pipeline Company, 1:21-CV-02522 (U.S. District Court, Northern District of Georgia (Atlanta) June 21, 2021).
- Facebook v Power Ventures, Inc, No. 17-16161 (Ninth Circuit Court of Appeals January 16, 2019).
- FBI. (2018, November 2). *The Morris Worm 30 Years Since First Major Attack on the Internet*. Retrieved from FBI News: <https://www.fbi.gov/news/stories/morris-worm-30-years-since-first-major-attack-on-internet-110218>
- FBI. (2021, October 10). *FBI History - Famous Cases and Criminals*. Retrieved from Morris Worm: <https://www.fbi.gov/history/famous-cases/morris-worm>
- Federal Bureau of Investigation. (2014, December 19). *Update on Sony Investigation*. Retrieved from FBI News: <https://www.fbi.gov/news/pressrel/press-releases/update-on-sony-investigation>
- Federal Trade Commission . (2010, March 9). *LifeLock Will Pay \$12 Million to Settle Charges by the FTC and 35 States That Identity Theft Prevention and Data Security Claims Were False*. Retrieved from FTC News and Events: <https://www.ftc.gov/news-events/press-releases/2010/03/lifelock-will-pay-12-million-settle-charges-ftc-35-states>

- Federal Trade Commission . (2020). *Federal Trade Commission 2020 Privacy and Data Security Update*. Washington, DC: FTC.
- Fidler, B. (2017). Cybersecurity governance: a prehistory and its implications. *Digital Policy, Regulation, and Governance*, 449-465.
- Fidler, B. (2020, June 11). (J. Winterton, Interviewer)
- Fischbein, E. (1982). Intuition and Proof. *For the Learning of Mathematics*, 9-18, 24.
- Fisher, K. (2010, February 23). *Risque website offers \$10 million for Sky Harbor name change*. Retrieved from ABC15:  
<https://web.archive.org/web/20100225054012/http://www.abc15.com/content/news/phoenixmetro/central/story/Risque-website-offers-10-million-for-Sky-Harbor/z4uPlkToaEy6HGHA12IqOQ.csp>
- Fowles, G., & Cassiday, G. (1999). *Analytical Mechanics*. Saunders College Publishing.
- Frank, M. (2017, June 9). *UA Med School Hosts Summit on Medical Device Hacking*. Retrieved from UA News: <https://uanews.arizona.edu/story/ua-med-school-hosts-summit-medical-device-hacking>
- Freedom Online Coalition . (2022). *Fact Sheet 2022*. Retrieved from Freedom Online Coalition: <https://freedomonlinecoalition.com/wp-content/uploads/2022/12/FOC-Factsheet-2022.pdf>
- Fruhlinger, J. (2020, February 12). *Equifax data breach FAQ: What happened, who was affected, what was the impact?* Retrieved from CSO Online:  
<https://www.csoonline.com/article/3444488/equifax-data-breach-faq-what-happened-who-was-affected-what-was-the-impact.html>
- FTC. (2017, September 27). *FTC Earns Prestigious International Award for AshleyMadison.com Data Breach Investigation*. Retrieved from Federal Trade Commission News: <https://www.ftc.gov/news-events/news/press-releases/2017/09/ftc-earns-prestigious-international-award-ashleymadisoncom-data-breach-investigation>
- Gallagher, S. (2015, June 8). Why the “biggest government hack ever” got past the feds. *Ars Technica*.
- GAO. (2019). *CYBERSECURITY Agencies Need to Fully Establish Risk Management Programs and Address Challenges*. Washington, D.C.: U.S. Government Accountability Office.
- Gartner. (2021, May 17). *Gartner Forecasts Worldwide Security and Risk Management Spending to Exceed \$150 Billion in 2021* . Retrieved from Gartner:



- <https://www.gartner.com/en/newsroom/press-releases/2021-05-17-gartner-forecasts-worldwide-security-and-risk-managem>
- Gatlan, S. (2023, February 10). *Clop ransomware claims it breached 130 orgs using GoAnywhere zero-day*. Retrieved from Bleeping Computer:  
<https://www.bleepingcomputer.com/news/security/clop-ransomware-claims-it-breached-130-orgs-using-goanywhere-zero-day/>
- Goel, V. (2017, February 21). *Verizon Will Pay \$350 Million Less for Yahoo*. Retrieved from The New York Times:  
<https://www.nytimes.com/2017/02/21/technology/verizon-will-pay-350-million-less-for-yahoo.html>
- Goel, V., & Lichtblau, E. (2017, March 15). *Russian Agents Were Behind Yahoo Hack, U.S. Says*. Retrieved March 2023, from The New York Times:  
[https://www.nytimes.com/2017/03/15/technology/yahoo-hack-indictment.html?\\_r=0](https://www.nytimes.com/2017/03/15/technology/yahoo-hack-indictment.html?_r=0)
- Goodin, D. (2017, April 14). *NSA-leaking Shadow Brokers just dumped its most damaging release yet*. Retrieved from Ars Technica:  
<https://arstechnica.com/information-technology/2017/04/nsa-leaking-shadow-brokers-just-dumped-its-most-damaging-release-yet/>
- Gorzoch, D. (1979, July 18). *Facsimile to Reduce Writer-to-Reader Time*. Retrieved from Defense Technical Information Center:  
<https://apps.dtic.mil/sti/pdfs/ADA086803.pdf>
- Graz University of Technology. (2018). Retrieved from Meltdown and Spectre:  
<https://spectreattack.com/>
- Greenberg, A. (2018, August 22). *The Untold Story of NotPetya, the Most Devastating Cyberattack in History*. Retrieved March 2023, from Wired:  
<https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/>
- Greenberg, A. (2020, May 12). *The Confessions of Marcus Hutchins, the Hacker Who Saved the Internet*. Retrieved from Wired:  
<https://www.wired.com/story/confessions-marcus-hutchins-hacker-who-saved-the-internet/>
- Greenberg, J. (2016, December 17). *Fact-checking the integrity of the vote in 2016*. Retrieved from PolitiFact: <https://www.politifact.com/article/2016/dec/17/fact-checking-claims-voter-fraud-2016/>
- Gregoire, C. (2015, August 20). *Ashley Madison Hack Could Have A Devastating Psychological Fallout*. Retrieved from The Huffington Post:

[https://www.huffpost.com/entry/ashley-madison-hack-psychological-fallout\\_n\\_55d4afcee4b07addcb44f5d4](https://www.huffpost.com/entry/ashley-madison-hack-psychological-fallout_n_55d4afcee4b07addcb44f5d4)

Greschler, G. (2023, May 1). *Breach blamed on Russian-linked hackers exposes San Jose healthcare group's sensitive data*. Retrieved from Mercury News: <https://www.msn.com/en-us/health/other/breach-blamed-on-russian-linked-hackers-exposes-san-jose-healthcare-group-s-sensitive-data/ar-AA1aBPJb>

Grieg, J. (2022, June 11). *FBI, DOJ say less than 25% of NetWalker ransomware victims reported incidents*. Retrieved from The Record: <https://therecord.media/fbi-doj-say-less-than-25-of-netwalker-ransomware-victims-reported-incidents/>

Groseclose, T., & Milyo, J. (2005). A measure of media bias. *The Quarterly Journal of Economics*, 1191-1237.

Hadagny, C., & Fincher, M. (2014, October 27). *Social Engineer Capture the Flag DEFCON 22 Report*. Retrieved from Social Engineer: [http://www.social-engineer.org/wp-content/uploads/2014/10/SocialEngineerCaptureTheFlag\\_DEFCON22-2014.pdf](http://www.social-engineer.org/wp-content/uploads/2014/10/SocialEngineerCaptureTheFlag_DEFCON22-2014.pdf)

Harcup, T., & O'Neill, D. (2016). What Is News? . *Journalism Studies*, 1470-1488.

Harwell, D., & MacMillan, D. (2020, December 15). *Investors in breached software firm SolarWinds traded \$280 million in stock days before hack was revealed*. Retrieved from The Washington Post: <https://www.washingtonpost.com/technology/2020/12/15/solarwinds-russia-breach-stock-trades/>

Harwell, D., & McMillan, D. (2020, December 15). *Investors in breached software firm SolarWinds traded \$280 million in stock days before hack was revealed*. Retrieved from The Washington Post: <https://www.washingtonpost.com/technology/2020/12/15/solarwinds-russia-breach-stock-trades/>

Hautala, L. (2018, January 8). *Spectre and Meltdown: Details you need on those big chip flaws*. Retrieved from CNet: <https://www.cnet.com/news/privacy/spectre-meltdown-intel-arm-amd-processor-cpu-chip-flaw-vulnerability-faq/>

Hayles, N. K. (1999). *Virtual Bodies in Cybernetics, Literature, and Informatics*. Chicago: University of Chicago Press.

Herzkovitch, S., & Tzezana, R. (2017, January 19). *CONNECTED DEVICES GIVE SPIES A POWERFUL NEW WAY TO SURVEIL*. Retrieved from Wired: <https://www.wired.com/2017/01/connected-devices-give-spies-powerful-new-way-surveil/>

- Heylighen, F., & Joslyn, C. (2001). Cybernetics and Second-Order Cybernetics. In R. A. Myers, *Encyclopedia of Physical Science & Technology (3rd Ed.)*. New York: Academic Press.
- hiQ Labs, Inc. v. LinkedIn Corp., No. 17-16783 (Ninth Circuit Court of Appeals April 18, 2022).
- Hutcherson, K. (2018, March 28). *Six days after a ransomware cyberattack, Atlanta officials are filling out forms by hand* . Retrieved from CNN: <https://www.cnn.com/2018/03/27/us/atlanta-ransomware-computers/index.html>
- Hutcherson, K. (2018, March 28). *Six days after a ransomware cyberattack, Atlanta officials are filling out forms by hand*. Retrieved from CNN: <https://www.cnn.com/2018/03/27/us/atlanta-ransomware-computers/index.html>
- ICANN. (n.d.). *ICANN's Historical Relationship with the U.S. Government*. Retrieved August 2022, from ICANN History Project: <https://www.icann.org/en/history/icann-usg>
- In Re: Ashley Madison Customer Data Security Breach Litigation, 4:15-MD-02669 (U.S. District Court, Eastern District of Missouri (St. Louis) November 20, 2017).
- Internet Free and Secure Initiative. (2021, April 24). *A human rights respecting definition of cybersecurity*. Retrieved from The Internet Free and Secure Initiative: <https://freeandsecure.online/definition/>
- Isadore, C. (2021, May 12). *Who owns the Colonial Pipeline? It's complicated*. Retrieved from CNN: <https://www.cnn.com/2021/05/12/investing/colonial-pipeline-ownership/index.html>
- Jackson, D. (2015, September 25). Obama, Xi vow cooperation on climate, cyber issues. *USA Today*. Retrieved from USA Todaya.
- Jaikaran, C. (2023). *CRS INSIGHT Prepared for Members and Committees of Congress INSIGHTi The National Cybersecurity Strategy—Going Where No Strategy Has Gone Before*. Washington, D.C.: Congressional Research Service.
- Jain, S., & Ropple, L. (2018, December 14). Stopping Data Breaches Will Require Help from Governments. *Harvard Business Review*.
- Johnson, D. B. (2021, May 19). *SolarWinds CEO expresses regret for 'blame the intern' defense during Orion hack investigation*. Retrieved from SC Magazine: <https://www.scmagazine.com/news/application-security/solarwinds-ceo-expresses-regret-for-blame-the-intern-defense-during-orion-hack>
- Johnson, D. B. (2021, May 19). *SolarWinds CEO expresses regret for 'blame the intern' defense during Orion hack investigation*. Retrieved from SC Magazine:

- <https://www.scmagazine.com/news/solarwinds-ceo-expresses-regret-for-blame-the-intern-defense-during-orion-hack>
- Kan, M. (2016, September 28). *The Yahoo hackers weren't state-sponsored, a security firm says*. Retrieved March 2023, from Network World:  
<https://www.networkworld.com/article/3125594/the-yahoo-hackers-werent-state-sponsored-a-security-firm-says.html>
- Kaplan, F. (2019, February 19). 'WarGames' and Cybersecurity's Debt to a Hollywood Hack. *The New York Times*.
- Kassner, M. (2015, February 2). *Anatomy of the Target data breach: Missed opportunities and lessons learned*. Retrieved from ZDNet:  
<https://www.zdnet.com/article/anatomy-of-the-target-data-breach-missed-opportunities-and-lessons-learned/>
- Keegan, R. (2014, December 6). *Sony hack 'unprecedented, damaging and unique' cyber security firm says* . Retrieved from The Los Angeles Times:  
<https://www.latimes.com/entertainment/envelope/cotown/la-et-ct-sony-hack-20141206-story.html>
- Kerr, O. [. (2021, June 25). "For those interested in the academic exercise of what Van Buren means for the CFAA, am I right that the...". [Tweet],  
<https://twitter.com/OrinKerr/status/1408498750813655042>: Twitter.
- Kerr, O. (2021, June 9). *The Supreme Court Reins In the CFAA in Van Buren*. Retrieved July 2022, from Lawfare: <https://www.lawfareblog.com/supreme-court-reins-cfaa-van-buren>
- King, I., Kahn, J., Webb, A., & Turner, G. (2018, January 8). *'It Can't Be True.' Inside the Chip Industry's Meltdown*. Retrieved from Bloomberg:  
<https://www.bloomberg.com/news/articles/2018-01-08/-it-can-t-be-true-inside-the-semiconductor-industry-s-meltdown>
- King, N., & Raja, V. (2012). Protecting the privacy and security of sensitive customer data in the cloud. *Computer Law and Security Review*, Vol. 28, Issue 3, pp. 308-319.
- Kline, R., & Pinch, T. (1996). Users as Agents of Technological Change: The Social Construction of the Automobile in the Rural United States. *Technology and Culture*, 763-795.
- Koerner, B. (2002, May). *From Russia With LØpht*. Retrieved from Legal Affairs: .  
[https://legalaffairs.org/issues/May-June-2002/feature\\_koerner\\_mayjun2002.msp](https://legalaffairs.org/issues/May-June-2002/feature_koerner_mayjun2002.msp)

- Kolevski, D., Michael, K., Abbas, R., & Freeman, M. (2020). Stakeholders in the cloud computing value-chain: A socio-technical review of data breach literature. *2020 IEEE International Symposium on Technology and Society (ISTAS)*, 290-293.
- Kolevski, D., Michael, K., Abbas, R., & Freeman, M. (2021). Cloud Data Breach Disclosures: the Consumer and their Personally Identifiable Information (PII)? *IEEE 3rd Conference on Norbert Wiener in the 21st Century*. Chennai: IEEE.
- Kolevski, D., Michael, K., Abbas, R., & Freeman, M. (Forthcoming). Cloud Computing Data Breaches in News Media: Disclosure of Personal and Sensitive Data. *IEEE*.
- Kuhn, T. S. (2012). *The Structure of Scientific Revolutios*. Chicago: University of Chicago Press.
- Ladyman, J., Lambert, J., & Wiesner, K. (2011, February 27). What is a complex system? . *European Journal for Philosophy of Science*.
- Launius, S. (2020, April 9). *Evaluation of Comprehensive Taxonomies for Information Technology Threats*. Retrieved from Cybersecurity and Information Systems Information Analysis Center: <https://csiac.org/articles/evaluation-of-comprehensive-taxonomies-for-information-technology-threats/>
- Lee, J., & Rosin, R. (1992, April/June). The Project MAC Interviews. *IEEE Annals of the History of Computing*, pp. 14-35.
- Lemos, R. (2003, November 13). *Bush unveils final cybersecurity plan*. Retrieved from CNet: <https://www.cnet.com/tech/tech-industry/bush-unveils-final-cybersecurity-plan/>
- Leonhardt, M. (202, February 6). *If you got an email about the \$117.5 million Yahoo data breach settlement, here are your options*. Retrieved March 2023, from CNBC: <https://www.cnbc.com/2020/02/06/what-to-do-if-you-got-email-from-yahoo-about-a-data-breach-settlement.html>
- Licklider, J. (1960). Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics*, 4-11.
- Lofgren, Z. (2013, June 20). *Rep. Zoë Lofgren Introduces Bipartisan Aaron's Law*. Retrieved from <https://lofgren.house.gov/media/press-releases/rep-zoe-lofgren-introduces-bipartisan-aarons-law>
- Maas, D. (2013, April 11). *What We're Up Against: Software Lobby SIIA Spends Big to Stop CFAA Reform*. Retrieved from Electronic Frontier Foundation : <https://www.eff.org/deeplinks/2013/04/what-were-against-software-lobby-siia-spends-big-stop-cfaa-adjustments>
- Magin, R. L. (2017, Jan/Feb 15). Bioengineering and Cybernetics. *IEEE Pulse*.

- Malatji, M., Von Solms, S., & Marnewick, A. (2019). Socio-technical systems cybersecurity framework. *Information and Computer Security*, 233-272.
- Mandiant. (n.d.). *Advanced Persistent Threats (APTs)*. Retrieved March 2023, from Mandiant: <https://www.mandiant.com/resources/insights/apt-groups>
- Manion, A. (2018, July 11). *Hearing on “Complex Cybersecurity Vulnerabilities: Lessons Learned from Spectre and Meltdown”*. Retrieved from U.S. Senate Committee on Commerce, Science, and Transportation: <https://www.commerce.senate.gov/services/files/48B3BB9C-570E-4C82-B85E-1B1F017A19AB>
- Marchant, G. (2011). The Growing Gap Between Emerging Technologies and the Law. In G. Marchant, *The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight* (pp. 19-33). Dordrecht, Germany: Springer.
- Markoff, J. (1999, December 20). An Internet Pioneer Ponders the Next Revolution. *The New York Times*.
- Marks, J. (2022, March 9). (J. Winterton, Interviewer)
- Martin, M. (2021, April). *Computer and Internet Use in the United States: 2018*. Retrieved from United States Census: <https://www.census.gov/content/dam/Census/library/publications/2021/acs/acs-49.pdf>
- Masuch, K., Greve, M., Trang, S., & Kolbe, L. (2022, January). Apologize or justify? Examining the impact of data breach response actions on stock value of affected companies? *Computers and Security*, 112.
- Maynard, A., & Garbee, E. (2019). Responsible innovation in a culture of entrepreneurship: a US perspective . In R. von Schomberg, & J. Hankins, *International Handbook on Responsible Innovation* (pp. 488-502). Cheltenham: Edward Elgar Publishing Ltd.
- Mazzucato, M. (2013). *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*. New York: Anthem Press.
- McAndrew, E. (2018, May 11). *The Hacked & the Hacker-for-Hire: Lessons from the Yahoo Data Breaches (So Far)*. Retrieved from The National Law Review: <https://www.natlawreview.com/article/hacked-hacker-hire-lessons-yahoo-data-breaches-so-far>
- McKenzie, A. A., & Walden, D. C. (1991). ARPANET, the Defense Data Network, and Internet. In F. E. Froelich, & A. Kent, *The Froelich/Kent Encyclopedia of Telecommunications* (pp. 341-376). New York: Marcel Dekker, Inc.

- McMillan, R. (2018, January 28). *Intel Warned Chinese Companies of Chip Flaws Before U.S. Government*. Retrieved from The Wall Street Journal: <https://www.wsj.com/articles/intel-warned-chinese-companies-of-chip-flaws-before-u-s-government-1517157430>
- McMillan, R., & Knutson, R. (2017, October 3). *Yahoo Triples Estimate of Breached Accounts to 3 Billion* . Retrieved from The Wall Street Journal: <https://www.wsj.com/articles/yahoo-triples-estimate-of-breached-accounts-to-3-billion-1507062804>
- Meijer, I. C., & Bijleveld, H. P. (2016). Valuable Journalism. *Journalism Studies*, 827-839.
- Michael, K., & Clarke, R. (2013). Location and tracking of mobile devices: überveillance stalks the streets. *Computer Law and Security Review*, 216-228.
- Miller, J. H., & Page, S. E. (2007). *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton and Oxford: Princeton University Press.
- Mimoso, M. (2016, September 29). *Yahoo Challenged on Claims Breach Was State-Sponsored Attack*. Retrieved from ThreatPost: <https://threatpost.com/yahoo-challenged-on-claims-breach-was-state-sponsored-attack/120975/>
- Monaco, L. O. (2021, October 6). *Op-Ed: America needs Congress's help to solve the ransomware threat*. Retrieved from CNBC: <https://www.cnbc.com/2021/10/06/deputy-ag-congress-must-create-standard-to-encourage-companies-to-report-cyberattacks.html>
- Monahan, B. (2010). *The Shock of the News: Media Coverage and the Making Of 9/11*. New York: New York University Press.
- Mondelez International Inc. v. Zurich American Insurance Co, 2018-L-011008 (Circuit Court, Cook County October 10, 2018).
- Moon, M. (2018, May 30). *Attacker involved in 2014 Yahoo hack gets five years in prison*. Retrieved March 2023, from Engadget: <https://www.engadget.com/2018-05-30-yahoo-hacker-sentence.html>
- Morganti, M., & Kling, A. (2021, May 4). Perverse Incentives: How We Disincentivized Vendor Transparency & How We Can Do Better. *Hack the Capitol* . Washington, DC.
- Moscaritolo, A. (2015, December 29). *Ashley Madison Adds 4 Million Users Since Hack*. Retrieved from PC Magazine: <https://www.pcmag.com/news/ashley-madison-adds-4-million-users-since-hack>

- in re U.S. Office of Personnel Management Data Security Breach Litigation, 1:15-mc-01394-ABJ (United States District Court for the District of Columbia October 26, 2022).
- Nakashima, E. (2015, July 9). *Hacks of OPM databases compromised 22.1 million people, federal authorities say*. Retrieved from The Washington Post: <https://www.washingtonpost.com/news/federal-eye/wp/2015/07/09/hack-of-security-clearance-system-affected-21-5-million-people-federal-authorities-say/>
- Nakashima, E., & Starks, T. (2023, January 5). *U.S. national cyber strategy to stress Biden push on regulation*. Retrieved from The Washington Post: <https://www.washingtonpost.com/national-security/2023/01/05/biden-cyber-strategy-hacking/>
- Nason, R. (2017). *It's not complicated: The art and science of complexity in business*. Toronto: University of Toronto Press.
- (2023). *National Cybersecurity Strategy*. Washington, D.C. : The White House.
- National Institute of Standards and Technology (NIST). (2018, December). NIST Special Publication 800-37 Revision 2. *Risk Management Framework for Information Systems and Organizations A System Life Cycle Approach for Security and Privacy*. <https://doi.org/10.6028/NIST.SP.800-37r2>.
- National Institute of Standards and Technology. (2018, April 16). *Cybersecurity Framework Version 1.1*. Retrieved from NIST Cybersecurity Framework: <https://www.nist.gov/cyberframework/framework>
- Neuberger, A. (2021, October 15). International Community Joins Forces as Ransomware Attacks Create Major Disruptions. (N. Schifrin, Interviewer)
- Neuberger, J. D. (2020, September 14). *Thoughtful Presentations of Terms of Use Crucial for Enforceability*. Retrieved from National Law Review: <https://www.natlawreview.com/article/thoughtful-presentations-terms-use-crucial-enforceability>
- Neumann, A., Statland, N., & Webb, R. (1977). Post processing audit tools and techniques. *Computer Science and Technology: Audit and Evaluation of Computer Security* (pp. 11-3 - 11-21). Washington, D.C.: U.S. Department of Commerce.
- Newman, A. (2015, August 31). *Ashley Madison Code Shows More Women, and More Bots*. Retrieved from Gizmodo: <https://gizmodo.com/ashley-madison-code-shows-more-women-and-more-bots-1727613924>



- Newman, L. H. (2017, October 3). *Yahoo's 2013 Email Hack Actually Compromised Three Billion Accounts*. Retrieved from Wired: <https://www.wired.com/story/yahoo-breach-three-billion-accounts/>
- Newman, L. H. (2018, August 9). *A New Pacemaker Hack Puts Malware Directly on the Device*. Retrieved from Wired: <https://www.wired.com/story/pacemaker-hack-malware-black-hat/>
- Newman, L. H. (2021, August 23). *38M Records Were Exposed Online—Including Contact-Tracing Info*. Retrieved from Wired: <https://www.wired.com/story/microsoft-power-apps-data-exposed/>
- Ng, A. (2018, July 11). *Congress: Please tell US about security vulnerabilities sooner*. Retrieved from CNet: <https://www.cnet.com/news/politics/congress-please-tell-us-about-security-vulnerabilities-sooner/>
- Nissim, K., & Wood, A. (2018). *Is Privacy Privacy?* Boston: Berkman Klein Center Research.
- Norton. (n.d.). *Norton Security Center*. Retrieved from What is Malware and How Can We Prevent It?: <https://us.norton.com/internetsecurity-malware.html>
- Novikov, D. (2016). *Cybernetics: from Past to Future*. Springer Verlag.
- Obar, J. A., & Oeldorf-Hirsch, A. (2018). The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication, and Society*, 128-147.
- O'Connor, N. E. (2021, August 16). *Katz, Marshall, and Banks*. Retrieved from Whistleblowers Accessing Company Documents Likely Will Not Be Prosecuted Under the Computer Fraud and Abuse Act: <https://www.kmblegal.com/whistleblower-blog/whistleblowers-accessing-company-documents-likely-will-not-be-prosecuted-under>
- Office of Inspector General. (1996). *Phaseout of the Automatic Digital Network*. Arlington: U.S. Department of Defense.
- Owaida, A. (2021, February 3). *Identity theft spikes amid pandemic*. Retrieved from WeLiveSecurity: [https://www.welivesecurity.com/2021/02/03/identity-theft-spikes-amid-pandemic/?utm\\_source=newsletter&utm\\_medium=email&utm\\_campaign=wls-newsletter-050221](https://www.welivesecurity.com/2021/02/03/identity-theft-spikes-amid-pandemic/?utm_source=newsletter&utm_medium=email&utm_campaign=wls-newsletter-050221)
- Owens, J. C. (2018, September 10). *The Equifax data breach, in one chart*. Retrieved from Marketwatch: <https://www.marketwatch.com/story/the-equifax-data-breach-in-one-chart-2018-09-07/>

- Panetta Warns of 'Cyber Pearl Harbor'*. (2022, April 26). Retrieved March 2023, from Association of the United States Army: <https://www.ausa.org/news/panetta-warns-cyber-pearl-harbor>
- Parformak, P., & Jaikaran, C. (2021, May 11). *Colonial Pipeline: The DarkSide Strikes*. Retrieved from Congressional Research Service: <https://crsreports.congress.gov/product/pdf/IN/IN11667>
- Parrish, A. (2021, August 20). *How to stop a hate raid*. Retrieved from The Verge: <https://www.theverge.com/22633874/how-to-stop-a-hate-raid-twitch-safety-tools>
- Pavković, N., & Perkov, L. (2011). Social-Engineering Toolkit - A systematic approach to social engineering. *2011 Proceedings of the 34th International Convention MIPRO* (pp. 1485-1489). Opatija: IEEE.
- Pepitone, J. (2015, June 25). *China Is 'Leading Suspect' in OPM Hacks, Says Intelligence Chief James Clapper*. Retrieved from NBC News: <https://www.nbcnews.com/tech/security/clapper-china-leading-suspect-opm-hack-n381881>
- Peters, J. (2013, February 7). *The Idealist*. Retrieved from Slate: <https://slate.com/technology/2013/02/aaron-swartz-he-wanted-to-save-the-world-why-couldnt-he-save-himself.html>
- Peters, J. (2017, December 20). *Should This Thing Be Smart? Coffee Mug Edition*. Retrieved from Slate: <https://slate.com/technology/2017/12/should-the-ember-ceramic-coffee-mug-be-smart.html>
- Peters, J. (2018, December 20). *Should This Thing Be Smart? Christmas Lights Edition*. Retrieved from Slate: <https://slate.com/technology/2018/12/twinkly-christmas-lights-smart-led-review.html>
- Pew Research Center. (2021, April 7). *Mobile Fact Sheet*. Retrieved from Pew Research Center: <https://www.pewresearch.org/internet/fact-sheet/mobile/>
- Pfefferkorn, R. (2021, September 7). *America's anti-hacking laws pose a risk to national security*. Retrieved from Brookings Institute: <https://www.brookings.edu/techstream/americas-anti-hacking-laws-pose-a-risk-to-national-security/>
- Pfefferkorn, R. (2022, February 4). *The EARN IT Act Is Back, and It's More Dangerous Than Ever*. Retrieved from The Center for Internet and Society, Stanford Law School : <https://cyberlaw.stanford.edu/blog/2022/02/earn-it-act-back-and-it%E2%80%99s-more-dangerous-ever>

- Philips, M., & Austad, S. (1996). Animal communication and social evolution. In C. Allen, & D. Jamison, *Readings in Animal Cognition* (pp. 257-267). Boston: MIT Press.
- Piccoli, G., & Wagner, E. (2003). The Value of Academic Research. *Cornell Hotel and Restaurant Administration Quarterly*, 29-38.
- Reuters. (2018, June 6). *Atlanta officials reveal worsening effects of cyber attack* . Retrieved from Reuters: <https://www.reuters.com/article/us-usa-cyber-atlanta-budget/atlanta-officials-reveal-worsening-effects-of-cyber-attack-idUSKCN1J231M?feedType=RSS&feedName=technologyNews>
- Rid, T., & Buchanan, B. (2015). Attributing Cyber Attacks. *Journal of strategic studies*, 4-37.
- Rinehart, R. (2020, February 13). *Faith in Elections in Relatively Short Supply in U.S.* Retrieved from Gallup: <https://news.gallup.com/poll/285608/faith-elections-relatively-short-supply.aspx>
- Rivero, N. (2021, May 10). *Hacking collective DarkSide are state-sanctioned pirates.* Retrieved from Quartz: <https://qz.com/2007399/the-darkside-hackers-are-state-sanctioned-pirates>
- Roberts, J. (2021, August 17). *Summons to Appear: NotPetya and the War Exclusion Clause.* Retrieved March 2023, from Johns Hopkins School of Advanced International Studies: <https://saisreview.sais.jhu.edu/notpetya-and-the-war-exclusion-clause/>
- Roberts, P. R., & ElRefai, M. H. (2019). The use of App-baed Follow-up of Cardiac Implantable Electronic Devices. *Cardiac Failure Review*.
- Rogin, J. (2012, July 9). NSA Chief: Cybercrime Constitutes the 'Greatest Transfer of Wealth in History'. *Foreign Policy*.
- Rogin, J. (2012, July 9). *NSA Chief: Cybercrime constitutes the “greatest transfer of wealth in history”*. Retrieved from Foreign Policy : [http://thecable.foreignpolicy.com/posts/2012/07/09/nsa\\_chief\\_cybercrime\\_constitutes\\_the\\_greatest\\_transfer\\_of\\_wealth\\_in\\_history](http://thecable.foreignpolicy.com/posts/2012/07/09/nsa_chief_cybercrime_constitutes_the_greatest_transfer_of_wealth_in_history)
- Romo, V. (2021, June 18). *How A New Team Of Feds Hacked The Hackers And Got Colonial Pipeline's Ransom Back.* Retrieved from National Public Radio: <https://www.npr.org/2021/06/08/1004223000/how-a-new-team-of-feds-hacked-the-hackers-and-got-colonial-pipelines-bitcoin-bac>
- Root, D., Kennedy, L., & Souzan, M. (2018, February 12). *Election Security in All 50 States: Defending America's Elections.* Retrieved March 2023, from The Center

- for American Progress: <https://www.americanprogress.org/article/election-security-50-states/>
- Rosati, P., Deeney, P., Cummins, M., van der Werff, L., & Lynn, T. (2019). Social media and stock price reaction to data breach announcements: Evidence from US listed companies. *Research in International Business and Finance*, 458-469.
- Rundle, J. (2023, March 6). *U.S. Government to Explore Cyber Insurance Backstop*. Retrieved from The Wall Street Journal: <https://www.wsj.com/articles/u-s-government-to-explore-cyber-insurance-backstop-ddc94c11>
- Ryals, P. (2017, March 27). *Museum of Computer Culture*. Retrieved from Recalling the AUTODIN: [www.computerculture.org/2012/02/recalling-the-autodin-part-i/](http://www.computerculture.org/2012/02/recalling-the-autodin-part-i/)
- Sanders, S. (2015, June 4). *Massive Data Breach Puts 4 Million Federal Employees' Records At Risk*. Retrieved from NPR: <https://www.npr.org/sections/thetwo-way/2015/06/04/412086068/massive-data-breach-puts-4-million-federal-employees-records-at-risk>
- Sanger, D. E., Krauss, C., & Perlroth, N. (2021, May 8). *Cyberattack Forces a Shutdown of a Top U.S. Pipeline*. Retrieved from The New York Times: <https://www.nytimes.com/2021/05/08/us/politics/cyberattack-colonial-pipeline.html>
- Sarian, R. (2021, July 1). *The Computer Fraud And Abuse Act Now Provides Less Protection From Insider Threats. Here's What Employers Need To Be Doing*. Retrieved from JD Supra: <https://www.jdsupra.com/legalnews/the-computer-fraud-and-abuse-act-now-1612458/>
- Schaffer, A., Marks, J., & Knowles, H. (2021, December 1). *Planned Parenthood Los Angeles says hack breached about 400,000 patients' information*. Retrieved from The Washington Post: <https://www.washingtonpost.com/nation/2021/12/01/los-angeles-planned-parenthood-hack/>
- Schukai, H. (2023, May 6). Personal Communication .
- Schwartz, E. (2018). *CISA: A good start, but challenges remain on security information sharing*. Retrieved from TechBeacon: <https://techbeacon.com/security/cisa-good-start-challenges-remain-security-information-sharing>
- Schwartz, M. (2017, March 3). *Verizon: Most breaches trace to phishing, social engineering*. Retrieved from Bank Info Security: <https://www.bankinfosecurity.com/interviews/most-breaches-trace-to-phishing-social-engineering-attacks-i-3516>

- Seal, M. (2015, March 2). *An Exclusive Look at Sony's Hacking Saga*. Retrieved from Vanity Fair: <https://www.vanityfair.com/hollywood/2015/02/sony-hacking-seth-rogen-evan-goldberg>
- Seals, T. (2020, December 16). *The SolarWinds Perfect Storm: Default Password, Access Sales and More*. Retrieved from ThreatPost: <https://threatpost.com/solarwinds-default-password-access-sales/162327/>
- Seckel, S. (2017, August 17). *ASU Now*. Retrieved from image title Creativity ASU, Amazon bring first-of-its-kind voice-technology program to campus: <https://asunow.asu.edu/20170817-asu-news-asu-amazon-dots-tooker-house>
- Sega Enterprises, Ltd. v. Accolade, Inc, 977 F.2d 1510 (9th Circuit Court of Appeals July 20, 1992).
- Segall, L. (2015, August 21). *Ashley Madison users now facing extortion*. Retrieved from CNN: <https://money.cnn.com/2015/08/21/technology/ashley-madison-users-extorted/>
- Segall, L. (2015, September 8). *Pastor outed on Ashley Madison commits suicide*. Retrieved from CNN: <https://money.cnn.com/2015/09/08/technology/ashley-madison-suicide/index.html>
- Selwyn, N. (2014). The Internet and Education. In M. Castells, D. Gelertner, J. Vásquez, E. Morozov, & M. Hyppönen, *Change: 19 Key Essays on How Internet Is Changing our Lives*. Madrid: Turner Libros.
- Shokrai, M. (2023, June 7). *The Case for a Federal Cyber-Insurance Backstop*. Retrieved from DarkReading: <https://www.darkreading.com/operations/the-case-for-a-federal-cyber-insurance-backstop>
- Sibit, J. M. (2018, October 12). *AUTODIN: The Air Force's first high speed data communications network*. Retrieved from Air Combat Command News: <https://www.acc.af.mil/News/Article/1660470/autodin-the-air-forces-first-high-speed-data-communications-network/>
- Singer, P. W., & Friedman, A. (2014). *Cybersecurity and Cyberwar: What Everyone Needs to Know*. New York: Oxford.
- Smith, A. (2001, February 9). *Bruce Springsteen Loses Cybersquatting Dispute*. Retrieved from The Register: [https://www.theregister.com/2001/02/09/bruce\\_springsteen\\_loses\\_cybersquatting\\_dispute/](https://www.theregister.com/2001/02/09/bruce_springsteen_loses_cybersquatting_dispute/)
- Southern Poverty Law Center. (Accessed 11 September 2022). *Andrew "weev" Auernheimer*. Retrieved from SPLC Extremist Files:

<https://www.splcenter.org/fighting-hate/extremist-files/individual/andrew-%E2%80%9Cweev%E2%80%9D-auernheimer>

- Spafford, E. H., & Antón, A. I. (2008). Controversies in Science and Technology. In E. H. Spafford, *The Balance Between Security and Privacy* (pp. 152-168). Mary Anne Liebart, Inc. .
- Spinde, T., Rudnitskaia, L., Mitrović, J., Hamborg, F., Granitzer, M., Gipp, B., & Donnay, K. (2021). Automated identification of bias inducing words in news articles using linguistic and context-oriented features. *Information Processing and Management*, Vol. 58, Issue 3.
- Starks, T. (2018, September 21). *The Cybersecurity 202: Trump administration seeks to project tougher stance in cyberspace with new strategy*. Retrieved from The Washington Post: <https://www.washingtonpost.com/news/powerpost/paloma/the-cybersecurity-202/2018/09/21/the-cybersecurity-202-trump-administration-seeks-to-project-tougher-stance-in-cyberspace-with-new-strategy/5ba3e85d1b326b7c8a8d158a/>
- Starks, T. (2023, March 27). *The latest mass ransomware attack has been unfolding for nearly two months*. Retrieved from The Washington Post: <https://www.washingtonpost.com/politics/2023/03/27/latest-mass-ransomware-attack-has-been-unfolding-nearly-two-months/>
- Stauss, D. (2021). *State privacy law tracker*. Retrieved from Husch Blackwell: <https://www.huschblackwell.com/2021-state-privacy-law-tracker>
- Stecklow, S., & Harney, A. (2019, December 24). *Exclusive: Malware broker behind U.S. hacks is now teaching computer skills in China*. Retrieved from Reuters: <https://www.reuters.com/article/us-china-usa-cyber-exclusive/exclusive-malware-broker-behind-u-s-hacks-is-now-teaching-computer-skills-in-china-idUSKBN1YS0UI>
- Temple-Raston, D. (2021, April 16). *A 'Worst Nightmare' Cyberattack: The Untold Story Of The SolarWinds Hack*. Retrieved March 2023, from National Public Radio: <https://www.npr.org/2021/04/16/985439655/a-worst-nightmare-cyberattack-the-untold-story-of-the-solarwinds-hack>
- The Cyberwire. (2022, May 10). *The Cyberwire*. Retrieved from Privacy Briefing: <https://thecyberwire.com/newsletters/privacy-briefing/4/90>
- The Ohio State University. (2018). *Choosing & Using Sources: A Guide to Academic Research*. Retrieved from Ohio State Libraries: <https://ohiostate.pressbooks.pub/choosingsources/>
- The United States Department of Justice. (2020, October 19). *Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and*

- Other Disruptive Actions in Cyberspace*. Retrieved March 2023, from Department of Justice Office of Public Affairs: <https://www.justice.gov/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and>
- The White House. (2021, May 12). *Executive Order on Improving the Nation's Cybersecurity* . Retrieved from <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>
- The White House. (n.d.). *Office of the National Cyber Director*. Retrieved June 2023, from The White House: <https://www.whitehouse.gov/oncd/>
- Thierer, A., & Szoka, B. (2009, August 12). Cyber-Libertarianism: The Case for Real Internet Freedom. *The Technology Liberation Front*.
- Tidy, J., & Molloy, D. (2021, October 6). *Twitch confirms massive data breach*. Retrieved from BBC: <https://www.bbc.com/news/technology-58817658>
- Tsvetanov, T., & Slaria, S. (2021, December). The effect of the Colonial Pipeline shutdown on gasoline prices. *Economics Letters*, 209.
- Turner, M. (2021, October 27). (J. Winterton, Interviewer)
- Turton, W., & Mehrotra, K. (2021, June 4). *Hackers Breached Colonial Pipeline Using Compromised Password*. Retrieved from Bloomberg: <https://www.bloomberg.com/news/articles/2021-06-04/hackers-breached-colonial-pipeline-using-compromised-password>
- Turton, W., & Mehrotra, K. (2021, June 4). *Hackers Breached Colonial Pipeline Using Compromised Password* . Retrieved from Bloomberg: <https://www.bloomberg.com/news/articles/2021-06-04/hackers-breached-colonial-pipeline-using-compromised-password>
- U.S. Congress. (n.d.). Retrieved May 2023, from United States Congressional Record : <https://www.congress.gov/search?q=%7B%22source%22%3A%22congreord%22%7D>
- U.S. Department of Justice. (2022, May 19). *19-48.000 - COMPUTER FRAUD AND ABUSE ACT*. Retrieved from Department of Justice Office of Public Affairs: <https://www.justice.gov/opa/press-release/file/1507126/download>
- U.S. Department of State. (2021, November 4). *DarkSide Ransomware as a Service (RaaS)*. Retrieved from Transnational Organized Crime Rewards Progra: <https://www.state.gov/darkside-ransomware-as-a-service-raas/>

- U.S. General Accounting Office. (1974). *Need to Consolidate Responsibility for Automatic Digital Network (AUTODIN) Terminals, Department of Defense*. Washington, DC: United States.
- U.S. Health and Human Services. (2023, February 22). *Clop Allegedly Targets Healthcare Industry in Data Breach*. Retrieved from HC3: Sector Alert: <https://www.hhs.gov/sites/default/files/clop-allegedly-targeting-healthcare-industry-sector-alert.pdf>
- U.S. Office of the President. (2017). *National Security Strategy of the United States*. Retrieved from <https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905.pdf>
- U.S. Senate Committee on the Judiciary. (2015, November 3). *Committee on the Judiciary*. Retrieved from Data Brokers – Is Consumers’ Information Secure?: [https://www.judiciary.senate.gov/meetings/data-brokers\\_is-consumers-information-secure](https://www.judiciary.senate.gov/meetings/data-brokers_is-consumers-information-secure)
- Uchill, J. (2021, August 21). *"What are the "wake-up" moments for cybersecurity policy since 2015?"*. Retrieved from Twitter: <https://twitter.com/JoeUchill/status/1429234014154526725>
- United States of America, Appellee, v. Robert Tappan Morris, Defendant-Appellant, 774 (U.S. Court of Appeals, Second Circuit March 7, 1991).
- United States Office of the President. (2017, August 1). *National Security Strategy of the United States of America*. Retrieved from <https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905.pdf>
- United States v Aaron Swartz, Crim. No. 1:11-cr-10260 (United States District Court of Massachusetts July 4, 2011).
- United States v Andrew Aurnheimer, 748 F.3d 525 (United States Court of Appeals, Third Circuit January 13, 2011).
- United States v David Nosal (The Supreme Court of the United States November 30, 2020).
- United States v. David Nosal, 676 F.3d 854 (United States Court of Appeals for the Ninth Circuit February 14, 2011).
- United States v. Gilberto Valle, Nos. 14–2710–cr, 14–4396–cr. (United States Court of Appeals, Second Circuit December 3, 2015).
- United States v. Ivanov, 175 F. Supp. 2d 367 (D. Conn 2001).



- UpGuard. (2021, August 23). *By Design: How Default Permissions on Microsoft Power Apps Exposed Millions*. Retrieved from UpGuard:  
<https://www.upguard.com/breaches/power-apps>
- US Food and Drug Administration. (2019, June 27). *FDA Warns Patients and Health Care Providers about Potential Cybersecurity Concerns with Certain Medtronic Insulin Pumps*. Retrieved from USFDA Division of Industry and Consumer Education : <https://www.fda.gov/medical-devices/safety-communications/cybersecurity-updates-affecting-medtronic-implantable-cardiac-device-programmers-fda-safety>
- Van Buren v United States, 940 F. 3d 1192 (The Supreme Court June 3, 2021).
- Varghese, S. (2020, December 15). *Backdoored Orion binary still available on SolarWinds website*. Retrieved from IT Wire: <https://www.itwire.com/business-it-news/security/backdoored-orion-binary-still-available-on-solarwinds-website.html>
- Vedantam, S. (2016, August 23). *Do You Read Terms Of Service Contracts? Not Many Do, Research Shows*. Retrieved from National Public Radio:  
<https://www.npr.org/2016/08/23/491024846/do-you-read-terms-of-service-contracts-not-many-do-research-shows>
- Verizon. (2022). *Data Breach Investigations Report*. Verizon.
- Vigna, G. (2019, October 10). *How to Think Like a Hacker*. Retrieved from DarkReading: <https://www.darkreading.com/vulnerabilities-threats/how-to-think-like-a-hacker>
- Vincent, J. (2021, August 24). *Check your permissions: default settings in Microsoft tool exposes 38 million user records online*. Retrieved from The Verge :  
<https://www.theverge.com/2021/8/24/22639106/microsoft-power-apps-default-permissions-settings-user-records-exposed-38-million-upgard>
- Vincent, J. (2021, November 5). *U.S. government offers \$10 million bounty for information on Colonial Pipeline hack*. Retrieved from The Verge:  
<https://www.msn.com/en-us/news/technology/us-government-offers-10-million-bounty-for-information-on-colonial-pipeline-hackers/ar-AAQlyJk>
- Vindu, G. (2016, December 14). *Yahoo Says 1 Billion User Accounts Were Hacked*. Retrieved from The New York Times:  
<https://www.nytimes.com/2016/12/14/technology/yahoo-hack.html>
- Visvanathan, S. (2002). The future of science studies. *Futures*, 91-101.
- Waddell, K. (2016, May 10). *The Computer Virus That Haunted Early AIDS Researchers*. Retrieved from The Atlantic:

- <https://www.theatlantic.com/technology/archive/2016/05/the-computer-virus-that-haunted-early-aids-researchers/481965/>
- Ward, M. (2015, August 20). *Ashley Madison: Who are the hackers behind the attack?* Retrieved from BBC News: <https://www.bbc.com/news/technology-34002053>
- Warren, S., & Brandeis, L. (1890). The Right to Privacy. *Harvard Law Review*, Vol. 4, No. 193.
- Weiner, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. Cambridge: MIT Press.
- Weiner, R., & Hawkins, D. (2018, June 19). *Hackers stole federal workers' information four years ago. Now we know what criminals did with it.* Retrieved from The Washington Post: [https://www.washingtonpost.com/local/public-safety/hackers-stole-feds-information-four-years-ago-now-we-know-what-criminals-did-with-it/2018/06/19/f42ff2b2-73d3-11e8-805c-4b67019fcfe4\\_story.html](https://www.washingtonpost.com/local/public-safety/hackers-stole-feds-information-four-years-ago-now-we-know-what-criminals-did-with-it/2018/06/19/f42ff2b2-73d3-11e8-805c-4b67019fcfe4_story.html)
- Weise, E. (2017, September 26). *A Timeline of Events Surrounding the Equifax Data Breach*. Retrieved from USA Today: <https://www.usatoday.com/story/tech/2017/09/26/timeline-events-surrounding-equifaxdata-breach/703691001/#>
- Well Go USA, Inc. v. Unknown Participants in Filesharing Swarm Identified by Hash B7FEC872874D0CC9B1372ECE5ED07AD7420A3BBB, 2012 WL 4387420 (Southern District of Texas, Houston Division September 24, 2012).
- Westcott, K., Loucks, J., Littmann, D., Whilson, P., Srivastava, S., & Ciampa, D. (2021). *Build it and they will embrace it: Consumers are preparing for 5G connectivity in the home and on the go.* Retrieved from The Deloitte Center for Technology, Media, and Communications: <https://www2.deloitte.com/us/en/insights/industry/telecommunications/connectivity-mobile-trends-survey.html>
- Whitacre, B., Gallardo, R., & Strover, S. (2014, November 1). Does rural broadband impact jobs and income? Evidence from spatial and first-differenced regressions. *The Annals of Regional Science*.
- White House. (2021, October 13). *Background Press Call on the Virtual Counter-Ransomware Initiative Meeting*. Retrieved from White House Briefing Room: <https://www.whitehouse.gov/briefing-room/press-briefings/2021/10/13/background-press-call-on-the-virtual-counter-ransomware-initiative-meeting/>
- White House. (2021, October 12). *Background Press Call on the Virtual Counter-Ransomware Initiative Meeting*. Retrieved from White House Briefing Room : <https://www.whitehouse.gov/briefing-room/press->

briefings/2021/10/13/background-press-call-on-the-virtual-counter-ransomware-initiative-meeting/

- White House. (2022, August 25). *OSTP Issues Guidance to Make Federally Funded Research Freely Available Without Delay*. Retrieved from White House Press Releases: <https://www.whitehouse.gov/ostp/news-updates/2022/08/25/ostp-issues-guidance-to-make-federally-funded-research-freely-available-without-delay/>
- Whittacker, Z. (2018, March 27). *Atlanta, hit by ransomware attack, also fell victim to leaked NSA exploits*. Retrieved from ZDNet: <https://www.zdnet.com/article/atlanta-hit-by-ransomware-attack-also-fell-victim-to-leaked-nsa-exploits/>
- Whittacker, Z. (2018, March 27). *Atlanta, hit by ransomware attack, also fell victim to leaked NSA exploits*. Retrieved from ZD Net: <https://www.zdnet.com/article/atlanta-hit-by-ransomware-attack-also-fell-victim-to-leaked-nsa-exploits/>
- Wilder, R. L. (1967). The Role of Intuition . *Science*, 605-610.
- Williams, J. (2020, December 16). *Durbin says alleged Russian hack ‘virtually a declaration of war’*. Retrieved from The Hill: <https://thehill.com/policy/cybersecurity/530461-durbin-says-alleged-russian-hack-virtually-a-declaration-of-war/>
- Winterton, J. (2015, July 16). How OPM Betrayed Me. *Slate*.
- Winterton, J. (2017, October 4). *Equifax: Continuing to Monitor Data-Broker Cybersecurity*. Retrieved from U.S. Senate Subcommittee on Privacy, Technology, and the Law: [https://www.judiciary.senate.gov/imo/media/doc/10-04-17 Winterton Testimony1.pdf](https://www.judiciary.senate.gov/imo/media/doc/10-04-17%20Winterton%20Testimony1.pdf)
- Winterton, J. (2017, October 4). *Equifax: Continuing to Monitor Data-Broker Cybersecurity*. Retrieved from U.S. Senate Subcommittee on Privacy, Technology, and the Law: [https://www.judiciary.senate.gov/imo/media/doc/10-04-17 Winterton Testimony1.pdf](https://www.judiciary.senate.gov/imo/media/doc/10-04-17%20Winterton%20Testimony1.pdf)
- Winterton, J. (2018, September 29). Keynote: Building a Better Hacker Future. *CactusCon*. Mesa, Arizona: <https://www.cactuscon.com/2018-keynote-jamie-winterton>.
- WIPO. (2001, January 25). *ADMINISTRATIVE PANEL DECISION Bruce Springsteen - v- Jeff Burgar and Bruce Springsteen Club Case No. D2000-1532*. Retrieved from WIPO Arbitration and Mediation Center: <https://www.wipo.int/amc/en/domains/decisions/html/2000/d2000-1532.html>

- WIPO. (2001, February 13). *ADMINISTRATIVE PANEL DECISION Celine Dion and Sony Music Entertainment (Canada) Inc. v. Jeff Burgar operating or carrying on business as Celine Dion Club Case No. D2000-1838*. Retrieved from WIPO Arbitration and Mediation Center:  
<https://www.wipo.int/amc/en/domains/decisions/html/2000/d2000-1838.html>
- Wolff, J. (2018). *You'll See This Message When It Is Too Late: The Legal and Economic Aftermath of Cybersecurity Breaches*. Cambridge: MIT Press.
- Wolff, J. (2018). *You'll See This Message When It's Too Late: The Legal and Economic Cost of Cybersecurity Breaches*. Cambridge: MIT Press.
- Wolff, J. (2020, March 30). *Ransomware Attacks are the Last Thing Hospitals Need Now*. Retrieved from Slate: <https://slate.com/technology/2020/03/ransomware-attacks-hospitals-coronavirus.html>
- Wolff, J. (2020, March 3). *Slate*. Retrieved from Ransomware Attacks are the Last Thing Hospitals Need Now: <https://slate.com/technology/2020/03/ransomware-attacks-hospitals-coronavirus.html>
- Wolff, J. (2021, October 10). (J. Winterton, Interviewer)
- Wolff, J. (2021, December 1). *How the NotPetya attack is reshaping cyber insurance*. Retrieved March 2023, from Brookings:  
<https://www.brookings.edu/techstream/how-the-notpetya-attack-is-reshaping-cyber-insurance/>
- Wolton, S. (2019). Are Biased Media Bad for Democracy? *American Journal of Political Science*, 548-562.
- Wood, L. (2022, October 27). *Mondelez, Zurich Settle NotPetya Dispute Before Trial Close*. Retrieved March 2023, from Law360:  
<https://www.law360.com/articles/1544141/mondelez-zurich-settle-notpetya-dispute-before-trial-close>
- Yeung, J. (2023, April 8). *The city without TikTok offers a window to America's potential future*. Retrieved from CNN Business:  
<https://www.cnn.com/2023/04/08/tech/hong-kong-tiktok-ban-intl-hnk-dst/index.html>
- Yu, E. (2009, September 30). *U.S. Government Finally Lets ICANN Go*. Retrieved from ZDNet: <https://www.zdnet.com/article/us-government-finally-lets-icann-go/>
- Zadeh, L. (1973). Outline of a New Approach to the Analysis of Complex Systems and Decision Processes. *IEEE Transactions on Systems, Man and Cybernetics*, 28-44.

Zetter, K. (2014, December 3). *Sony Got Hacked Hard: What We Know and Don't Know So Far*. Retrieved from Wired: <https://www.wired.com/2014/12/sony-hack-what-we-know/>

Zetter, K. (2014, December 3). Sony Got Hacked Hard: What We Know And Don't Know So Far. *Wired*.

Zorabedian, J. (2019, February 1). *Data Breach Fatigue Makes Every Day Feel Like Groundhog Day*. Retrieved from Security Intelligence: <https://securityintelligence.com/data-breach-fatigue-makes-every-day-feel-like-groundhog-day/>

## APPENDIX A

### DESCRIPTION OF EVENTS IN PRIMARY DATA SET

**Incident:** Target<sup>29</sup>  
**Date of incident:** 2013  
**Type:** Data breach  
**Number of records:** 70M  
**Attacker:** Unknown  
**Suspected intent:** Financial  
**Classification:** Personal

The 2013 Target data breach was one of the largest and most significant cyberattacks in history, affecting millions of customers and leading to widespread concerns about the security of personal and financial information held by retailers and other organizations.

The breach occurred during the holiday shopping season, beginning on November 27, 2013, and lasting until December 15, 2013. During this period, cybercriminals gained unauthorized access to Target's computer network and installed malware on the point-of-sale (POS) systems in its stores. The malware was designed to steal credit and debit card information from customers making purchases at the affected stores.

As a result of the breach, sensitive data from approximately 40 million credit and debit cards used in Target stores were compromised. The stolen information included names, card numbers, expiration dates, and, in some cases, the three-digit security codes (CVV/CVC) on the back of the cards. Additionally, personal information, such as email addresses and phone numbers, of approximately 70 million Target customers were also exposed.

The scale and impact of the Target data breach were significant, leading to widespread public concern about the security of personal information in retail

---

<sup>29</sup> Written in part by OpenAI. "ChatGPT (GPT-3.5) - AI Language Model." Accessed July 29, 2023. <https://www.openai.com/>.

transactions. The company faced heavy criticism for its handling of the incident, including delays in detecting and responding to the breach and insufficient security measures to protect customer data.

In the aftermath of the breach, Target took several measures to mitigate the damage and improve its cybersecurity practices. The company offered free credit monitoring services to affected customers and invested heavily in enhancing its data security and incident response capabilities.

The 2013 Target data breach highlighted the importance of robust cybersecurity measures, regular system monitoring, and proactive risk management. It also sparked discussions about the need for stronger data protection regulations and greater accountability for organizations that handle sensitive customer information – although little progress has been made in these areas. The Target breach is binned in the Personal Impacts category, due to the nature of the data exposed and lack of lasting impact on the company.



**Incident:** Bowman Dam<sup>30</sup>  
**Date of incident:** March 2016  
**Type:** Data breach  
**Number of records:** 70M  
**Attacker:** Iran  
**Suspected intent:** Financial  
**Classification:** Personal

In 2016, the Bowman Dam, a small dam located in Rye Brook, New York, United States, became the target of a cyberattack conducted by Iranian hackers (Cohen, 2021). The attack aimed to gain unauthorized access to the dam's control systems and infrastructure, raising concerns about the vulnerability of critical infrastructure to cyber threats.

Fortunately, the hackers' attempts to breach the dam's systems were not successful. The gate controls of the dam were offline when the dam was attacked (Berger, 2016). However, the incident shed light on the growing landscape of cyberwarfare and state-sponsored cyber-espionage activities. The attack on the Bowman Dam represented the potential risks to essential facilities and systems, emphasizing the importance of implementing robust cybersecurity defenses for critical infrastructure. In response to the incident, the United States charged individuals associated with the Iranian hacking group for their involvement in cyberattacks targeting not only the Bowman Dam but also various banks. None of the individuals has been apprehended by the U.S., and a lack of extradition treaty with Iran makes a future capture unlikely (CBS News, 2016). However, the charges serve as an open acknowledgement of U.S. ability to identify the source of internet-based attacks.

---

<sup>30</sup> Written in part by OpenAI. "ChatGPT (GPT-3.5) - AI Language Model." Accessed July 29, 2023. <https://www.openai.com/>.

This incident was one of the U.S.'s first indicators that critical infrastructure was vulnerable to cyberattack. For this reason, it is binned in the Political category of primary impacts. The Bowman Dam incident underscored the need for investment in and vigilance toward cybersecurity practices to safeguard essential facilities and sensitive systems from potential disruption and damage. The incident also highlights the necessity of international cooperation and collaboration to address and mitigate such cyber threats effectively, although these activities are challenging.

**Incident:** Sony Pictures Entertainment (SPE) hack  
**Date of incident:** November 2014  
**Type:** Data breach  
**Number of records:** N/A  
**Attacker:** North Korea (Zetter, 2014)  
**Suspected Intent:** Political activism (Keegan, 2014)  
**Classification:** Political

In a world where cyber-attacks happen nearly every day, the Sony Pictures Entertainment (SPE) incident stands out for many reasons. From the leering red skull graphic that announced the attack on the company's monitors (see Figure 3) to the exposure of celebrity squabbles and internal politics, North Korea's attack on SPE was theatrical, brazen, and shocking. But importantly, it was also the first time the federal government intervened on behalf of a private company.

In 2014, SPE released previews for a new movie, 'The Interview', wherein North Korean dictator Kim Jong Un (played by Randall Park) was soundly lampooned then set ablaze in a cartoonishly gruesome helicopter attack. As one might expect, the movie was perceived by the North Korean government as a grave insult, and after some (rather typical) saber-rattling rhetoric on the topic, they escalated their retaliation in a new way. North Korean military hackers broke into Sony's network, exfiltrated troves of sensitive and embarrassing information, then took control over Sony's desktop machines to display a threatening message (see Figure 4, (Zetter, Sony Got Hacked Hard: What We Know and Don't Know So Far, 2014)). This message didn't make specific financial demands, but instead made vague references to "our request" and prior "warnings" and threatened to release SPE's "secrets... to the world".

Five days later, several journalists received links and passwords to the “secret” information alluded to in the message (Wolff, You'll See This Message When It's Too Late: The Legal and Economic Cost of Cybersecurity Breaches, 2018). This included social security numbers, spreadsheets of salary information (which exposed huge gender and racial disparities in pay), and employee performance reviews. The folders also contained emails from the top levels of SPE, criticizing the company’s talent and even making racist references to then-president Barack Obama (Seal, 2015).

In the following days, more documents were released publicly in an attempt to force Sony to abandon the movie. Sony initially caved to these demands and cancelled the movie’s release, until the Obama administration unexpectedly stepped in, and strongly encouraged them to reconsider. The White House’s position was that acquiescing to an adversarial foreign dictator would be a security issue for the entire United States - not just the SPE company. "The cyber attack against Sony Pictures Entertainment was not just an attack against a company and its employees. It was also an attack on our freedom of expression and way of life,” explained U.S. Secretary of Homeland Security Jeh Johnson in a December 2014 statement. The FBI concurred, adding that “North Korea’s actions were intended to inflict significant harm on a U.S. business and suppress the right of American citizens to express themselves” (Federal Bureau of Investigation, 2014). Secretary Johnson continued to encourage the private sector to work closely with the Department of Homeland Security to share information on cyber threats and protect U.S. interests.

This taxonomy focuses not on the intent of the attackers, but where the harm is concentrated for each event. For the SPE breach, the effects impact all three categories to

a degree. Clearly, individuals whose personal information was exposed absorbed some harm, whether that information was a social security number, a salary, or an embarrassing email. One could potentially consider the entertainment industry as “infrastructure”; the way SPE does business was changed significantly due to this event. But the significant intervention of the federal government, on the grounds that the attack on SPE was in fact an attack on democratic principles, is compelling enough to list this event in the Political category. However, it remains a good example of the “complex layers of victims... different types of harm each of them suffered, and the lack of clarity around who was responsible for mitigating those harms” (Wolff, You'll See This Message When It's Too Late: The Legal and Economic Cost of Cybersecurity Breaches, 2018) that indicate the gaps and conflicts in incentives which this research aims to illuminate.

**Incident:** Ashley Madison data breach  
**Date of incident:** July 2015  
**Type:** Data breach  
**Number of records:** 36 million  
**Attacker (if known):** Impact Team hacking group  
**Suspected Intent:** Activism  
**Classification:** Personal

The Ashley Madison hack of 2015 is a perfect example of a “hactivist” group trying to drive cybersecurity towards one outcome, but due to misunderstanding the system’s incentives, ending up with the opposite outcome. The Ashley Madison company claimed to discreetly and securely connect individuals who were looking to have affairs. A hacking group called ‘Impact Team’ infiltrated the databases of Ashley Madison’s parent company, Avid Life Media (ALM), siphoning off user data, then threatened to publicly shame the company - not for its infidelity-based business model, but for its poor security practices. (There is no record of Impact Team’s activities prior to the Ashley Madison breach; indicating they may have formed specifically to carry out this activity (Ward, 2015).) ALM had publicly touted its “unprecedented” password protection and an extra-secure “full delete” option that would instantly and permanently remove all account data (for a price), but Impact Team found none of these to be true. The hactivists called out ALM’s lack of security basics and threatened to dump their user data, including names and credit card numbers of users who had paid for the “full delete” option, attempting to convince customers to leave the company and put it out of business (Wolff, 2018).

Impact Team’s plan backfired spectacularly. The company, whose security measures were overpromised and insufficient, suffered almost negligible costs (financial

and reputational). Ashley Madison's brand recognition grew, more traffic was driven to their website due to the news coverage of the breach, and the company hit the 60 million member mark in 2019 (Dodgson, 2019). A class action lawsuit was settled in 2017 for \$11.2M, down from the original filing of \$576M – a small percentage of the company's profits that year. The customers who had no responsibility for the site's security, however, suffered significant costs due to the breach. Individuals – some of which had been signed up as a joke - suffered upheaval in their personal and professional lives, including attempts at extortion (Segall, Ashley Madison users now facing extortion, 2015). It seems odd that Impact Team didn't consider (or care much about) the individual human consequences of "naming and shaming" the people that used Ashley Madison's website in their attempt to attack the company.

Because ALM had promised a certain level of security and failed, the Federal Trade Commission (FTC) was able to form a case against them. This case resulted in a \$1.6M fine and a consent decree that required the company implement a more rigorous security program. (\$1.6M may seem significant, but as a comparison, Ashley Madison offered \$10M to the City of Phoenix in 2010 to change the name of Sky Harbor International Airport to 'The Ashley Madison International Airport' for a period of five years (Fisher, 2010).) The FTC won a prestigious international award for their cross-border collaboration on the investigation of the Ashley Madison breach (FTC, 2017). The FTC's engagement may seem like a mere slap on the wrist, but worse, it could actually be a disincentive to promote security. The FTC is a user advocacy group which engages when companies are dishonest about their practices, whether it be the ability to prevent identity theft (Lifelock) or the labeling of "dolphin-safe tuna" (50 CFR part 216 Subpart

H). The FTC did not get involved in the Ashley Madison breach because of ALM's lack of security; they engaged because ALM promised certain security features that they did not deliver. Instead of an incentive to create robust security for users, the threat of FTC action may be acting as a disincentive to promise and enact specific security measures for customers.

As a side note, it turned out that not only were some of ALM's security practices fictitious, so were millions of potential female clients it offered. Data analysis of the Impact Team dump described a website "*like a science fictional future where every woman on Earth is dead, and some Dilbert-like engineer has replaced them with badly-designed robots*" (Newman A. , 2015).

The hackers' intention – to embarrass and shut down the company – ended up driving more business to the company, and harming individuals who were in no way responsible for the site's security practices. Standard individual remediations like credit monitoring were almost completely irrelevant to compensate for the very personal type of harm inflicted on the victims of this event. Therefore, this event is categorized as having primarily Personal impacts.



**Incident:** U.S. Office of Personnel Management (OPM) data breach  
**Date of incident:** July 2012 – April 2015 (Wolff, You'll See This Message When It's Too Late: The Legan and Economic Cost of Cybersecurity Breaches, 2018)  
**Type:** Data breach  
**Number of records:** 22.1 million (Nakashima, Hacks of OPM databases compromised 22.1 million people, federal authorities say, 2015)  
**Attacker:** Chinese intelligence organization (Sanders, 2015)  
**Suspected Intent:** Espionage (Pepitone, 2015)  
**Classification:** Political

One common feature of data breach disclosures is the increase in number affected after the initial discovery. OPM initially reported that 49,000 names and social security numbers were affected, but in a matter of weeks that number had risen dramatically – not just in the number of records (topping out at 22.1 million) but more importantly, in the types of data that were lost. An SF-86 form is the paperwork used to adjudicate a person's eligibility for a government security clearance. It includes not just name and SSN, but years of prior addresses, contact information and interviews of friends and family, information on hobbies, and deeply personal information used to assess a person's trustworthiness and susceptibility to foreign blackmail (Winterton, How OPM Betrayed Me, 2015). Sensitive non-personal information, such as security documents and network architecture information, was also lost in the breach (Chaffetz, Meadows, & Hurd, 2016). A fractured maze of government oversight and years of technical debt (outdated systems and procedures) were no match for a sophisticated nation-state attacker that obtained over 21 million records of cleared personnel and their families (Gallagher, 2015). It's estimated that every SF-86 processed by OPM since the year 2000 was siphoned off during OPM's exposure, which lasted at least 33 months, possibly longer.

Chinese hackers, acting on behalf of the Chinese government, were initially suspected, due to the type of malware used in the attack. In a rare turn of events, one Chinese hacker was arrested at a Los Angeles airport for his connections with the malware, although the OPM data breach was not mentioned in the court filings (Stecklow & Harney, 2019). He pleaded guilty and served 18 months in a U.S. federal prison, although the Chinese Ministry of Foreign Affairs claimed to have “no understanding” of this – or any other Chinese cyberattack.

The fact that OPM’s networks were insecure was not a surprise to anyone familiar with their systems, as they had racked up considerable deficiencies during Federal Information Security Management (FISMA) audits going back to 2007 (Esser, 2015). These deficiencies remained unmitigated for years, indicating a dangerous lack of incentive for government organizations to respond to FISMA audits. In fact, an audit by the U.S. Government Accountability Office (GAO) in 2021 recommended that OPM “establish a process for conducting an organization-wide cybersecurity risk assessment” - OPM agreed but responded that “due to resource challenges and that it plans to revisit this effort in the second quarter of FY 2021” (GAO, 2019) . As of November 2021, this action remains categorized as “Open” – and therefore unresolved - on the GAO dashboard<sup>31</sup>. Even though the OPM data breach reported in 2015 received considerable attention, and resulted in a few high-profile resignations, the incentive to create basic protections, such as an organizational risk assessment, remain absent.

---

<sup>31</sup> As of November 2021, the dashboard of GAO’s recommendations can be found here: [https://www.gao.gov/products/gao-19-384#summary\\_recommend](https://www.gao.gov/products/gao-19-384#summary_recommend)

The OPM breach was a failure of technology, but also a failure to understand the human, legal and social aspects of cybersecurity, which led to irreparable harm. One reason OPM's networks weren't immediately quarantined after the breach discovery was the fear that "senators will come for us", as OPM is responsible for hosting federal job search and employment websites (Chaffetz, Meadows, & Hurd, 2016). The full extent of the damage from this breach is still not clearly understood. A few fraudulent consumer loans have been attributed to OPM's data loss (Weiner & Hawkins, 2018), but given that they occurred four years after the breach, that claim is far from conclusive. The only remediation offered by OPM was a year of credit monitoring services for those individuals impacted – a clearly misaligned response, given the true level of the breach's impact, but what other *post facto* action would be more meaningful? One must assume that the bounty of information detailing the United States' cleared workforce has been absorbed by a foreign adversary, possibly combined with other sensitive data breaches such as Equifax and Ashley Madison, providing this adversary with intelligence value for decades to come. This event has been binned in the Political category.

**Incident:** Anthem<sup>32</sup>  
**Date of incident:** 2014  
**Type:** Data breach  
**Number of records:** 78.8M  
**Attacker:** Unknown  
**Suspected intent:** Financial  
**Classification:** Infrastructure

In February 2015, Anthem Inc, the second-largest health insurance company in the United States, reported a data breach that affected approximately 78.8 million individuals. The breach was discovered when Anthem's security team noticed an increase in database queries and detected unauthorized access to their information systems. The stolen information included names, dates of birth, Social Security numbers, medical IDs, addresses, email addresses, and employment information. This data breach was one of the largest healthcare data breaches in history and has raised significant concerns regarding the security of personal information in the healthcare industry.

Anthem Inc. responded to the data breach by providing two years of credit monitoring and identity theft protection to all affected individuals. (Victims were also offered \$50 in lieu of credit monitoring.) However, there are many ways to use stolen health care information that fall outside the purview of credit monitoring. The Anthem data breach was one of the first events to highlight the black-market value of health data. These data can be used more flexibly than stolen credit card numbers – for criminal activities such as falsified prescriptions or insurance fraud.

The protection of personal health records in the U.S. is covered by the Health Insurance Portability and Accountability Act (HIPAA). Because Anthem did not properly

---

<sup>32</sup> Written in part by OpenAI. "ChatGPT (GPT-3.5) - AI Language Model." Accessed March 14, 2023. <https://www.openai.com/>.

protect these sensitive data, they were fined \$16M USD by the U.S. Department of Health and Human Services, Office for Civil Rights (OCR) and were required to undertake substantial corrective actions, such as updated password policies and network monitoring to identify intrusions (U.S. Department of Health and Human Services, 2015). Anthem also paid \$115M USD in 2017 to settle a class-action lawsuit, the largest settlement for a data breach at the time (Pierson, 2017). This incident is categorized in the Infrastructure category.

**Incident:** Cyberattacks on election security  
**Date of incident:** 2015-2016  
**Type:** Data breach  
**Number of records:** N/A  
**Attacker:** Russian intelligence organizations  
**Suspected intent:** Espionage, election interference  
**Classification:** Political

Across much of the U.S., many electronic voting and tabulation systems are over a decade old. Some machines are no longer in production, so election officials often resort to platforms like eBay to find replacement parts, which in itself presents a security risk. Similarly, registration databases are often running on unsupported software, leaving them susceptible to intrusion due to the lack of regular security updates (Brennan Center for Justice, 2019). Voter registration databases were breached but there is no evidence that any data were changed, nor were the vote tallies impacted by this event (Greenberg J., 2016). What did change is the important realization of electoral systems as critical infrastructure with unique cybersecurity needs – not just the vote tabulation machines but supporting technologies such as state and local databases of voter information.

Since the attacks in 2016, all 50 states have improved their election security posture, many through cooperation with federal-level entities such as DHS or the National Guard (Root, Kennedy, & Souzan, 2018). While it would be preferable to not have had *any* intrusions into the U.S. electoral system, the response from state and local governments resulted in some meaningful improvements to election infrastructure, which aggregate to a stronger national security position.

It may not be possible to know the full intentions behind the electoral hacks in 2016, but it is possible that one of the primary motivations was not just to exploit the

vulnerabilities in machines, but to exploit human vulnerabilities as well. Specifically, human trust is required to make democracy functional, and trust in election systems has been damaged significantly in the past several years (Rinehart, 2020). The 2016 election system hacks are not the only avenue by which voter trust has been eroded; the problem is broader than cybersecurity and continues to evolve. But from the individual perspective, these attacks were powerful and effective, undermining many peoples' trust in democracy. This is an example of how the impact of cybersecurity events can extend beyond technology to the human and social components of the system, and how difficult it is to remediate once that happens.

Due to the suspected intent of those who target election systems, and the public trust in such systems, this event has been categorized as a Political incident.

**Incident:** Yahoo!<sup>33</sup>  
**Date of incident:** 2014  
**Type:** Data breach  
**Number of records:** 3 billion  
**Attacker:** Russia  
**Suspected intent:** Financial  
**Classification:** Personal

In 2014, Yahoo experienced one of the largest series of data breaches in history, with hackers stealing personal information from an aggregated 3 billion user accounts. The breach, which was not disclosed until 2016, exposed user data, including names, email addresses, phone numbers, birthdates, encrypted passwords, as well as security questions and answers (some unencrypted). For this reason, this event is binned in the Personal impacts category. Yahoo stated that the breach was likely conducted by state-sponsored actors and in 2017, the FBI officially charged four individuals with the breach and subsequent criminal activity using data from the breach (Goel & Lichtblau, 2017). This was the first time Russian intelligence officers had been formally charged by the U.S. federal government with computer abuse crimes.

The only individual arrested was a Canadian hacker named Baratov, who was not involved in breaching Yahoo's databases, but had been paid by the Russian hackers to use the data in a series of subsequent crimes (mostly involving spamming). Baratov is assumed to have made \$1.1M USD in the deal, but was charged with 5 years in prison and ordered to pay \$2.25M in restitution to the victims. Baratov maintains he did not know he was working for Russian spies, and has expressed regret for the incident (Moon, 2018).

---

<sup>33</sup> Written in part by OpenAI. "ChatGPT (GPT-3.5) - AI Language Model." Accessed March 14, 2023. <https://www.openai.com/>.



The Yahoo data breach had significant financial implications, with Verizon Communications, Inc. negotiating to purchase Yahoo in 2016 for \$4.8 billion. However, after the data breach was revealed, Verizon reduced its offer by \$350 million, citing concerns about the impact of the breach on Yahoo's business and reputation (Goel, Verizon Will Pay \$350 Million Less for Yahoo, 2017). The breach also resulted in legal action, with multiple class-action lawsuits filed against Yahoo alleging that the company failed to take adequate measures to protect user data (Leonhardt, 202). In addition to the \$80M USD class action lawsuit, Yahoo received a \$35M USD fine from the U.S. Securities and Exchange Commission for not disclosing the chain of breaches to its investors – the first application of the 1934 Securities Exchange Act to address cyber incidents (McAndrew, 2018).

**Incident:** Equifax  
**Date of incident:** May 2017  
**Type:** Data breach  
**Number of records:** 143.5 million  
**Attacker:** China (allegedly)  
**Suspected intent:** Espionage  
**Classification:** Political

In May of 2017, hackers infiltrated the Equifax network, creating approximately 30 points of intrusion. In late July, the intrusion was detected. It wasn't until September that victims of the breach were notified. Another important point in the timeline is the discovery that the vulnerability used to break into Equifax's sensitive databases could've been patched back in March of the same year – prior to the intrusion (Weise, 2017).

A U.S. Senate hearing was held in October 2017, in the Subcommittee for Science, Technology, and the Law, to address large-scale data breaches. Former Equifax CEO Richard Smith testified at this hearing, and claimed several times that Equifax was a security-centric company, and this breach happened due to a single person who did not install software patches in a timely fashion.

Shortly after the hearing, Smith retired with a payout of \$90M USD (which translates to \$0.63 per victim of the data breach). The breach victims were offered a year of free credit monitoring.

Credit monitoring companies are of particular concern from a security perspective because consumers can't opt out. An individual cannot stop their personal data from being collected by these companies, and therefore security practices (or lack thereof) cannot serve as a force in the marketplace. And when private companies are hacked, it's not just the individual and the company that are jeopardized; there's a national security

implication as well. Consider the value to a foreign adversary of 134.5 million credit records – the number of records lost in the Equifax data breach in 2017. An individual’s credit history tells a lot about them – their financial security, their patterns of life, what they’re willing to go into debt over. Scale that up by 134.5 million, and you get the same picture of a nation – the financial security and patterns of life of our entire population, and our strengths and vulnerabilities (Winterton, 2017). This is concerning in and of itself, but combined with other data breaches, is dire. Foreign hackers exfiltrated 21.5 million security clearance forms from the Office of Personnel Management – not just social security numbers, but exhaustive family information, financial details, and areas of technical expertise (Winterton, 2015). Combine these two data breaches, and a foreign adversary could have a very detailed picture of our national security enterprise and how it might be exploited. As Senator Jeff Flake said in a July 2015 hearing on the Experian data breach, “What may not be sensitive as one item may become sensitive in the aggregate”. Given the history of data breaches, across multiple sectors, we are facing the potential consequences of an aggregate-of-aggregates, a detailed multi-dimensional data set that affects individuals to nation states and all the levels in between. This is why the Equifax incident has been binned in the Political category.

**Incident:** Facebook  
**Date of incident:** 2018/2019  
**Type:** Data breach  
**Number of records:** >500M  
**Attacker:** Unknown  
**Suspected intent:** Unknown  
**Classification:** Personal Impacts

There are a wide variety of personal impact incidents. Some incidents, like the Target breach, expose credit card numbers and addresses, but others include much more personal data, such as photos, messages, and personal connections that span decades (Bell, 2023). Facebook is an example of the latter type of event.

Sometime prior to 2019 – Facebook has not been specific about the timing – information from over half a billion Facebook accounts was scraped from the platform, including full names, phone numbers, locations, and email accounts. The breach was discovered when the entire batch of data was posted to an online hacking forum in early 2021. Facebook decided not to inform the users that had been impacted, saying that it “did not have complete confidence in which users” were affected, but also that the data were already public and that it “was not an issue that users could fix themselves” (Bowman, 2021).

After a data breach has been made public, the breached company often sends an email to their customer base about how they care about their customers, that they take security seriously, and that they will do better in the future. These missives are often interpreted as theater, but Facebook did not even go through the motions, opting instead for an indifferent shrug. While it’s true that an impacted user could not remediate the impacts of a breach from a technical perspective, they should be made aware so they have

the opportunity to be more vigilant, to assess their personal threat profile and decide what actions they may need to take. Research has shown that an apology for a data breach negatively impacts investor behavior, but a justification may have a positive impact (Masuch, Greve, Trang, & Kolbe, 2022). Regardless, some notification should have been made. Some users may not care, but others may be in a vulnerable situation – with more than half a billion accounts compromised, it’s not hard to imagine that this group includes activists, journalists, people who are targets of abuse or stalking. It’s also worth noting that Facebook allows anyone over the age of 12 to create an account, so a percentage of the profiles exposed may well be from minors. Not giving these people an opportunity to increase their situational awareness shows a flagrant disregard for their well-being. Given that the data breach included phone numbers, the line “we can’t inform victims because we aren’t sure who was impacted” reads more like “we won’t inform victims because we don’t care enough to set up a cheap automated message system”. This comes after Facebook’s 2019 settlement with the FTC for \$5B USD over two specific abuses of their user population. First, Facebook requested user phone numbers, ostensibly to verify their account, but instead used them for advertising. Second, Facebook used 60 million users for facial recognition and tracking experiments without notifying them. (Facebook stated via their blog that the FTC settlement "is not only about regulators, it's about rebuilding trust with people".)

This incident came after the Cambridge Analytica scandal in which the social media company harvested data from users for undisclosed psychological profiling and subsequent political advertising. It was a challenge to pick a single Facebook event in the cybersecurity domain for this analysis, but losing data on over half a billion users and

responding with a shrug ended up being an irresistible choice. This incident is binned in the Personal Impacts category.

**Incident:** SolarWinds<sup>34</sup>

**Date of incident:** March 2020 – December 2020 (estimated)

**Type:** Malware

**Number of records:** 18,000 customers were impacted, of the 300,000 potential victims, but many were critical industry or government organizations (Harwell & McMillan, 2020)

**Attacker:** APT29 (Russian foreign intelligence service)

**Suspected intent:** Reconnaissance, espionage

**Classification:** Political

Like PowerApps, the SolarWinds event affected a centralized software platform that few people outside of IT had heard of, impacting a large web of industry and government entities. Unlike PowerApps, it was an intentional and sophisticated piece of malware created by a foreign adversary, and as such impacted a much more concerning group of targets, including several U.S. federal government agencies, including the Department of Homeland Security, the Treasury Department, and the Commerce Department, as well as private companies such as Microsoft and FireEye. The hackers were able to monitor and steal data from these networks for several months without detection (Cimpanu, 2020).

In a Congressional hearing, which took place in February 2021, SolarWinds executives attributed the security problem to an unnamed intern, who was allegedly using a weak password ("SolarWinds123"), which was left exposed on a GitHub server for years. At the 2021 RSA Conference, Sudhakar Ramakrishna, who became SolarWinds' CEO in January 2021, acknowledged that the comments made by him and the former CEO, Kevin Thompson, during a congressional hearing were "not appropriate." The blame-shifting onto the intern faced backlash from the cybersecurity community, who

---

<sup>34</sup> The March 14 2023 version of ChatGPT contributed to this section.

criticized the executives for not taking responsibility for their high-level budgetary and operational failures. Third-party forensic examiners later found that the leaked password likely had no role in the security breach.

It is worth noting that Kevin Thompson, who was CEO during the hack, left the company just before the breach was publicly disclosed and is among several SolarWinds executives accused of selling millions of dollars in company stock before the revelation. The company also faced a class-action lawsuit brought by SolarWinds stockholders alleging deception about their cybersecurity risks and practices. The SEC investigation is likely responsible for the prolonged dip in SolarWinds stock price (Johnson, SolarWinds CEO expresses regret for ‘blame the intern’ defense during Orion hack investigation, 2021).

This event is binned in the Political Impacts category, due to the suspected intent behind the attack and the prominent position of the victims.



**Incident:** Colonial Pipeline attack  
**Date of incident:** May 2021  
**Type:** Ransomware  
**Number of records:** N/A  
**Attacker (if known):** DarkSide (Russian-speaking cybercrime group)  
**Suspected Intent:** Financial  
**Classification:** Political

The first ransomware attack, back in 1989, traveled via floppy disk, and demanded that the computer's owner pay \$189 USD to recover access to their files (Waddell, 2016). Since then, ransomware attackers have shifted focus from individual computer users to larger and more lucrative targets in the infrastructure sector, such as hospitals, cities, and schools, leveraging the global connectivity of the internet and the anonymity of cryptocurrency to scale up their enterprises. As disruptive and pervasive as these attacks were, they didn't garner much interest from the federal government – that is, until the ransomware attack on Colonial Pipeline. Colonial transports 45% of the East Coast's fuel from the Gulf Coast (Sanger, Krauss, & Perlroth, 2021), and with the pipeline out of commission for 5 days, fuel reserves dried up and gas prices rose across the country (Turton & Mehrotra, Hackers Breached Colonial Pipeline Using Compromised Password, 2021).

The federal government responded. Whether it was the geographic extent that made a difference, the disgruntled populace, or the hot-button issue of fuel prices is unclear (Turner, 2021) (Wolff, 2021), but ransomware suddenly became a topic of unprecedented popularity at the federal level. The White House convened an international Counter-Ransomware Initiative, a virtual meeting including 30 countries to discuss “everything from efforts to improve national resilience, to experiences addressing the

misuse of virtual currency to launder ransom payments, our respective efforts to disrupt and prosecute ransomware criminals, and diplomacy as a tool to counter ransomware” (White House, 2021) - although notably neither China nor Russia were included in the talks.

In July, 2021, the U.S. leveraged the ‘Rewards for Justice’ program, established in 1984 to fight international terrorism, to offer a \$10M bounty for the Colonial Pipeline hackers (or anyone engaging in “malicious cyber activities against US critical infrastructure”) – a significant step farther than the federal government had engaged in any prior ransomware engagement (Vincent, U.S. government offers \$10 million bounty for information on Colonial Pipeline hack, 2021).

Interestingly, the Colonial Pipeline attackers offered up an apology note, claiming ***“We are apolitical... Our goal is to make money and not creating problems for society.”*** (Clark, 2021) The group also promised to vet their customers’ targets more carefully in the future (DarkSide runs a “Ransomware-as-a-Service” model, providing ransomware packages to affiliates which then perform the intrusion and share the payment with the provider (U.S. Department of State, 2021).)

Ransomware has shifted categories over the years. Initially, this type of attack was focused on individual harm, but accelerated to larger infrastructure-based targets through the late 2010’s. The attack on Colonial Pipeline, however, reflects another transition – for a potential myriad of reasons, from the political to the financial, ransomware now is viewed as a national security threat, so this particular incident is binned in the Political category.

**Incident:** Ransomware  
**Date of incident:** Various  
**Type:** Ransomware  
**Number of records:** N/A  
**Attacker:** Various  
**Suspected intent:** Often financial  
**Classification:** Infrastructure

Ransomware is a type of malicious software (malware) that encrypts a victim's files or locks them out of their computer or network, rendering the data inaccessible. The attackers behind ransomware demand a ransom, usually in cryptocurrency, in exchange for providing the decryption key or restoring access to the system. This digital extortion tactic has become increasingly prevalent in recent years, targeting individuals, businesses, and even government institutions worldwide.

For example, in March 2018 the city of Atlanta was hit by a ransomware attack (Whittacker Z. , 2018). The attack paralyzed the city for over a week – online services like utility or parking payments were frozen and court proceedings were halted as city officials scrambled to restore backups or implement manual methods as a fallback (Hutcherson, Six days after a ransomware cyberattack, Atlanta officials are filling out forms by hand , 2018). Forensic research on this incident showed that the city had been a victim of two ransomware attacks – one sophisticated, and one simple - because it simply hadn't patched its networks properly.

Ransomware typically spreads through various means, such as phishing emails, malicious downloads, or exploiting vulnerabilities in software and systems. Once it infiltrates a system, it quickly encrypts the victim's files, making them unusable without the unique decryption key held by the attackers. Victims are often faced with a difficult

decision: either pay the ransom and hope to regain access to their data or refuse to comply and risk losing valuable information. Ransomware attacks have had significant financial and operational impacts on individuals and organizations, across a wide variety of targets. Critical functions of several hospitals have been severely impacted when their computer systems were encrypted and held hostage by hackers, who demanded a ransom before de-encryption keys were provided. According to Josephine Wolff, assistant professor of cybersecurity policy at Tufts, more than 1,000 health-related organizations have been hit with ransomware since 2016, costing the health care system more than \$157 million (Wolff, 2020).

Given these impacts, ransomware has been binned primarily in the Infrastructure category. However, its effects on individuals and national security should not be underestimated. The overall scoring of ransomware in the three categories of harm reflects this broad and multi-faceted impact.

**Incident:** MS PowerApps  
**Date of incident:** May 2021  
**Type:** Data exposure  
**Number of records:** 38 million  
**Attacker:** None  
**Suspected Intent:** N/A  
**Classification:** Infrastructure

The Microsoft PowerApps incident is a clear example of how security problems can propagate across a wide spectrum of applications using the same core software. It's also a good example of how a poorly designed user interface, rather than insecure code, can lead to security problems.

Microsoft PowerApps isn't a familiar name to most people. It's a tool that allows software developers to quickly spin up new apps – both the front-facing site for users, and the back-end database to record and store user data, hosted in the cloud. But PowerApps is the engine behind a huge number of applications that manage things from flight records (American Airlines), medical information (COVID vaccine appointment scheduling), drug screening information including social security number (J.B. Hunt), and information of minor students (NYC schools) (UpGuard, 2021). Some of the data should be public, like the addresses of vaccination clinics, but some should be protected, like the names, addresses, and vaccination status of patients. PowerApps' default setting was to make data publicly accessible, and although their documentation explained how to change that setting to private, security researchers from UpGuard found thousands of databases that had been inadvertently left open to public sharing (Newman L. H., 38M Records Were Exposed Online—Including Contact-Tracing Info, 2021). UpGuard submitted a vulnerability report to Microsoft, but the Microsoft analyst responded that

this configuration was “considered to be by design” – putting the responsibility on the software users - and closed the case (UpGuard, 2021). As responsible stewards, UpGuard started to reach out to the affected companies, although nearly 50 separate corporate and government entities were involved (Vincent, 2021).

After some of the more severe cases were disclosed, Microsoft revisited the problem and changed PowerApps’ default setting to private, rather than public. As of March 2023, there is no evidence of the data being misused, but there is also no way to know who may have accessed sensitive data, and for what purposes, before the change was made. This incident highlights the challenges in securing highly distributed software systems. It is also an example of how a design choice – to use public sharing as default, rather than private – was convenient for application developers but ultimately dangerous for corporate users of those applications, as well as the individual end-users whose data had been exposed. This incident has been binned in the Infrastructure category.

**Incident:** NotPetya  
**Date of incident:** May 2017  
**Type:** Malware  
**Number of records:** N/A  
**Attacker:** Russia  
**Suspected intent:** Attack on Ukrainian systems  
**Classification:** Infrastructure

The uncontrolled spread and wide-ranging consequences of the NotPetya malware is a good example of why cybersecurity issues can't accurately be described by geographic boundaries. Like ransomware, this malicious code travels through a victim's computer network, encrypting each machine it can access. Unlike ransomware, however, there is no way to pay a ransom and recover the machines – they are permanently and irretrievably locked. The NotPetya malware was designed by Russian military hackers specifically to target Ukrainian accounting software, one part of a continuing Russian campaign to damage Ukrainian infrastructure (Greenberg A. , The Untold Story of NotPetya, the Most Devastating Cyberattack in History, 2018). (This attack was part of the extended pre-kinetic conflict before Russia's physical invasion of Ukrainian territory on February 24, 2022.)

As discussed extensively in Chapter 2, however, computer networks aren't segmented along geographic boundaries. Ukrainian computers connect to machines all over the world, and many of them share the same software configuration that allowed the NotPetya malware to take hold. NotPetya moved extremely quickly, knocking out significant infrastructure in Ukraine (banks, hospitals) and moving across the globe in a matter of hours. The malware seized machines in hospitals in Pennsylvania, halted global operations of the shipping company Maersk, and even reached back to infect the Russian

state-owned oil company Rosneft. Overall, the extent of NotPetya's damage, as well as the speed at which it inflicted the damage, were both astonishing.

One of NotPetya's unintended victims was Mondelez International, a global food company headquartered in the U.S., that makes Oreos and Triscuits, among other snack products. The company's operations were impacted for weeks, resulting in an estimated \$100M USD of damages. Mondelez filed an insurance claim, but it was rejected by their insurance company, who claimed the NotPetya malware was "hostile and warlike action" and therefore excluded from their coverage (Wood, 2022). Mondelez sued the insurance company for breach of contract (Mondelez International Inc. v. Zurich American Insurance Co, 2018). The companies went to court but settled for undisclosed terms on the final day of the two-week trial. Pharmaceutical company Merck & Co also sued their insurance company after their NotPetya claim was denied, claiming the malware was "an act of war or terrorism purportedly excluded from coverage under the Policies" ( Merck & Co. Inc. vs. Ace American Insurance Co. et al, 2022). The case was settled in favor of Merck for \$1.4B USD.

The NotPetya malware had global cross-sector impact, with an estimated worldwide cost of \$10B USD. The event also raised interesting policy questions, many of them around definitions of cyberwar (Roberts J. , 2021) and the potential ramifications of war-centric communication about cyber events (Wolff, 2021). Josephine Wolff, cybersecurity professor at Tufts University, suggests that when officials describe cybersecurity events with dramatic language, such as former Secretary of Defense Leon Panetta's "cyber Pearl Harbor" (Panetta Warns of 'Cyber Pearl Harbor', 2022) or Senator Dick Durbin's description of the SolarWinds campaign as "virtually a declaration of war"



(Williams, 2020), they are creating incentives for insurance companies to exclude cybercrimes from cyberinsurance policies. <so what?>

The origins of the NotPetya malware add an interesting angle to the story. The malware includes an exploit called Eternal Blue, a powerful hacking tool that had been stolen from the U.S. National Security Agency (NSA), and dumped online by a group calling themselves Shadow Brokers (Goodin, 2017). The path of the malware was not just from Russia to Ukraine, with countries across the globe as collateral damage, but from the U.S. to the cybercrime world at large, then adopted by Russia and built into the NotPetya global menace.

It is unusual for specific individuals to be charged in global cyber events, but in 2020 the U.S. Department of Justice charged six Russian Intelligence officers with a spate of cybercrimes, “wantonly causing unprecedented damage to pursue small tactical advantages and to satisfy fits of spite” (The United States Department of Justice, 2020), including the global NotPetya malware infection. (For its part, the Kremlin claims that U.S. attribution of NotPetya as a Russian offense are part of a “Russophobic campaign that is not based on any evidence.” (Wood, 2022)) It’s unlikely that the named Russian officers will be arrested, but the detailed attribution at the individual level does make a powerful statement about U.S. intelligence gathering abilities in cyberspace. It is also interesting to note that the U.S. DOJ cites several companies as being integral to their investigation, including Google, Twitter, Facebook, and “some private sector companies”.

The 2023 White House Cybersecurity Strategy includes the concept of a federal cyberinsurance “backstop” – federal funds that would provide relief to insurers in the

case of “catastrophic events”. Similar insurance backstop programs exist around natural disasters such as hurricanes, allowing insurers to provide more flexible policies without taking on irretrievable risk (Rundle, 2023).

For its wide ranging effects and origin story, this event is binned in the Political Impacts category.

## BIOGRAPHICAL SKETCH

Jamie Winterton is the Senior Director of Research Strategy for ASU's Global Security Initiative, where she designs interdisciplinary research concepts for defense and security missions. Jamie oversees GSI's large-scale research centers on cybersecurity, human/AI teaming, and disinformation within GSI, while creating and maintaining the organization's five-year strategic plan and financial projections. In October 2017 Jamie provided expert testimony on data breaches to the US Senate Subcommittee for Technology, Privacy, and the Law. Jamie chairs the university's DARPA Working Group and is a Faculty Fellow at the Center for Law, Science and Innovation at the Sandra Day O'Connor College of Law. Jamie has discussed cybersecurity issues in several media outlets, including the Washington Post, NPR, USA Today, the Christian Science Monitor, Slate, KCBS, KTAR, and the Associated Press.

Prior to joining ASU in August 2014, Jamie worked as a staff scientist for Lockheed Martin, where she developed and directed projects in electro-optical and radar processing/analysis for multiple military and government organizations. Jamie's work in optical characterization of materials, high-fidelity physics-based 3D modeling and simulation, and exploitation of synthetic aperture radar (SAR) data has been recognized for both innovation and mission utility. Jamie received her Bachelor's degree in Physics from Arizona State University and her Master's degree in Physics from the University of Massachusetts, Amherst.