

Cyberhealth and Informational Wellbeing

By John Michael Thornton
Darwin College
April 2019

This dissertation is submitted for the degree of Doctor of Philosophy

Preface

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text

It does not exceed the prescribed word limit for the relevant Degree Committee.

Abstract

Title: Cyberhealth and Informational Wellbeing

Author: John Michael Thornton

In this dissertation, I present a new framework for conceptualizing the digital landscape inspired by the field of public health. I call this framework Public Cyberhealth. This framework is an alternative to the dominant cybersecurity paradigm, which frames cyberspace as a digital battleground. I argue that the *philosophy of public health* can be useful for thinking about the normative justification for—and ethical limits on—government intervention in cyberspace, while *public health policy and institutions* can serve as examples of how to manifest these higher principles (e.g. the WHO, ethical review boards). This Public Cyberhealth framework takes seriously non-malicious threats to network robustness and resilience (e.g. human error, buggy code, natural disasters), highlights the impact of network threats and interventions on health and wellbeing, and is more thoughtful about protecting individual rights compared to the dominant cybersecurity lens typically used by policymakers and IT professionals.

In addition to defining this alternative framework, I demonstrate how it may be used in three contexts. First, I explore how thinking about the digital landscape like a public health expert can help one to understand the role public goods play in maintaining robust digital networks. Second, I explore how this framework can help one to create policies which adequately account for how digital technologies impact health. And third, I define a theory of “informational wellbeing,” which seeks to capture the myriad of ways in which digital information and its use, control, accuracy, and accessibility impact personal wellbeing. The Public Cyberhealth framework is not only a useful and coherent way of thinking about technology policy, but also reveals interesting and surprising things about the nature of health, wellbeing, and identity in the digital age.

Acknowledgments

I am grateful to the Department of History and Philosophy of Science, the University of Cambridge, and Darwin College for the opportunity to create this work.

My deepest appreciation goes to my supervisor, Dr Stephen John, for his guidance throughout this process. I thank him for his patience, encouragement, and meticulous feedback. I would also like to thank my advisor Professor Tim Lewens and Dr Anna Alexandrova for their insightful comments on specific chapters.

Lastly, I thank my parents for always encouraging my intellectual curiosity and my wife Isabel for being my partner in this and every adventure. I dedicate this thesis to them.

Contents

Preface.....	3
Abstract.....	4
Acknowledgments.....	5
Contents.....	6
List of Tables	7
Introduction	8
0.1 Structure	16
0.2 A Note on Scope	19
Chapter 1: Public Goods for Cyberhealth	21
1.1 Externalities, Public Goods, and Networks	23
1.2 The Containment of Malware.....	29
1.3 Obligation and Public Goods for Cyberhealth	48
1.4 Conclusion	63
Chapter 2: Two Levels of Abstraction	66
2.1 Levels of Abstraction	67
2.2 The Cybersecurity LoA.....	72
2.3 The Public Cyberhealth LoA	82
2.4 Two Applications of Public Cyberhealth	88
2.5 Conclusion	102
Chapter 3: Health and Cyberhealth	104
3.1 Health and Public Health—Privileged Categories.....	106
3.2 Poor Cyberhealth and Public Health	111
3.3 Definitions of Disease	125
3.4 Poor Cyberhealth as Pathology	128
3.5 Conclusion	143
Chapter 4: Informational Wellbeing	144
4.1 The Capability Approach to Wellbeing	148
4.2 Informational Capabilities and Wellbeing.....	153
4.3 A Theory of Informational Wellbeing.....	163
4.4 Value of the Theory	166
4.5 Generalizability	172
4.6 Informational Wellbeing as a Fundamental Capability	174
4.7 Conclusion	189
Conclusion.....	191
Bibliography.....	197

List of Tables

Table 1	74
Table 2	83
Table 3	93
Table 4	112
Table 5	167
Table 6	168

Introduction

“[A] problem well put is half-solved...The way in which the problem is conceived decides what specific suggestions are entertained and which are dismissed; what data are selected and which rejected; it is the criterion for relevancy and irrelevancy of hypotheses and conceptual structures.”¹ – John Dewey

Today was a typical Wednesday. I woke up to my iPhone’s alarm and then read the news on half a dozen websites before getting up to make coffee. While I ate breakfast, I listened to music played on an Amazon Echo. As I finished breakfast, the gas company called about installing a smart meter. They verified my identity using my email address and phone number and sent the confirmation to my email account. As I checked my email, Gmail automatically reminded me to follow-up with my supervisor about an email I sent the previous week, and I noticed my accountant had sent over my tax return to be digitally signed. I paid for lunch with a contactless ID card and then bought a coffee with a digital wallet at the art museum. While walking around the museum, I took a couple of pictures with my phone, which were immediately uploaded to cloud storage, and I read the biography of the artist John Everett Millais on Wikipedia. Despite having deleted my Facebook account a few months back, never using Twitter or other social media, and engaging in the decidedly Victorian activity of a walk in a museum, I had interacted with over a dozen different digital devices and services before afternoon tea. Not to mention the legion of CCTV cameras capturing my stroll down the street and the half dozen or so state intelligence agencies who may have been tracking my banal online activity.

It is hard to overstate the importance of digital networks and digital information to life in the 21st century. For many the internet is where they socialize, fall in love, express themselves, and learn about the world. As of 2017, nearly half the

¹ John Dewey, *Logic: The Theory of Inquiry* (New York: Henry Holt & Co., 1938), 108.

world's population has internet access.² Furthermore, even many of those who do not personally use the internet still rely on the digital networks that underpin the global economy, voting systems, and critical infrastructure, such as telecommunications, the electrical grid, banking and finance, national defence, the oil and gas industries, transportation services, the water supply, and emergency services.³ In the coming decades, we will likely rely on digital technologies to an even greater degree as smart homes and cities, autonomous vehicles, biometric identification, and personal robots become more commonplace. Robust and resilient digital networks and devices are essential to the stability of modern life. In this dissertation, I will explore the myriad of ways in which these networks impact our lives and propose a new framework for conceptualizing the digital landscape inspired by the philosophy of public health.

As I will use the terms 'digital network,' 'robustness,' and 'resilience' throughout this dissertation, it will be helpful to have a clear definition of each term up front. For the purpose of this work, a 'digital network' is defined as *any network which transmits digital information between nodes*. A 'node' can be a person using a digital technology, such as a computer or smartphone, or it can be a digital technology which operates in a largely autonomous fashion, such as a server or sensor. Meanwhile, 'robustness' is defined as *the degree to which networks and devices are able to withstand malicious (e.g. malware), accidental (e.g. human error, buggy code), and natural threats (e.g. hurricanes)*. A robust network will also be a secure network by this definition, capable of repelling adversary-based attacks and keeping information safe. Lastly, by 'resilience' I mean *a network's ability to recover from a successful cyberattack, damaging accident, or environmental problem*. Often this will entail having adequate back-up systems or the ability to bypass inoperable or insecure systems.

Generally speaking, the primary goal of a digital network is to allow digital information to be accurately and confidentially transferred between nodes to fulfil a given purpose (e.g. visiting a website, sharing a picture, paying a bill). Digital

² International Telecommunications Union, "ICT Facts and Figures 2017," International Telecommunications Union, July 2017, <https://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2017.pdf> (accessed Jan. 13, 2019).

³ Fred Kaplan, *Dark Territory* (New York: Simon and Schuster Paperbacks, 2016), 41.

networks can fail to fulfil this goal for a number of reasons. Network infrastructure (including individual nodes) can be infected with malware, buggy software can malfunction, individuals can give away their passwords in phishing scams, or network infrastructure can be physically disrupted or destroyed—recent examples of the latter include sharks biting through transatlantic cables,⁴ a Georgian grandmother cutting off Armenia’s internet while scavenging for copper,⁵ and the disruption of Puerto Rico’s networks in the wake of Hurricanes Maria and Irma in 2017.⁶ Sometimes these failures are small in scale, for instance a bug may only affect a single user. In other cases, these failures can crash substantial portions of the internet, as in the case of the 2016 DDoS attack on part of the Domain Name System (DNS),⁷ or shutdown critical infrastructure, as was the case with the WannaCry ransomware attack in 2017 that led to the temporarily closure of many NHS facilities.⁸ However, it is important to note that human errors or buggy code can be just as devastating as a skilled hacker. A few months after the 2016 Mirai botnet attack on the DNS, a fat-finger mistake by an

⁴ Robert McMillan, “Sharks Want To Bite Google’s Undersea Cables,” *Wired*, Aug. 15, 2014, <https://www.wired.com/2014/08/shark-cable/> (accessed Dec. 19, 2018).

⁵ Tom Parfitt, “Georgian Woman Cuts Off Web Access to Whole of Armenia,” *The Guardian*, April 6, 2011, <https://www.theguardian.com/world/2011/apr/06/georgian-woman-cuts-web-access> (accessed Nov. 6, 2018).

⁶ Nick Thieme, “After Hurricane Maria, Puerto Rico’s Internet Problems Go from Bad to Worse,” *PBS*, Nov. 23, 2018, <https://www.pbs.org/wgbh/nova/article/puerto-rico-hurricane-maria-internet/> (accessed Nov. 6, 2018).

⁷ Berkeley Lovelace Jr. and Antonio José Vielma, “Friday’s Third Cyberattack on Dyn ‘Has Been Resolved,’ Company Says,” *CNBC*, Oct. 21, 2016, <https://www.cnbc.com/2016/10/21/major-websites-across-east-coast-knocked-out-in-apparent-ddos-attack.html> (accessed Feb. 15, 2019).; Ethan Chiel, “Here Are the Sites You Can’t Access Because Someone Took the Internet Down,” *Splinter*, Oct. 21, 2016, <https://splinternews.com/here-are-the-sites-you-cant-access-because-someone-took-1793863079> (accessed Feb. 15, 2019).

⁸ Comptroller and Auditor General, “Investigation: WannaCry cyber attack and the NHS,” National Audit Office, HC 414 Session 2017–2019, April 25, 2018, <https://www.nao.org.uk/wp-content/uploads/2017/10/Investigation-WannaCry-cyber-attack-and-the-NHS.pdf> (accessed Dec. 20 2018).

Amazon.com employee overloaded the company's popular cloud services and disrupted hundreds of thousands of websites and internet connected products, such as smart light bulbs and thermostats.⁹

Historically, philosophers have played a relatively minor role in defining how we think and talk about the digital landscape. Rather, the way we understand digital technologies, information security, network failure, and the various obligations of states, corporations, and individuals in cyberspace has largely been defined by computer scientists, engineers, defence agencies, novelists, filmmakers, online communities, and lawyers.¹⁰

One of the most significant conceptual frameworks we use to understand these technologies can be called the cybersecurity framework, or cybersecurity lens. At a high level, this lens frames the digital landscape as a battleground between good guys and bad guys.¹¹ The good guys are law enforcement, your own country's cyber defence forces, cybersecurity experts, and corporate IT departments. The bad guys, meanwhile, are cybercriminals, hackers, and other states' intelligence agencies.¹² In this hostile landscape, companies, individuals, and states are generally responsible for securing their own digital castle or homestead. In this sense, the framework can

⁹ Amazon.com, "Summary of the Amazon S3 Service Disruption," *Amazon.com*, <https://aws.amazon.com/message/41926/> (accessed Feb. 15, 2019).; Elizabeth Weise, "Massive Amazon Cloud Service Outage Disrupts Sites," *USA Today*, 28 February 2017, <https://www.usatoday.com/story/tech/news/2017/02/28/amazons-cloud-service-goes-down-sites-scramble/98530914/> (accessed April 15, 2017).; Darrell Etherington, "Amazon AWS S3 Outage Is Breaking Things For A Lot Of Websites And Apps," *TechCrunch*, Feb. 28, 2017, <https://techcrunch.com/2017/02/28/amazon-aws-s3-outage-is-breaking-things-for-a-lot-of-websites-and-apps/> (accessed April 2, 2017).

¹⁰ I do not mean to imply that there is no philosophy being written on these topics. There is certainly philosophical work on privacy, personal information, and obligations in cyberspace. However, generally speaking there are far fewer philosophers working on these topics than lawyers, computer scientists, etc.

¹¹ Kaplan, *Dark Territory*.

¹² Panayotis A. Yannakogeorgos and Adam B. Lowther, *Conflict and Cooperation in Cyberspace: The Challenge of National Security*, eds. Panayotis A. Yannakogeorgos and Adam B. Lowther (Boca Raton: Taylor & Francis, 2014).

loosely be described as ‘feudal.’ In order to secure their digital information and networks, states often employ offensive cyber capabilities, and a proposed bipartisan piece of legislation in the United States would empower corporations to do the same.¹³ Often these offensive capabilities are deployed with little regard for the broader effects on the network.¹⁴ While this way of conceptualizing the digital landscape emerged out of the Cold War and the need to protect military (and later corporate) secrets, it is not limited to cybersecurity experts.¹⁵ The popularity of this framework is in part due to its prevalence in popular media, including dozens of films, such as *Tron* (1982), *War Games* (1983), *Sneakers* (1992), *Hackers* (1995), *The Net* (1995), *Enemy of the State* (1998), *The Matrix* trilogy (1999, 2003, 2003), *Swordfish* (2001), *Live Free or Die Hard* (2007), and *Skyfall* (2012).

In some contexts, the cybersecurity lens is useful and appropriate. There *are* cybercriminals trying to steal personal information, military secrets, and corporate intellectual property, and *one* way to prevent that loss is by identifying those threats, arresting the perpetrators, and using offensive cyber weapons to disrupt adversaries’ systems. But this cybersecurity lens—inspired by law enforcement, criminal justice, and military intelligence—can also act as a set of blinders, leading one to underappreciate aspects of network robustness and resiliency which do not fit this adversarial narrative. For example, the cybersecurity framework does not address network failures caused by natural disasters, human error, or buggy code; the framework’s feudal notion of responsibility makes it ill-suited for mobilizing the collective action needed to respond to large-scale malware outbreaks; and it tends to downplay or ignore the myriad (but sometimes subtle) ways that poor network robustness and resiliency impact personal wellbeing. Together, these limitations suggest the cybersecurity lens is too narrow to serve as an overarching way for policymakers to conceptualize the digital landscape.

While the cybersecurity lens may not be up to the task, there are reasons to believe that an overarching framework for conceptualizing our relationship to digital

¹³ Active Cyber Defense Certainty Act, HR 4036, 115th Congress, 1st session (2017), <https://www.congress.gov/115/bills/hr4036/BILLS-115hr4036ih.pdf>.

¹⁴ Peter Trim and David Upton, *Cyber Security Culture: Countering Cyber Threats Through Organizational Learning and Training* (Farnham: Gower Publishing, 2013).

¹⁵ Kaplan, *Dark Territory*.

information and networks is needed to create effective, consistent, and justifiable technology policies. Without an overarching guide, technology policy has often been crafted as an ad hoc reaction to the latest security crisis. In turn, this has often led to conflicting cybersecurity strategies, spotty protection of individual rights, and ineffective international collaboration on matters of mutual concern.¹⁶ In the proper context, the cybersecurity lens is very useful, but all too often it is treated as the one true way to conceptualize the digital landscape.

As mentioned above, the alternative framework I will present in this dissertation is inspired by the philosophy of public health. Although an approach inspired by the philosophy of public health may seem odd at first, the language of epidemiology and public health has been used to describe digital networks for decades. In 1993 David Chess, Jeffrey Kephart, and Steve White of the IBM Thomas J. Watson Research Center fleshed out the first biological analogy for self-replicating computer viruses as a way to think about using the tools of epidemiology to improve the health of computer networks.¹⁷ Brent Rowe, Tony Lentz, and Michael Halpern of RTI International expanded on this analogy to create an entire framework for comparing types of cyberattacks to their biological counterparts, borrowing ideas from public health like communicability, risk behaviours, and environmental exposures.¹⁸ This effort to systematize threats was aimed at helping cybersecurity experts categorize types of risk, suggest potential prevention strategies, and understand individuals' risk preferences.¹⁹

¹⁶ Elaine Sedenberg and Deirdre Mulligan, "Public Health as a Model for Cybersecurity Information Sharing," *Berkeley Technology Law Journal* 30, no. 3 (2015): 6.; Nazli Choucri, Stuart Madnick, and Priscilla Koepke, "Institutions for Cyber Security: International Responses and Data Sharing Initiatives," Working Paper Cybersecurity Interdisciplinary Systems Laboratory, MIT, October 2016, <http://web.mit.edu/smadnick/www/wp/2016-10.pdf> (accessed May 29, 2017).

¹⁷ Jeffrey O. Kephart, Steve R. White, and David M. Chess, "Computers and Epidemiology," *Spectrum IEEE* 30, no. 5 (1993): 20.

¹⁸ Brent Rowe, Michael Halpern, and Tony Lentz, "Is a Public Health Framework the Cure for Cyber Security?," *Cross-Talk* 25, no. 6 (2012): 31-32.

¹⁹ *Ibid.*, 30.

More recently, Robert May and Alun Lloyd have explored the similarities between how viruses spread in human and computer networks. Focusing on the work of physicists Romualdo Pastor-Satorras and Alessandro Vespignani, May and Lloyd identified that epidemic spreading in scale-free networks like the World Wide Web bore significant similarities to the spread of infection in sexual-partner networks and suggested future study of computer networks for those seeking to manage epidemics.²⁰ These similarities have proved useful in helping to develop more flexible and automated cyber defences based on epidemiological strategies to combat infection.²¹ One specific application is Scott Charney's development of device health certificates as a way to encourage digital herd immunity.²² While these comparisons of malware to biological viruses are useful and revealing, in this work I will deepen this analogy in new and illuminating ways by demonstrating that a public health inspired approach is not only useful for describing malware, but can also serve as an overarching guide to create technology policy which promotes health and wellbeing, while protecting individuals' rights.

I call this alternative approach "Public Cyberhealth." By cyberhealth I mean *the robustness and resiliency of a network, be it a home network or the entire internet*. Whereas cybersecurity only refers to problems caused by adversaries,²³ cyberhealth takes into account a network's ability to withstand and recover from buggy code, natural disasters, and human error, in addition to malicious threats. In contrast to the cybersecurity framework described previously, Public Cyberhealth addresses both malicious and non-malicious cyber threats; highlights the ways in which cyber threats and interventions impact health, wellbeing, and individuals' rights; and uses the

²⁰ Alun Lloyd and Robert May, "How Viruses Spread Among Computers and People," *Science* 292 (2001): 1316.

²¹ United States Department of Homeland Security, "Enabling Distributed Security in Cyberspace," United States Department of Homeland Security, 2011, <https://www.dhs.gov/xlibrary/assets/nppd-cyber-ecosystem-white-paper-03-23-2011.pdf>.

²² Scott Charney, "Collective Defense: Applying the Public-Health Model to the Internet," *Security & Privacy IEEE* 10, no. 2 (2012): 55.

²³ P. W. Singer and Allan Friedman, *Cybersecurity and Cyberwar* (Oxford: Oxford University Press 2014), 34.

philosophy of public health to understand the normative justification for—and the ethical limits on—government interventions in cyberspace. Furthermore, this approach is not only a useful and coherent way of thinking about technology policy, but also reveals interesting and surprising things about the nature of health, wellbeing, and identity in the digital age.

0.1 Structure

The four chapters in this dissertation work together to demonstrate the coherence and utility of the Public Cyberhealth framework. First, in Chapter 1, I will demonstrate one application of the public health approach—the provisioning of public goods. While public goods are by no means unique to public health, the public goods which are relevant to promoting robust digital networks are similar to those which are relevant to public health, including monitoring programs, regulatory regimes, and herd immunity. Specifically, using the example of the Conficker computer worm, I will demonstrate how the containment of infectious diseases suggests useful ways to think about the containment of malware. The first similarity between these two public goods is that they are *Global Public Goods* (GPGs)—public goods which require global cooperation and produce global benefits. GPGs are unlike other public goods in that they require states to work together in a context where no one state possesses the authority to compel another state to act. A prime example is the World Health Organization’s monitoring of infectious diseases. The second similarity between these public goods is that they are *participatory public goods*. Unlike many public goods, a participatory public good requires the beneficiaries to participate in the creation of the good beyond merely paying their taxes. The paradigmatic participatory public good is herd immunity. In the public health context herd immunity requires one to get vaccinated, while in the digital context one must patch one’s devices. In both cases, once a certain percentage of the network is protected, it is significantly more difficult for infections to spread. While the containment of malware and the containment of communicable disease are not perfectly analogous, the philosophy of public health can help us think through how to fairly and adequately produce the public goods which promote cyberhealth.

Having demonstrated the utility of the public health approach for addressing a classic cybersecurity problem like malware, in Chapter 2 I will more formally and

comprehensively define the Public Cyberhealth framework using the method of Levels of Abstraction (LoA). This method, popularized by Luciano Floridi,²⁴ is based on the concept that whenever we consider a given system, such as digital networks, we highlight certain variables, observables, and behaviours that are relevant to our goal and downplay those we deem irrelevant. The variables, observables, and behaviours we use to model the system comprise a ‘level of abstraction’ or LoA. This method is useful for clearly defining a framework’s purpose, outlining the assumptions of one’s framework, and for comparing the utility and coherence of competing frameworks.²⁵ In this chapter, I will use the method of LoA to compare the competing frameworks of Cybersecurity and Public Cyberhealth, demonstrating the inadequacy of the former and the promise of the latter.

In Chapters 3 and 4, I will then further flesh out the Public Cyberhealth framework by exploring what it means to conceptualize the digital landscape with health and wellbeing front and centre. First, in Chapter 3, I will explore the ways in which poor cyberhealth impacts health. Demonstrating that poor cyberhealth has the ability to significantly impact health, 1) strengthens the normative justification for viewing the digital landscape through a public health inspired LoA and 2) strengthens the justification for governments to invest more heavily in the public goods for cyberhealth discussed in Chapter 1. In the first half of the chapter, I will outline a number of straightforward ways in which poor cyberhealth can impact our health, including the insecurity of critical infrastructure (e.g. dams, water treatment, emergency services, etc.), medical devices (e.g. insulin pumps, defibrillators), and hospital infrastructure. Despite their importance to health, these systems and devices are often running old operating systems and remain unpatched. I will then explore how a lack of reliable network access can exacerbate existing health inequalities caused by poverty and geographic isolation; poor and rural populations without network access may receive a lower level of healthcare, have less access to their doctors, and may miss important public health information.

²⁴ Luciano Floridi, *The Philosophy of Information* (Oxford: Oxford University Press, 2011).; Luciano Floridi, *The Ethics of Information* (Oxford: Oxford University Press, 2013).

²⁵ Luciano Floridi, *The Philosophy of Information*.

In the second half of Chapter 3, I will then look at more unusual cases where cyberhealth should itself be considered a *constituent* part of what it means to be healthy. I will argue that when an artificial component (e.g. a digital pacemaker) is coupled to a biological system, one should assess the functioning of this coupled system when determining whether an individual suffers from a pathology. This may require us to consider an individual with a properly functioning pacemaker as essentially disease-free, while simultaneously considering a malware infected pacemaker a pathology in and of itself. These cases raise interesting questions about bodily integrity, the regulation of medical devices, models of technology ownership, and what it means to be healthy in the 21st century. While the philosophy of public health and biology do not provide a single clear way to answer these questions, they provide the vocabulary and theory to explore these topics in nuanced ways which are sensitive to how policy, interventions, and inaction impact health, wellbeing, and rights.

Having demonstrated the significant ways in which poor cyberhealth can impact health, in Chapter 4 I will explore how one captures the impact of poor cyberhealth (i.e. poor network robustness and resiliency) on the broader notion of wellbeing. In this chapter, I will propose a theory of “informational wellbeing,” based on Amartya Sen and Martha Nussbaum’s capability approach to wellbeing which seeks to capture the myriad ways in which digital information and its use, control, accuracy, and accessibility impact personal wellbeing. I will specifically focus on Nussbaum’s version of the capabilities approach, which is centred around one’s ability to achieve ten fundamental capabilities, such as the ability to live a natural life span, the ability to be healthy, and the ability to have bodily integrity. I will argue that an individual has a high degree of informational wellbeing to the extent they have achieved the ‘informational capabilities and functionings’ (e.g. the ability to control access to their personal information, etc.) which are necessary to achieve fundamental human capabilities (e.g. the ability to live a normal lifespan). Using my theory, one can assess how a given technology policy or intervention impacts a person or group’s wellbeing. This is useful for assessing the success of existing policies in wellbeing terms and for creating new policies and products which promote wellbeing.

Capturing the impact of poor cyberhealth on wellbeing is particularly relevant for the discussion of public goods for cyberhealth discussed in Chapter 1. Not only can a theory of informational wellbeing help one to understand the value of such

goods—strengthening the justification for their production—but it can help one to set thresholds for determining when a given good has been adequately produced. For example, one can determine that a state’s digital voting infrastructure is sufficiently secure when it no longer prevents one from being able to achieve the fundamental human capability of being able to participate politically.

Finally, I will argue that in addition to being instrumental to the fundamental capabilities listed by Nussbaum, in at least two ways the concept of informational wellbeing should itself be thought of as a fundamental capability on par with the other capabilities on Nussbaum’s list. First, in a logical extension of my argument in Chapter 3, I will argue that if one accepts that digital devices can be considered a part of someone’s body, then the proper functioning of these devices (and the networks they use) partly constitutes an individual’s health status. Second, I will argue that informational wellbeing is essential to the ability to define one’s self, an ability which underlies Sen and Nussbaum’s eudaimonism. Using Daniel Dennett’s narrative theory of the self, I will illustrate the important role personal information plays in defining the self. Then, I will argue that when the self is conceived of informationally, having an adequate degree of informational wellbeing partly constitutes the ability to define one’s self. Both cases suggest that eudaimonists like Sen and Nussbaum should consider having an adequate level of informational wellbeing to be central to what it means to flourish.

0.2 A Note on Scope

Nearly any of these chapters could have been expanded into its own dissertation. There is a danger in this breadth. Many of the topics I write about—wellbeing, definitions of disease, public goods, cybersecurity—have been debated in vast bodies of literature. Inevitably some important questions and arguments will be glossed over, while others may be entirely ignored. However, this breadth helps to define the landscape and demonstrate the viability of a public health approach as a full-blooded alternative to the existing cybersecurity lens. Ultimately, I have tried to strike a balance. While the topics I have chosen exhibit the broad utility of such an approach, they also work as part of a cohesive argument.

Finally, note that most of these chapters follow a pattern of, first, presenting a relatively uncontroversial but significant point, then, a more controversial and more

interesting point. While this was not initially intended, I think it is appropriate given that digital technologies are relatively new, ubiquitous, and often ignored by philosophers. As such, there is value in both pointing out the straightforward but significant impact these technologies have on our lives and speculating about how these new technologies may be changing what it means to be a human in ways we are only just beginning to recognize. In Chapter 3, for instance, I begin by highlighting the health impacts of insecure critical infrastructure, such as nuclear power plants, and conclude with a discussion about whether a broken pacemaker should be considered a pathology. My hope is that even if you doubt aspects of my more controversial arguments, I will nonetheless succeed in demonstrating the value of thinking about the digital landscape in public health terms.

The stakes are high. Given the growing ubiquity of digital technologies, the explosion of digital personal information, the rapid development of artificial intelligence, and the increasing sophistication of biotechnologies, the potential for digital technologies to impact health and wellbeing will only grow in the coming years. While conceiving of cyberspace as a battlefield is understandable, it should not be the default for policymakers or product designers when there is an alternative framework that puts the promotion of health and wellbeing front and centre.

Chapter 1: Public Goods for Cyberhealth

Public goods are central to public health policies and interventions. Some of these public goods include the containment of infectious diseases, public education campaigns, herd immunity, fundamental health research, regulatory regimes, the draining of malarial swamps, information sharing programs, and disease surveillance.¹ Public goods are characterized by being non-rivalrous and non-excludable.² When a good is non-rivalrous, my consumption of the good does not diminish your ability to benefit from the good. For example, when a malarial swamp is drained, the benefit I receive does not diminish the benefit you receive. Meanwhile, when a good is non-excludable, an individual who does not participate in creating the good cannot easily be prevented from gaining the benefit. For example, one cannot prevent an unvaccinated individual from benefitting from herd immunity even though they have not contributed to the production of the good.³ Public goods can be contrasted with private goods, such as a slice of cake. Private goods are rivalrous and excludable. If there is one slice of cake at a cafe and I buy it, you can no longer purchase it, and a café can withhold a slice of cake from you until you pay. Identifying when a good is non-excludable and non-rivalrous is valuable because goods that have these characteristics tend to be underproduced by private markets. As

¹ Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager, eds., *Global Public Goods for Health: Health Economic and Public Health Perspectives* (Oxford: Oxford University Press, 2003).; Inge Kaul, Isabelle Grunberg, and Marc Stern, eds., *Global Public Goods* (New York; Oxford: Oxford University Press, 1999), 264-304.

² John G. Head, *Public Goods and Public Welfare* (Durham, N.C.: Duke University Press, 1975), 69.; John Rawls, *A Theory of Justice*, revised edition (Cambridge, MA: Harvard University Press, 1971, 1999), 235.

³ Angus Dawson, “Herd Protection as a Public Good: Vaccination and our Obligations to Others,” in *Ethics, Prevention, and Public Health*, eds. Angus Dawson and Marcel Verweij (Oxford: Oxford University Press, 2007), 163.

a result, public goods are often produced with some form of public assistance (e.g. subsidies, tax incentives, direct provisioning).⁴

Despite the similarities between communicable diseases and malware discussed in the Introduction,⁵ the theory of public goods plays a relatively minor role in cybersecurity policy and practice compared to its prevalence in the field of public health. As a result, many non-excludable and non-rivalrous goods, such as the containment of malware, are treated as if they were private goods. This, in turn, leads to those goods being underproduced compared to the level that would be best for society on the whole. In this chapter, I will argue that the similarities between the mitigation of malware and the containment of communicable disease 1) suggest that the theory of public goods should play a more significant role in how we think about and maintain robust and resilient digital networks, and 2) that public goods for public health are the best example of how to provision said goods. Specifically, I will argue that one can learn valuable lessons about how to adequately and equitably provision public goods for cyberhealth by studying historical attempts to provision public goods for public health (e.g. the creation of the World Health Organization) and by analysing the problem through the lens of the philosophy of public health.

In Section 1, I will provide a more detailed overview of public goods, describe their role in public health, and argue for the importance of public goods for maintaining the cyberhealth (i.e. robustness and resiliency) of large digital networks like the internet. In Section 2, I will focus on one specific public good for cyberhealth—the containment of malware. Using the example of the Conficker computer worm, I will argue that the containment of malware, like the containment of communicable diseases, is both a Global Public Good (GPG) and a participatory public good. I will argue that these similarities suggest that the containment of communicable diseases can provide useful insight into how best to design policies and institutions to contain malware. Finally, in Section 3, I will discuss what obligations states, individuals, and corporations have to contribute to the production of public goods for cyberhealth.

While some researchers, such as Deirdre Mulligan and Fred Schneider, have previously argued that cybersecurity should be thought of as a public good, this

⁴ Head, *Public Goods and Public Welfare*.

⁵ Lloyd and May, “How Viruses Spread Among Computers and People.”

analysis generally has not gone beyond a superficial acknowledgement that the benefits of cybersecurity are to some degree non-rivalrous and non-excludable.⁶ However, this superficial analysis is insufficient for determining the best way to provision public goods for cyberhealth, as public goods are a diverse set of outcomes and services which avoid easy categorization—the category of public goods includes goods as diverse as lighthouses and the containment of TB. Some public goods are presumptively beneficial, meaning almost no one would choose to live without them (e.g. national defence), while others are discretionary (e.g. fireworks displays). Some public goods are best created by direct government provisioning, while others may be best created by changing individuals’ and companies’ incentives through tax breaks and regulations. And some public goods are intrinsically “public” (e.g. broadcast television), while others can be made excludable in certain contexts (e.g. information). In order for public goods to be a useful category in the development of technology policy, one must go beyond the simple economic definition.

Having demonstrated the value of a public health inspired approach for addressing a classic cybersecurity threat like malware, in Chapter 2 I will comprehensively and formally define the Public Cyberhealth framework and argue that it can help policymakers create coherent, justified, and ethical technology policies in a wide variety of contexts. While beginning with a case study and then introducing the theory may seem a bit backwards, this more concrete discussion will help to provide context for the more abstract discussion to come.

1.1 Externalities, Public Goods, and Networks

To understand public goods, it is first important to understand the concept of an externality. Externalities, also referred to as external economies or spillover effects, can be defined as, “an event which confers an appreciable benefit (inflicts an appreciable damage) on some person or persons who were not fully consenting parties in reaching the decision or decisions which led directly or indirectly to the event in

⁶ Deirdre K. Mulligan and Fred B. Schneider, “Doctrine for Cybersecurity,” *Dædalus* 140, no. 4 (2011): 70-92.

question.”⁷ Colloquially we can think of externalities as side-effects which result from a given action. For example, when an individual with the flu does not wash their hands, other people who they meet will be more likely to catch flu. In this example, the spread of the flu is a negative externality of one’s personal decision to not wash one’s hands. Meanwhile, if someone gets a flu shot, not only is that individual’s risk of getting the flu lowered, but others in their network also receive some small degree of protection. The protection received by others in the network is a positive externality.

Public goods are special extreme cases of positive externalities where 1) there is essentially no additional cost to expand the beneficial side-effect to another person (i.e. the good is non-rivalrous)⁸ and 2) there is often no easy way to prevent someone from gaining the beneficial side-effect (i.e. it is non-excludable). Recognizing when a good is non-rivalrous and non-excludable is important because goods with these characteristics are typically underproduced by private markets compared to the socially optimal level—this inefficient distribution of goods is a kind of market failure.⁹ Private markets fail to efficiently produce public goods as in an open market, it is “grossly unrealistic,” to use John Head’s phrase, to expect individuals to voluntarily pay for a benefit that they will receive for free.¹⁰ This, in turn, reduces the incentive for affected parties (such as individuals or companies) to produce said

⁷ James E. Meade, *The Theory of Economic Externalities: The Control of Environmental Pollution and Similar Social Costs* (Geneva: Sijthoff-Leiden, 1973), 15.

⁸ Deborah Spar, “The Public Face of Cyberspace,” in *Global Public Goods*, eds. Inge Kaul, Isabelle Grunberg, and Marc Stern (New York; Oxford: Oxford University Press, 1999), 350.

⁹ John O. Ledyard, “Market Failure,” in *The New Palgrave Dictionary of Economics*, eds. Palgrave Macmillan (London: Palgrave Macmillan, 2008), https://doi.org/10.1057/978-1-349-95121-5_1052-2; It is worth noting that market failure does not imply none of the good will be produced—toll roads could be produced by private companies—but rather that the good will be produced inefficiently and social benefit will be unrealized.

¹⁰ Head, *Public Goods and Public Welfare*, 170.

good.¹¹ For example, let us assume a company announced a plan to drain a malarial swamp for a community at a cost of \$100,000. While the benefit the community will receive on the whole is worth that cost, each individual may not be adequately incentivized to contribute their fair share to the cost because they will receive the good for free, assuming the \$100,000 cost is eventually met and the company drains the swamp. Accepting a given benefit, without contributing to its production is called free-riding.¹²

Another classic public good for public health that suffers from free-riding is herd immunity. As more people in a population get vaccinated, the ambient protection unvaccinated individuals receive increases. As this ambient protection increases, so does the incentive to free-ride (i.e. not get personally vaccinated).¹³ In general, when too many people free-ride, the public good in question will often cease to be produced at the socially optimal level. In the case of herd immunity, this failure is particularly stark as the relationship between passive protection and the vaccination rate of a population is not linear.¹⁴ I will explore one's obligations to contribute to public goods in greater depth in Section 1.3.

Many goods that are valuable to public health efforts are non-excludable and non-rivalrous, including disease surveillance, fundamental research, public health campaigns, herd immunity, and regulations. In each case, private markets will tend to underproduce the good in question compared to the level which would be best for society. When private markets fail to adequately provision a good, the state (if it

¹¹ Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager, preface to *Global Public Goods for Health: Health Economic and Public Health Perspectives*, eds. Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager (Oxford: Oxford University Press, 2003), IX.

¹² George Klosko, *The Principle of Fairness and Obligation*, (Lanham: Rowman & Littlefield Publishers, 1992), 42.

¹³ Paul Fine, Ken Eames, and David L. Heymann, "Herd Immunity," *Clinical Infectious Diseases* 52, no. 7, April 1, 2011: 911–916.
<https://doi.org/10.1093/cid/cir007>.

¹⁴ Ibid.

chooses to intervene) can either directly provision said good or change the underlying incentives of the market to boost production and/or consumption.

At the national level, the production of public goods can be encouraged via a number of governmental mechanisms. Some of these mechanisms include providing subsidies to individuals to lower the cost of consumption, tying one good to another (e.g. only vaccinated children can attend public school), mandating technical solutions, enacting taxes or fines, changing default options, or even physical coercion. To promote herd immunity, for instance, states deploy a number of these strategies. In the United States, children must be vaccinated to attend public schools, public education campaigns encourage vaccination, the state incentivizes manufacturers to produce vaccines, and the Vaccine Injury Compensation Program is intended to reassure individuals that they will receive assistance in the case of a negative side-effect.¹⁵ At the international level, only some of these tools may be available as there is no world government with the authority to unilaterally enforce penalties. I will discuss this issue in more depth in Section 1.2, when I discuss global public goods (GPGs).

It is important to note that not all public goods are equally non-excludable and non-rivalrous. In fact, there are very few (if any) “pure” public goods—national defence may come closest. While the degree of excludability and rivalrousness of a given good will impact the type and degree of government involvement needed to correct market failure (e.g. direct provisioning vs. incentives), all goods which possess these traits are, to some degree, undersupplied by private markets compared to the level that would be socially optimal.¹⁶

¹⁵ Lee Ventola, “Immunization in the United States: Recommendations, Barriers, and Measures to Improve Compliance,” *Pharmacy and Therapeutics* 41, no. 7 (2016): 426–436.; Health Resources and Services Administration, “National Vaccine Injury Compensation Program,” Health Resources and Services Administration, October 2018, <https://www.hrsa.gov/vaccine-compensation/index.html>.

¹⁶ David Woodward, and Richard Smith, “Global Public Goods and Health: Concepts and Issues,” in *Global Public Goods for Health: Health Economic and Public Health Perspectives*, eds. Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager (Oxford: Oxford University Press, 2003), 4-5.; Head, *Public Goods and Public Welfare*, 80-81.; Inge Kaul, Isabelle Grunberg, and Marc A. Stern,

1.1.1 Public Goods for Cyberhealth

While digital networks can be entirely built and managed by private markets, many of the elements which promote cyberhealth (i.e. network robustness and resiliency) have the characteristics of public goods discussed above. Some of these public goods include network protocols, standards, encryption algorithms, cybersecurity knowledge, and the containment of malware. Some of these goods have traditionally been produced by public institutions (standards), while others have historically been treated as essentially private goods (the containment of malware). In each case, however, these goods are to varying degrees non-rivalrous and non-excludable.

One textbook example of a public good for cyberhealth is the creation of the Data Encryption Standard (DES). In 1972, the United States' National Bureau of Standards (NBS)—now the National Institute of Standards and Technology (NIST)—determined the government needed an encryption algorithm for encrypting unclassified but sensitive material.¹⁷ The NBS requested proposals for a cipher meeting a rigorous set of criteria, and in 1974 IBM submitted a cipher which—after being tweaked by the NSA—was accepted. While the DES has since been replaced by stronger algorithms, it served as a national standard used across a wide range of industries for decades.¹⁸ Once the cipher was publically released, anyone could take advantage of it, and one individual's use did not diminish another's ability to use the cipher. As such, the DES was non-excludable and non-rivalrous.

It is more controversial to declare cybersecurity itself a public good due to it being historically produced as if it were a private good. By cybersecurity I mean the state of a network being able to repel malicious attacks and preserve the confidentiality, integrity, and availability of the digital information that is transmitted and stored on the network. However, upon closer examination, the benefits of cybersecurity are generally non-rivalrous and non-excludable, even as some of the

“Introduction,” in *Global Public Goods*, eds. Inge Kaul, Isabelle Grunberg, and Marc A. Stern (Oxford: Oxford University Press, 1999), XX.

¹⁷ Gallagher, Link, and Rowe, *Cyber Security*, 144.

¹⁸ Ibid.

goods used to create a secure network are private goods (e.g. hardware).¹⁹ As is the case with public goods like herd immunity, the fact that the benefits of cybersecurity are non-rivalrous and non-excludable can lead to its underproduction. Describing the externalities which lead to the underproduction of cybersecurity by private markets, Gallagher, Link, and Rowe argue that any cybersecurity improvement a company or individual makes, especially of a proactive variety, will create “social benefits in excess of private benefits.”²⁰ Examples of investments or behaviours which generate positive externalities include intrusion detection systems, automated security patching, and practicing safe browsing. When your computer is secure, there is one fewer node which can pass on malware to me, and one fewer machine which can be drafted into a botnet which could attack critical infrastructure. In both instances, I am gaining a benefit for which I am not paying. When determining how much one should invest in cybersecurity, an individual or company tends to consider their personal benefits and costs, while ignoring these broader social benefits. While these externalities are not in and of themselves public goods, in aggregate they lead to a safer digital environment which produces benefits which are non-rivalrous and non-excludable. As digital networks underpin nearly all forms of critical infrastructure (e.g. water treatment, dams, banks, emergency services), the social benefits of cybersecurity are distributed widely. Even if one does not personally use the internet, one will likely benefit from the increased security of these critical networks.

Having said this, rather than focus on cybersecurity itself—which is too large a topic for one chapter—for the remainder of the chapter I will focus on one important aspect of cybersecurity, the containment of malware and its similarity to the

¹⁹ Johannes Bauer and Michel Van Eeten, "Cybersecurity: Stakeholder Incentives, Externalities, and Policy Options," *Telecommunications Policy* 33, no. 10 (2009): 706-19.; Arben Asllani, Charles Stephen White, and Lawrence Ettkin. “Viewing Cybersecurity as a Public Good: The Role of Governments, Businesses, And Individuals,” *Journal of Legal, Ethical and Regulatory Issues* 16, no. 1 (2013): 7-14.

¹⁹ Gallagher, Link, and Rowe, *Cyber Security*, 248.; Software occupies a middle ground between private and public goods. It can be thought of as a club good, like internet access itself. Club goods are non-rivalrous, but can be made excludable if the creator chooses, such as by requiring one to purchase a license.

²⁰ Gallagher, Link, and Rowe, *Cyber Security*, 248.

containment of communicable disease. While the public health approach is useful for addressing a wide range of issues related to network robustness and resiliency, this example will help demonstrate that even if one is only interested in archetypal matters of cybersecurity, there is value to using a public health lens.

1.2 The Containment of Malware

The containment of malware is particularly relevant for this discussion as 1) it is extremely important to robustness and resiliency of networks, 2) it displays the traits of a public good, and yet 3) it has largely been treated as if it were a private good capable of being adequately produced by private markets. As such, it is a likely candidate to be underproduced compared to the socially optimal level. Additionally, its similarity to the containment of communicable diseases suggests a public health approach can serve as a valuable template for how to adequately and equitably provision the good.

Rather than speak in generalities, it will be helpful to focus on one specific malware threat—the Conficker worm. While there have been many more recent and harmful global malware threats (e.g. WannaCry, Mirai), the Conficker case is particularly well-documented, and thus it is a good example for understanding the roles the private and public sectors have traditionally played in containing malware.²¹ My expanded treatment of this case study can be found in the cybersecurity policy book *Rewired*.

1.2.1 The Conficker Infection

On October 23, 2008, during the eighth annual meeting of the International Botnet Task Force, Microsoft released an out-of-band emergency security patch. The patch fixed a Windows vulnerability which could allow malware to spread between unprotected machines without any user interaction.²² While releasing an emergency

²¹ Michael Thornton, “Containing Conficker,” in *Rewired*, eds. Ryan Ellis and Vivek Mohan (Hoboken: Wiley, 2019).

²² The Rendon Group, “Conficker Working Group: Lessons Learned,” The Rendon Group, (2010),

patch cast a spotlight on the vulnerability, Microsoft had already seen the flaw exploited in the wild. On November 22nd, a month after the patch's release, a new piece of highly-contagious malware—the Conficker worm—was first detected. In response, Microsoft issued a security alert recommending people immediately patch their systems.

For the most part, Conficker A (as it would come to be called) simply hid among a computer's background activity. However, when it was time to call home for instructions, the worm would contact 250 pseudo-randomly generated domains spread out across 5 Top Level Domains (TLDs).²³ Behind any of those domains, the creators of the worm could be waiting to issue commands. A few weeks later a more sophisticated variant called Conficker B appeared which could propagate via thumb drives, disable Windows Automatic Update, block certain DNS look-ups, and call domains from eight TLDs.²⁴ While individually these strategies were not new, it was unusual for so many features to be packed into a single piece of malware. More than one researcher described it as “elegant.”²⁵ By the end of 2008, SRI International estimated between 1-1.5 million computers were infected.²⁶ Over the following five months, three additional versions of the worm would be introduced. At its peak in 2009, Conficker infected between 5 and 13 million machines.²⁷ While the worm's purpose was not clear, a botnet of that size could be used to disrupt critical infrastructure, including large parts of the internet.

While Microsoft's release of an emergency patch was a sign that the vulnerability was particularly dangerous, in general the cybersecurity community (and especially governments) were slow to recognize the scope of the problem. While in late 2008 the worm was being discussed with increased frequency on a number of

http://www.confickerworkinggroup.org/wiki/uploads/Conficker_Working_Group_Lessons_Learned_17_June_2010_final.pdf (accessed Feb. 20, 2019).

²³ Common TLDs include .com, .org., .gov., etc.

²⁴ Dave Piscitello, “Conficker Summary and Review,” ICANN, 2010, <https://www.icann.org/en/system/files/files/conficker-summary-review-07may10-en.pdf>.

²⁵ Ibid., 5.

²⁶ The Rendon Group, “Conficker Working Group: Lessons Learned,” 16.

²⁷ Ibid., 10.

cybersecurity e-mail lists,²⁸ until early 2009 there was little organized activity within the private sector to control the spread of the worm. Governments, meanwhile, were entirely absent from the discussion. The security firm Qualys estimated that two months after the emergency patch was released 30 per cent of computers running Windows remained unpatched.²⁹ The absence of government involvement should have raised red flags, given that the benefit of containing Conficker would be non-rivalrous and non-excludable—one need not contribute to this effort to reap the benefit, and one person's benefit would not diminish the benefit to others.

A small number of security experts, who would later call themselves the Conficker Working Group (CWG), did notice that Conficker threatened the internet at large. Shortly after the worm's appearance, they began to study the worm and devise ways to control it. Early members of the all-volunteer CWG, many of who knew each other from conferences and social media, included representatives of Microsoft, SRI International, and several companies which managed TLDs, as well as a number of independent security researchers and academics. Relatively quickly it was discovered that the domain names which could be used for command and control communications were not random. By running the domain name generation algorithm for a future date, the group could identify the domains that would be called and register the names themselves (often with *personal* credit cards) before the worm's creators could use them for passing the botnet instructions. When infected computers called these domains, the CWG redirected the traffic to designated sinkhole servers³⁰ which were then used to map the spread of the infection.

This strategy worked reasonably well until the introduction of the Conficker C variant in late February 2009. In reaction to the CWG's sinkholing project, the creators of Conficker C designed the new version to generate a list of 50,000 domains every day from among a list of 116 TLDs. Unlike in the past, this list included not only general TLDs (e.g. .org, .biz) but over 100 country level domains (e.g. .cn, .fr). Each day an infected machine would attempt to contact 500 domains from this list of

²⁸ Ibid., 16.

²⁹ Ibid., 4.

³⁰ A sinkhole server is a server used by researchers or law enforcement to capture traffic intended for another source. When a machine attempts to call a given domain, the DNS server will reroute that call to the designated sinkhole server.

50,000. In order for the sinkholing strategy to continue to work, the organizations and companies managing these TLDs would need to work together to block 50,000 domain names a day, forever. In some countries this practice was of dubious legality, and in some cases domains were already owned and operated for legitimate purposes. Additionally, the strategy relied on the International Corporation for Assigned Names and Numbers (ICANN) agreeing to waive its fees for registering the domains—something that had never before been asked. Despite these difficulties, the CWG—now numbering hundreds of volunteers—was ultimately able to convince the relevant stakeholders to cooperate by leveraging personal connections, although a number of the companies that managed TLDs dragged their feet.

This work went on with essentially no government involvement. Despite repeatedly trying to raise the alarm in Washington, the CWG was largely rebuffed. While leaders of the CWG eventually were able to meet with U.S. government officials, the U.S. government generally failed to understand the risk and was ill prepared to take any kind of active role in mitigation efforts. One member of the working group would later sum up the U.S. government’s role as “zero involvement, zero activity, zero knowledge.”³¹

While the Conficker worm ultimately infected between 5-13 million machines, the catastrophe that was feared never came to pass. Having said this, it is unclear how much credit the CWG deserves. While the sinkholing effort was generally successful at keeping the creators of the worm from taking control of the botnet (a few domains did slip through), later variants of the worm possessed the ability to pass on instructions via peer-to-peer connections. Even if the sinkholing project was perfect, the creators of Conficker could have circumvented the CWG’s efforts by using this slower payload delivery system. Additionally, there was nothing stopping the creators from simply upping the number of domains that needed to be blocked. Would the organizations which manage TLDs voluntarily block 100,000 domains a day? What about 500,000? Whether the heat became too much, or the effort too costly, the creators of the Conficker botnet never implemented their master plan (if there was one to begin with).

Before dissecting the flaws in the CWG’s approach, it is worth noting their successes. The CWG’s greatest successes were rallying an unprecedented degree of

³¹ Ibid.

private sector collaboration and gaining cooperation from ICANN and the TLDs without any enforcement authority. Despite at times conflicting incentives, twenty private companies and a number of non-profits were mostly able to work together effectively. Decisions on information sharing and strategy were generally made by consensus and decisions on when to talk to the press were typically made as a group. While some key members were accused of sharing information with other stakeholders or the press for selfish reasons,³² these instances were relatively rare and ultimately not fatal to the overall project. CWG members cited the informal organization and lack of a hierarchy as major factors in keeping the group together. Additionally, the fact that many members knew each other via social media helped facilitate trust.³³

In the Lessons Learned report, the leaders of the CWG listed the following as failures or downsides to their model: remediation efforts (they did little to remove Conficker from infected machines), communication with ISPs, collaboration and information sharing with the U.S. government, public relations, a lack of accountability, a lack of a tasking authority, and balancing inclusion of stakeholders with efficiency.³⁴ The last three entries on the list are perhaps the required cost of the informal organization and reliance on social networks that contributed to the group's successes. In regard to government collaboration, members specifically mentioned U.S. government representatives being willing to take information (including plagiarizing CWG slides) without providing any information or resources in return.

When thinking about the CWG or other private sector responses as a model for future large-scale malware control efforts, the most important questions are: How repeatable are the successes? and How fixable are the failures? First, I will consider the successes.

While in the case of Conficker the private sector collaboration was impressive, the fact that companies often have conflicting incentives means there is always the risk that a collaboration effort like the CWG will break down as underlying incentives shift. Additionally, the incentives which brought private companies together for Conficker may be slightly different and less persuasive in the case of other large-scale

³² Mark Bowden, *Worm* (London: Grove Press UK, 2011) 232.

³³ The Rendon Group, "Conficker Working Group: Lessons Learned."

³⁴ Ibid., 34-36.

malware threats which nonetheless need attention. One fundamental conflict is between security companies, which sell remediation tools, and Microsoft, which seeks to patch vulnerable machines before they become infected. Companies will also always have strong incentives to talk up their own efforts in the media. This caused tensions within the CWG and surely would in future efforts as well.³⁵ Additionally, there is no guarantee that TLDs would voluntarily and universally support future large-scale domain registration.³⁶ While the sinkholing strategy was mostly successful in the case of Conficker, multiple TLDs dragged their feet and balked at the initial ask.³⁷ As of 2018, there remains no enforcement mechanism to ensure compliance. On the whole, while the CWG's successes were impressive and admirable within the context of this specific threat, it is unlikely that the same level of cooperation and collaboration can be counted on to control future threats—more dangerous threats may not be able to be contained by voluntary efforts, and less dangerous (or more highly targeted) threats may not sufficiently inspire broad collaboration.³⁸

I am equally pessimistic that the failures of the CWG's model can be easily fixed. Without a leadership structure, it is almost impossible to effectively and repeatedly assign tasks and hold people accountable. Likewise, without organizational permanence it is difficult to build better working relationships with ISPs and governments. While loose networks of technical experts may be great at solving complex engineering problems and will surely be needed to address future threats, they are ill-suited to coordinating an ongoing international crisis response effort, let alone several simultaneously.

In the Lessons Learned report, several participants suggested that if only the CWG had two to three full-time administrative resources, many of the flaws of the CWG's structure could be fixed, but this ignores the fundamental issue that states cannot responsibly leave the protection of critical infrastructure to volunteers. The ICANN post-mortem specifically mentioned that one cannot rely on a similar calibre of volunteer the next time around and questioned the group's ability to potentially

³⁵ Ibid., 23.

³⁶ Piscitello, "Conficker Summary and Review," 14.

³⁷ Bowden, *Worm*, 230.

³⁸ One notable systematic change was that ICANN now has a formal process for waiving fees for malware control efforts.

deal with two threats simultaneously.³⁹ One member reinforced this concern, saying the only reason the Zeus Trojan spread so widely was that everyone was focused on Conficker.⁴⁰ In *Worm*, Mark Bowden characterizes the members of the CWG as the X-Men—outsiders who possess almost supernatural skills and swoop in to save the day. The problem with the X-Men is that sometimes they save the planet and sometimes they start a civil war.

Given that the benefit being produced is non-excludable and non-rivalrous, there are good reasons to believe that, in the long run, a public health approach to infections like Conficker will be superior to the ad hoc volunteer approach of the CWG. Specifically, a public health approach will be more reliable and more likely to act in the public interest. In the following section, I will describe how the containment of communicable diseases can serve as a template for how to adequately provision public goods for cyberhealth like the containment of malware. Then in Section 1.3, I will discuss what obligations states, corporations, and individuals have to contribute to the production of public goods for cyberhealth.

1.2.2 Global Public Goods and Malware

One of the primary reasons to think that a public health approach to the containment of malware may be useful is that both the containment of malware and the containment of communicable diseases belong to a special class of public goods whose benefits are global. While researchers and companies in the United States took the lead in the containment, all countries benefitted from the containment. Even if the CWG had wanted to restrict the benefit to countries who had contributed to the effort, they would have been unable to do so. As the benefit of containing Conficker is non-rivalrous, non-excludable, and transcends national borders, it is what is called a global public good (GPG). These traits are not unique to Conficker—the containment of all large-scale malware outbreaks will have these traits. In the context of global health policy, Kaul, Grunberg, and Stern define GPGs as those whose:

Benefits are quasi universal in terms of countries (covering more than one group of countries), people (accruing to several, preferably all, population groups), and generations (extending to both current and future generations, or

³⁹ Piscitello, “Conficker Summary and Review,” 12.

⁴⁰ The Rendon Group, “Conficker Working Group: Lessons Learned,” 41-42.

at least meeting the needs of current generations without foreclosing development options for future generations).⁴¹

Classic examples of GPGs include the protection of the ozone layer, efforts to combat climate change, and, most relevant to this work, the containment of communicable diseases.

In some ways, GPGs are not qualitatively different from national public goods. As William Nordhaus says, “They are only ones where the effects spill widely around the world and for a long time to come.”⁴² However, while most public goods are produced at the national level by national governments with the power to coerce, global public goods must be produced by the wilful collaboration of states. We can see this need for global collaboration in the case of Conficker. Even if the United States demanded that all TLDs managed within its borders participated in the sinkholing project, the worm could simply call a domain outside the United States for commands.

As mentioned in Section 1, most public goods can be produced using a variety of tools (e.g. subsidies, direct provisioning, standards) as long as the national government has the will to do so, but in the international context the proliferation of actors and the lack of a global government make it much harder to achieve a consensus on appropriate action. The lack of a central authority also makes it difficult to coerce free-riding countries to contribute their fair share.⁴³ Tools like withholding aid create additional ethical problems. As Nordhaus describes the problem, “there is

⁴¹ Inge Kaul, Isabelle Grunberg, and Marc A. Stern, “Defining Global Public Goods,” in *Global Public Goods: International Cooperation in the 21st Century*, eds. Inge Kaul, Isabelle Grunberg, Marc A. Stern (New York ; Oxford: Oxford University Press, 1999), 2-3.

⁴² William Nordhaus, “Paul Samuelson and Global Public Goods,” in *Samuelsonian Economics and the Twenty-First Century*, eds. Michael Szenberg, Lall Ramrattan, Aron Gottesman (Oxford: Oxford University Press, 2006), doi:10.1093/acprof:oso/9780199298839.003.0006.

⁴³ Woodward, David, and Richard Smith, “Global Public Goods and Health: Concepts and Issues,” in *Global Public Goods for Health: Health Economic and Public Health Perspectives*, eds. Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager (Oxford: Oxford University Press, 2003), 18.

no legal mechanism by which disinterested majorities, or supermajorities short of unanimities, can coerce reluctant free-riding countries into mechanisms that provide for global public goods.”⁴⁴ He calls this the ‘Westphalian dilemma,’ as it is inherent in our understanding of sovereign states. States must consent to any form of coercion by joining international agreements and organizations which establish binding responsibilities, but the non-excludability of any public good means there are substantial incentives for states to free-ride. Lastly, GPGs are also different from typical issues of international concern like tariffs or border issues in that they often require joint facilities as well as internal national policies to converge.⁴⁵

One can imagine a spectrum where public goods that can be adequately produced at the national level are on the left and public goods which can only be adequately produced by the global community on the right. Public roads may fall all the way to the left, while all the way on the right sits protection of the ozone layer. Somewhere in between are public goods which can be produced at the national level, but those national level solutions will be insufficient to bring about the optimal outcome. Generally, those on the right, like the protection of the ozone layer are intrinsically global, while those which fall somewhere in the middle are historically national public goods which have become global due to the opening of borders and the increasing interconnectedness of modern societies. The containment of communicable diseases and the containment of malware fall into this middle ground. While global collaboration may be required to achieve an optimal outcome, an individual state would still receive some benefit from vaccinating its population or protecting its essential networks even in the face of international apathy.⁴⁶

We can see this dynamic in the eradication of a disease like smallpox. In the United States, vaccination efforts controlled smallpox at the national level, such that the last naturally occurring outbreak occurred in 1949.⁴⁷ However, only through an

⁴⁴ Nordhaus, “Paul Samuelson and Global Public Goods.”

⁴⁵ Kaul, Grunberg, and Stern, “Introduction,” XXV.

⁴⁶ The degree of benefit will vary depending on the effectiveness of the intervention (i.e. vaccination effectiveness), the porousness of borders, and the nature of the specific threat.

⁴⁷ “Smallpox,” Centers for Disease Control and Prevention, July 12, 2017, <https://www.cdc.gov/smallpox/index.html> (accessed Dec. 30, 2018).

aggressive international effort organized by the World Health Organization (WHO), could the United States ensure the disease would not be reintroduced. While the United States reaped substantial benefits from its national vaccination campaign, the socially optimal outcome could only be achieved through a global eradication campaign consisting of diligent monitoring and vaccination campaigns. That is to say, securing the “national” public good, required the production of a global public good, which could only be produced through global agreement.

Similarly, in the case of Conficker, if a given state had required individuals and companies to install Microsoft’s security patch, individuals and companies in that state would have gained protection from certain harms. For instance, these patched machines could not be drafted into the Conficker botnet, nor could the worm be used to install other malware on these machines. However, if the worm remained uncontained in other countries, the country with a high patch rate could still be the victim of DDoS attacks from this international botnet. As an example, one could look to the Mirai botnet’s attack on Liberia’s internet in 2016. While relatively few of the devices that comprised the Mirai botnet were in Liberia, a DDoS attack on the country’s internet service providers was able to shutdown the country’s internet for several days.⁴⁸

As the containment of malware exists somewhere in the middle of the spectrum between national public goods and “pure” global public goods (e.g. protection of the ozone layer), one could imagine two paths forward. Either states can erect stronger borders and treat the containment of malware as a national public good, or they can keep open internet borders and engage in greater levels of international collaboration. While both paths are logistically difficult, the first may also be of limited effectiveness unless a state is willing to fully cut itself off from the internet, as illustrated by the case of Liberia and Mirai mentioned previously. As even states like China, with notoriously tight control on internet activity, remain connected to the internet, I will focus on the second option—increased international collaboration.

⁴⁸ Nicky Woolf, “Massive Cyber-Attack Grinds Liberia's Internet to a Halt,” *The Guardian*, Nov. 3, 2016, <https://www.theguardian.com/technology/2016/nov/03/cyberattack-internet-liberia-ddos-hack-botnet> (accessed Dec. 29, 2018).

Since Conficker there have been some tentative steps towards greater international collaboration on digital matters of mutual concern (e.g. the Global Alliance against Child Sexual Abuse Online), but these efforts have typically been relatively small scale and created in an ad hoc manner. As states have generally chosen this open borders approach in the face of communicable diseases, public health efforts, such as the creation of the WHO and the International Health Regulations, are a valuable historical guide for how to produce public goods at the global level.

1.2.3 The WHO and Disease Monitoring

In broad strokes, the public health approach to containing communicable disease at the global level involved building consensus among states as to what constituted a shared risk, creating international institutions to monitor a limited number of diseases, and then using those institutions, such as the WHO, to continue to build consensus and broaden the mandate over time. According to Mark Zacher, the modern notion of disease surveillance originated in 1897 at the International Sanitary Conference when participating countries recognized the need for some kind of global disease surveillance.⁴⁹ As the benefit of disease surveillance is non-excludable and non-rivalrous, private market solutions were ill-suited to the task, and while individual states could monitor diseases within their own borders, these efforts were insufficient to mitigate the threat since diseases could easily cross borders.

In 1902 the newly formed Pan-American Sanitary Bureau was charged with collecting and sharing information on disease outbreaks, and a year later the International Sanitary Convention was adopted which called for the creation of a new international organization to monitor diseases. The organization that was subsequently formed was called the Organisation Internationale Publique which along with the Health Organization of the League of Nations was a precursor to the WHO, established in 1948.⁵⁰ With the introduction of the International Sanitary Regulations in 1951, the WHO began to require states to report cases of designated diseases to the

⁴⁹ Mark Zacher, “Global Epidemiological Surveillance: International Cooperation to Monitor Infectious Disease,” in *Global Public Goods*, eds. Inge Kaul, Isabelle Grunberg, and Marc Stern (New York; Oxford: Oxford University Press, 1999).

⁵⁰ *Ibid.*, 266.

organization within twenty-four hours.⁵¹ These regulations were renamed the International Health Regulations in 1969. While originally only concerning four diseases, the IHR was revised in 2005 to require states to notify the WHO of all events that may constitute a public health emergency of international concern and to respond to requests for verification of information regarding such events.”⁵²

Over time the mandate of the WHO has continued to expand to encompass the production of numerous public goods that would likely be underproduced if production was left in the hands of either private markets or national health services acting in their own (national) self-interest. In addition to disease surveillance, the WHO serves as a centre for research, crisis coordination, training, standards for readiness and response, and information sharing.⁵³ While GPGs need global collaboration to be produced, it is important to note that states play a fundamental role in the provisioning of GPGs. While the WHO provides logistical support and directly provisions some of these goods, the organization works in close collaboration with national public health agencies, such as the United States’ Centers for Disease Control and Prevention (CDC), European Centre for Disease Prevention and Control (ECDC), Chinese Center for Disease Control and Prevention (China CDC), and Public Health England.

The cyberhealth equivalent may involve building an international institution that is equivalent to the WHO—a World Cyberhealth Organization. As mentioned previously, many of the types of public goods which the WHO provides are also critical to cyberhealth, including fundamental research, crisis coordination, response and readiness standards, training, ongoing monitoring programs, and information sharing. While in this chapter I have focused on malware, many of these public goods could also mitigate the harmful network impacts of environmental threats or human error. A World Cyberhealth Organization would have ongoing relationships with the organizations, companies, and states that manage TLDs and internet infrastructure (e.g. DNS, security certificates, transmission lines, routing equipment), and would possess coercive tools not available to an ad-hoc volunteer group like the CWG. A

⁵¹ Ibid., 272.

⁵² “Frequently Asked Questions About the International Health Regulations,” World Health Organization, 2009, <http://www.who.int/ihr/about/FAQ2009.pdf>.

⁵³ World Health Organization, “Programmes and Projects.”

World Cyberhealth Organization would also address many of the organizational problems of the CWG by having the authority to delegate tasks and hold people and organizations accountable. This accountability would also extend to the organization itself as governmental and inter-governmental organizations ultimately—if indirectly—fall under the purview of elected leaders. In the case of the WHO, the governing body is comprised of representatives from all WHO member states.⁵⁴ If the CWG had failed in its mission, there would have been no justifiable reason or mechanism to hold them accountable for their failure. Meanwhile, many of the strengths of the CWG could still be preserved by a World Cyberhealth Organization. If an issue exceeded the technical capabilities of the organization, specialized working groups of independent and private sector security experts could still be convened, but those groups would have institutional resources and legitimate authority to implement their solutions.

Such an organization could also deploy preventative strategies, which were unavailable to the CWG. In the public health context, one example of this approach is the WHO's standards on the prevention of drug use and non-communicable disease. In the case of drug use, the WHO encourages interventions targeting pregnant women, early childhood education, addressing mental health disorders, keeping children in school, mentoring programs, and media campaigns, rather than simply treating individuals after they have developed an addiction.⁵⁵ Meanwhile, in regards to non-communicable diseases, the WHO recommends surveillance, reduction of risk factors, and the promotion of health across the life course as the most effective way of reducing premature death and disability.⁵⁶ In the case of Conficker, a preventative

⁵⁴ World Health Organization, "World Health Assembly," World Health Organization, <https://www.who.int/mediacentre/events/governance/wha/en/> (accessed Feb. 27, 2019).

⁵⁵ World Health Organization, "International Standards on Drug Use and Prevention," 2nd edition, *World Health Organization*, 2018, https://www.unodc.org/documents/prevention/standards_180412.pdf (accessed Dec. 13, 2018).

⁵⁶ Director General, "Global Strategy for the Prevention and Control of Noncommunicable Diseases," *World Health Organization*, A53/14, March 22, 2000, http://apps.who.int/gb/archive/pdf_files/WHA53/ea14.pdf (accessed Dec. 13, 2018).

approach would have been preferable to the CWG's reactive approach. While forcing people to patch their systems may have been overly burdensome,⁵⁷ public education campaigns and tighter security standards could have sufficiently mitigated the risk—a botnet of half a million machines poses far less of a risk to critical infrastructure than one of 10 million.

While it may seem like the conversation has shifted from the value of public goods to the value of prevention strategies, the two issues are closely related, for two reasons. First, when faced with threats like malware or contagious disease, individual actors acting in isolation (individuals, corporations, states) will tend to underinvest in preventative strategies (compared to the socially optimal level) because these strategies generate substantial externalities. As such, individual actors will not receive the full benefit of their actions. And second, private markets are often unable to implement preventative strategies to address global threats, as these strategies are ineffective without widespread international collaboration. As a result, preventative strategies can often only be deployed as part of a public response.

While there is no World Cyberhealth Organization today, one may want to point to national institutions like the United States' National Cybersecurity and Communications Integration Center (NCCIC) as evidence that this public health approach to cybersecurity is already being explored at the national level. After all the NCCIC—part of the department of Homeland Security—is tasked with performing many of the functions I described previously (e.g. monitoring, crisis response, information sharing). However, while the creation of the NCCIC in 2009 suggests there is some recognition in the United States that the government should be providing public goods for cyberhealth, the organization is too small to effectively fulfil these functions. For example, as of 2016 the NCCIC only had seven cybersecurity advisors on staff to advise the private sector on the security of critical infrastructure—up from one in 2009.⁵⁸ With such a small staff, it is not surprising that the cybersecurity of critical infrastructure is generally left in the hands of the private sector, despite these private companies often having conflicting incentives (e.g.

⁵⁷ For further discussion on proportionality see 2.4.2.

⁵⁸ "FY 2018 Budget in Brief," United States Department of Homeland Security, <https://www.dhs.gov/sites/default/files/publications/DHS%20FY18%20BIB%20Final.pdf>.

updating their systems eats into profits). Additionally, as part of the security apparatus, the NCCIC has its own conflicts of interest, at least from an international point of view. For instance, in certain cases it may be in the strategic interest of the United States to withhold certain information that would benefit the world community, such as the existence of certain vulnerabilities that have strategic value to intelligence agencies. While someone operating within a public health mindset would be predisposed to share such information, one operating within a security framework is more likely to keep such information close to the vest.

While the history of public health suggests a few cyber threats of common concern can serve as the basis for greater collaboration over time, today there is little agreement about what those common threats may be. While some types of network risks, such as spam, botnets, and ransomware negatively impact most countries, other risks such as intellectual property theft have asymmetric costs—the United States and Europe tend to be victims, while China benefits⁵⁹ by using stolen intellectual property to bypass expensive research and development.⁶⁰ Part of the problem is that we often conflate cybercrime with asymmetric costs (e.g. IP theft, military spying) with those that negatively impact all states (e.g. ransomware, spam). As a result, China and Russia have generally been wary of joining international enforcement efforts. Singer and Friedman describe the issue saying, “The parallel would be treating the actions of a prankster with fireworks, a bank robber with a revolver, an insurgent with a roadside bomb, and a state military with a cruise missile as if they were all the same phenomenon simply because their tools all involved the same chemistry of gunpowder.”⁶¹ However, there are signs that greater collaboration is possible. The past two-decades has seen the creation of international and regional CERTs (Computer Emergency Response Teams), including AP-CERT (Asia Pacific), TF-

⁵⁹ Paulo Shakarian, Jana Shakarian, Andrew Reuf, *Introduction to Cyber-Warfare* (Waltham: Syngress, 2013), 114-153.

⁶⁰ Bryan Krekel, “Capability of the People’s Republic of China to Conduct Cyber Warfare and Computer Network Exploitation,” The US-China Economic and Security Review Commission, 2009, <http://www.thing.net/~rdom/ucsd/CyberwarUSChina2009.pdf>.

⁶¹ Singer and Friedman, *Cybersecurity and Cyberwar*, 68.

CSIRT (Europe), and CERT/CC CERT (International);⁶² the formation of the U.S. E.U. Working Group on Cybersecurity and Cybercrime; and the Global Alliance against Child Sexual Abuse Online—a collaboration of 50 states working to reduce and combat child pornography.⁶³ There is even evidence that traditional holdouts like Russia may begin to feel the cost of inaction as they become more reliant on digital networks. After Russian critical infrastructure, including governmental offices and railroads were disabled by WannaCry, Frants Klintsevich, the deputy chairman of the Russian Senate’s defence committee said, “Humanity is dealing here with cyberterrorism...It’s an alarming signal, and not just a signal but a direct threat to the normal functioning of society, and important life-support systems.”⁶⁴

The success or failure of voluntary arrangements like the WHO and IHR frequently has much to do with the nature of the production of the public good in question. In cases where the benefits or harms are additive, the benefit of a country to free-ride is substantial. An example of a harm which is additive is the emission of greenhouse gases. Whether or not Panama cuts its emissions has little bearing on the overall levels of greenhouse gases, as they are only a minor producer. As such, they have a greater incentive to free-ride. For goods which have weakest link characteristics—where the system is only as strong as the weakest link—there is substantially less incentive to free-ride, as the good will only be produced if each country holds up its end of the bargain.⁶⁵ These traits may help explain why climate change treaties like the Kyoto Protocol have generally failed, while countries have generally cooperated with the WHO in its work to contain communicable diseases—a public good which demonstrates weakest link characteristics. Conficker, like the containment of contagious disease, is a type of weakest link problem. In both cases,

⁶² Choucri et al., “Institutions for Cybersecurity.”

⁶³ “Fact Sheet: U.S.-EU Cyber Cooperation.” The White House, March 26, 2014, <https://obamawhitehouse.archives.gov/the-press-office/2014/03/26/fact-sheet-us-eu-cyber-cooperation> (accessed May 18, 2017).

⁶⁴ Andrew Kramer, “Russia, This Time the Victim of a Cyberattack, Voices Outrage,” *New York Times*, May 14, 2017, <https://www.nytimes.com/2017/05/14/world/europe/russia-cyberattack-wannacry-ransomware.html> (accessed May 18, 2017).

⁶⁵ Nordhaus, “Paul Samuelson and Global Public Goods.”

national efforts to patch computers or vaccinate individuals help to mitigate the spread of the epidemic, but free-riding countries make it difficult to fully eliminate the threat, which may then emerge at a later time. In the case of Conficker, if a single TLD refused to sinkhole the relevant domains, the entire containment effort may have been useless. While not every piece of malware follows this pattern, perhaps countries can at least agree to work together to contain Conficker-like threats as a starting point.

While other global organizations could serve as useful models of international collaboration, public health is a particularly promising model given that many of the public goods which are relevant to public health and cyberhealth can only be secured through the combined efforts of individuals, companies, NGOs, and states. The containment of polio, for example, involved an international monitoring regime, national health systems, the WHO, private vaccine manufacturers, public and private research, and individuals willing to be vaccinated. Global public goods which involve so many actors with potentially competing rights and incentives are relatively rare. Many that do exist come from the realm of public health—Polio eradication,⁶⁶ TB control,⁶⁷ antimicrobial drug resistance,⁶⁸ environmental protection.⁶⁹ These

⁶⁶ Bruce Aylward, Arnab Acharya, Sarah England, Mary Agocs, and Jennifer Linkins, “Polio Eradication,” in *Global Public Goods for Health: Health Economic and Public Health Perspectives*, eds. Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager (Oxford: Oxford University Press, 2003).

⁶⁷ Jim Yong Kim, Aaron Shakow, Arachu Castro, Chris Vanderwarker, and Paul Farmer, “Tuberculosis Control,” in *Global Public Goods for Health: Health Economic and Public Health Perspectives*, eds. Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager (Oxford: Oxford University Press, 2003).

⁶⁸ Richard Smith and Joanna Coast, “Antimicrobial Drug Resistance,” in *Global Public Goods for Health: Health Economic and Public Health Perspectives*, eds. Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager (Oxford: Oxford University Press, 2003).

⁶⁹ Anthony McMichael, Colin Butler, and Michael Ahern, “Global Environment,” in *Global Public Goods for Health: Health Economic and Public Health Perspectives*, eds. Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager (Oxford: Oxford University Press, 2003).

examples, both the successes and failures, provide insight into how best to juggle individual rights, incentives, and obligations in the global context.

The concept of GPGs is significant as it highlights certain difficulties which are inherent to goods which cross global boundaries. As with national public goods, calling something a global public good does not in any way remove the collective action problems, but the concept can be helpful as a tool for analysis and advocacy.⁷⁰ By more accurately describing the ways in which states' interests are interconnected and by highlighting that national programs are often inadequate to address global problems, the concept can spur wealthier nations to invest in the production of the global public good in question, not out of altruism, but the more reliable motivation of self-interest.⁷¹ While there are a number of non-health related global public goods, global public goods for public health typically must balance a similarly diverse set of actors and factors as malware mitigation, and represents a historical example of the international community coming together over time to contain an emergent threat to global stability.

1.2.4 Participatory Public Goods

In addition to being GPGs, the containment of malware and the containment of communicable diseases are also what I will call *participatory public goods*. Whereas the term 'global public good' is used in the literature on public goods, 'participatory public good' is my own term. A participatory public good is *a public good which can only be produced thorough the active participation of the beneficiaries beyond the mere provision of financial resources*. As with GPGs, participatory public goods are also somewhat rare. Most paradigmatic public goods (both national public goods and global public goods) can be produced by a state or group of states without the active participation of the majority of the beneficiaries. Examples include monitoring programs, lighthouses, national defence, and standards; in each of these cases, all that is required for beneficiaries to enjoy the public good is that they pay their taxes so

⁷⁰ Smith and Coast, "Antimicrobial Drug Resistance," 84.

⁷¹ Lincoln Chen, Tim G. Evans, and Richard A. Cash, "Health as a Global Public Good," in *Global Public Goods*, eds. Inge Kaul, Isabelle Grunberg, and Marc Stern (New York; Oxford: Oxford University Press, 1999), 299.

others can provide that good. In Section 1, I introduced perhaps the most well-known and important participatory public good—herd immunity.⁷²

The containment of malware is like herd immunity in that individuals must actively participate in the production of the good for it to be produced. In the public health context one must get vaccinated, while the containment of malware requires individuals to install security patches. As in the case of infectious disease, when a high percentage of the population has patched a digital vulnerability, it becomes much harder for worms to spread throughout a digital network.⁷³ The only reason Conficker was able to spread as it did was that over 30 per cent of machines running Windows remained unpatched at the time of the worm’s appearance. As the CWG had no authority to coerce or incentivize individuals to patch their systems, they did not seek to raise the patch adoption rate. A public health inspired approach to containing malware could have nipped the problem in the bud by encouraging or requiring individuals and companies to update their devices. While ‘digital herd immunity’ may not stop highly targeted pieces of malware, it can significantly curtail the growth of some large botnets⁷⁴ like the ones formed by Conficker or Mirai—the botnet which attacked the DNS in 2016. It is important to note that Conficker is not an edge case, individuals play an active role in containing many forms of malware.

Using herd immunity as an example, at the national level the state could employ a range of incentives to encourage individuals and corporations to participate in the containment of malware.⁷⁵ Similar to states limiting access to public schools to vaccinated children, internet access (or certain parts of the web) could be limited to

⁷² Fine, Eames, and Heymann, “Herd Immunity.”

⁷³ It is worth noting that in the last section when I called Conficker a ‘weakest link problem,’ this was related to infected machines calling domains for commands (only one domain needed to slip through), not the actual spread of the infection.

⁷⁴ Meng Zhang, Guohua Song, Lansun Chen, “A State Feedback Impulse Model For Computer Worm Control,” *Nonlinear Dynamics* 85 (2016), <https://doi.org/10.1007/s11071-016-2779-0>.

⁷⁵ While the containment of malware is a GPG, as mentioned in 1.2.1, global strategies often must be enacted by national institutions.

those with essential security patches.⁷⁶ For individuals who cannot update their system for various reasons, exemptions could be given. Additionally, in recognition of the potential downsides of updating one's device, the government could create an equivalent of the National Vaccine Injury Compensation Program in the United States⁷⁷ or the Vaccine Damage Payment scheme in the United Kingdom⁷⁸ to pay out benefits to those who are significantly harmed by installing a required patch. In the case of the production of patches, the state could incentivize or simply mandate that companies support programs for a certain number of years. The former is broadly similar to government subsidies to pharmaceutical companies manufacturing low-margin vaccines.

Thus far, I have argued that public goods play an important role in network robustness and resilience and that the containment of malware is similar in many ways to the containment of communicable disease. As such, the containment of communicable disease can serve as a valuable example for states as they develop policies to adequately and justly contain certain kinds of malware, such as the Conficker worm. Up to now, I have mostly been focused on the practical problems of coordinating the provisioning of public goods for cyberhealth, given that these goods have two features which are not shared with canonical public goods—they are global and participatory. Having discussed what could be done to adequately provision these goods, in the next section I will turn to the obligation states, individuals, and corporations have to contribute to these goods.

1.3 Obligation and Public Goods for Cyberhealth

As Section 1 explained, the category of public goods contains a wide array of goods, everything from national defence to the eradication of polio to fireworks displays. While all of these goods may be underproduced by private markets compared to the

⁷⁶ Note that these national policies can still be part of an overarching global strategy, as discussed in the previous section.

⁷⁷ “National Vaccine Injury Compensation Program,” Health Resources & Services Administration, October 2018, <https://www.hrsa.gov/vaccine-compensation/index.html> (accessed Dec. 13, 2018).

⁷⁸ Vaccine Damage Payments Act 1979, Chapter 17, https://www.legislation.gov.uk/ukpga/1979/17/pdfs/ukpga_19790017_301114_en.pdf

socially optimal level, states, corporations, and individuals surely do not have an obligation to always correct this market failure, i.e. there is no obligation for a state to produce or subsidize fireworks displays. Determining precisely which public goods states, corporations, and individuals may have an obligation to help provide will be heavily dependent on the political theory of a specific state, but there are two qualities of a public good that generally influence this calculation—the degree to which the good is essential and the degree to which it is excludable. First, I will look at the obligations of the state, and then I will turn to individuals and corporations.

1.3.1 Obligations of the State

While some public goods are clearly always discretionary, such as the aforementioned fireworks display, others are necessary to be able to live a minimally decent life and typically are considered among the primary responsibilities of the state. The least objectionable of these goods is national defence—generally, even libertarians believe that states have a responsibility to protect citizens from external threats. However, even if one is a libertarian, there are good reasons to believe a state has an obligation to ensure adequate production of certain public goods for cyberhealth that are essential to a state's ability to wage defensive war. The most relevant of those goods is the cyberhealth of critical infrastructure—one cannot wage defensive war effectively if chemical plants, banks, dams, transportation, and manufacturing facilities are crippled.

While national defence is the least objectionable of state obligations, most political philosophers accept that states have an obligation to provide at least some other basic functions. In addition to national defence, the most basic might be protection from a hostile environment and those which enable the satisfaction of basic bodily needs (e.g. clean water).⁷⁹ The creation of these public goods is often central to

⁷⁹ George Klosko, "Presumptive Benefit, Fairness, and Political Obligation," *Philosophy & Public Affairs* 16, no. 3 (1987): 241-259. Goods like clean water will sometimes be considered public goods and sometimes common goods or club goods depending on the degree of rivalrousness and excludability of the good in question. In a country with very limited water resources, clean water is closer to an exhaustible common good. Having said this, for the purposes of this argument rivalrousness is of

a state's legitimacy within social contract theories, such as the those developed by James Buchanan,⁸⁰ Gordon Tullock,⁸¹ and Michael Moeler.⁸² In Buchanan's theory, national defence and the protection of rights are the most basic responsibilities—forming the basis of the 'protective state.' But once states enter their post-constitutional stage, one of the primary responsibilities of the state is to produce public goods for the benefit of society; the production of public goods is the basis of what Buchanan calls the 'productive state.'⁸³ As digital networks underpin nearly all forms of critical infrastructure, regardless of the specific public goods one feels a state has an obligation to produce there are reasons to think that states will also have an obligation to ensure the adequate production of public goods for cyberhealth. For instance, as I will argue in more depth in Chapter 3, as critical infrastructure and medical devices increasingly rely on digital networks, cyberhealth becomes increasingly important to an individual's ability to access healthcare. This was starkly demonstrated when hurricanes Irma and Maria destroyed Puerto Rico's internet and telecommunications infrastructure preventing many from getting urgent assistance. To get assistance, individuals needed to fill out a form online or over the phone, although

secondary importance to excludability, and thus these distinctions are of minor importance.

⁸⁰ James Buchanan, *The Limits of Liberty: Between Anarchy and Leviathan* (Chicago: University of Chicago Press, 1975), 68.

⁸¹ James Buchanan and Gordon Tullock, *The Calculus of Consent: Logical Foundations of Constitutional Democracy* (Ann Arbor: University of Michigan Press, 1962).

⁸² Michael Moeler, *Minimal Morality: A Multilevel Social Contract Theory* (Oxford: Oxford University Press, 2018).

⁸³ Buchanan, *The Limits of Liberty: Between Anarchy and Leviathan*, 68.; While Buchanan's commitment to contracts leads him to generally refrain from listing which goods should be publically produced, he lists the following goods as examples of those which most benefit from public interference in the market: herd immunity, national defence, soil erosion, college education, and public parks. James M. Buchanan, *The Demand and Supply of Public Goods* (Chicago : Rand McNally & Company, 1968).

for many people these services remained inoperable for many months.⁸⁴ While relatively few people directly died from the storms, thousands ultimately unnecessarily died over the following weeks due to broken infrastructure.⁸⁵ I will also argue in Chapter 3 that when digital devices are closely coupled to a biological system, as in the case of digital pacemakers and the circulatory system, the cyberhealth of these devices in part constitutes what it means to be healthy. In these cases, ensuring an environment with an adequate degree of cyberhealth is broadly akin to states ensuring the ability to live in a disease-free environment. As in the case of public health, states should be concerned with eliminating the most serious threats, not necessarily eliminating all malware.

Despite the importance of cyberhealth to these fundamental responsibilities of the state, traditionally many states, including the United States, have allowed the private companies which manage critical infrastructure to determine which investments in robustness and resiliency are worthwhile. Singer and Friedman describe the problem in the United States saying:

The CEO of one cybersecurity firm told us...that ‘The most critical of the critical infrastructure are the biggest laggards in cybersecurity.’ While much attention has been paid to securing areas like finance, where the incentives are more in alignment for regulation and investment, other areas of even more core importance and danger like water control, chemical industry, or the ports have almost none. In 2013, for instance, a study we helped guide of six major American ports found only one had any proper level of cybersecurity, due to the fact that the Coast Guard and Department of Transportation officials, who are in charge of regulating and protecting the ports, had literally no power or expertise in the area.⁸⁶

⁸⁴ Oliver Milman, “Six weeks after Hurricane Maria, Puerto Ricans Still Waiting for Help from Fema,” *The Guardian*, 9 November 2017, <https://www.theguardian.com/world/2017/nov/09/six-weeks-after-hurricane-maria-puerto-ricans-still-waiting-for-help-from-fema> (accessed Dec. 14 2018).

⁸⁵ Amy Sherman, “Fact-checking the Death Toll Estimates from Hurricane Maria in Puerto Rico.”

⁸⁶ Singer and Friedman, *Cybersecurity and Cyberwar*, 202.

Whereas the dominant ‘feudal’ approach to cybersecurity described in the introduction treats the security and robustness of these networks as predominantly a private good, a public health inspired approach empowers the state to meet its obligations by providing the normative justification for intervening in cyberspace. While technological ignorance is an obstacle for a state to overcome, it is not an excuse for failing to meet its most uncontroversial obligations.

In this section, I wanted to demonstrate that there are good reasons to think states have a responsibility to promote public goods for cyberhealth even if one holds a rather limited view of state power. However, throughout the rest of this dissertation, I will assume that states have a responsibility to provide a number of goods, including access to healthcare, access to education, safe environments in which to live, roads, and clean water. In this chapter, I explored one justification for states to provide these goods—the correction of market failure. In the following chapters, I will explore additional justifications for—and limits on—state action in cyberspace. In Chapter 2, I will explore the role Mill’s harm principle plays in defining spheres of public and private responsibility; in Chapter 3 I will explore the way poor cyberhealth impacts health; and in Chapter 4 I will explore the myriad of ways in which poor cyberhealth impacts personal wellbeing.

1.3.2 Obligations of the Individual

The indispensability of a public good is also important to consider when determining the obligations of human individuals to contribute to the production of public goods. This is particularly relevant for participatory public goods like the containment of malware that require individuals to perform certain actions beyond merely paying their taxes, as these goods typically place a higher burden on the individual. For example, to contribute to herd immunity one needs to get vaccinated. This in turn may require one to go to a pharmacy, clinic, or doctor’s office; pay for a vaccine; endure some physical discomfort; and face possible negative side-effects. As this is typically more onerous than paying one’s taxes, it would seem the benefits of a given participatory public good must be correspondingly higher in order for an individual to have an obligation to contribute to its production. In this section, I will explore what obligations an individual might have to contribute to the promotion of public goods for cyberhealth.

The general question of whether an individual is obligated to contribute to public goods has been looked at by many theorists including John Rawls and Robert Nozick, but I find George Klosko's argument centred on 'presumptively' beneficial public goods to be most convincing. By presumptively beneficial goods, Klosko means goods which it can be presumed all people in a community want regardless of "what their rational plans are in detail."⁸⁷ As such, Klosko argues they are "public analogues of Rawls' primary goods."⁸⁸ After laying out Klosko's argument, I will apply it specifically to public goods for cyberhealth.

Klosko's argument that individuals have an obligation to help produce certain public goods is based on the principle of fairness—the idea that "those who benefit from the cooperative efforts of others have an obligation to cooperate as well."⁸⁹ For example, if one wants to take advantage of a well being dug in one's neighbourhood, then the principle of fairness would suggest that one has an obligation to contribute to the well's creation. And if one does not contribute, then it is reasonable for others to exclude one from the benefit.⁹⁰ While applying the principle of fairness to excludable goods is relatively "trouble-free,"⁹¹ it becomes trickier in the case of non-excludable goods (like the containment of disease and malware) because one cannot choose whether or not one will receive the benefit.

The primary challenge to applying the principle of fairness to public goods is what Klosko calls the "limiting argument,"⁹² which is that the principle of fairness does not apply in the case of non-excludable goods unless, as Rawls argues, an individual has "voluntarily accepted the benefits."⁹³ Nozick, another proponent of the limiting argument, goes further than Rawls in his classic exploration of whether an individual has an obligation to participate in a neighbourhood public address

⁸⁷ Klosko, "Presumptive Benefit, Fairness, and Political Obligation," 246-247.

⁸⁸ Ibid., 246.

⁸⁹ Ibid., 242.

⁹⁰ Ibid., 243.

⁹¹ Ibid., 243.

⁹² Ibid., 244.

⁹³ Examples of Rawls' primary goods include rights, wealth, and the social bases for self-respect. (Rawls, *A Theory of Justice*, 111-12.)

system.⁹⁴ In this example, everyone in a neighbourhood is assigned a day to create entertainment for the PA system. Nozick argues that even if one accepts the benefit, say by opening up one's window and listening to the day's entertainment, one has no obligation to produce one's own program of entertainment on one's appointed day unless one supported the creation of the scheme. Nozick's argument is based on the presumption that individuals should decide for themselves if they will be forced to have their liberty curtailed, and that putting an obligation on an individual is no small matter. Klosko accepts that the limiting argument may apply in Nozick's example as the public good is discretionary—perhaps beneficial, but not presumptively so. However, he argues that when the public good is 1) presumptively beneficial, 2) worth the cost, and 3) non-excludable, then one does have political obligations to contribute regardless of whether one has “voluntarily accepted the benefits.”⁹⁵

One of Klosko's primary examples is a resident of a country which is surrounded by hostile neighbours. The threat is such that the country has instituted mandatory military service. As the individual resident cannot be excluded from the protection, would never choose to live without the benefit of national defence, and the benefit is worth the burden on the average citizen, Klosko argues there is a political obligation to serve.⁹⁶ If any criterion above is not met, then the obligation disappears. For instance, there is no obligation to contribute to a hopeless defence effort (e.g. the defence of the Alamo), as the expected benefit does not outweigh the expected cost.⁹⁷

Klosko's two other examples are less extreme circumstances. The second scenario is a city with unhealthy levels of air pollution caused by automobiles. To mitigate the problem, the city enacts restrictions on automobile use and requires automobiles to be modified to curb air pollution. Meanwhile, the third scenario

⁹⁴ Robert Nozick, *Anarchy, State, and Utopia* (Oxford: Basil Blackwell, 1974), 93-95.

⁹⁵ Klosko, “Presumptive Benefit, Fairness, and Political Obligation,” 249.; Rawls, *A Theory of Justice*, 111-12.

⁹⁶ Klosko, “Presumptive Benefit, Fairness, and Political Obligation,” 249.

⁹⁷ As an aside, it is worth noting the irony that despite Klosko's primary example coming from the arena of national defence, the cybersecurity mindset described in the introduction typically has treated cybersecurity as a private good, which individuals have no obligation to support. This is just one of a number of incoherencies within the cybersecurity framework that I will explore in greater depth in Chapter 2.

concerns an area beset by drought that enacts water restrictions on personal use to protect crops and avoid a famine.⁹⁸ As was the case in the first scenario, in each of these cases Klosko argues an obligation exists to participate in the production of the good, given that breathing clean air and avoiding famines are presumptively beneficial goods, the contributions are worth the cost, and the goods are essentially non-excludable.

Klosko's argument tracks with our intuitions that there is no obligation in Nozick's PA system example, but that there is an obligation to contribute to goods like national defence, which almost no one would choose to live without. While Nozick is right that there is no obligation to contribute to the neighbourhood entertainment, his reasoning is incomplete. It is not merely that the individual did not agree to participate, but that the benefit was the wrong kind of benefit to establish an obligation.

Klosko's argument may make it seem as if individuals have an obligation to contribute to relatively few public goods, as most public goods are not presumptively beneficial. However, Klosko argues that the obligation extends to goods that are essential to the production of presumptively beneficial goods, even if in other contexts those "access goods" may be discretionary. By access goods I mean goods that are necessary to being able to achieve other goods (e.g. vaccines are access goods for herd immunity). I will first apply this argument to non-participatory public goods and then turn to participatory public goods.

In regards to non-participatory public goods, Klosko argues that discretionary goods like highways, railroads, airports, bridges, communication technologies, and harbours are practically indispensable to national defence.⁹⁹ As such, it is reasonable to assume that the individual's obligation to support the presumptively beneficial good of national defence persists even when that good is packaged with other goods of a discretionary nature—presuming the cost does not get too high relative to the benefits. As other presumptively beneficial public goods, such as clean water, likewise rely on a host of access goods, Klosko's argument suggests individuals have an obligation to support a rather large number of seemingly discretionary government services.

⁹⁸ Ibid., 250.

⁹⁹ Klosko, *The Principle of Fairness and Political Obligation*, 88.

If we turn to cyberhealth, there are a number of public goods which seem essential to the production of presumptively beneficial goods like national defence or water treatment. As mentioned above, the cybersecurity of critical infrastructure, including military and government networks seems to be one. Vulnerabilities to those networks would hamper the ability of a state to wage defensive war. For similar reasons, I would add maintaining robust internet infrastructure and encryption standards to that list. Further research would be necessary to determine if containing malware more broadly makes the cut, but there are reasons to think it should. While most of the time malware is more of a nuisance than a serious threat to critical infrastructure, large-scale botnets, such as those created by Conficker and Mirai, do have the ability to disrupt critical infrastructure. In 2007, it is suspected that Russian forces used botnets to launch DDoS attacks against Estonia during a time of military tension between the two countries, shutting down parts of Estonia's internet for several weeks.¹⁰⁰

Where the containment of malware diverges from encryption standards or more robust infrastructure, is in being a participatory public good. As participatory public goods are typically more onerous for individuals to contribute to than non-participatory public goods, they are more likely to fail Klosko's second criterion. Having said this, the benefits of containing malware are such that I believe they typically exceed the costs. In order to actively participate in the mitigation of malware, individuals should patch their personal devices, practice safe browsing habits, and install real-time malware protection; on the whole these steps do not seem particularly burdensome compared to the benefits of both avoiding malware infections and contributing to the security of critical infrastructure. Additionally, states can craft policies to help people fulfil their participatory obligations, including educating or training individuals, subsidizing security software, and creating standards to simplify and standardize the process of installing security updates. There is an analogy in the case of states promoting herd immunity. In the UK, for instance,

¹⁰⁰ Damien McGuinness, "How a Cyber Attack Transformed Estonia," *BBC News*, April 27, 2017, <https://www.bbc.co.uk/news/39655415> (accessed Feb. 28, 2019).

the state provides most vaccines for free to lower the burden of participation.¹⁰¹ Furthermore, as updating one's system and practicing safe browsing habits also produces significant individual benefit, it seems reasonable to put the burden of proof on the objector to demonstrate that the costs exceed the benefit. As the costs can be significantly decreased by opting all users into automatic patching and real-time malware monitoring, it seems reasonable that if individuals have an obligation to support the production of cyberhealth via tax payments (as I argued previously), then they likely also have an obligation to contribute via these slightly more burdensome steps, although where the line should be drawn is open to debate. Perhaps, automatic patching is worthwhile, but insisting individuals practice safe browsing habits is not, given the cost to one's freedom. My aim here is not to resolve these specific issues, but to show how thinking about cybersecurity in terms of public goods provides us with a way of thinking about how individuals' enjoyment of the benefits of the internet can generate reciprocal obligations.

1.3.3 Corporate Obligations

In this section, I will assume that it makes sense to talk about corporations as having obligations.¹⁰² However, even if one does not accept this view, the arguments that follow may still be useful for thinking about the kinds of regulations that states may reasonably impose upon corporations. In addition to the obligations discussed above, which were grounded in the principle of fairness, in the corporate case I will also argue there are obligations derived from the general moral principle to not cause harm to others.¹⁰³ I will refer to this principle as the harm to others principle.¹⁰⁴ As with the

¹⁰¹ Public Health England, "Complete Routine Immunisation Schedule," GOV.UK, May 7, 2014, <https://www.gov.uk/government/publications/the-complete-routine-immunisation-schedule> (accessed Feb. 28, 2019).

¹⁰² Peter French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984).

¹⁰³ This general principle includes a collection of obligations related to beneficence and non-maleficence. Following Dawson's lead, I will refer to these obligations collectively as the 'harm to others principle.' Angus Dawson, "Herd Protection as a Public Good: Vaccination and our Obligations to Others," in *Ethics, Prevention, and*

notion of corporate obligations in general, I will be assuming that it makes sense to ground corporate obligations in notions of fairness and harm prevention, while acknowledging that this is controversial. While the harm to others principle may have some relevance in the case of human individuals and their obligations to contribute to the mitigation of malware, it is more relevant in the case of corporate obligation as a human individual's contribution to poor cyberhealth will only trivially harm others. By comparison, corporations' poor security choices have the ability to cause substantial harm. For example, the DDoS attack on Dyn, Inc. (part of the DNS) was largely possible because of the poor security practices of a few Chinese device manufacturers.

My argument in this section is based on Angus Dawson's argument that individuals have an obligation to get vaccinated against serious diseases in order to protect others from harm. I will first introduce this argument, then I will adapt it for the case of malware, and finally I will suggest what this argument and the principle of fairness may mean for three types of businesses. I will only focus on the mitigation of malware in this section, as I will assume that like human individuals corporations have a clear obligation to contribute to at least some non-participatory public goods for cyberhealth through the payment of taxes, given the presumptively beneficial nature of national defence (or at least the stability which national defence provides).

1.3.3.1 Vaccination and the Harm to Others Principle

Dawson argues that when a state of herd immunity does not exist, then an individual has a moral obligation get vaccinated against serious diseases based upon the general moral principle to not cause harm to others. The essential formulation of Dawson's argument, quoted directly, is:

1. Contagious diseases that might result in (more than trivial) harm can be passed on to others through non-intentional action.
2. Such a risk of harm can be reduced through vaccination of any potential source individual in advance (where a relevant vaccine exists).

Public Health, eds. Angus Dawson and Marcel Verweij (Oxford: Oxford University Press, 2007), 166-167.

¹⁰⁴ This is the way Angus Dawson refers to this collection of principles. I will use this term as well, given his argument is the basis for my argument in this section.

3. We have a general moral obligation not to cause harm to others through our own actions and inactions.
4. Given 1 and 2, an individual can reduce the risk of causing (non-trivial) harm to others through vaccination for (serious) contagious disease.

Conclusion: Given 3 and 4, we are morally obligated to have vaccinations for (serious) contagious diseases (where available).¹⁰⁵

The similarities between malware and infectious disease, make this an appealing argument for thinking about corporate obligations to contain malware. Specifically, like communicable diseases malware can spread between devices, the harm malware can cause is not trivial, and effective mitigation strategies exist.

Adapting Dawson's argument for the purposes of mitigating malware, one gets the following:

1. Malware that might result in (more than trivial) harm can be passed on to others through non-intentional action.
2. Such a risk of harm can be reduced if the device has certain reasonable security features (e.g. up to date patches, strong passwords).
3. We have a general moral obligation to not cause harm to others through our own actions and inactions.
4. Given 1 and 2, corporations can reduce the risk of causing (non-trivial) harm to others by patching their own devices and adding reasonable security features to the devices they design and manufacture (if applicable).

Conclusion: Given 3 and 4, technology producers have a moral obligation to patch their own devices and add adequate security features to their products.

This argument should be appealing to those who accept that there is a general moral obligation to not harm others. If the people being harmed were the one's *knowingly* buying products with poor security features, then the harm to others principle would likely not apply, but poor security practices can harm those who did not consent to being exposed to harm. In fact, due to the risk to critical infrastructure, poor security practices can even harm those who lack internet access.

While this argument suggests a moral obligation for corporations to protect their devices and ensure the devices they produce have adequate security features, two caveats must be mentioned. First, Dawson argues that this obligation falls away when

¹⁰⁵ Ibid., 168.

a state of herd immunity exists as getting vaccinated will add no additional benefit, but may cause harm to the individual getting vaccinated (e.g. negative side-effects).¹⁰⁶ And second, Dawson argues there is no obligation to get vaccinated against non-communicable diseases like tetanus.¹⁰⁷ In the case of malware, however, neither of these caveats apply.

First, we can dismiss the ‘no additional benefit’ argument because in the case of malware, herd immunity is much harder to achieve. As infected machines can pass on infections or launch attacks as part of a botnet from anywhere on the globe, local herd immunity is insufficient to mitigate the threat. In 2017, for instance, ten years after Conficker first appeared, there were over two million new infections worldwide.¹⁰⁸ These new infections from all over the globe still could make up a single botnet capable of attacking targets anywhere. Given that in the case of malware we are always in a sub-herd immunity state, Dawson’s ‘no additional benefit’ argument does not come into play—there is always an additional benefit to protecting one’s devices.¹⁰⁹ Second, unlike vaccines which target specific diseases, the malware protection strategies mentioned above work against a broad array of malware threats. Some very targeted malware attacks may be analogous to tetanus, but (in general) the same techniques are used to block most forms of malware (i.e. general security patches, safe browsing habits, real-time monitoring). One could imagine a parallel might be if there were a single all-purpose vaccine that worked against all types of disease. Based on Dawson’s argument, it seems plausible that there would be a moral obligation to get this vaccine because it generates a substantial public benefit by containing serious infectious diseases. The fact that it *also* would protect one from tetanus would not matter.

Having said this, the specific ways in which a corporation might be obligated to contribute to the mitigation of malware will be dependent on the type of business. First, let us consider a large accounting firm that has 100,000 employees, but does not

¹⁰⁶ Ibid., 177.

¹⁰⁷ Ibid., 165.

¹⁰⁸ Patrick Howell O’Neill, “Conficker Worm Still Spreading Despite Being Nearly 10 Years Old,” *Cyberscoop*, Dec. 8, 2017, <https://www.cyberscoop.com/conficker-trend-micro-2017/> (accessed Feb 27, 2019).

¹⁰⁹ Dawson, “Herd Protection as a Public Good,” 171.

make technology products.¹¹⁰ In this case, the corporation's obligations to mitigate malware are not fundamentally different from 100,000 individuals with internet connected computers. Like an individual, the corporation should keep their computers up to date, encourage their employees to practice safe browsing, and use real-time malware detection systems. While this firm may be more likely to be a victim of a targeted malware attack than the 100,000 individuals, most of the harm associated with such an attack is isolated to the firm and its customers.

The second type of business is a technology company making consumer technology products. For example, let us assume that they make internet connected webcams, the kind of products that were drafted into the Mirai botnet used to attack the DNS system in 2016.¹¹¹ This company will have all the obligations of the first in regards to their own machines, plus obligations to ensure the devices they manufacturer are adequately secure based on the harm to others principle.

The third type of business is one that manages critical infrastructure, such as a nuclear power plant. While these businesses are not exempt from the obligation to mitigate malware derived from the principle of fairness and harm to others principle, the more important obligation is certainly to protect their own network's integrity given that the failure of critical infrastructure is one of the main ways in which malware can cause significant harm. The obligation to sufficiently invest in their own network security, like the obligation to add adequate security features to manufactured devices, can be justified by the harm to others principle. Given the direct and significant harm that can result from the failure of critical infrastructure networks, this obligation is stronger than in the case of the device manufacturer discussed previously.

While in this section I have used the 'harm to others principle' to consider what obligations corporations might have to contribute to malware mitigation, another

¹¹⁰ Large firms have easily over 100,000 employees. For instance, EY Global (Ernst & Young) had over 270,000 employees in 2018. "EY Global Chairman and CEO Mark Weinberger to step down effective July 1, 2019," EY Global, December 3, 2018, https://www.ey.com/en_gl/news/2018/12/ey-global-chairman-and-ceo-mark-weinberger-to-step-down-effective-july-1-2019 (accessed April 4, 2019).

¹¹¹ Lovelace Jr. and Vielma, "Friday's Third Cyberattack on Dyn 'Has Been Resolved,' Company Says."

way to look at this issue is through the related notion of Mill's harm principle.¹¹² The harm to others principle speaks to individuals' obligations, while Mill's harm principle guides the appropriate use of state power. Given that corporations can harm others through their poor cyberhealth practices, states may have an obligation to prevent this harm through punishment or regulation. I will revisit the question of the role of the state in cyberspace in 2.4.1, at which point I will explore the use of Mill's harm principle in greater depth.

Having established that there are potentially obligations for businesses to mitigate malware and invest in cyberhealth, or at least grounds to legally require them to do so, a major obstacle to corporate investment in cybersecurity has been that we currently lack a good understanding of the specific value of various security strategies and the costs of cyberattacks. Challenges to estimating costs include identifying an appropriate time horizon, monetizing qualitative impacts, quantifying the risk, and determining the social discount rates applied to monetized future events.¹¹³ While these challenges should not stop corporations from investing in tried and true methods, such as keeping software up to date and following security standards, they may still lead to underinvestment overall compared to what their obligations require. While some companies may over-invest in cybersecurity, the evidence suggests that the vast majority underinvest and act as quasi free-riders.¹¹⁴

To help correct this underinvestment, it would be sensible for states to invest public funds in analytic tools to help measure the likelihood and costs of network threats (e.g. malicious attacks, human error, natural disasters) and the effectiveness of various types of cyberhealth interventions. These analytic tools could also help policymakers quantify the benefit companies receive from public cyberhealth, which could bolster the case for corporate obligations derived from the principle of fairness and serve as a basis for establishing appropriate regulations.

In addition to these analytical tools, legislation can help ensure that companies do reap the costs of their insecurity by requiring public disclosure of breaches and

¹¹² John Stuart Mill, *On Liberty*, in *Utilitarianism and Other Writings*, 1859, ed. Mary Warnock (Glasgow: Collins, 2003), 94-95.

¹¹³ Bauer and Van Eeten, "Cybersecurity," 714.

¹¹⁴ Howard Kunreuther and Geoffrey Heal, "Interdependent Security," *Journal of Risk and Uncertainty* 26, no. 2 (2003): 231-49.

holding companies liable for the damage their insecurity caused to other nodes in the network. Holding the Chinese device manufacturers responsible for the harm caused to Dyn, Inc. and the thousands of websites affected by that attack would help shift these companies' cost-benefit analyses to favour proactive strategies over reactive strategies. While reputational harm alone can help some of these externalities be internalized, Bauer and Van Eeten have found the “feedbacks were too weak, localized, or too slow to move agents' behavior swiftly towards more efficient social outcomes.”¹¹⁵

1.4 Conclusion

In this chapter, I argued that the practices, philosophy of, and history of public health can serve as useful guides for thinking about the public goods which bolster cyberhealth. Specifically, I focused on one classic cybersecurity issue—the mitigation of malware. First, in Section 1, I provided an overview of public goods and introduced a number of public goods for public health, including most importantly the containment of infectious diseases. Next, I argued that many of the kinds of public goods which are valuable to public health also are valuable for promoting cyberhealth, including surveillance programs, information sharing programs, standards, and fundamental research. In Section 2, I then explored one specific public good for cyberhealth—the containment of malware—using the example of the Conficker worm. I argued that the containment of malware was similar to the containment of communicable disease in two primary ways. First, both are Global Public Goods—goods whose benefit spreads across international borders and which often can only be adequately produced via international efforts. The lack of a world government with coercive powers means that the provision of these goods relies on building international consensus and institutions which states are willing to grant authority and power. While public goods like the containment of malware and the containment of communicable diseases can be produced at the national level, these efforts will be limited in their success given the ability of malware and diseases to spread across borders. While acknowledging the differences between malware and communicable diseases, I argued that the creation of the WHO and the International

¹¹⁵ Bauer and Van Eeten, “Cybersecurity,” 714.

Health Regulations could serve as an example of how to build international consensus over time around the containment of malware.

The second main similarity is that the containment of malware and communicable diseases cannot be achieved merely through the payment of taxes, but rather individuals must actively participate in the production of the good. This requirement raised interesting questions as to what obligations individuals and companies have to contribute to the containment of malware. Using George Klosko's response to the limiting argument, I argued that the principle of fairness suggests that individuals and corporations do generally have an obligation to contribute to public goods for cyberhealth if they are presumptively beneficial or enable presumptively beneficial public goods like national defence. Apart from the principle of fairness, I argued that corporations have additional obligations based on the 'harm to others principle.' First, I introduced Angus Dawson's argument that individuals have an obligation to get vaccinated against serious infectious diseases based on the general obligation to avoiding harming others. Then, I applied this argument to the containment of malware, arguing that corporations have an obligation to patch their systems and ensure that any technology products they make have adequate security protections. In cases where the costs are potentially substantial and the benefits amorphous, I suggested a sensible first step would be for governments to invest more heavily in developing analytical tools and to change legislation to help internalize the costs of cybersecurity failures. These actions should at least allow companies and individuals to more accurately assess their own risk and adjust their cybersecurity investment accordingly.

Finally, it is worth saying that the reason a public good is provided may be as important as whether it is provided at all, as many techniques that could be employed to secure networks would also destroy privacy and jeopardize notions of the open web. In this regard, public health is a more benevolent model than law enforcement or economics. This will be a central focus of Chapter 2. A tool like network monitoring, for instance, can gather information in an anonymous, minimized, and decentralized way which protects individual privacy (as it is in the public health context),¹¹⁶ or it can be used as a mechanism for crushing political dissent. When the motivation for

¹¹⁶ Sedenberg and Mulligan, "Public Health as a Model for Cybersecurity Information Sharing."

improving cybersecurity is merely economic or strategic, then often privacy and personal freedom end up sacrificed at the altar of security.

In this chapter, I have presented a series of parallels and analogies between cyberhealth and public health that demonstrate the potential value of a public health inspired approach as an overarching framework for guiding technology policy. Having demonstrated one context in which such an approach is useful, in the next chapter, I will more formally define the public cyberhealth framework. This formalization will be useful for distinguishing Public Cyberhealth from the dominant cybersecurity lens, and for demonstrating that this approach is a cohesive and reasonably comprehensive way to conceptualize the digital landscape as a whole.

Chapter 2: Two Levels of Abstraction

In Chapter 1, I demonstrated the value of thinking about a paradigmatic cybersecurity issue, the mitigation of malware, in public health terms. Using the example of the Conficker computer worm, I argued that a public health approach, grounded in the theory of public goods, was superior to the existing paradigm, which generally treated malware mitigation as a private good. I then argued that the philosophy of public health can help one to understand the obligations of states, individuals, and corporations to contribute to public goods for cyberhealth, while public health institutions can serve as blueprints for how to provision said public goods.

While one may see the value of the public health approach for thinking about public goods for cyberhealth, one might rightfully ask whether a public health lens has broader utility for thinking about technology policy and our relationship to digital information. After all, it is not unusual that a specific learning from one field is useful in another. I was once at a conference where a data scientist described how the behaviour of trout helped him to better understand the movement of retail customers on a store floor, but no one is arguing that angling is the appropriate lens to understand all aspects of retail strategy. In this chapter, I will argue that the public health approach is not only useful for mitigating malware but can be a cohesive and reasonably comprehensive way to conceptualize our relationship to digital information. I will demonstrate this by formalizing the framework using the method of Levels of Abstraction (LoA), a method for clearly defining the variables, observables, and behaviours that comprise a framing device.

First, in Section 1, I will introduce the method of LoA and explain its utility. Then, in Section 2, I will define what I call the Cybersecurity LoA. This will be a formal statement of what I have referred to informally as the cybersecurity mindset or lens. As a reminder, in the Introduction, I described this lens as focused on malicious-attacks, and I argued that those using this approach tend to characterize cyberspace as a kind of battlefield. Once I have defined the main features of this LoA, I will evaluate its cohesiveness and utility. I will argue that the LoA's limited scope make it inadequate as an overarching framework for creating technology policy, and that its

internal inconsistencies undermine its primary goal of keeping digital information and networks secure.

In Section 3, I will then formally define the public health inspired alternative approach, which I will call “Public Cyberhealth.” In contrast to the Cybersecurity LoA, the Public Cyberhealth LoA is designed to address both malicious and non-malicious threats to network robustness and resiliency, capture the impact of technology policies and interventions on health and wellbeing, and identify potential ethical conflicts. I will argue that this alternative framework corrects a number of the incoherencies of the Cybersecurity LoA, and not only is a better approach to thinking about cybersecurity (as seen in Chapter 1) but shows greater promise as an overarching framework for creating technology policies which improve individuals’ lives. Lastly, in Section 4, I will expand upon the discussion from 1.3 and explore how the Public Cyberhealth LoA can be used to understand the normative justification for—and ethical limitations on—government interventions in cyberspace, a necessity for crafting consistent and justified technology policies.

The Public Cyberhealth LoA is intended to be useful to a number of actors, including (but not limited to) policymakers seeking to craft consistent, effective, and just technology policies; creators of technology products who wish to safeguard or improve their customers’ wellbeing; and ISPs aiming to improve the reliability of their networks while protecting individuals’ rights. As an alternative way of viewing the digital landscape, the Public Cyberhealth LoA can be used in a variety of ways. In some cases, the value may simply be in causing one to question the assumptions of the existing security paradigm, while in other cases the LoA may highlight a previously overlooked impact or suggest the applicability of a specific public health tool. While I will explore some of these applications in this dissertation, one should not assume these examples to be comprehensive.

2.1 Levels of Abstraction

The method of LoA has its roots in computer science, but has most recently been developed and popularized by Luciano Floridi.¹ It is based on the idea that whenever one tries to answer questions about a given system, one highlights certain relevant variables, observables, and behaviours while ignoring those that are deemed

¹ Luciano Floridi, *The Philosophy of Information*.

irrelevant. In this sense it is simply a way of more formally describing the colloquial notion of framing a problem. As Floridi says, “[the method of LoA] should not be confused with some neo-Leibnizian dream of a *calculemus* approach to philosophical problems.”² Rather than get too bogged down in theory, I will demonstrate how the method of LoA works by applying it to one issue which is frequently seen through different lenses—illegal drug use.

As illegal drug use is a complicated societal issue it can be described in different ways depending on one’s goals. Two common frames applied to the problem are law enforcement and public health.³ If one is a police officer, one’s goal in understanding illegal drug use in a society may be to disrupt the drug economy by arresting drug users and sellers.⁴ In order to achieve this goal, one will highlight certain variables, observables, and behaviours of the system in question, while ignoring others. One might focus on variables such as drug users (who they are, where they live, etc.), sellers (who they sell to, their criminal connections), growers, and manufacturers; and one might focus on behaviours of the system such as how an influx of new drugs impacts the existing market. This collection of variables, observables, and behaviours could be called the “Law Enforcement LoA” for understanding and responding to illegal drug use. Meanwhile, if one is a public health expert, one’s goal might be to limit deaths by overdose and improve the health of drug users. In pursuit of this goal, one might focus on variables like treatment options for individuals, training for paramedics, the impact of drug use on families, life

² Ibid., 79.

³ Hilgunn Olsen, “Open Drug Scenes and Police Strategies in Oslo, Norway,” *Journal of Scandinavian Studies in Criminology and Crime Prevention* (2017): 141-156.; Douglas N. Husak, “Drugs, Crime and Public Health: A Lesson From Criminology,” in *Criminal Law, Philosophy and Public Health Practice*, eds. A. M. Viens, John Coggon, and Anthony Kessel, 42-61 (Cambridge: Cambridge University Press, 2013), <https://doi.org/10.1017/CBO9781139137065.003>.

⁴ For the purpose of this exercise, I will risk being overly simplistic. Surely, in many locations police have a number of goals, including the health of drug users.

expectancy, and health care costs;⁵ and one might focus on behaviours such as how an influx of new drugs burdens the healthcare system or affects health outcomes. This collection of variables, observables, and behaviours could be called the “Public Health LoA” for understanding illegal drug use. The law enforcement officer and the public health expert are both describing the *same* system of illegal drug use within a society, but their divergent goals lead them to focus on *different aspects* of the system. To one using the Law Enforcement LoA, a drug user may be a ‘criminal.’ Meanwhile, to one using the Public Health LoA, the same person is a potential ‘patient.’ Which is the correct designation depends on the questions one is trying to answer. Often, states use both of these frames simultaneously as part of dual-track policies.⁶

I will be engaging in a similar exercise in this chapter by formally defining the variables, observables, and behaviours of the cybersecurity and public cyberhealth frameworks, which can be used to conceptualize the digital landscape. Using the method of LoA to more formally define these approaches is useful for 1) clearly identifying the goal of a given framework, 2) spelling out one’s assumptions, 3) comparing competing frameworks, and 4) helping one build more useful models of the system in question.⁷ In particular, the method of LoA can be useful for forcing one to consider the implicit assumptions of dominant mindsets like the cybersecurity mindset. While the language of cybersecurity is ubiquitous, it is of course just as much of a LoA as the public health inspired alternative I am proposing.

Formalization is most useful when dealing with smaller, more easily quantifiable problems, such as the selling of a used car.⁸ Sprawling concepts like *society*, for instance, may be simply too complicated to be usefully described using

⁵ World Health Organization, “Management of Substance Abuse: Terminology & Classification,” World Health Organization, https://www.who.int/substance_abuse/terminology/en/ (accessed Mar. 1, 2019).

⁶ Olsen, “Open Drug Scenes and Police Strategies in Oslo, Norway.”

⁷ This process is even more essential when one frame is entrenched as the dominate paradigm.

⁸ For smaller problems, there are often fewer variables, and one can create models that provide more predictable outcomes, such as the appropriate price for a used vehicle.

the method of LoA.⁹ Digital networks and their impact on human wellbeing fall somewhere in the middle. For more complex systems, like digital networks, properly speaking one will often describe them using what is called a Gradient of Abstraction (GoA), which is an interlocking group of LoAs that each describe a piece of the overall system. Technically, when I speak of the Cybersecurity LoA and Public Cyberhealth LoA, I will be speaking about Gradients of Abstraction comprised of practitioner LoAs (i.e. the way a IT professional might view the digital landscape), strategic LoAs (i.e. the way a CTO or organization head might view the digital landscape), legal LoAs, etc. However, this can quickly become very complex and difficult to effectively illustrate. As such, in this dissertation, I will present a simplified version of the Cybersecurity and Public Cyberhealth GoAs. As I will be using the simplified version, I will continue to use the term ‘level of abstraction’ rather than the cumbersome ‘gradient of abstraction.’¹⁰

2.1.1 Evaluating LoAs

While one can frame a given system in any number ways using the method of LoA, not all LoAs are equally useful or reliable. For instance, if one wanted to test out whether a new policy was going to improve health outcomes for users of illicit drugs, then one would not want to use the Law Enforcement LoA which lacked the relevant variables and observables to measure health impacts. As LoAs are used to build models of systems, which in turn can be used to test theories about that system, one can evaluate LoAs on their utility and coherence. Using these criteria one can both assess a LoA on its own merits and compare it to other competing LoAs describing the same system. Below I will describe what utility and coherence mean in this context. I will then use these concepts to assess the Cybersecurity LoA and the Public Cyberhealth LoA in Sections 2.2 and 2.3 respectively.

⁹ Floridi, *The Philosophy of Information*, 79.

¹⁰ While the term Levels of Abstraction might suggest a hierarchical structure, the method of LoA does not assume or require that the system in question be structured hierarchically or that the levels used to model the system relate hierarchically.

2.1.1.1 Utility

When evaluating the utility of a LoA one can speak of internal and external utility. I define internal utility as *the degree to which a LoA is effective at achieving its stated purpose*. In the case of the Public Health LoA for illegal drug use, for instance, we can ask if using the LoA improves health outcomes. Meanwhile, I define external utility as *the degree to which the stated purpose of a LoA is useful to the broader goals of society*. A Law Enforcement LoA for combating illegal drug use may result in many people going to jail, but does mass incarceration create more problems than it solves, all things considered?

2.1.1.2 Coherence

Coherence, meanwhile, can be broken down into three subcategories, logical coherence, operational coherence, and inter-LoA coherence, although only the latter two are relevant for this exercise, as both LoAs we will look at are logically coherent. By operational coherence I mean that *the various observables and behaviours of the LoA work together to efficiently achieve one's aim*. This definition is similar to, and inspired by, Hasok Chang's definition of pragmatist coherence—"a harmonious fitting-together of actions that leads to the successful achievement of one's aims."¹¹ Operational coherence differs from internal utility in that internal utility is about the end result, while operational coherence speaks to the process of arriving at that end result. For instance, as I will explore in more depth in the next section, one could argue that the Cybersecurity LoA does achieve its goal of securing digital information and networks, but does so in an inefficient, and at times self-defeating, manner. Having said this, frequently the two concepts cannot be fully separated, as a lack of operational coherence typically reduces the utility of a LoA.

Finally, by inter-LoA coherence I mean *how well a LoA works with other established LoAs*. In isolation, Ptolemy's geo-centric 'LoA' for modelling the movement of celestial objects can be used to make reasonably accurate predictions about when certain celestial phenomenon will occur.¹² However, the heliocentric

¹¹ Hasok Chang, "Pragmatic Realism," *Humanities Journal of Valparaíso*, no. 8 (2016): 112.

¹² Stanley E. Babb, Jr., "Accuracy of Planetary Theories, Particularly for Mars," *Isis* 68, no. 3 (1977): 426-434.

model of Copernicus was ultimately more compatible with other LoAs used to describe the physical world including Newtonian physics.¹³ While one might interact with a LoA in relative isolation, no LoA is an island entire of itself. Inter-LoA coherence is certainly not sufficient to determine the quality of an LoA, but it is a useful check when used in conjunction with the other criterion mentioned previously.

2.2 The Cybersecurity LoA

Before defining the Public Cyberhealth LoA, I will first define the Cybersecurity LoA, which will serve as a point of comparison. As the cybersecurity lens is the dominant way we tend to conceive of cyberspace, defining this lens in more formal terms will hopefully reveal the assumptions inherent in this approach. While it may seem like cybersecurity is the natural way to discuss threats to network robustness, I hope to demonstrate that it is in fact a very specific and idiosyncratic way of conceptualizing cyberspace. One challenge of defining the Cybersecurity LoA, however, is that there is not one single cybersecurity framework or LoA. Someone working at the NSA might conceive of the problem of informational security differently than a computer scientist at the University of Cambridge; in my experience the latter are more concerned with protecting privacy and individual rights than the former. Therefore, there is a real danger that any attempt to formalize a single cybersecurity approach will be overly reductive or merely a straw man.

While the Cybersecurity LoA I will define in this section cannot represent all of the diversity within the cybersecurity community, I believe the variables and behaviours I will describe are broadly representative of how cybersecurity practitioners and policymakers in the United States, United Kingdom, and Europe think about cyberspace and information security.¹⁴ The variables, observables, and behaviours I have selected come from a review of cybersecurity literature, the public statements of technology policymakers, my personal experience working in the

¹³ Roy Porter, *The Scientific Revolution in National Context* (Cambridge, UK: Cambridge University Press, 1992).

¹⁴ While I will not go into it in this dissertation, one should note that describing the digital landscape as a battlefield not only shapes the behaviours of states, but also the behaviour of cybercriminals.

technology industry, and discussions with computer scientists, lawyers, and practitioners at technology conferences in the United States and United Kingdom.¹⁵

2.2.1 The Cybersecurity LoA

The first step in defining a LoA is to define the purpose of the LoA. It is this purpose that, in theory, dictates which variables, observables, and behaviours are highlighted and which are ignored. While I will complicate this claim a bit later by arguing that the Cybersecurity LoA is heavily influenced by the pre-existing LoAs of criminal justice and military intelligence, the purpose of the Cybersecurity LoA is 1) *to stop adversaries from gaining unauthorized access to digital information, networks, or devices* and 2) *to bring those who commit these illegal acts to justice*. At the core of this idea is ‘the adversary’—a malicious actor. According to Singer and Freidman, if there is no adversary, then technically speaking there is no cybersecurity threat.¹⁶ Natural disasters, human error, or poorly written code would not be considered cybersecurity threats, per se, although each could contribute to the susceptibility of a network to malicious attacks.

Connected to the variable of ‘the adversary,’ are a variety of typed variables inspired by, or taken wholesale from, the domains of criminal justice and military intelligence, including tools and strategies to identify, capture, prosecute, and deter attackers. It is important to note that the Cybersecurity LoA did not spring forth fully fledged but evolved along with the threat of information theft. The primary eras of this evolution include the Cold War, the rise of digital corporate espionage in the 1990s, and the emergence of cybercrime targeting private citizens in the early 2000s.¹⁷ Each era saw new variables and observables added to address the emerging threat of malicious attacks. In the table below I have outlined a simplified version of the Cybersecurity LoA.

¹⁵ Singer and Friedman, *Cybersecurity and Cyberwar*.; Yannakogeorgos and Lowther, *Conflict and Cooperation in Cyberspace: The Challenge of National Security*; Kaplan, *Dark Territory*.; Gallagher, Link, and Rowe, *Cyber Security*.; Trim and Upton, *Cyber Security Culture: Countering Cyber Threats through Organizational Learning and Training*.

¹⁶ Singer and Friedman, *Cybersecurity and Cyberwar*, 34.

¹⁷ Kaplan, *Dark Territory*.

Table 1

Simplified Cybersecurity LoA

Variables	Valid Values	Invalid or Minimized Values
Adversary	<ul style="list-style-type: none"> • Cybercriminal • State • Hacktivist • Advanced Persistent Threat (may or may not be state-aligned) 	
Type of Threat	<ul style="list-style-type: none"> • Worm • Trojan • Phishing • Other adversary based attacks 	<ul style="list-style-type: none"> • Non-adversary based threats (natural disasters, human error, fragile infrastructure)
Cost of Attack	<ul style="list-style-type: none"> • Financial value of information • Strategic value of information • Reputational harm • Financial cost of destruction of infrastructure 	<ul style="list-style-type: none"> • Impact on health and wellbeing, • Impact on other network nodes (negative externalities) • Impact on individuals' rights
Defensive Capabilities	<ul style="list-style-type: none"> • Private network monitoring • Firewalls • Employee/Personal education 	<ul style="list-style-type: none"> • Public defence efforts (threat monitoring, public education campaigns)
Cost of Defensive Efforts	<ul style="list-style-type: none"> • Financial cost of preventative strategies • Employee/personal time 	<ul style="list-style-type: none"> • Health and wellbeing costs • Impact on other nodes on the network (negative externalities) • Impact on individuals' rights
Offensive Capabilities	<ul style="list-style-type: none"> • Hacking back • Infiltrating adversary networks • Pre-emptive cyberattacks 	
Cost of Offensive Efforts	<ul style="list-style-type: none"> • Financial • Employee time • Provocation of additional attacks 	<ul style="list-style-type: none"> • Health and wellbeing costs, • Impact on other nodes on the network (negative externalities) • Impact on individuals' rights, • Militarization of cyberspace
Ethical Considerations	<ul style="list-style-type: none"> • Proportionality of attacks 	<ul style="list-style-type: none"> • Impact on individuals' rights • Obligations to others on the network
Insurance	<ul style="list-style-type: none"> • Insurance for financial costs 	<ul style="list-style-type: none"> • Insurance for non-financial harms (wellbeing, health)
Possible Role of State (specifics depend on circumstances)	<ul style="list-style-type: none"> • Limited information sharing • Limited assistance (CERTs) • Liability protection • Law enforcement (prosecution, extradition, punishment) 	<ul style="list-style-type: none"> • Strong regulations • Robust network monitoring • Crisis coordination • Mandatory information sharing

Key Behaviours:

1. Adversary must be present for there to be a cybersecurity threat
2. Some (but not all) offensive capabilities may be limited to state actors
3. Owner of network or information is responsible for its security

While this is a necessarily simplified picture, cybersecurity practitioners, policymakers, and strategists implicitly (and sometimes explicitly) use these variables, observables, and behaviours to make sense of the digital landscape. This includes estimating the likelihood and cost of various attacks, identifying potential vulnerabilities, devising defences, and planning counterattacks. While the Cybersecurity LoA can be used to devise a myriad of strategies, each are generally composed of the building blocks listed in the table above. For example, the Conficker Working Group might not seem to fit the feudal portrait I have painted. However, given the nature of the threat, an ad hoc coalition of hosting companies, security companies, internet governance organizations (e.g. ICANN), and Microsoft is exactly the kind response one should expect within a framework which limits state power and emphasizes financial and reputational harms. In the next sub-section, I will evaluate this LoA. Then, in Section 3, I will define the Public Cyberhealth LoA, which I am proposing as an alternative.

2.2.2 Evaluating the Cybersecurity LoA

While the Cybersecurity LoA is useful as a way of conceptualizing malicious attacks, I will argue that 1) its lack of operational coherence and limited focus on financial harm undermine its supposed goal of protecting digital information and networks, and 2) its narrow focus on malicious attacks limits its external utility as an overarching way to conceptualize the digital landscape. In Chapter 1, I explored some of these issues in the specific context of Conficker. In this section, I will explore these issues as they relate to cybersecurity and the digital landscape more generally, beginning with a lack of operational coherence and internal utility.

2.2.2.1 Internal Utility

As a reminder, internal utility is an LoA's ability to achieve the goal of the LoA. In the case of the Cybersecurity LoA, I defined the goal as: 1) to stop adversaries from gaining unauthorized access to digital information, networks, or devices and 2) to bring those who commit these illegal acts to justice. The Cybersecurity LoA fails to meet this goal as effectively as possible due to a lack of operational coherence and an incomplete accounting of the harms of cyberattacks.

The Cybersecurity LoA lacks coherence in at least three primary ways. First, as discussed at length in Chapter 1, while cybersecurity (and security more generally) displays the characteristics of a public good, the Cybersecurity LoA treats cybersecurity as a private good to be supplied by the owner of the information, network, or device in question. While one can produce goods which exhibit the characteristics of being non-excludable and non-rivalrous via private markets, this production will be inefficient and will lead to the underproduction of the good in question compared to the socially optimal level.¹⁸ National security is the paradigmatic public good; therefore, it is particularly odd that cybersecurity—which is certainly a part of national security—is treated as predominately a private good. This incoherence can be seen clearly in the example of the Conficker worm discussed in Chapter 1. As the Conficker worm posed a threat to critical infrastructure all over the world, it was in the world community’s interest to contain the threat. However, working within the cybersecurity mindset, states left the problem to be dealt with by an ad hoc group of volunteers. Viewing the problem through the Cybersecurity LoA, the United States government treated the vulnerability as Microsoft’s problem; Microsoft released a patch, but left it up to individuals whether or not it would be installed; and many individuals did not feel the need to install the patch as Conficker posed little risk to their own devices. By framing shared network problems as private problems to be resolved largely through private actions, the LoA undermines its own goal of securing information, networks, and devices from malicious attacks. Individuals and less wealthy states are left particularly vulnerable to attack. For example, while the Mirai botnet temporarily disrupted access to many popular websites in the United States and Europe, it was able to almost completely shut down Liberia’s internet for several days.¹⁹ Here the feudal analogy may once again be evocative; in the case of an invasion, those behind the castle walls may be safe, while those outside are left to defend themselves with little more than pitchforks.

The second way in which the Cybersecurity LoA lacks coherence is by over-emphasizing the importance of traditional law enforcement strategies, such as the identification, extradition, and prosecution of cyber criminals. While the Cybersecurity LoA generally downplays the role of the state, one area in which

¹⁸ Head, *Public Goods and Public Welfare*, 80-81.

¹⁹ Woolf, “Massive Cyber-Attack Grinds Liberia's Internet to a Halt.”

governments are expected to play a role is in the investigation and prosecution of cybercrimes after they have been committed. However, traditional law enforcement strategies are often of limited use in cyberspace.²⁰ Apprehending and prosecuting individuals requires the ability to positively identify cyber attackers and extradite them to the country where the crime was committed. Identification is time and resource intensive,²¹ legal notions of responsibility online are frequently fuzzy,²² and false flag operations are common.²³ As a result, the accused often has plausible deniability and most cyberattacks are never investigated. This is especially true of the types of criminals who target individuals. While the cybercriminals who attack major corporations may be brought to justice, the attacks which impact individuals are almost never investigated, as they do not justify the substantial cost of cyber forensics.²⁴ More intensive and invasive network monitoring could improve attribution, but these strategies would certainly jeopardize individual rights. In the extreme, such policies may force individuals to give up the anonymity which helps enable fundamental rights like the freedom of speech and association.²⁵

Even in cases where positive identification can be made, often law enforcement cannot arrest the perpetrator as they fall outside of their jurisdiction. For example, the largest prosecution offices in Texas only reported having prosecuted a handful of individuals for cybercrimes between 2012-2017,²⁶ despite there being over

²⁰ Mulligan and Schneider, "Doctrine for Cybersecurity," 8-9.

²¹ Nick Selby, "Local Police Don't Go After Most Cybercriminals. We Need Better Training," *Washington Post*, April 21 2017, <https://www.washingtonpost.com/posteverything/wp/2017/04/21/local-police-dont-go-after-most-cybercriminals-we-need-better-training> (accessed Jan. 10, 2019).

²² Mulligan and Schneider, "Doctrine for Cybersecurity."

²³ Andy Greenberg, "Russian Hacker False Flags Work—Even After They're Exposed," *Wired*, February 27, 2018, <https://www.wired.com/story/russia-false-flag-hacks/> (accessed March 22, 2019).

²⁴ Nick Selby, "Local Police Don't Go After Most Cybercriminals. We Need Better Training."

²⁵ Mulligan and Schneider, "Doctrine for Cybersecurity," 75.

²⁶ Selby, "Local Police Don't Go After Most Cybercriminals. We Need Better Training."

21,000 cybercrime incidents in the state in 2017 alone.²⁷ Additionally, in many cases, the most sophisticated cyberattacks are carried out by states themselves, which cannot be effectively punished through traditional legal approaches. As a result, while the law enforcement and criminal justice approaches to information security are appropriate given the very real and substantial criminal activity in cyberspace, traditional law enforcement strategies are often ineffective in the digital domain even at stopping straightforward criminal behaviour.

I do not mean to imply that traditional law enforcement should play *no* role in cyberhealth. Rather, it needs to be downplayed compared to other more effective preventative strategies, even if we are concerned solely with reducing cybercrime. If anything, my appreciation of the role of law enforcement in cyberspace has only grown since I have been working on this dissertation. An example of this value can be seen in the 2017 joint effort by the FBI and Dutch National Police to shutdown two of the largest dark web marketplaces, AlphaBay and Hansa. On these marketplaces hacking tools were sold alongside, drugs, weapons, and other black-market goods. Before the FBI shutdown AlphaBay, Dutch National Police took control of Hansa, but allowed it to continue to operate for a period of time. This allowed them to monitor the illegal activity on the site for several weeks and capture the activities of all the new users fleeing from AlphaBay.²⁸ However, this type of law enforcement action is very unusual, and traditional law enforcement techniques are generally of limited use for *preventing* cybercrimes and large-scale malware outbreaks.

The final aspect of the Cybersecurity LoA that undermines operational coherence and internal utility is the state-sanctioned development and use of offensive cyber capabilities, which in the long run undermine defensive efforts. Unlike in the Public Cyberhealth LoA I will discuss next, within the Cybersecurity LoA there is often little distinction between defensive and offensive capabilities. One of clearest

²⁷ FBI's Internet Crime Complaint Center, "2017 Internet Crime Report," FBI's Internet Crime Complaint Center, May 7, 2017, https://pdf.ic3.gov/2017_IC3Report.pdf (accessed March 12, 2019).

²⁸ Samuel Gibbs and Lois Beckett, "Dark Web Marketplaces AlphaBay and Hansa Shut Down," *The Guardian*, 20 July 2017, <https://www.theguardian.com/technology/2017/jul/20/dark-web-marketplaces-alphabay-hansa-shut-down> (accessed Nov. 26, 2018).

cases where short-term security benefits were valued over long-term cyberhealth, was the United States' and Israel's development of the Stuxnet worm used to destroy Iranian nuclear centrifuges.²⁹ Stuxnet was a sophisticated piece of software targeting the centrifuges' SCADA industrial control systems—a cyber sniper shot that would have required many months or years of planning and code development.³⁰ Once the worm broke into the wild, overnight the number of people who could develop such a weapon grew exponentially. Today, most countries possess some offensive cyber capabilities,³¹ yet such a development was not inevitable as evidenced by the international community's collective efforts to place significant limits on other classes of weapon (e.g. biological and chemical weapons). The use of state-sanctioned offensive cyberweapons has made global collaboration more difficult at a time when it is needed more than ever to overcome the challenges related to identification and extradition discussed above.

In addition to a lack of operational coherence, the internal utility of Cybersecurity LoA is further diminished by the LoA's downplaying of externalities and non-financial harms. Not accounting for certain harms is not incoherent, as LoAs are rarely meant to be entirely comprehensive, but ignoring these harms does lead to states, corporations, and individuals to undervalue cybersecurity investments, relative to the level that would be best for society as a whole. Even when these externalities are acknowledged, they often do not factor into an individual's, corporation's, or state's cost benefit analyses, as each is responsible for their own security; within the Cybersecurity LoA you truly are not your brother's keeper. When one only looks at one's own economic costs, often the most sensible choice from a financial perspective is to either insure against losses or simply hope that one does not suffer a devastating attack. By failing to take externalities into account, states, corporations, and individuals collectively underinvest in cybersecurity compared to the socially optimal level.

²⁹ Brian Orend, "Fog in the Fifth Dimension: The Ethics of Cyber-War," in *The Ethics of Information Warfare*, edited by Luciano Floridi and Mariarosaria Taddeo (Cham: Springer, 2014), 6-7.

³⁰ Singer and Friedman, *Cybersecurity and Cyberwar*, 98

³¹ Ibid.

In addition to downplaying externalities, those using the Cybersecurity LoA also typically fail to account for non-financial harms, such as the impact of cyberattacks and cybersecurity interventions on health and wellbeing. This undercounting of costs can exacerbate the underinvestment in cybersecurity. If the product being protected is a medical device, some non-financial impacts may be considered. However, as we will see in Chapter 3, in many cases medical devices have some of the weakest cybersecurity of networked devices. 41 per cent of the machines infected with Conficker in 2017—nearly ten years after the initial infection—were in the healthcare sector.³²

Note that throughout this dissertation when I speak of ‘financial harms’ or ‘financial costs,’ I am referring to relatively direct monetary costs associated with cyberattacks or interventions. These include costs related to IP theft, employee salaries, software and hardware expenses, monetary theft, and destruction of property. While one can account for health and wellbeing impacts in financial terms, within the Cybersecurity LoA this is typically not done. As such, while recognizing that for policy purposes one might place dollar values on health and wellbeing impacts, I will continue to refer to health and wellbeing impacts as non-financial harms.

While the Cybersecurity LoA could be modified to account for non-financial harms, in this chapter I am seeking to define the Cybersecurity LoA as it is used in practice. I suspect that the reason non-financial impacts have largely been overlooked is that there is a shortage of tools for measuring how various cyber threats or interventions impact health or wellbeing. In Chapter 4, I will attempt to remedy one of those gaps by defining a theory of informational wellbeing that can be used to measure how personal wellbeing is impacted by digital information and its use, control, accuracy, and accessibility.

2.2.2.2 External Utility of the Cybersecurity LoA

Whereas internal utility considers the LoA’s ability to fulfil its goal, external utility considers whether that goal is useful in the context of a society’s broader needs.³³ In this case, the Cybersecurity LoA is also somewhat lacking. This could be an

³² O’Neill, “Conficker Worm Still Spreading Despite Being Nearly 10 Years Old.”

³³ Admittedly, this distinction can sometimes be hazy. The previous discussion related to unaccounted harms could also be discussed as a matter of external utility.

expansive discussion, but I would prefer to not disappear down that rabbit hole. For my purposes, I will simply accept that one goal of most societies or states in the 21st century is to effectively keep digital information secure and digital networks up and running. While the Cybersecurity LoA can be useful for addressing certain adversary-based threats, it is less useful for addressing the threats associated with natural disasters, human error, bugs, poor product design, and bad technology policy, all of which harm network robustness and resiliency. Recent examples include a glitch in a Federal Aviation Association computer which grounded half the planes in the US in 2011,³⁴ the damage to Puerto Rico's networks following hurricane's Irma and Maria in 2017,³⁵ and the fat finger mistake which brought down Amazon's S3 cloud system (and with it a number of the Web's most popular sites).³⁶

The Cybersecurity LoA is also of limited use for conceptualizing technology matters unrelated to network failure, such as the value of network access and the ethics of technology use. In regards to network access, while the Cybersecurity LoA can help one understand how adding a node to a network makes it less secure, it cannot help one to understand the physical, psychological, economic, and social cost of not having access to digital networks.³⁷ I will look at this issue in greater depth in Chapters 3 and 4 as part of my elaboration of the Public Cyberhealth LoA. In regard to the second point, let us consider a question like whether there is a moral obligation to stop using Facebook. Matthew Liao argues that if Facebook is leading to the destruction of certain democratic norms (e.g. spreading fake news) or harms wellbeing, one may have a responsibility to leave the service. He argues that even if one does not actively engage in spreading fake news or other destructive behaviours, by contributing to Facebook's analytics and bolstering its user count one may be an

³⁴ Singer and Friedman, *Cybersecurity and Cyberwar*, 35.

³⁵ Thieme, "After Hurricane Maria, Puerto Rico's Internet Problems Go from Bad to Worse."

³⁶ Amazon.com. "Summary of the Amazon S3 Service Disruption."

³⁷ In the previous section I discussed how failing to account for these types of harm undermines the internal utility by leading to an underinvestment in cybersecurity. Here I am making the point that by failing to include these types of variables, the LoA cannot be used to think about technology issues beyond security.

accessory to harmful behaviour.³⁸ With its focus on cyberattacks and information security, the Cybersecurity LoA is not useful for thinking through this type of dilemma as the relevant variables—such as impacts on wellbeing and rights—are downplayed or entirely absent from the LoA. By contrast, the public health inspired LoA with its greater focus on health, wellbeing, and individual rights is better suited for conceptualizing our relationship to digital information and devices on the whole.

In this section, I described the Cybersecurity LoA and its main features, including a focus on the adversary, a focus on strategic and financial costs, and the use of offensive cyber capabilities in certain circumstances. I then assessed the utility and coherence of the LoA. I argued that its lack of operational coherence diminished its internal utility (i.e. its ability to protect information and networks from malicious attacks), while its narrow focus limited its value outside of the context of adversary-based attacks. In the next section, I will outline the Public Cyberhealth LoA and then illustrate a number of its advantages over the cybersecurity alternative.

2.3 The Public Cyberhealth LoA

While the goal of the Cybersecurity LoA is:

1) To stop adversaries from gaining unauthorized access to digital information, networks, or devices and 2) To bring those who commit these illegal acts to justice

the goal of the Public Cyberhealth LoA is:

To promote cyberhealth (i.e. network robustness and resilience) as part of broader societal efforts to promote health and wellbeing.

Whereas the Cybersecurity LoA arose from the domains of military intelligence and criminal justice, the Public Cyberhealth LoA makes use of the vocabulary, philosophy, and tools of public health. As such, the Public Cyberhealth LoA downplays the importance of the adversary and business interests, while taking seriously non-malicious points of failure (bugs, human accidents, natural disasters,

³⁸ Matthew Liao, “Do You Have a Moral Duty to Leave Facebook?,” *The New York Times*, November 24, 2018,

<https://www.nytimes.com/2018/11/24/opinion/sunday/facebook-immoral.html> (accessed Jan. 10, 2019).

bad product design) and the impact of poor cyberhealth on health and wellbeing. In Chapter 1, I explored how this LoA can be used to address traditional cybersecurity issues like malware. In this section, I will demonstrate how this way of thinking can be used to frame the digital landscape more broadly.

Unlike the Cybersecurity LoA, which was a formalization of the dominant way of conceptualizing the digital landscape, the Public Cyberhealth LoA is an original contribution of this thesis, albeit one inspired both by the field of public health and researchers like Deirdre Mulligan, Fred Schneider, and Elaine Sedenberg who have suggested the value of a public health approach to technology policy. This is the first attempt to formally define such an approach, and as such my intention is not to be exhaustive, but to capture the most important observables and behaviours. Bolded line items in the table below are those which appear in the Public Cyberhealth LoA but not in the Cybersecurity LoA.

Table 2

Simplified Public Cyberhealth LoA

Variables	Valid Values	Invalid or Minimized Values
Type of Threat	<ul style="list-style-type: none"> • Human error • Fragile infrastructure • Natural disasters • Buggy code • Adversary based attacks 	<ul style="list-style-type: none"> • Adversary based attacks lose relative importance <i>only</i> in the sense that other threats are now made more prominent than in the Cybersecurity LoA
Cost of Threat	<ul style="list-style-type: none"> • Impact on wellbeing³⁹ • Impact on health⁴⁰ • Impact on other network nodes (externalities) • Impact on individuals' rights • Financial value of information stolen • Strategic value of information stolen • Reputational harm • Financial cost of Destruction of Infrastructure 	<ul style="list-style-type: none"> • Financial costs and reputational harm lose relative importance due to addition of health, wellbeing, and rights considerations.
Stakeholders	<ul style="list-style-type: none"> • Individuals • Communities • States • Corporations • Non-Profits 	

³⁹ To be discussed in Chapter 4: Informational Wellbeing.

⁴⁰ To be discussed in Chapter 3: Health and Cyberhealth.

Defensive Capabilities	<ul style="list-style-type: none"> • Public network monitoring • Public education campaigns • Infrastructure improvements • Herd immunity • Private network monitoring • Firewalls • Employee/Personal education 	<ul style="list-style-type: none"> • Public defensive capabilities supplement private capabilities
Cost of Defensive Efforts	<ul style="list-style-type: none"> • Wellbeing costs • Health costs • Impact on other nodes on the network (negative externalities) • Impact on individuals' rights • Financial cost of preventative strategies • Employee/personal time 	
Offensive Capabilities	<ul style="list-style-type: none"> • Very few, if any, offensive responses are acceptable 	<ul style="list-style-type: none"> • Hacking back • Infiltrating adversary networks • Pre-emptive cyberattacks
Cost of Offensive Efforts	<ul style="list-style-type: none"> • Militarization of cyberspace • Provocation of additional attacks • Wellbeing costs • Health costs • Impact on other nodes on the network (negative externalities) • Impact on individuals' rights 	<ul style="list-style-type: none"> • Financial • Employee time
Ethical Considerations	<ul style="list-style-type: none"> • Protection of individuals' rights • Local values • Obligations to others on the network • Proportionality 	
Insurance	<ul style="list-style-type: none"> • Insurance for health impacts • Insurance for financial loss 	
Possible Role of State (specifics depend on circumstances)	<ul style="list-style-type: none"> • Production of public goods (information sharing, network monitoring, production of basic research, public education) • Crisis coordination • Protection of individual rights • Law enforcement (prosecution, extradition, punishment) 	<ul style="list-style-type: none"> • Liability protection • Law enforcement actions lose relative importance as preventative strategies gain in importance

Key Behaviours:

1. Impacts on health and wellbeing are given greater weight compared to business or military interests
2. Use of the philosophy of public health, where applicable, to understand proportionality, engagement with local stakeholders, and the normative justification for government interventions

At a high level, the Public Cyberhealth LoA attempts to: 1) address a host of threats to network robustness and resiliency (e.g. accidents, buggy code, cyberattack, natural disasters, etc.), 2) take seriously the effect on human health and wellbeing of network failure, interventions, and technology policy, 3) downplay or even delegitimize the use of offensive cyber capabilities, and 4) explicitly consider the impact of both network failure and interventions on individual rights. It meets these goals by considering a broader array of variables and observables and using the philosophy of public health to think through the normative justification for—and ethical limits on—intervening in cyberspace. While these changes are significant, it is important to note that many typed variables exist in both LoAs (e.g. Type of Threat, Defensive Capabilities). This similarity allows one to switch back and forth between the two LoAs (to some degree) for the purpose of debate or analysis. For example, in Chapter 1, I initially described the Conficker infection using language more at home in the Cybersecurity LoA, and then I critiqued the response using the Public Cyberhealth LoA.

In fact, much of the Cybersecurity LoA exists (in an altered form) within the Public Cyberhealth LoA. Just as public health experts must have a plan for dealing with biological weapons, the public cyberhealth expert must be prepared to deal with cyberattacks. This being said, the strategies one would devise if using the Public Cyberhealth LoA would likely look quite different from the strategies devised by a General or CTO using the Cybersecurity LoA. As Sedenberg and Mulligan argue, “adversarial considerations are simply less relevant when dealing with prevention and management orientations—in contrast to deterrence oriented strategies that are focused on intent—because harms manifest, and protections work, regardless of intent.”⁴¹ Even in cases where traditional tools of cybersecurity and criminal justice are the best tools for the job, the Public Cyberhealth LoA forces one to explicitly consider the impact on individual rights, health, and wellbeing; engage with relevant stakeholders; and think about how one’s actions or inactions will affect others on the network. In taking these variables into account, certain strategies such as security backdoors to assist law enforcement and many offensive capabilities will be

⁴¹ Sedenberg and Mulligan, “Public Health as a Model for Cybersecurity Information Sharing,” 1705.

delegitimized, while other strategies, such as automatic patching, stronger security standards, and public education gain prominence.⁴²

While it is hard to assess the utility and coherence of the Public Cyberhealth LoA until it is used, one can make a few initial assessments. First, it does not possess the same incoherencies as the Cybersecurity LoA explored in Section 2.2. In particular, 1) it does not try to solve shared problems with private solutions, but recognizes a potentially broader role for states to provide public goods for cyberhealth, and 2) by downplaying the use of offensive cyber weapons, it discourages short term strategies that harm information security and human wellbeing on the whole and in the long run. Additionally, by highlighting impacts to health, wellbeing, and individual rights, it helps one to more fully account for costs and benefits while determining the appropriate level of cyberhealth investment. In Chapter 3, I will explore some of the health impacts of poor cyberhealth, and in Chapter 4, I will explore how to measure the impact of technology policies, interventions, and network threats on personal wellbeing using a version of the capabilities approach to wellbeing.

Second, in theory at least, the Public Cyberhealth LoA has greater external utility compared to the Cybersecurity LoA. While the latter only focused on adversary-based attacks, the Public Cyberhealth LoA is also useful for thinking about non-malicious threats, including natural disasters, human error, and buggy code. However, perhaps more importantly, the Public Cyberhealth LoA is useful for guiding technology policy beyond how to respond to network failure. After all, public health experts not only respond to acute epidemics, but conduct or support scientific research, collect and share health data, and address harmful social behaviours like smoking, overeating, and a lack of exercise. Similarly, the Public Cyberhealth LoA can be used to guide technology policy on a wide range of issues. Elaine Sedenberg and Deirdre Mulligan have demonstrated one such application of the public health inspired approach—the collection and sharing of cybersecurity information. Their

⁴² Offensive cyber capabilities represent a broad spectrum of tools. While some forms of beaconing may be acceptable within public cyberhealth for attribution purposes, retaliatory attacks would likely not be acceptable, as in the long run this tit for tat strategy of escalating attacks reduces the security of the network as a whole.

work is a specific application of the kind of approach I seek to formalize and generalize in this dissertation.

While Sedenberg and Mulligan do not explicitly use the concept of the ‘Public Cyberhealth LoA,’ they essentially design an information sharing scheme using the variables, observables, and behaviours listed in the table above. Specifically, they emphasize the protection of individual rights, the potential impacts on wellbeing, and the variety of stakeholders who may be impacted by various approaches.⁴³ By studying the information sharing systems used by public health institutions, policymakers, and researchers, they derived four principles which can guide the development of cybersecurity information sharing systems: “expert and collaborative data governance, reporting minimization and decentralization, earliest feasible de-identification, and limitations on use.”⁴⁴ In addition to these principles, they recommend that cybersecurity information should be made available for public use and that cybersecurity information sharing practices should emphasize ethical research.⁴⁵ These principles contrast with how cybersecurity information is currently shared. In the United States, which is the focus of their work, there are few restrictions on what types of cybersecurity information should be shared (including personally identifiable information), few restrictions on how shared information is used, and groups representing users and privacy advocates are often in a reactive role, rather than part of the governance process.⁴⁶

While Sedenberg and Mulligan are seeking to improve cybersecurity, the Public Cyberhealth LoA, as a generalization of their approach, can be used in a variety of contexts, including product design, data protection regulations, and infrastructure robustness. In Chapter 1, I explored one of these uses—thinking through the provisioning of public goods for cyberhealth. In Section 2.4, I will further explore how it can help us think about the normative justification for government interventions in cyberspace and the ethical limits on those interventions. And in

⁴³ Sedenberg and Mulligan, “Public Health as a Model for Cybersecurity Information Sharing.”

⁴⁴ *Ibid.*, 1692.

⁴⁵ *Ibid.*, 1730-1736.

⁴⁶ *Ibid.*

Chapters 3 and 4, I will use the Public Cyberhealth LoA to explore how cyberhealth impacts health and personal wellbeing.

One may be tempted to argue that as the two LoAs do not have precisely the same goal, they should simply be deployed in different contexts; the Cybersecurity LoA can be used for adversary-based threats, while the Public Cyberhealth LoA can speak to other aspects of technology policy. While I believe these two frames can exist side by side in some form, it would be a mistake to think of them as non-overlapping magisteria. In particular, many of the tools and strategies of cybersecurity undermine the strategies and goals of cyberhealth. For instance, the development of offensive cyber capabilities by those working in cybersecurity has contributed to a cyber arms race around the world, with many of the tools eventually falling into the hands of cyber criminals. Similarly, as stated before, treating other state actors as adversaries harms the collaboration needed for international monitoring programs and regulatory regimes. For these reasons, adopting the Public Cyberhealth LoA may also necessitate fundamental changes to the Cybersecurity LoA.

2.4 Two Applications of Public Cyberhealth

The exercise of formally defining LoAs is useful for clearly defining a framework's purpose, outlining the assumptions of one's framework, and for comparing different frameworks. Using this exercise, I highlighted the Cybersecurity LoA's lack of operational coherence and limited utility as an overarching guide for understanding the digital landscape. I then argued that the Public Cyberhealth LoA avoids some of the incoherencies of the Cybersecurity LoA and is more broadly useful as a framework for guiding technology policy.

In this section I will further explore the “cash-value” of such an approach—to use William James' term—by exploring two applications of the Public Cyberhealth LoA beyond the production of public goods discussed in Chapter 1 and the information sharing scheme developed by Sedenberg and Mulligan described in Section 2.3. Specifically, I will explore how using the Public Cyberhealth LoA to think about the digital landscape can help one to 1) distinguish which problems deserve public interventions and which are best left to the private sector, and 2) ensure that policies and interventions do not unnecessarily infringe upon individuals' rights.

2.4.1 Spheres of Public and Private Cyberhealth

As public health policies typically are enacted by governments⁴⁷ and exist within legal systems, they must be justifiable within the context of a ‘defensible political theory.’⁴⁸ The previous chapter articulated one important justifiable role of the state: to provide public goods necessary for sustaining a minimally decent life. In this section, we turn to consideration of the proper role and limits of states in other areas of health policy. By combining these various responsibilities and limits one can define a sphere of *public cyberhealth*, within which there is a normative justification for governments to promote network robustness and resiliency, and a sphere of *private cyberhealth*, which is best left to private individuals and the private sector to manage. First, I will look at how these spheres are defined in public health, and then I will apply this same way of thinking to the digital landscape.

2.4.1.1 Shared and Overlapping Problems

In much of the Global North the political system which bounds the proper use of governmental action is some form of liberal democracy, which, depending on the country, to a greater or lesser extent follows Millian notions of limited government. Within Millian liberalism, problems which justify public health interventions are generally those where one person’s health status can adversely affect the health of another.⁴⁹ These problems can be called ‘shared’ problems.⁵⁰ An archetypical shared

⁴⁷ While non-state actors like the Gates Foundation can perform some public health services, ultimately only governments can justifiably coerce populations, physically if needed.

⁴⁸ John Coggon. *What Makes Health Public?: A Critical Evaluation of Moral, Legal, and Political Claims in Public Health* (Cambridge: Cambridge University Press, 2012), 266.

⁴⁹ Coggon, *What Makes Health Public*, 25.

⁵⁰ Note I define ‘shared’ problems more narrowly than Jennings, as I am focused on the liberal context, while he uses it in the context of civic republicanism. Bruce Jennings, “Public Health and Civic Republicanism,” in *Ethics, Prevention, and Public*

problem is infectious disease—if I come in contact with someone with the flu, I am more likely to get the flu. In cases where the harm one individual poses to others is direct and substantial (e.g. Ebola), a narrow application of Mill’s harm principle—which essentially states that a state can only limit a person’s freedom of action to prevent harm to others⁵¹—can justify intrusive governmental action like quarantine. In cases where the threat is less extreme, like chicken pox, there is less justification for an intrusive governmental response. In these cases, governments may still address the problem, but only through less intrusive means, such as public education.

Shared problems can be contrasted with ‘overlapping’ problems. Overlapping health problems are those which we all might be concerned about, but where your health status does not influence my health status. For example, we all might care about weight management, but generally your weight will not impact my weight. While the default in liberal democracies is to leave overlapping problems to individuals and private markets, there are two types of overlapping health problems which are normally thought to justify government intervention—overlapping health problems which generate substantial negative externalities and external threats which harm or could harm a large number of people (e.g. natural disaster). An example of the first category might be widespread obesity. While weight management is an overlapping problem, widespread obesity can strain health systems, hurt the economy, and normalize unhealthy behaviours like eating fast food or drinking soda, all of which can indirectly harm others.⁵² In these cases, interventions can be justified under a softer version of Mill’s harm principle. However, as the risk posed to the general

Health, eds. Angus Dawson and Marcel Verweij (Oxford: New York: Clarendon Press; Oxford University Press, 2007).

⁵¹ John Stuart Mill, *On Liberty*, in *Utilitarianism and Other Writings*, 1859, ed. Mary Warnock (Glasgow: Collins, 2003), 94-95.

⁵² Youfa Wang, May A. Beydoun, Lan Liang, Benjamin Caballero, Shiriki K. Kumanyika, “Will All Americans Become Overweight or Obese? Estimating the Progression and Cost of the US Obesity Epidemic,” *Obesity* 16, no. 10 (2008): 2323-2330. Note that the distinction between shared and overlapping problems is often not hard-and-fast (as can be seen in this example of obesity). Problems will often lie along a spectrum between these two poles. Having said this, I believe the concepts are a useful heuristic for thinking about whether a state’s interventions are justifiable.

population by obese individuals is indirect, non-urgent, and relatively minor, only relatively unobtrusive interventions (e.g. nudging, public education) are justifiable.⁵³ Examples of the second category, meanwhile, include threats like natural disasters and environmental hazards. As discussed in Chapter 1.3, states are generally justified in addressing these threats as part of their responsibility to ensure individuals' ability to live a minimally decent life.⁵⁴

2.4.1.2 Application to Digital Landscape

Turning to the digital landscape, many of these same justifications can be used to define spheres of public and private cyberhealth problems, which can serve as the basis for consistent and justifiable technology policies. One of the simplest examples of a shared network problem is a computer worm like Conficker, which I discussed in Chapter 1. Conficker spread from computer to computer and could pass commands (in its later iterations) via peer to peer connections. This infection model looks very similar to communicable diseases and would be a good candidate to be framed as a public problem which justifies some liberty encroaching measures (e.g. requiring one to patch one's devices) even within a stronger version of Millian liberalism. Not only is one's own device at a greater risk of infection if it is closely connected to infected devices, but large numbers of infected devices could be wielded in a botnet that could endanger critical infrastructure. As a result, even unconnected individuals could be harmed. Note that while these issues are related to the 'harm to others principle' discussed in 1.3.2, they are not identical. The 'harm to others principle' is relevant to one's individual obligations, while Mill's harm principle governs the proper use of state power.

While computer worms may be a shared problem, the short lifespan of laptop batteries looks like an overlapping concern. While everyone with a laptop might have the concern, my battery's lifespan will not affect your battery's lifespan. However, like obesity, if there was such an 'epidemic' of dying batteries that there were broad economic consequences, then perhaps liberty-encroaching interventions could be

⁵³ I will explore the concept of proportionality more in the following subsection when I discuss ethical limits on interventions.

⁵⁴ As mentioned in 1.3, the specific responsibilities of the state will depend on a state's political system and the specifics of the threat.

justifiable under a more expansive, softer version of Millian liberalism, or as a way to correct a particularly pernicious market failure as discussed in Chapter 1.

While the above example of an epidemic of dying batteries may seem a bit ridiculous, the example of obesity (a paradigmatic overlapping problem) being treated as a public health problem suggests that in highly connected networks very few ailments which affect a large segment of the population will not lead to some form of harm for the broader population. One case where this is relatively easy to see is in the case of national health services. As healthcare funding is coming from a collective pool all citizens pay into, the unhealthy life choices of one individual does in some minor way negatively impact all other taxpayers. In the context of the internet, which is defined by its interconnectedness, something similar seems to happen. While prior to the internet the security of one's camera or thermostat was an overlapping problem, today the insecurity of these devices can lead to attacks like the DDoS attack on the DNS system in 2016 that negatively affected (albeit minimally) hundreds of millions of people.

It might seem a bit weird to think of the cybersecurity of one's thermostat as being a matter of public concern, but as more devices of a previously discrete nature become network connected, problems that were once overlapping in nature become shared. A parallel might be to think about a person on a desert island. If they have measles, it is not a public health issue as they are not connected to any other people. But if you drop that person in London, that person's health status becomes a shared problem. By including variables for a broader array of harms and by placing particular emphasis on externalities, the Public Cyberhealth LoA is better suited to capturing how nodes affect one another compared to the Cybersecurity LoA, which tends to downplay externalities and the impacts of network threats on health and wellbeing.

If we combine this analysis with the discussion of public goods from Chapter 1, one can define a sphere of public cyberhealth problems, which governments are justified in addressing, and a sphere of private cyberhealth problems, which are not within the proper purview of state action.

Table 3

Public Cyberhealth	Private Cyberhealth
Shared Problems (e.g. infectious malware)	
Overlapping Problems with Substantial Negative Externalities (e.g. widespread internet outages)	Overlapping Problems with Minimal Externalities (e.g. broken computer, targeted hacking)
Production of Public Goods (e.g. basic research, standards, network monitoring, national defence)	Production of Private Goods (e.g. buying a better router, a more secure computer, etc.)

While I have only sparingly used the language of the method of Levels of Abstraction in this section, the above framework is firmly grounded in the Public Cyberhealth LoA, which takes into account a variety of threats, a variety of harms, and emphasizes the externalities which arise in connected systems. Someone using the Cybersecurity LoA would find it difficult to conceptualize many of these types of problems given that LoA's focus on adversary-based threats and limited concept of harm. When one focuses on financial and strategic harms and downplays externalities, most problems will seem "private" in nature and hence outside the proper scope of state action.

2.4.2 Ethical Limits on Interventions

In addition to helping one to understand the normative justification for government intervention in cyberspace, the Public Cyberhealth LoA also helps one think about ethical conflicts which may arise in the course of these interventions by 1) helping one to identify potential ethical conflicts and 2) suggesting solutions to those conflicts from the field of public health ethics.

2.4.2.1 Identifying Conflicts

Those using the Cybersecurity LoA are often blind to potential ethical conflicts because their conceptual framework simply does not include (or at least minimizes) variables like stakeholders, the impact on individual rights, externalities, and the impact of interventions and policies on health and wellbeing. By highlighting these variables, those using the Public Cyberhealth LoA are capable of identifying ethical conflicts which exist but have previously gone unrecognized or underappreciated.

One simple example is the case of updating the software that runs robotic prosthetics. While someone using the Cybersecurity LoA may treat the device like any other (albeit one of greater importance than an Xbox), one using the Public Cyberhealth LoA will be more likely to recognize the individual patient as an important stakeholder and recognize that both device insecurity and device updates introduce interesting questions about bodily integrity, device ownership, and consent. As I will explore in more depth in Chapter 3, empirical research suggests that such devices can be incorporated into one's sense of one's body.⁵⁵ In these cases, is it ethical to simply stop supporting a product that is a part of someone's body if it becomes unprofitable? Is it ethical to develop new models with more advanced features, if that means risking the stability of the old models? Is it ethical to automatically push an update that fixes a security vulnerability, but also changes the prosthetic's functionality? Should one share information generated by this device, and, if so, with whom and for what purpose?⁵⁶

As some of the specific questions raised above regard technologies that are still on the horizon, they have not been fully considered in public health literature; however, publications like the WHO and USAID's *Standards for Prosthetics and Orthotics* demonstrate the kind of thoughtfulness that is needed to ethically develop technology products. This set of standards defines a comprehensive list of stakeholders, provides guidance for working with local populations, and outlines users' rights.⁵⁷

In next section, I will explore how two specific public health tools can help one think about the ethics of digital interventions. First, I will consider The Intervention Ladder, a tool developed by the Nuffield Council on Bioethics to think

⁵⁵ Abbe Brown, Shawn H. E. Harmon, Rory O'Connor, Sita Popat and Sarah Whatley, "Body Extension And The Law: Medical Devices, Intellectual Property, Prosthetics And Marginalisation (Again)," *Law, Innovation and Technology* 10, no. 2 (2018): <https://doi-org.ezp.lib.cam.ac.uk/10.1080/17579961.2018.1526853>.

⁵⁶ While I will not be definitively answering the ethical questions raised above, I will return to some of these themes in Chapter 3, as I explore how cyberhealth impacts health.

⁵⁷ WHO and USAID, *WHO Standards for Prosthetics and Orthotics*, World Health Organization (Geneva: World Health Organization, 2017).

about proportionality. Then, I will turn to ethical review boards and their role in public health institutions.

2.4.2.2 Public Health Strategies for Managing Ethical Conflicts

As public health interventions may need to infringe upon individuals' rights like personal freedom and privacy,⁵⁸ public health philosophers and policymakers have developed a number of strategies to weigh the effectiveness of a specific intervention, the severity of the problem, and the impact on individuals' rights.⁵⁹ One way to stop the spread of a disease would be to lock everyone in their homes, but such a solution would be, in all but the most extreme scenarios, ethically unacceptable.

While there is not a one to one comparison between the ethical conflicts within the two fields—quarantining a machine is not the same as quarantining a person—digital interventions which benefit the population may still infringe upon personal rights in a similar manner. In the last section, I mentioned how updating health devices can raise questions related to bodily integrity, and digital quarantining does restrict one's freedom of association and potentially one's freedom of movement—especially as virtual reality becomes more commonplace. For those with relatively few avenues to connect to the internet, this restriction could cause significant social and economic harm. In this section, I will outline how two tools from public health can help one balance effective solutions with the protection of rights in the digital context—the Intervention Ladder and ethical review boards.

⁵⁸ For example, quarantine restricts freedom of movement and association, mandatory vaccinations may violate someone's bodily integrity or restrict their freedom of choice, and surveillance programs may share sensitive information about a person's health history, including their sexual partners or drug use.

⁵⁹ Note, other fields also have a rich history of thinking about proportionality. For example, Just War theory can help one understand proportionality in the context of war. However, outside of that context, Just War theory is ill-suited to thinking about proportional response to cyber attacks. Thomas W. Simpson, "The Wrong in Cyberattacks," in *The Ethics of Information Warfare*, eds. Luciano Floridi and Mariarosaria Taddeo (Heidelberg: Springer, 2014), 144.

2.4.2.2.1 *The Intervention Ladder*

The first strategy is the Nuffield Council on Bioethics' Intervention Ladder which is designed to help one to think through "acceptability and justifiability" of various public health policies, and has been used to create public health policy in a wide variety of contexts including food labelling standards and transportation.⁶⁰ While the Intervention Ladder is just one public health policy tool, and not universally used or accepted, it formalizes a general process of thinking about proportionality that is a hallmark of public health policy and philosophy. The way the ladder works is by providing a spectrum of actions from the least intrusive to the most. The higher up the ladder, "the stronger the need for justification and sound evidence for implementation."⁶¹ From the Nuffield Council on Bioethics' report on public health, the steps on the Intervention Ladder are:

- 1) Do nothing or simply monitor the current situation.
- 2) Provide information. Inform and educate the public, for example as part of campaigns to encourage people to walk more or eat five portions of fruit and vegetables per day.
- 3) Enable choice. Enable individuals to change their behaviours, for example by offering participation in a NHS 'stop smoking' programme, building cycle lanes, or providing free fruit in schools.
- 4) Guide choices through changing the default policy. For example, in a restaurant, instead of providing chips as a standard side dish (with healthier options available), menus could be changed to provide a more healthy option as standard (with chips as an option available).
- 5) Guide choices through incentives. Regulations can be offered that guide choices by fiscal and other incentives, for example offering tax-breaks for the purchase of bicycles that are used as a means of travelling to work.
- 6) Guide choice through disincentives. Fiscal and other disincentives can be put in place to influence people not to pursue certain activities, for

⁶⁰ "Intervention Ladder Informs Lords Behaviour Change Report," Nuffield Council on Bioethics, July 19, 2011, <http://nuffieldbioethics.org/news/2011/intervention-ladder-informs-lords-behaviour-change-report> (accessed Dec. 7, 2018).

⁶¹ John Krebs, "The Importance of Public Health Ethics," *Bulletin of the World Health Organization* 86, no. 8 (2008): 577-656.

example through taxes on cigarettes, or by discouraging the use of cars in inner cities through charging schemes or limitations of parking spaces.

- 7) Restrict choice. Regulate in such a way as to restrict the options available to people with the aim of protecting them, for example removing unhealthy ingredients from foods, or unhealthy foods from shops or restaurants.
- 8) Eliminate choice. Regulate in such a way as to entirely eliminate choice, for example through compulsory isolation of patients with infectious diseases.⁶²

The Intervention Ladder does not suggest any specific solutions for a particular problem, but rather lays out a set of options to facilitate thinking about what constitutes a proportional response. By assessing a suite of options rather than simply accepting the first that addresses the problem, one is more likely to find a politically and ethically acceptable balance between personal rights and population health. Additionally, it is a useful reminder that one should use the least serious intervention, even if a stronger reaction might be ethically justified. For example, hypothetically let us assume a tax on soda is justifiable given the health impacts of soda consumption. Ignoring the financial reasons a state may have to impose such a tax, if incentivizing someone to choose healthier options is equally effective, then a state should use that option first.

While there is not space in this chapter for a full case study, I would like to sketch out how one might apply the Intervention Ladder to a cyberhealth issue by considering the problem of unprotected personal computers (PCs). In 2014, PCs with no anti-malware software (20% of PCs worldwide) were six times more likely to be infected than machines that ran up-to-date monitoring.⁶³ Not only will individuals with infected machines be at greater risk of identity theft and other personal harms, but unprotected devices can be easily drafted into large botnets that can be used to endanger critical infrastructure, as was the case with the Conficker worm. Individuals may not update their devices because they fear the update may damage their

⁶² Nuffield Council on Bioethics, *Public Health: Ethical Issues*, XIX.

⁶³ Dennis Batchelder, et al., *Microsoft Security Intelligence Report*, volume 18. (2015), <http://www.microsoft.com/security/sir/archive/default.aspx>, 79-80.

machines, they may be unaware of the seriousness of being unprotected, or they may be unable to pay for anti-malware software or the bandwidth to download patches.⁶⁴

Using the Intervention Ladder as a guide, depending on the level of risk assessed, the following interventions could be applied in increasing level of intrusiveness:

- 1) Do nothing.
- 2) Educate the public about cybersecurity without infringing on their autonomy.
- 3) Provide all new PC owners with the option to enable anti-malware software.
- 4) Nudge individuals towards protection by changing the default to opt-in new PC owners to anti-malware software and automatic updates.
- 5) Subsidize anti-malware software and create a fund analogous to the Vaccine Injury Compensation Program to pay for damage caused by patches.
- 6) Charge unprotected users more for internet access.
- 7) Internet Service Providers (ISPs) could restrict unprotected PCs to only allow access to verified safe websites.
- 8) ISPs could fully restrict unprotected machines until they install anti-malware software.

As with the example of the soda tax, one should always run through the various options and typically deploy the least intrusive response that will get the job done, even if a more intrusive option may also be justifiable.

A cyberhealth Intervention Ladder would likely need different steps than the one developed by the Nuffield Council. For example, in the digital context, charging unprotected individuals more for internet access may be more burdensome than restricting their access to safe websites. In the public health context, however, quarantines are more burdensome than fines in that they restrict an individual's freedom of movement and association. However, even in its current form, the Intervention Ladder is a useful tool for thinking through how to balance risks, responses, and personal rights in cyberspace.

⁶⁴ Rowe, Halpern, and Lentz, "Is a Public Health Framework," 30-38.

This public health approach, which assesses a wide variety of interventions and seeks to find the right balance between efficacy and ethical costs, is fundamentally different from the approach employed by those using the Cybersecurity LoA described in Section 2.2. As cybersecurity is treated as a private problem by those using the Cybersecurity LoA, companies choose the types of interventions which are in their self-interest, area of expertise, and legal authority. While an ISP can shut-off someone's internet access, they cannot auto-enable application security updates, and they might have little direct interest in doing so even if they could. Singer and Friedman report that 27 per cent of internet providers do not attempt to track outbound attacks, and half of those that do track them take no action to mitigate these attacks.⁶⁵ With limited capabilities and incentives which frequently diverge from the public interest, private companies are ill-suited to identifying and implementing the cyberhealth solution which best balances effectiveness and personal rights. Having said this, as companies are well-suited to balancing effectiveness and cost, they may play an important role in controlling the cost of public cyberhealth policies. For instance, if policymakers determine that ISPs should shut down the internet access of spammers and other sources of malware, the ISP itself is likely better suited to creating a cost-effective solution to meet that mandate than the policymaker. In many cases, specific cyberhealth measures, such as patching proprietary code, will only be implementable by private companies.

It bears repeating that the Intervention Ladder is merely one formal articulation of the type of thinking that public health experts engage in as part of their day to day work within public health institutions. It is this type of thinking—characterized by an awareness of ethics and a thoughtfulness about proportionality—that the Public Cyberhealth LoA seeks to encourage by explicitly highlighting individual rights, non-financial forms of harm, and the impact of externalities on the broader network.

2.4.2.2.2 *Ethical Reviews*

Beyond the Intervention Ladder, this thoughtfulness about ethics can also be seen in the use of ethical reviews in public health institutions. The World Health

⁶⁵ Singer and Friedman, *Cybersecurity and Cyberwar*, 175-177.

Organization's Ethics Review Committee, for instance, formally reviews all WHO funded research which involves human subjects and provides guidance to member nations on internal ethical issues pertaining to public health.⁶⁶ Similarly, the Centers for Disease Control and Prevention's Public Health Ethics Unit seeks to raise awareness within the organization of ethical problems which public health interventions can cause and integrate this way of thinking into everyday work.⁶⁷ It may seem trite to say the way to balance ethics and effectiveness considerations is to think about how one balances ethics and effectiveness considerations, but this basic level of ethical review used by public health institutions is almost entirely absent from the technology sector and technology policy. While researchers and developers at companies like Google's Deep Mind have begun to think about the ethical issues which arise specifically in the field of artificial intelligence, this type of analysis is reserved for special cases. While runaway, superintelligent AIs get quite a bit of attention, everyday ethical concerns are often ignored. In fact, it is reasonable to think that the focus on 'killer robots' is a way of diverting attention from more everyday concerns, such as whether or not a product helps or harms customers, whether customers have consented to certain practices (e.g. Facebook running experiments on users),⁶⁸ and whether customers are adequately informed if their data is compromised or purposefully shared with third parties. While the ethics of these practices may be occasionally discussed in the public sphere, they are less commonly discussed by those actually developing technology products on a day to day basis. The lack of discussion about these topics can, in part, be attributed to the downplaying of externalities, rights, and wellbeing in the Cybersecurity LoA. If one's framework does

⁶⁶ World Health Organization, "Research Ethics Review Committee," World Health Organization, <http://www.who.int/ethics/review-committee/en/> (accessed May 16, 2017).

⁶⁷ Centers for Disease Control and Prevention, "Public Health Ethics," Centers for Disease Control and Prevention, May 10, 2015, <https://www.cdc.gov/od/science/integrity/phethics/> (accessed May 16, 2017).

⁶⁸ Vindu Goel, "Facebook Tinkers With Users' Emotions in News Feed Experiment, Stirring Outcry," *New York Times*, June 29, 2014, <https://www.nytimes.com/2014/06/30/technology/facebook-tinkers-with-users-emotions-in-news-feed-experiment-stirring-outcry.html>

not include these variables, it is easy to assume that ethical reviews are unnecessary and merely a hindrance to technological progress. One can point to the recent collapse of Google's Advanced Technology External Advisory Council as evidence that technology companies only pay lip service to many ethical concerns. The board, which was dissolved one week after being founded, had no authority to stop projects, was only to meet four times a year, and included a number of members with questionable qualifications—one of which seemed to be on the board primarily as a way of currying favour with conservative lawmakers.⁶⁹

Lastly, because of the direct impact public health interventions often have on individuals, there is more of a culture of involving a broad number of stakeholders in discussions surrounding specific interventions. This can take place at the individual level with informed consent procedures, at the local level with public health workers engaging communities around the treatment of HIV, and at the national level in a forum like the WHO. By contrast, the technology industry typically prefers to make decisions under a veil of secrecy rather than with meaningful public discussion, and frequently purposefully obscures their intent through impenetrable terms and conditions. These everyday capitalistic practices may have been acceptable and even appropriate when most technology problems looked more like overlapping problems, but as the internet has grown into the essential connective tissue of modern life these overlapping problems are increasingly of shared concern.⁷⁰ While these practices are not a result of the Cybersecurity LoA, per se, it is worth highlighting as an example of how the public cyberhealth approach has broad applicability for reframing how we think about technology policy and corporate practice beyond the narrow scope of network failure.

⁶⁹ Kelsey Piper, "Exclusive: Google Cancels AI Ethics Board in Response to Outcry," *Vox.com*, April 4, 2019, <https://www.vox.com/future-perfect/2019/4/4/18295933/google-cancels-ai-ethics-board> (accessed April 5, 2019).

⁷⁰ For example, in the past if an individual's data was taken or given away without consent only that individual may have been harmed. Today, that data may be used to help sway elections, as was the case with Cambridge Analytica's involvement in Brexit and the 2016 US presidential election. As such, a previously overlapping problems becomes one of shared concern.

While it may seem a bit unlikely that a new framing device can replace a LoA as ubiquitous as cybersecurity, the history of public health once again provides some reason to think such a shift is possible. One historical parallel is the private control of water sources and the 1854 cholera outbreak in London. At the time of the cholera outbreak, many water sources were provided by private companies with minimal oversight. While the dominant theory of the time attributed cholera to miasma or bad air, by mapping incidents of the disease John Snow identified that a specific water pump on Broad Street was the likely source. He also identified that the company supplying the water was using water from sewage polluted sections of the Thames.⁷¹ As people's understanding of the importance of clean water grew, governments took on a larger role in the oversight of water quality. The people of London did not simply wait for the water companies to become public health campaigners and clean-up their act on their own. As the internet and other digital technologies become an ever more essential component of modern life—underpinning critical infrastructure and social interactions—our notions of corporations' private rights and public responsibilities must be continually re-evaluated. The alternative is to accept that the digital landscapes in which we live and the security of the critical infrastructure on which we rely, will be designed to maximize profits rather than human flourishing.

2.5 Conclusion

In Chapter 1, I argued that the philosophy of public health and the history of public health institutions could help one to understand 1) the importance of public goods to network robustness and resilience and 2) the obligations of states, individuals, and corporations to participate in the production of said goods. In this chapter, I argued that this usefulness was not an anomaly, and that there is substantial value in viewing many technology matters through a public health lens.

Using the method of levels of abstraction, I formally defined what I called the Cybersecurity LoA and the Public Cyberhealth LoA. The Cybersecurity LoA represented the dominant way corporations and states tend to think about the digital landscape. It was characterized by a focus on adversary-based risks and financial costs, a very limited view of state action, and the use of offensive capabilities. The

⁷¹ Judith Summers, *Soho: A History of London's Most Colourful Neighborhood* (London: Bloomsbury, 1989), 113-117.

Public Cyberhealth LoA, meanwhile, considered adversary and non-adversary-based threats, highlighted non-financial harms and individual rights, downplayed the use of offensive capabilities, and was characterized by a greater awareness of externalities. I argued that not only is the Public Cyberhealth LoA more operationally coherent than the Cybersecurity LoA, but that it could also be used as a framework for guiding technology policy more generally.

Lastly, I illustrated how one might use the Public Cyberhealth LoA to understand the normative justification for—and ethical limits on—government interventions in cyberspace. First, expanding on the discussion from 1.3.1, I explored how thinking like a public health expert can help one to define spheres of public and private cyberhealth, which could be used to determine whether governments were justified in addressing a specific digital problem. Then, I demonstrated how the Public Cyberhealth LoA surfaces ethical concerns that may be ignored by those using the Cybersecurity LoA and specifically argued for the value of the Intervention Ladder and ethical reviews for thinking about proportionality.

In the chapters to come, I will continue to develop the Public Cyberhealth LoA by exploring the role digital technologies and information play in our health and wellbeing. These chapters will strengthen the justification for using a public health inspired lens and demonstrate what it means to take health and wellbeing into account when designing technology policies and products.

Chapter 3: Health and Cyberhealth

In Chapters 1 and 2, I introduced and then formalized an alternative framework for conceptualizing the digital landscape inspired by the philosophy of public health. I argued that the philosophy of public health could be useful for thinking about the normative justification for and ethical limits on government intervention in cyberspace, while public health policy and institutions could serve as examples of how to manifest these higher principles (e.g. the WHO, the Intervention Ladder, ethical review boards). This Public Cyberhealth LoA takes seriously non-malicious threats to network robustness and resilience, highlights the impacts of network threats and interventions on health and wellbeing, and is more thoughtful about protecting individual rights compared to the dominant cybersecurity lens typically used by policymakers and IT professionals.

In this chapter and the next, I will flesh out this framework by exploring in greater depth what it means to think about the digital landscape with human health and wellbeing front and centre. This continues the work begun in Chapter 2 of demonstrating that the Public Cyberhealth LoA is a full-blooded alternative to the cybersecurity lens. First, in this chapter I will explore how poor cyberhealth impacts health. Then, in Chapter 4, I will define a theory of ‘informational wellbeing’ that policymakers can use to assess how digital information and its use, control, accessibility and accuracy impact personal wellbeing. Together these chapters strengthen the justification for using the Public Cyberhealth LoA by revealing the extent to which technology policy and digital threats can impact health and wellbeing. The greater these impacts, the stronger the argument is for using a LoA which explicitly considers these variables when constructing technology policy and designing technology products. Additionally, I will argue that beyond affecting how we think about and assess technology policy, the Public Cyberhealth LoA also suggests we should reassess how we define the very concepts of health and personal wellbeing.

In this chapter, I will put broader wellbeing to the side and focus on health. In Section 1, I will outline why identifying certain cyberhealth issues as health or public health issues is important, and I will define what qualifies something as a health or public health issue. This will be an expansion of the description of public health issues introduced in 2.4.1. In Section 2, I will then describe the somewhat straightforward ways in which poor cyberhealth is a health issue or matter of public health. First, I will consider how the poor cyberhealth of critical infrastructure and medical devices can impact health. Then, I will touch on the ways in which internet access is becoming increasingly important to good health outcomes, and how poor network connectivity can exacerbate unjust health inequalities.

In Sections 3 and 4, I will then explore more interesting (and controversial) cases where the poor cyberhealth of devices and networks can be the causal basis for disease. These cases arise when digital components become coupled to biological systems, as in the case of digital pacemakers. I will argue that for the purpose of determining if one is healthy, we should consider these coupled devices to be part of an individual's body. As such, when a pacemaker works properly and one's symptoms disappear, we should say one is healthy, and if a piece of malware reduces the functional efficiency of the pacemaker (and one's circulatory system), then we should consider this drop in functional efficiency to be a separate pathology from the underlying condition which necessitated the pacemaker to begin with. I will argue that identifying certain cyberhealth issues as pathologies is valuable for two reasons. First, it potentially affects how one allocates funding for public health and cyberhealth, representing one concrete way in which viewing the digital landscape through the Public Cyberhealth LoA differs from the dominant cybersecurity lens. And second, it helps one to re-examine the meaning of familiar concepts, like health, disease, and bodily integrity in the digital age.

It is worth noting that this novel conception of health is quite different from the arguments I made in Chapters 1 and 2 for applying the philosophy and tools of public health to technology policy. In previous chapters, the argument for applying the tools and philosophy of public health to technology policy was based on an analogy between digital networks and human networks—i.e. the 'health' of a digital network is in some ways similar to the health of a population. In Sections 3 and 4 of this chapter, the argument for using the philosophy of public health to craft technology policy does not rely on analogy—I will argue a pacemaker should be

treated as part of one's body and a hacked pacemaker is a disease comparable to a torn ligament or TB. These two approaches are separable. One does not need to adopt this conception of health to think it is a good idea to use the Intervention Ladder to think about proportionality, and conversely one might think it makes sense to treat a pacemaker as a part of one's body but think the cyberhealth of networks is not similar enough to population health to justify the creation of a cyber WHO. I do not think the separability is a weakness but rather is a sign that viewing the digital landscape through a public health lens has benefits in a variety of contexts.

3.1 Health and Public Health—Privileged Categories

Health is generally considered central to personal and collective wellbeing¹ and is essential to being able to function in the world and pursue goals and opportunities.² Sudhir Anand has argued that this importance is recognized across cultures and time.³ Given this significance, states and individuals rightly treat health issues seriously and public health—the “efforts of society as a whole to improve the health of the population and prevent illness”⁴—is a core function of any reasonably well-functioning modern state. A right to health is even included in both the Universal Declaration of Human Rights⁵ and the constitution of the WHO.⁶ While the inclusion of a right to health in these documents is a very contentious issue,⁷ the mere fact that it is considered appropriate by many is a sign of health's significance. In Chapter 1

¹ Amartya Sen, *Commodities and Capabilities*, (Amsterdam: North-Holland, 1985).

² Sudhir Anand, “The Concern for Equity in Health,” in *Public Health, Ethics, and Equity*, eds. Sudhir Anand, Fabienne Peter and Amartya Sen (Oxford: Oxford University Press, 2004), 18.

³ Anand, “The Concern for Equity in Health,” 17.

⁴ Nuffield Council on Bioethics, “Public Health: Ethical Issues,” (London: Nuffield Council on Bioethics, 2007), V.

⁵ United Nations, *Universal Declaration of Human Rights*, (1948): <http://www.un.org/en/universal-declaration-human-rights/>.

⁶ World Health Organization, *Basic Documents*, 48th edition, (World Health Organization, 2014), 1.

⁷ Jonathan Wolff, *The Human Right to Health*, (New York: W.W. Norton & Co., 2012).

(Section 1.3.1), I included access to healthcare in a list of public goods which I believe states have an obligation to provide, given its importance to one's ability to live a minimally decent life.

Despite this importance, within the dominant Cybersecurity LoA described in Chapter 2, relatively little attention is given to the health impacts of poor network robustness and resiliency. When companies, states, and individuals fail to account for these potentially significant health effects, they are prone to underinvest in network robustness and resiliency compared to the socially optimal levels. By identifying these potential health impacts, one can identify a subset of cyberhealth problems which may deserve a greater level of funding, research, and regulation. Historical parallels include our shifting understanding of the health risk of nuclear fallout in the 1950s,⁸ air pollution in the 19th century,⁹ and smoking in the mid-twentieth century.¹⁰ Before the latter was recognized as causing various serious health problems, it was just a leisure activity, like reading a book. After those health issues were identified, governments introduced regulations, launched education campaigns, and allocated funds for treatment and research.

Some aspects of cyberhealth are already recognized as affecting health but are conceptualized using the Cybersecurity LoA (e.g. the security of medical devices). Given the limited role of states within this framework, these devices are under-regulated compared to their potential to impact health. Other cyberhealth issues, meanwhile, are like smoking prior to the link with lung cancer; the potential for these problems to impact health has yet to be studied in depth (e.g. reliable access to the internet). Others still have little to no impact on human health, for example, a targeted cyberattack to steal the IP of a clothing company. However, bundling the three types of cyberhealth problems together as strictly matters of IT security makes it difficult to create nuanced and effective technology policy and to determine proportional responses to specific threats.

⁸ “Fact on the Fall-Out,” *The Washington Post*, December 16, 1954, pg. 20.

⁹ Peter Thorsheim, *Inventing Pollution*, (Athens, OH: Ohio University Press, 2006).

¹⁰ K. Michael Cummings and Robert N. Proctor, “The Changing Public Image of Smoking in the United States: 1964–2014,” *Cancer, Epidemiology, Biomarkers & Prevention* 23, no. 1 (2014): 32-36.

Furthermore, if some issues of cyberhealth are matters of health or public health, then it is particularly appropriate to apply a public health inspired framework to these problems. First, the Public Cyberhealth LoA places more focus on the harm-sufferer and less on the adversary. And second, the Public Cyberhealth LoA will be more sensitive to possible ethical concerns. For instance, returning to the example from Chapter 2, to one using the Cybersecurity LoA a robotic prosthetic infected with malware may be seen as just a technical problem to be solved with an automatic update, whereas a public health frame would be more likely to flag issues related to consent and bodily integrity. I will return to this issue in Section 3.4.

Finally, this exercise is important not only because it forces us to re-evaluate the nature of digital problems, but because it encourages a re-evaluation of what counts as a health problem—this point will be explored in Section 3.4. As individuals rely ever more on digital devices and networks, traditional boundaries between humans and our environment need to be reassessed. If we say a person with a slow heartbeat is unhealthy, what is that person’s health status when that heartbeat is corrected by a pacemaker? And then what if that device malfunctions or is maliciously compromised? Our traditional notions of health and disease may provide an answer, but I will argue that these answers are outdated and inadequate given the increasing sophistication and ubiquity of networked biotechnologies.

3.1.1 Health and Public Health Issues

Having demonstrated that there is value in identifying which cyberhealth issues are health or public health issues, the question then becomes: What is a health or public health issue? In Section 3.3, I will explore definitions of disease in depth, but for now it is sufficient to rely on a more intuitive definition:

Something can be classified as a ‘health issue’ if it is either [a] typically a proximate cause of a harmful biological condition (e.g. environmental hazards, occupational hazards, unhealthy behaviours, disease vectors) or [b] prevents an existing harmful biological condition from being fixed (e.g. a lack of roads; a shortage of hospitals, medical equipment, or shortage of medical professionals; a lack of insurance).

Within the context of cyberhealth, the former may include a cyberattack leading to contaminated water supplies, while the latter may include ransomware which blocks medical professionals from accessing a hospital's computer network, as was the case with the 2017 WannaCry attack on the NHS. It is worth noting that this definition of a 'health issue' is rather limited, in that it does not include distal causes of poor health (e.g. poverty). In this chapter, I want to illustrate that even within the bounds of this modest definition, the health impacts of network robustness and resiliency are potentially significant.

One subcategory of 'health issues' that is particularly relevant for this discussion of cyberhealth is that of 'public health issues.' Public health is often described as some variation of the "efforts of society as a whole to improve the health of the population and prevent illness."¹¹ While I introduced what makes something a public health issue in Chapter 2.4.1, it is worth expanding upon that definition a bit here, as in the last chapter the scope of the discussion was limited by the task at hand. By 'public health issue' I will mean an issue which should concern public health policymakers because it has the ability to impact the collective health of a population. While different states may have different views on what makes something a public health issue, in broad strokes, public health issues generally fall within one or more of the following categories: shared health issues, overlapping health issues at scale, and certain kinds of health inequalities.

The descriptions I use below should not be treated as a comprehensive picture of the field of public health, but rather as a useful set of categories for the discussion that follows. Additionally, as I wrote about the distinction between shared and overlapping health issues in Chapter 2, I will only briefly summarize them here.

3.1.1.1 Shared Health Issues

Shared health issues are those where one person's health status affects the health status of others in the population. Paradigmatic cases include contagious diseases like TB or the flu. Governmental interventions in these types of cases is generally justifiable in liberal democracies under some version of Mill's harm principle (see

¹¹ Nuffield Council on Bioethics, "Public Health: Ethical Issues," V.

Section 2.4.1.1). Examples of specific public health policies to address shared health issues include vaccination campaigns, education programs, and quarantines.

3.1.1.2 Overlapping Health Issues at Scale

The second category is overlapping health issues which affect a very large number of people. While one person being obese is not a matter of public health, tens of millions of obese people might be a matter of public concern. When a sizeable portion of the population is obese, the health system may become overburdened. This may require the general population to pay higher taxes and insurance premiums to support the health system and wait longer for appointments, at least if one accepts Klosko's argument that individuals have an obligation to contribute to presumptively beneficial public goods (see Section 1.3.2). Additionally, if obesity becomes common enough, harmful behaviours like frequent soda consumption and poor eating habits may be normalized, which increase the likelihood of obesity in others.

A subcategory of overlapping health issues worth specifically highlighting due to its relevance in the cyberhealth debate is public safety. Examples include hazardous workplaces, crime, and environmental pollution.¹² In this chapter, the most relevant public safety concern is the fragility of certain forms of critical infrastructure (e.g. nuclear plants, chemical plants, dams). While some threats to public safety are treated as public health issues, many others are not. However, these distinctions are due primarily to the way bureaucracies have historically carved up responsibilities and should not be used to make a normative claim about what should fall inside and outside the bounds of public health. Crime for instance is usually seen as separate from public health, but, as a result, many of the health effects of crime on a community go unaddressed (see discussion on LoAs and illegal drug use in 2.1).

¹² Stephen John, "Why 'Health' Is Not a Central Category for Public Health Policy," *Journal of Applied Philosophy* 26, no. 2 (2009): 129-143.

3.1.1.3 Promoting Equal Access to Healthcare and Reducing Health Inequality

Finally, public health policies may be targeted at reducing certain kinds of health inequalities between populations.¹³ These policies may be targeted at improving access to healthcare for certain disadvantaged groups or on closing the health gap between different groups in a society.¹⁴ Examples include school lunch programs and opening health clinics in underserved areas.¹⁵ It is important to note that not all health inequalities are morally significant. I will discuss this topic in greater depth in Section 3.2.3.

3.2 Poor Cyberhealth and Public Health

If one sorts cyberhealth problems into the categories above, there are four contexts in which poor cyberhealth may appropriately be considered a public health issue: critical infrastructure, medical devices, hospital infrastructure, and a lack of access to the internet. While hospitals can also be considered a part of critical infrastructure, I am choosing to discuss hospital infrastructure separately from other forms of critical infrastructure due to its high potential to impact health directly. It is important to remember throughout this section that labelling certain cyberhealth problems as public health issues does not mean the Cybersecurity LoA is irrelevant in these cases. In many cases, we will be looking at instances of overlapping magisteria where both the Cybersecurity LoA and the Public Cyberhealth LoA have value.

¹³ Yukiko Asada, *Health Inequality*, (Toronto: University of Toronto Press, 2007), 4.

¹⁴ Department of Health and Social Care, *Our Healthier Nation: A Contract for Health*, Cm 3852, Feb. 9, 1998:
https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/265721/title.pdf.

¹⁵ Donald Acheson, *Independent Inquiry into Inequalities in Health Report*, (The Stationary Office, 1998):
https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/265503/ih.pdf.

3.2.1 Failure of Critical Infrastructure

The first and most significant way in which poor cyberhealth can adversely affect health is the fragility and insecurity of the networks which underpin critical infrastructure. While the United States identifies sixteen critical infrastructure sectors ranging from banking to wastewater treatment, the following seven forms of critical infrastructure are typically viewed as relevant to health or public health: the chemical sector, dams, emergency services, food and agriculture, healthcare and public health, nuclear reactors, and water and wastewater.¹⁶ Each of these sectors are underpinned by digital networks, and when those networks fail or are compromised there may be substantial health impacts. The following table summarizes the health risks associated with the poor cyberhealth of critical infrastructure:

Table 4

Sector	Cyberhealth Risk	Public Health or Health Issue
Chemical	<p>Example Vulnerabilities:</p> <ul style="list-style-type: none"> • Stuxnet style attack on SCADA system causing infrastructure damage and leak of chemicals into the environment. • Natural disaster damages digital infrastructure • Bad code/poor testing <p>Known Threat:</p> <ul style="list-style-type: none"> • In 2007, at the Idaho National Laboratory, the US government demonstrated for journalists how they were helpless to stop hackers from destroying a mock chemical plant.¹⁷ 	<ul style="list-style-type: none"> • Instrumental effect on health due to environmental contamination • Public safety

¹⁶ Department of Homeland Security, “Critical Infrastructure Sectors,” Department of Homeland Security, July 11, 2017, <https://www.dhs.gov/critical-infrastructure-sectors> (accessed Jan. 5, 2018). The complete list of sectors is: chemical sector, commercial facilities sector, communications sector, critical manufacturing sector; dams sector; defence industrial base sector; emergency services sector; energy sector; financial services sector; food and agriculture sector; government facilities sector; healthcare and public health sector; information technology sector; nuclear reactors, materials, and waste sector; transportation systems sector; and water and wastewater sector.

¹⁷ Singer and Friedman, *Cybersecurity and Cyberwar*, 37.

Dams	<p>Example Vulnerabilities:</p> <ul style="list-style-type: none"> • Cyberattack taking control of system and opening dams or damaging infrastructure causing flooding. • Natural disaster damages digital infrastructure • Bad code/poor testing <p>Known Threat:</p> <ul style="list-style-type: none"> • Iranian hackers attempt to take control of small Bowman Avenue Dam in New York. The dam was offline at the time. It is believed they thought they were attacking the much larger Arthur R. Bowman dam in Oregon.¹⁸ 	<ul style="list-style-type: none"> • Instrumental effect on health due to flooding • Public safety
Emergency Services	<p>Example Vulnerabilities:</p> <ul style="list-style-type: none"> • Ransomware and other cyberattacks shutting down health centres, disrupting emergency dispatch systems. • Natural disasters disrupting digital networks. • Bad code/poor testing <p>Known Threat:</p> <ul style="list-style-type: none"> • WannaCry ransomware in 2017 briefly shut down a number of UK health centres. • Hurricane Maria knocking out networks on Puerto Rico. 	<ul style="list-style-type: none"> • Instrumental effect on health by limiting access to healthcare • Reduces capacity to respond to outbreaks
Food and Agriculture	<p>Example Vulnerabilities:</p> <ul style="list-style-type: none"> • Cyberattack taking control of system and contaminating food supply. • Bad code/poor testing of software used for food purity and logistics 	<ul style="list-style-type: none"> • Instrumental effect on health by contaminating food supply • Public safety
Health Care and Public Health	<p>Example Vulnerabilities:</p> <ul style="list-style-type: none"> • Ransomware shutting down health centres and locking medical records, insecurity of hospital equipment. • Cyberattacks could give access to hospital systems to unauthorized users. • Bad code/poor testing <p>Known Threat:</p> <ul style="list-style-type: none"> • WannaCry (2017), see above. Isolated ransomware attacks somewhat frequently block individual health centres' access to medical files. 	<ul style="list-style-type: none"> • Instrumental effect on health by limiting access to healthcare • Reduces capacity to respond to outbreaks • Potentially increases inequality of access to healthcare/ health outcomes

¹⁸ Joseph Berger, "A Dam, Small and Unsung, Is Caught Up in an Iranian Hacking Case," *The New York Times*, March 25, 2016, https://www.nytimes.com/2016/03/26/nyregion/rye-brook-dam-caught-in-computer-hacking-case.html?_r=0 (accessed Nov. 17, 2017).

Nuclear Facilities	<p>Example Vulnerabilities:</p> <ul style="list-style-type: none"> • Natural disaster damaging digital control systems leading to meltdown and environmental contamination. • Cyberattack on SCADA system leading to meltdown and environmental contamination. • Bad code/poor testing <p>Known Threats:</p> <ul style="list-style-type: none"> • Stuxnet (2005-2010), the jointly built American/Israeli worm attacked SCADA system controlling centrifuges. Ultimately destroyed roughly a fifth of Iran's centrifuges.¹⁹ • Department of Homeland Security and FBI issued a report in 2018 outlining Russian state actors' attacks on nuclear power plants, water facilities, and other forms of critical infrastructure in the US. Hackers were able to infiltrate the systems and conduct reconnaissance on the workings of the Industrial Control Systems.²⁰ 	<ul style="list-style-type: none"> • Instrumental effect on health due to environmental contamination, loss of electricity • Public safety
Water and Wastewater	<p>Example Vulnerability:</p> <ul style="list-style-type: none"> • Stuxnet style attack on SCADA system causing infrastructure damage, • Cyberattackers taking control of system • Bad code/poor testing <p>Known Threats:</p> <ul style="list-style-type: none"> • Department of Homeland Security and FBI issued a report in 2018 outlining Russian state actors' attacks on nuclear power plants, water facilities, and other forms of critical infrastructure in the US. Hackers were able to infiltrate the systems and conduct reconnaissance on the workings of the Industrial Control Systems. 	<ul style="list-style-type: none"> • Instrumental effect on health due to water contamination. • Public safety

¹⁹ William J. Broad, John Markoff and David E. Sanger, "Israeli Test on Worm Called Crucial in Iran Nuclear Delay," *The New York Times*, January 15, 2011, <http://www.nytimes.com/2011/01/16/world/middleeast/16stuxnet.html> (accessed Nov. 17, 2017).

²⁰ US-CERT, "Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors," US-CERT, March 16, 2018, <https://www.us-cert.gov/ncas/alerts/TA18-074A> (accessed May 25, 2018).

The cyberhealth threats to critical infrastructure are diverse. Common threats include targeted cyberattacks (e.g. Stuxnet), untargeted cyberattacks (e.g. WannaCry), human error, and natural disasters (e.g. Hurricane Maria in Puerto Rico). These in turn can 1) create unsafe environments (e.g. chemical or nuclear leaks), 2) limit access to healthcare (e.g. disrupting healthcare infrastructure), 3) limit a state's ability to respond to an outbreak (e.g. disruption of emergency services), and 4) limit access to basic biological needs (e.g. food and water contamination).

One major technical vulnerability worth highlighting is the vulnerability of the supervisory control and data acquisition systems (SCADA) that manage many types of critical infrastructure. SCADA systems may stay in place for decades, can be very expensive or difficult to update, and in many cases were never meant to be connected to an external computer (let alone the entire internet).²¹ The first successful cyberattack on a SCADA system was the Stuxnet worm (discovered in 2010), which destroyed nearly a fifth of Iran's nuclear centrifuges. Similar attacks could threaten other sectors. Jan Kallberg and Rosemary Burk describe a scenario where taking control of a dam's SCADA system could allow one to open the floodgates and overwhelm dams and reservoirs downstream.²² In areas where dam systems are near dense population centres (e.g. Pennsylvania, West Virginia, Yunnan Province, Hubei Province), flooding could be severe and deadly.

Contributing to the danger of this type of risk is the privatization of critical infrastructure. In the United States, 90 per cent of critical infrastructure is in the hands of private corporations.²³ In the UK that number is around 85 per cent.²⁴ According to Singer and Friedman, "Several major American power companies have told Congress that they judge the known loss of revenue needed to take plants offline for just a few

²¹ Jan Kallberg and Rosemary A. Burk, "Cyberdefense as Environmental Protection—The Broader Potential Impact of Failed Defensive Counter Cyber Operations," in *Conflict and Cooperation in Cyberspace*, eds. Panayotis Yannakogeorgos and Adam Lowther (Boca Raton, London, Paris: Taylor and Francis, 2014).

²² *Ibid.*, 270.

²³ Singer and Friedman, *Cybersecurity and Cyberwar*, 15.

²⁴ Charlie Edwards, *National Security for the Twenty-first Century*, (London: Demos, 2007), 64.

hours to upgrade their cyber systems is greater than any unknown cyber risks.”²⁵ Unfortunately, this narrow financial cost/benefit approach to thinking about the problem has led to upgrades in sectors where the financial benefit is clear, such as in banking—bank hacks can directly lead to substantial monetary losses, and customers will take their money elsewhere if it is not secure. Meanwhile, the most critical types of infrastructure for human health are often the least well prepared for modern network threats.²⁶ While the private sector argues that private business will always know best how to protect their own infrastructure, Singer and Friedman note that the same arguments were deployed by the shipping industry prior to the Titanic and the nuclear industry before Three Mile Island.²⁷ In addition to often overlooking potential health impacts, corporations frequently do not adequately account for (or value) the positive externalities associated with proactive investments in cyberhealth, nor the negative externalities generated by underinvestment, as discussed in Chapter 1. As a result, allowing companies to determine what counts as adequate cyberhealth investments will often lead to underinvestment compared to what would be best for society on the whole.

While many of the threats listed above may look like paradigmatic security threats rather than health issues (e.g. vulnerable chemical plants), three things are important to note. First, as previously noted, calling something a public health issue does not mean it cannot also be simultaneously addressed using the Cybersecurity LoA. Second, historic bureaucratic distinctions are frequently fluid and somewhat arbitrary, such that what seems like a security issue today may be clearly understood as a matter of public health in the future.²⁸ And third, while I am identifying threats to health and collective health, this is not supposed to imply that public health policy must swoop in with new regulations. Rather, the threats should be identified and measured, and then appropriate policies should be created as needed. The overarching point of this dissertation is not to create a bunch of new public health regulations, but rather to understand what risks are posed by poor cyberhealth and to think about how

²⁵ Singer and Friedman, *Cybersecurity and Cyberwar*, 209.

²⁶ *Ibid.*, 202.

²⁷ *Ibid.*

²⁸ Dorothy Porter, *Health, Civilization and the State: A History of Public Health from Ancient to Modern Times*, (London: Routledge, 1997).

public health experts might frame and address the problem. Sometimes the answer might be to monitor the situation, sometimes it might be to intervene, and sometimes it may be to do nothing.²⁹

3.2.2 Insecurity of Medical Devices and Hospital Infrastructure

The second significant way in which poor cyberhealth can impact health or be considered a matter of public health is the security and robustness of medical devices and other hospital technologies. In this section, I will primarily focus on “external” sensors, monitoring equipment, hospital digital networks, and electronic medical records. I will only mention internal devices, such as pacemakers, in passing as I will discuss these devices in greater depth in Section 4.

Despite the importance of medical devices and hospital digital infrastructure to health, their cybersecurity is notoriously poor. According to May Wang, the Chief Technology Officer of Internet of Things (IoT) security firm ZingBox, “For the past three years the healthcare sector has been hacked even more than the financial sector. And more and more hacking incidents are targeting medical devices.”³⁰ The main problems affecting the cyberhealth of these devices are: 1) the ubiquity of devices, 2) proprietary software gives little visibility into potential security flaws, 3) many devices run out-of-date software. While cybersecurity experts do acknowledge these problems, the Cybersecurity LoA’s downplaying of health impacts and treatment of cybersecurity as a private good has not led to effective risk mitigation strategies in this sector.

The potential health risks associated with the poor cyberhealth of hospitals and medical devices are significant. Hospitals in the United States tend to average between ten and fifteen connected devices per hospital bed, and large hospital system can have several thousand beds.³¹ A 2017 survey of IoT search engine Shodan, showed over 30,000 healthcare related devices connected to the internet—three per cent of these devices were still running Windows XP, which Microsoft stopped

²⁹ See 2.4.2.2.1 The Intervention Ladder.

³⁰ Lily Hay Newman, “Medical Devices Are the Next Security Nightmare,” *Wired*, Mar. 2, 2017, <https://www.wired.com/2017/03/medical-devices-next-security-nightmare> (accessed Nov. 20 2017).

³¹ *Ibid.*

issuing security updates for in 2014.³² One effect of this outdated software is that old threats continue to plague healthcare devices. As mentioned in Chapter 2, in 2017, nearly ten years after Conficker appeared, there were over 2.5 million new Conficker infections. 41 per cent—over 1 million infections—were machines being used in the healthcare industry.³³ While Conficker is a relatively benign piece of malware by modern standards, new threats could exploit the same vulnerabilities.

We should be concerned about insecure healthcare devices because they can directly impact an individual's health and serve as an insecure gateway to the rest of a hospital's network. As Anthony James, vice-president of TrapX describes the problem:

Most of these [healthcare] facilities have no clue, because no one [at the facilities] is monitoring their healthcare devices for the presence of an attacker. No one is thinking about a CT scanner or an MRI machine and seeing a launchpad for a broader attack.³⁴

Researchers have demonstrated the ability to hack insulin pumps to alter doses of insulin;³⁵ pacemakers to run down batteries and alter heartbeat;³⁶ temperature control on refrigeration devices which hold medicines and samples; CT scanners; Bluetooth enabled defibrillators; and infusion pumps, which control morphine, chemotherapy

³² Ibid.

³³ O'Neill, "Conficker Worm Still Spreading Despite Being Nearly 10 Years Old."

³⁴ Lily Hay Newman, "Medical Devices Are The Next Security Nightmare."

³⁵ Jim Finkle, "J&J warns diabetic patients: Insulin pump vulnerable to hacking," *Reuters*, Oct. 4, 2016, <https://www.reuters.com/article/us-johnson-johnson-cyber-insulin-pumps-e/jj-warns-diabetic-patients-insulin-pump-vulnerable-to-hacking-idUSKCN12411L> (accessed Dec. 3, 2017).

³⁶ St. Jude had to push out a emergency patch to over 500,000 devices in summer of 2017 after the discovery of this vulnerability. [Alex Hern, "Hacking Risk Leads to Recall of 500,000 Pacemakers Due to Patient Death Fears," *The Guardian*, Aug. 31, 2017, <https://www.theguardian.com/technology/2017/aug/31/hacking-risk-recall-pacemakers-patient-death-fears-fda-firmware-update> (accessed Dec. 3, 2017).]

drugs, and antibiotics.³⁷ While many of the most devastating consequences require a dedicated, malicious actor, these devices can also simply malfunction due to bad code or a faulty update, or be collateral damage in poorly targeted cyberattacks. For this reason, the Public Cyberhealth LoA is superior to the Cybersecurity LoA for thinking about the robustness and resiliency of these devices more broadly. While the Public Cyberhealth LoA can be used to address malicious threats, it is also useful for addressing these non-malicious threats.

While I have singled out healthcare devices, many of the same problems exist for any set of networked devices. Smart homes, autonomous vehicles, and the Internet of Things all have the potential to impact physical health in meaningful ways, but the connection to health is more tenuous than in the case of healthcare devices which deserve a special level of scrutiny.

3.2.3 Network Access and Health Inequalities

The final issue I will consider is unreliable or sporadic access to ICT networks. This cyberhealth issue can both directly lead to poor health outcomes or contribute to unjust health inequalities. In some cases, this sporadic access is due to straightforward cyberhealth issues, such as faulty or inadequate hardware. In other cases, it is that someone lacks access to ICTs as a result of their geography or socio-economic status. Health inequalities that arise from socio-economic status are the paradigmatic example of unjust health inequality.³⁸ While lacking access to the internet is not a matter of network robustness or resilience, per se, the Public Cyberhealth LoA is still well-suited for conceptualizing this issue due to 1) its inclusion of variables and observables related to impacts on health, wellbeing, and individual rights, and 2) its use of the philosophy of public health, which contains rich discussion of health inequalities. This suitability can be contrasted to the Cybersecurity LoA, which has limited use beyond addressing malicious threats. How precisely network access fits into the Public Cyberhealth LoA will become clearer in Chapter 4 once I have defined my theory of informational wellbeing.

³⁷ Kim Zetter, "It's Insanely Easy to Hack Hospital Equipment," *Wired*, April 25, 2015, <https://www.wired.com/2014/04/hospital-equipment-vulnerable> (accessed Dec. 3, 2017).

³⁸ Asada, *Health Inequality*.

The first way that limited network access can harm one's health is by limiting the feature set of networked biotechnologies. Devices like digital pacemakers and internal defibrillators not only regulate heart rhythm, but also serve as data gathering tools which allow physicians to remotely monitor "metrics of device integrity (e.g. battery status, lead impedance), programming issues (e.g. disabling of ventricular fibrillation therapy, insufficient safety margins for sensing or capture), or medical data (e.g. arrhythmias, indication of lung fluid accumulation)."³⁹ As such, individuals who are unconnected to ICTs may receive a lower standard of care if this information is delayed in reaching their doctor.

In addition to limiting the feature set of biotechnologies, a lack of reliable connectivity can also impact an individual's ability to communicate with their doctor (e.g. email, video consultations) and ability to access public health information. While lack of network access is probably a relatively minor contributory factor to the emergence of health problems, it may be a substantial factor in determining whether patients receive adequate levels of care. Assuming these individuals lack network access due to their socio-economic status and not simply because they are choosing to live off the grid, then this inequity would be a good candidate for being considered unjust and a possible target for public health interventions.⁴⁰

In fact, there is good evidence that socio-economic status is largely to blame for being unconnected to ICTs. In the United States, 87 per cent of those who earn over \$75,000 a year have access to broadband at home, compared to 45 per cent who earn less than \$30,000. Rural communities are also disproportionately unconnected.⁴¹ An electrophysiologist in Western North Carolina estimated that approximately 10-25 per cent of his patients who receive digital pacemakers or internal defibrillators are

³⁹ Haran Burri and David Senouf, "Remote Monitoring and Follow-Up of Pacemakers and Implantable Cardioverter Defibrillators," *Europace* 11, no. 6 (2009): 701–709.

⁴⁰ Asada, *Health Inequality*, 38.; This also implicitly assumes there is some kind of right to health (or right to the social basis of health). This is also the basis of my assumption in Chapter 1, that states have an obligation to provide certain public goods required for one to be able to live a minimally decent life.

⁴¹ Pew Research Center, "Internet/Broadband Fact Sheet," Pew Research Center, Feb. 5, 2018, <http://www.pewinternet.org/fact-sheet/internet-broadband/> (accessed May 23, 2018).

unconnected from all forms of ICT networks at home.⁴² As many of these individuals lack network access for similar reasons, such as poverty or living in an area of poor network infrastructure, in specific geographic pockets these percentages will be significantly higher.

The connectivity gap also impacts health providers. In the United States, 1 per cent of small providers lack broadband access, but an estimated 7 per cent of small providers in rural communities remain unconnected.⁴³ Some services, such as telemedicine, require higher and more reliable broadband, which may require providers to have Dedicated Internet Access (DIA)—a special class of internet access that is often several times as expensive as mass-market products. In rural areas, DIA can be three times as expensive as in urban areas. This has led providers in some rural parts of the United States to transport medical records by thumb-drive rather than via digital networks.⁴⁴ In the United States, the Rural Health Care Program provides subsidies to help close the connectivity gap, but the program has been underutilized with only a fraction of the annual spending limit being distributed.⁴⁵ Together these various pieces of data suggest that socio-economic inequalities are one of the main causes of the connectivity gap.

It must be acknowledged that not all health inequalities are unjust or require a public health intervention, some differences in health are simply differences. For example, people who engage in risky leisure activities might be more likely to get

⁴² This fact was relayed to me in conversation.

⁴³ Kate Samuels, et al., “Closing the Rural Health Connectivity Gap: How Broadband Funding Can Improve Care,” USC-Brookings Schaeffer On Health Policy, April 1, 2015, <https://www.brookings.edu/blog/usc-brookings-schaeffer-on-health-policy/2015/04/01/closing-the-rural-health-connectivity-gap-how-broadband-funding-can-improve-care/> (accessed May 26, 2018).

⁴⁴ Steve Lohr, “Digital Divide Is Wider Than We Think, Study Says,” *New York Times*, Dec. 4, 2018, <https://www.nytimes.com/2018/12/04/technology/digital-divide-us-fcc-microsoft.html> (accessed Dec. 5, 2018).

⁴⁵ Kate Samuels, et al., “Closing the Rural Health Connectivity Gap: How Broadband Funding Can Improve Care.”

injured, but this health inequality would not be considered unjust.⁴⁶ Within the philosophy of public health there are numerous theories about what makes a health inequality unjust, of which I will mention three. The first approach argues that health inequalities are unjust if they are the result of socio-economic status.⁴⁷ The second approach argues health inequalities are unjust if they are the result of factors outside of an individual's control (e.g. a skydiver getting injured is not an unjust health inequality).⁴⁸ And a third approach argues that health inequalities are unjust if an intervention exists to solve the problem which is not being deployed.⁴⁹ Assuming unjust health inequalities should be a target of public health interventions, each of these approaches may lead one to adopt different policies in regard to the connectivity gap. The third approach might suggest states should provide network access to everyone (if that is possible), while the second approach would make room for people to choose to live off the grid, come what may. While the Public Cyberhealth LoA does not suggest one straightforward public solution to closing the connectivity gap, by linking issues of internet access to debates about health inequalities one can create more thoughtful and consistent technology policy.

3.2.4 Quantifying the Public Health Impact of Poor Cyberhealth

While the poor cyberhealth of critical infrastructure, hospitals, and medical devices clearly have the ability to impact health, it is difficult to assess how significant the risk to health is with the tools which are currently available. As mentioned in Chapters 1 and 2, we currently suffer from a lack of research into the likelihood of

⁴⁶ Allen Buchanan, Dan Brock, Norman Daniels, Daniel Wikler, "Introduction," in *From Chance to Choice*, eds. Allen Buchanan, Dan Brock, Norman Daniels, Daniel Wikler (Cambridge: Cambridge University Press, 2000), 17-18.

⁴⁷ Paula Braveman, "Health Disparities and Health Equity: Concepts and Measurement," *Annual Review of Public Health* 27 (2006): 167-194, <https://doi.org/10.1146/annurev.publhealth.27.021405.102103>.

⁴⁸ Julian Le Grand, *Equity and Choice: An Essay in Economics and Applied Philosophy*, (London: HarperCollins Academic, 1991).

⁴⁹ Emmanuela Gakidou, Christopher Murray, and Julio Frenk, "Defining and Measuring Health Inequality: An Approach Based on the distribution of Health Expectancy," *Bulletin of the World Health Organization* 78 (2000): 42-54.

cyberattacks, the costs of various network failures, and the effectiveness of interventions. Additionally, while the potential health effects of the connectivity gap are widely acknowledged, there have been no studies thus far that quantify this impact.

One particular challenge of assessing the health impacts of poor cyberhealth is the interrelated nature of network threats. For example, the failure of one critical infrastructure sector could cause others to fail (e.g. damaged cell networks would affect emergency services). Or in the case of medical devices, how does the risk of the least secure device jeopardize the security of more secure devices on the same network? For each of the vulnerabilities listed in the preceding sections, millions or even billions—in the case of critical infrastructure—of people worldwide are at a slightly higher risk of ill-health than they would be if these vulnerabilities were addressed, but how that risk translates into quantifiable health outcomes is beyond the scope of this work and unfortunately has yet to be tackled by other researchers. Such analyses will be essential for not only determining efficient and effective levels of government funding and regulations, but for ensuring state interventions are proportional and just.

A further challenge, apart from measuring the potential health effects of network failure, is measuring the impact of living in the state of risk caused by poor cyberhealth. Living in a state of risk or vulnerability can adversely impact wellbeing and lead to inefficient investment of resources (as individuals or companies must guard against future shocks).⁵⁰ As Jonathan Wolff and Avner De-Shalit put it, “exceptional risk and vulnerability is itself a disadvantage, whether or not the feared event ever actually happens.”⁵¹

3.2.5 Recommendations

As it is difficult to measure the potential health impacts of poor cyberhealth with existing tools, it is hard to determine what counts as a proportional response. Therefore, my policy recommendations are conservative. States should focus funding efforts on fixing the most egregious security lapses (e.g. updating or phasing out

⁵⁰ Stefan Dercon, “Risk, Poverty, and Vulnerability in Africa,” *Journal of African Economies* 14, no. 4 (2005): 483-488.

⁵¹ Jonathan Wolff and Avner De-Shalit, *Disadvantage*, 9.

devices using out-of-date operating systems), creating standards and testing procedures for new devices and networked systems that have yet to be deployed, and increasing network access for individuals (in part because this has positive effects on wellbeing beyond strictly improving access to healthcare).⁵² Focusing on new devices avoids the tricky issue of requiring owners of critical infrastructure and manufacturers of devices to make costly changes with little guarantee that they have meaningfully reduced their risk. Additionally, it is cheaper to improve the robustness of new devices than to attempt to modify old hardware and software, which is prone to breaking and perhaps is already being phased out. New devices can also be sold for more money to offset security investments, while patching old devices only costs a manufacturer. All of these steps should be filtered through institutional ethical review processes, as discussed in Chapter 2, to ensure that the pursuit of cyberhealth is not unnecessarily or unacceptably harming personal rights.

Lastly, one issue which greatly contributes to device and network vulnerability is a lack of visibility into proprietary software. For instance, while the United States' Food and Drug Administration began considering cybersecurity as part of the medical device approval process in 2013, testing is the responsibility of the device manufacturer.⁵³ In the European Union, meanwhile, manufacturers have been left to develop standards for medical device IT security. The cybersecurity of devices is only mentioned in passing in the EU's current medical device regulations (adopted in 2017).⁵⁴ Third party—or government—code validation and penetration testing would help ensure that devices were adequately protected. While frequently those using the “feudal” Cybersecurity LoA are hesitant to expose proprietary code for fear that vulnerabilities and valuable IP will become public, in practice dedicated hackers can

⁵² The connection between network access and wellbeing will be revisited in Chapter 4.

⁵³ U.S. Food & Drug Administration, “Cybersecurity,” U.S. Food & Drug Administration, <https://www.fda.gov/medicaldevices/digitalhealth/ucm373213.htm> (accessed March 4, 2019).

⁵⁴ Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC: <http://data.europa.eu/eli/reg/2017/745/oj>.

almost always discover these vulnerabilities. It is far superior from the standpoint of society at large that the vulnerabilities are found by friendly eyes and able to be addressed proactively.

While it is possible that one may come to similar recommendations without the Public Cyberhealth LoA, the Public Cyberhealth LoA highlights the urgency of these improvements and provides a justification for a more robust government response (see the discussion of public goods in Chapter 1). Additionally, the Public Cyberhealth LoA is useful for thinking through issues like the expansion of network access, which may weaken cybersecurity in a narrow sense, but nonetheless improve a more equitable distribution of health and wellbeing.

3.3 Definitions of Disease

In Section 2, I explored a few of the more straightforward ways in which poor cyberhealth can impact health. In the second half of this chapter, I will argue that when a digital technology is closely linked to a biological system (e.g. a digital pacemaker), poor cyberhealth should be considered itself a pathology or disease. Calling a cyberhealth issue a disease is potentially significant for at least three reasons. First, there may be a moral claim to treatment for some or all diseases, (depending on how one defines the term). This is a vast departure from the traditional way of conceiving of cyberhealth issues as hindrances to business or state strategic interests (see the Cybersecurity LoA in Section 2.2). Second, if poor cyberhealth can be a disease, then this strengthens the justification for using the Public Cyberhealth LoA, which being grounded in the philosophy of public health is better suited to conceptualizing ethical questions related to bodily-integrity, health inequality, and the moral right to treatment than the Cybersecurity LoA. And third, it may affect how we think about the distinction between biotech treatments and enhancements. This, in turn, has ramifications for deciding which biotechnologies states and insurance companies are willing to pay for.

Before discussing how a cyberhealth issue can be a constitutive part of a disease state, it is necessary to first define what is meant by the term disease or pathology (throughout this section, I will use these terms interchangeably, following the practice of one of two main theorists I will discuss in this section, Christopher

Boorse).⁵⁵ As there is not one agreed upon definition, I will present two of the most commonly used definitions below, which I will reference throughout the next section. Broadly speaking, definitions of disease fit into two categories—naturalist definitions, which define disease as some form of biological dysfunction, and hybrid definitions, which define disease as a *harmful* biological dysfunction, where harmful is a sociocultural designation.⁵⁶ In Section 4, I will argue that within both accounts of disease, certain cyberhealth issues should be considered as constitutive parts of pathologies.

3.3.1 Boorse's Biostatistical Theory

The first definition of disease I will employ is Boorse's influential biostatistical theory. The essential formulation of Boorse's theory, directly quoted, is as follows:

- 1) The *reference class* is a natural class of organisms of uniform functional design; specifically, an age group of a sex of a species
- 2) A *normal function* of a part or process within members of the reference class is a statistically typical contribution by it to their individual survival and reproduction
- 3) A *disease* is a type of internal state which is either an impairment of normal functional ability, i.e. a reduction of one or more functional abilities below typical efficiency, or a limitation on functional ability caused by environmental agents.
- 4) *Health* is the absence of disease.⁵⁷

If a person breaks their ankle, we would say they have a pathology according to BST because their ability to walk is far below typical efficiency for their reference group and being able to walk is important to one's ability to survive. Boorse argues that the term disease is value-free. As such, not all diseases according to BST need be considered harmful by society or the person with the pathology. For example, BST typically classifies homosexuality as a disease, as it generally reduces one's

⁵⁵ Christopher Boorse, "A Rebuttal on Health," in *What Is Disease?*, eds. James M. Humber and Robert F. Almeder (Totowa, NJ: Humana Press, 1997), 7.

⁵⁶ Jerome C. Wakefield, "The Concept of Mental Disorder: Diagnostic Implications of The Harmful Dysfunction Analysis," *World Psychiatry* 6, no. 3 (2007): 149–156.

⁵⁷ Boorse, "A Rebuttal on Health," 7-8.

functional ability to reproduce.⁵⁸ The fact that in many countries homosexuality is not generally considered a harmful or undesirable state does not factor into this determination.

3.3.2 Wakefield's Harmful Dysfunction Approach

The second definition of disease I will consider is Jerome Wakefield's 'harmful dysfunction' approach—one of the most influential hybrid accounts of disease. While Boorse argues that disease is a value free term, Wakefield's approach explicitly combines value judgments with scientific assessments of functionality. Wakefield describes his 'harmful dysfunction' approach saying, "a disorder is a harmful dysfunction, where 'harmful' is a value term, referring to conditions judged negative by sociocultural standards, and 'dysfunction' is a scientific factual term, referring to failure of biologically designed functioning."⁵⁹

There are two primary differences between Boorse and Wakefield's accounts. The first difference is in how they define dysfunction. In BST dysfunction is a departure from the species typical contribution of a part or process to an individual's capacity to survive and reproduce. In contrast, Wakefield defines dysfunction in relation to the evolutionary purpose of a part. The second difference, meanwhile, is that Wakefield argues that a dysfunction is only a disease if it is *considered* harmful. 'Considered' is the important word here. Within BST, dysfunctions are harmful in the sense that they reduce an individual's survivability or reproducibility but need not be considered harmful by the individual or the culture to count as diseases (e.g. homosexuality, being on birth control). In contrast to BST, Wakefield would argue that in cultures where homosexuality is not generally considered harmful, it should not be considered a disease—even if the trait is, in an evolutionary sense, a dysfunction. While this may seem like a preferable conclusion to some, it is important to reiterate that within BST disease is a value-free term; i.e. to say an individual is

⁵⁸ Christopher Boorse, "On the Distinction Between Disease and Illness," *Philosophy and Public Affairs* 5, no. 1 (1975): 63.

⁵⁹ Wakefield, "The Concept of Mental Disorder: Diagnostic Implications of The Harmful Dysfunction Analysis."; While Wakefield tends to use the term disorder rather than disease, I will treat the two terms as synonymous. As Wakefield and Boorse compare their respective theories to the other, I believe this is reasonable.

diseased is not to imply any further moral claim. Additionally, Wakefield's approach also may lead one to some surprising conclusions. Tim Lewens⁶⁰ and Rachel Cooper⁶¹ have argued in different works that depression may not be a disease within Wakefield's approach, as it may not qualify as a failure of biologically designed functioning. I mention this example merely to dissuade one of the simplistic notion that Wakefield's approach is on its face clearly preferable to BST.

By arguing that diseases are only those biological dysfunctions which are considered harmful, Wakefield imbues the term disease with an ethical salience that Boorse's value-free conception lacks. As such, a 'disease' by Wakefield's definition is more likely to carry a claim to care than a 'Boorsean disease,' although the strength of this claim varies considerably depending on the degree of harm and the cause of the pathology. Lastly, it is worth noting that in most cases, these two approaches lead one to the same designation. Wakefield and Boorse would both agree that someone with malaria, a broken hip, or a torn ACL have a pathology. With these two definitions in mind, in the next section I will argue that under certain conditions we should consider poor cyberhealth a pathology. While this claim may seem rather odd, I will argue that both Boorse and Wakefield's accounts of disease can accommodate such a claim with only minor, independently plausible, changes.

3.4 Poor Cyberhealth as Pathology

The underlying assumption of both Wakefield and Boorse's accounts of disease is that the constitutive causal basis of dysfunction must be some organic part or biological system. As such, when determining if a person has a dysfunction, one assesses the functioning of organic parts or biological systems without counting the contribution of artificial components or tools. For example, when determining if one is myopic, one assesses the functional efficiency of one's vision without glasses and contact lenses even though these devices play a role in one's ability to see on a day to day basis. While that distinction may be appropriate in the case of eyeglasses (I will revisit this later in the chapter), the assumption that one should not count the

⁶⁰ Tim Lewens, *The Biological Foundations of Bioethics*, (Oxford: Oxford University Press, 2015), 188.

⁶¹ Rachel Cooper, "Disease," *Studies in History and Philosophy of Biological and Biomedical Sciences* 33 (2002): 263-282.

contribution of artificial parts when determining if someone has a disease is complicated by the intimate integration of artificial and biological parts in many modern medical interventions.

To motivate this claim, I will now work through the following (somewhat complicated but plausible) example. Consider the following:

Jim has bradycardia (a slow heart rhythm) due to sinus node dysfunction. When Jim's bradycardia is symptomatic it causes fatigue, weakness, and can lead to fainting.⁶² Within both Boorse and Wakefield's accounts of disease, Jim has a pathology. Jim receives a digital pacemaker which corrects his heart rhythm when it is too slow, relieving his symptoms. The digital pacemaker can run without maintenance for 15 years, and Jim can resume all normal physical activities, including hobbies like mountain biking and hiking. As with other patients with bradycardia who have digital pacemakers, Jim's life expectancy is normal. After a number of years, Jim's digital pacemaker is infected with malware, leading to a drop in the functional efficiency of his circulatory system. He has a second intervention to replace the device, and he returns to his active lifestyle.

This example highlights the oddness of only considering biological parts when determining dysfunction and, by extension, disease. When Jim's pacemaker is working, he seems healthy; he can pursue an active lifestyle, and his lifespan is expected to be normal. Additionally, the pacemaker can only be separated from his circulatory system via surgical intervention. Based on this example, I will argue for the following three claims:

- 1) For the purpose of determining if Jim is healthy, we should count the contribution of Jim's pacemaker towards the functional efficiency of his circulatory system.
- 2) We should consider Jim to be disease-free when his pacemaker is working properly, assuming that a) the functional efficiency of Jim's circulatory system is typical for his reference class and b) that his circulatory system adequately performs its 'biologically designed' purpose of circulating blood.

⁶² Mayo Clinic Staff, "Bradycardia," Mayoclinic.org, Aug. 23, 2017, <https://www.mayoclinic.org/diseases-conditions/bradycardia/symptoms-causes/syc-20355474> (accessed Dec. 6, 2017).

- 3) When Jim's pacemaker is infected with malware, we should consider this a distinct disease or pathology from the underlying sinus node dysfunction.

I will defend these claims by addressing a series of objections, and then I will discuss the potential ramifications for health and cyberhealth policy.

Objection: A malfunctioning pacemaker cannot itself be considered a constitutive part of a disease state because it is not a *biological* dysfunction.

Response: When we speak of a *biological function*, we should separate two senses of the word biological. The first sense refers to a function that biological creatures normally must perform in order to go about their life, such as pumping blood, moving around, eating and digesting, thinking, seeing, etc. The second sense refers to a function that is being performed solely by organic parts—the pumping of blood is being performed by the heart as opposed to a heart-lung machine. I suggest that only the first of these senses should be relevant to diagnosing someone with a disease—in determining if Jim is healthy, we should care that his heart is beating at the appropriate rate and not that it is being regulated by a pacemaker. This can be intuitively understood in the case of less technologically sophisticated devices, such as an artificial hip. Consider the following example:

Barbara fractures her hip. She is in pain and cannot walk. By any common definition Barbara has a pathology. Barbara has surgery to fix her hip. The surgery entails replacing part of the hip socket and the upper portion of the femur with artificial components. Once Barbara recovers from surgery, her new hip performs at least as well as her old one, if not better.

While Barbara has a pathology when her hip is broken, once she has recovered from her hip surgery and regained her ability to walk, she should be considered disease-free. Lewens, for one, argues that artificial hips are indeed often thought of as cures.⁶³ However, this is only the case if we take into account her artificial hip when measuring her functional efficiency. If we only consider her organic parts when evaluating her functional efficiency, she is in fact worse-off than when she had a broken hip, as now she is also missing the top half of her femur and a sizeable portion of her hip socket. Measuring functionality in this way would be a ridiculous thing to

⁶³ Lewens, *Biological Foundations of Bioethics*, 180.

do. For clinicians—and most everyone else—the determining factor as to whether Barbara is diseased is whether or not she can perform the biological function of walking, not the artificial or organic nature of her hip. In all clinically important senses, the artificial parts are now simply a part of her ambulatory system. If Barbara broke her artificial hip, I suspect that most people would simply say she broke her hip.

As with an artificial hip, we should count the contribution of Jim's pacemaker when measuring the functional efficiency of his circulatory system, given that it is the functioning of this system that matters for Jim's survivability and ability to reproduce and not the performance of each individual part. I admit that the two cases are not identical. For instance, one might argue that 1) the pacemaker is an *addition* to the circulatory system rather than a direct one-in-one-out replacement of a dysfunctional part, and 2) the original underlying part-dysfunction remains in the case of the pacemaker. However, if what one ultimately cares about is an organism's ability to survive and reproduce—as is the case in Boorse's account—then these differences are immaterial.⁶⁴ This argument is similar to Lewens' argument for a pluralistic naturalism, which I will return to later in this section.⁶⁵

In the case of Wakefield's hybrid account, we should say Jim is disease free for two reasons. The first reason is that even if the heart's sinus node remains dysfunctional, Jim's condition (having a pacemaker) should no longer be considered harmful given that he can live an active life of normal length. And second, it is not exactly clear that "part" dysfunctions, per se, qualify as a dysfunction within Wakefield's account. Wakefield argues that dysfunction, "refers to failure of an internal mechanism to perform one of its naturally selected functions," and he defines internal mechanism as, "a general term to refer both to physical structures and organs as well as to mental structures and dispositions."⁶⁶ If the mechanism in question is treated as the circulatory system as a whole, rather than simply the problematic sinus

⁶⁴ As Boorse's goal is to describe how pathologists use the term disease or pathology, he could maintain that Jim has a disease despite the fact that in practice Jim is performing at a statistically normal level. In most cases, including the clinical context, we should accept that Jim is healthy.

⁶⁵ Lewens, *The Biological Foundations of Bioethics*, 179.

⁶⁶ Wakefield, "The Concept of Mental Disorder: Diagnostic Implications of The Harmful Dysfunction Analysis."

node, then there is no dysfunction when Jim's pacemaker is working as designed. The corollary to this is that when the pacemaker malfunctions (for whatever reason), such that it meaningfully reduces the functionality of the circulatory system, we should think of this as a different dysfunction than the underlying sinus node dysfunction.

Objection: While one may count the contribution of an artificial hip when measuring one's ability to walk, this is not the case for many other types of medical devices, such as glasses. If a person, David, uses glasses to correct his near-sightedness, he still has a disease. The glasses merely mitigate the symptoms of that disease. If David's glasses break or are smudged, we do not think David has a new disease.

Response: Glasses are substantially different from the case of the artificial hip or the pacemaker because glasses are not as integrated into the visual system of the near-sighted individual. If one pictures a spectrum of integration, on one side you have devices like glasses which I will call "tools," and on the other side there are technologies like pacemakers which once installed become a "part" of a given biological system. While there may not be a clear threshold between tools and parts, one can identify paradigmatic cases on either side. Paradigmatic tools include glasses and crutches. Meanwhile, paradigmatic parts include pacemakers, cochlear implants, and intraocular lenses used in cataract surgery. In between these poles would be devices such as wheelchairs and oxygen delivery systems.

One helpful set of criteria for determining which devices should be considered a part of a biological system and which should be thought of as tools has been developed by Andy Clark and David Chalmers as part of their work on the concept of the 'extended mind.' Below, I will outline their framework and then modify it for non-cognitive biological systems.

In brief, Clark and Chalmers's theory of extended mind says that a person's mind does not need to be defined only by the mental activity which occurs inside their skull.⁶⁷ Instead, what makes something a mind is that it is performing a cognitive task,

⁶⁷ Andy Clark and David Chalmers, "The Extended Mind," *InterAction* 8, no. 1 (2016): 48-64, <https://search.proquest.com/docview/1808003977?accountid=9851> (accessed March 28, 2019).

such as remembering, reasoning, or observing the world.⁶⁸ In arguing this, Clark and Chalmers implicitly distinguish between the two sense of biological that I described previously. Their classic example of an ‘extended mind’ is of a man named Otto and his notebook.⁶⁹ Otto (who has Alzheimer’s disease) and Inga (who does not) are both going to the Museum of Modern Art in New York. Upon deciding to go to the museum, Inga searches through her memory to recall where the museum is located. Otto, meanwhile, checks his notebook (which he always has with him) for the address. They both find the information and successfully make it to the museum.

Clark and Chalmers argue that the two instances of address retrieval “are entirely analogous.”⁷⁰ Inga’s memory is stored solely ‘inside’ her brain, while Otto’s is distributed between his brain and his notebook. As Clark and Chalmers say, “The information in the notebook functions just like the information constituting an ordinary non-occurrent belief; it just happens that this information lies beyond the skin.”⁷¹ Clark and Chalmers argue that the notebook and Otto’s brain are a *coupled system*. They describe a coupled system saying:

All the components of the system play an active causal role, and they jointly govern behavior in the same sort of way that cognition usually does. If we remove the external component the system’s behavioral competence will drop, just as it would if we removed part of its brain. Our thesis is that this sort of coupled process counts equally well as a cognitive process, whether or not it is wholly in the head.⁷²

Returning to my case, a device should be considered a part of a biological system if the system and device form a coupled system. For Clark and Chalmers, the key criteria for this coupling are that the constituent parts are 1) constantly available, 2) the information is easily and directly accessible, and 3) once received the information is readily endorsed.⁷³ Within Clark and Chalmers’ framework, not all notebooks are

⁶⁸ Andy Clark and David J. Chalmers, “The Extended Mind,” *Analysis* 58, no. 1 (1998): 7–19.

⁶⁹ Ibid.

⁷⁰ Ibid.

⁷¹ Ibid., 13.

⁷² Ibid., 8-9.

⁷³ Ibid., 18.

part of minds (some are merely what I called tools), but Otto's is because he constantly keeps it with him and accepts its contribution as if it came from his brain. While one can lose a notebook or it can contain some inaccurate information, Clark and Chalmers argue that these limitations are not fundamentally different from the brain which can be injured, contain faulty memories, and can become temporarily inaccessible through inebriation or sleep.

For now, put to the side whether or not a notebook can be part of a mind. While I find it convincing, it is controversial. I would argue that the basic idea behind Clark and Chalmers' coupling criteria is less controversial and more intuitive when applied to 'non-cognitive' functions, such as the ability to see or walk. First, here is a modified version of the coupling criteria for non-cognitive systems:

An artificial component is coupled to a biological system if:

- 1) It is contributing to the function of a biological system.
- 2) It is constantly available.
- 3) Its contribution to the functioning of the system is readily and directly provided.
- 4) The contribution is automatically endorsed/accepted by the system in question.

Applying these criteria to David and Barbara, one can say that David's glasses and vision system are not coupled, while Barbara's artificial hip and her ambulatory system are coupled. While David's glasses contribute to the functioning of this vision system, they are not always available (e.g. can be easily lost or stolen, prescription needs adjusting) and the contribution is not always readily provided (e.g. smudged lenses, glare). Barbara's artificial hip, meanwhile, is always available,⁷⁴ it directly and readily offers its contribution to her ambulatory system, and Barbara's ambulatory system automatically accepts the contribution. As a result, we should (and I would say generally do) think of Barbara's artificial hip as just another part of her body, but we do not and should not consider David's glasses a part of his body—despite their value, they remain a tool. Other forms of vision interventions, meanwhile, would pass the criteria of being closely coupled to the biological vision system. As mentioned previously, during cataract surgery the biological lens is removed and replaced with

⁷⁴ While Barbara's artificial hip could break, it is as reliable (at least) as her non-artificial hip.

an artificial lens. As with Barbara's hip, the lens is always available, performs reliably, and its contribution is automatically endorsed. Intuitively we accept that the new artificial lens is a part of one's vision system—almost no one with cataract surgery goes around talking about their bionic eye.⁷⁵

Returning to the case of Jim and his digital pacemaker, using Clark and Chalmers' criteria for coupling, it seems like we should consider Jim's pacemaker a part of his circulatory system—an “extended heart” in a manner of speaking. The device is always available, its contribution is reliable and automatically endorsed, and it is performing a heart-like function. If we accept the pacemaker as essentially a part of Jim, then 1) as long as it is working we should consider Jim disease-free, and 2) if a computer virus or other malfunction affects the performance of the device, we should consider this condition as much a dysfunction as his original sinus node dysfunction.

The fact that Jim's pacemaker is capable of transmitting data via digital networks does not change the fact that it is coupled to his circulatory system. Imagine Barbara's hip had sensors that sent her doctor data on her activity levels, or imagine that the cataract patient's artificial lens could measure glucose levels—a feature that was developed by Google and Novartis before ultimately being abandoned due to inconsistent results.⁷⁶ These additional features—assuming the base components are still readily available, reliable, and perform tasks associated with walking and seeing respectively—should not alter our fundamental belief that the hip and lens are now simply part of a person's functional systems. However, this should also be true if the core functionalities of the device depend on digital networks. In these cases, we should think of the network and external computing resources as also part of the coupled system. While the network enabled features of a digital pacemaker probably do not rise to this level, it is easy to imagine devices that would. For example,

⁷⁵ It is worth noting that these intuitions may not extend to how we think about the mind, but this may be that the workings of the mind are more mysterious than joints or the heart. Perhaps with greater clarity into the workings of the brain, our intuitions may change.

⁷⁶ Jihun Park, et al., “Smart Contact Lenses With Integrations Of Wireless Circuits, Glucose Sensors, And Displays,” *Science Advances* 4, no. 1 (2018): DOI: 10.1126/sciadv.aap9841.

imagine contact lenses with facial recognition capabilities that compensate for an individual with face blindness. Likely, this function would require cloud computing resources in order to work properly. In this case, the cyberhealth of the network and the shared computing resources are partly constitutive of one's health status.

There is certainly something a bit odd about the idea that a shared resource like a cloud server could be a part of multiple people's coupled systems. While we do not usually think of body parts as being shared, it is not without precedent. Conjoined twins can share a single liver, heart, pelvis, spine, part of the intestine and occasionally even brain tissue.⁷⁷ Yet, we recognize conjoined twins as separate people despite their shared resources. The case of cloud resources may also seem different because the parts are physically distant, whereas the pacemaker or hip are "internal." Again, this is not normally how we think of bodies working, but I do not think it should affect whether we treat the resources as being part of a coupled system as long as the contribution is reliable, always available, and readily accepted. One could surgically implant a small computer into a person to perform sophisticated feats of computing like facial recognition, but it just seems like a worse medical and technical solution than letting the server sit in a warehouse.

As an aside, while the digital pacemaker case still might *feel* a bit different than the hip, I chalk this up mostly to the terminology involved. An artificial hip is called a hip, something each of us naturally have two of. In contrast, a 'pacemaker' sounds more like it belongs on a racetrack than inside a human body. If the pacemaker was instead called an artificial heart or artificial sinus node, I think we would feel more comfortable accepting it as part of Jim.

Objection: While the clinician or philosopher may consider Jim to be disease-free when his pacemaker is working properly, Jim may still think of himself as having a pathology. He may even want the pacemaker removed or the network connected features turned off despite the physical benefit. It should be the patient who decides whether or not the artificial device or the shared computing resources it uses are considered a part of their body.

⁷⁷ Mayo Clinic Staff, "Conjoined Twins," [Mayoclinic.org](https://www.mayoclinic.org/diseases-conditions/conjoined-twins/symptoms-causes/syc-20353910), March 7, 2018, <https://www.mayoclinic.org/diseases-conditions/conjoined-twins/symptoms-causes/syc-20353910> (accessed March 6, 2019).

Response: The objection above conflates two different questions. The first question is whether one should consider artificial parts and their associated functions—especially parts and functions which rely on digital networks—as part of biological systems (e.g. the circulatory system) for the purpose of disease diagnosis. The second question is whether one should consider those parts and functions to be part of a person’s body apart from the diagnosis of disease.

In regard to the first question, an individual’s feelings about whether or not the artificial parts should be considered a part of their body is irrelevant for determining if the individual has a disease within both Boorse and Wakefield’s accounts. In both Boorse and Wakefield’s account of disease, if there is no dysfunction, then there is no disease. While Jim may be experiencing harm in the form of mental distress at the idea of having a pacemaker, the mental distress is not related to a physical dysfunction—his circulatory system is both performing its evolutionarily designed function (relevant to Wakefield)⁷⁸ and performing at typical efficiency (relevant to Boorse).

This leads us to the second question of whether or not we should consider artificial parts to be part of one’s body apart from the context of disease diagnosis. The argument I presented in this section was not intended to provide an answer to this question. While this question is largely beyond the scope of this particular chapter, empirical research suggests that whether or not people do in fact consider these devices to be a part of their body is highly context dependent. While some children who depend on medical devices incorporate these devices into their self-presentation, others try to conceal the device and pass as ‘normal.’⁷⁹ In the case of prosthetics, the degree to which individuals think of the device as embedded, or a part of the bodily assemblage, depends on both the purpose of the prosthetic (e.g. functional replacement, aesthetic addition, rehabilitative) and external factors (e.g. appearance,

⁷⁸ In the case of Wakefield, even if one argued there was a still dysfunction, this likely would not matter given that the condition of having a pacemaker is generally not thought of as being harmful.

⁷⁹ Susan Kirk, “How Children and Young People Construct and Negotiate Living with Medical Technology,” *Social Science & Medicine* 71, no. 10 (2010): 1796-1803.

capabilities, who controls the device).⁸⁰ One famous example is how Stephen Hawking came to accept his ‘robotic’ voice as part of his identity and refused to adopt more natural sounding voice synthesizers.⁸¹

Given these empirical findings, one could imagine that some networked features may be more easily incorporated into one’s bodily identity if they use cloud computing resources rather than a cumbersome physical device one must constantly lug around. This research also suggests potentially new ways of thinking about the ownership and control of shared computing resources. While today cloud computing resources are typically owned and controlled by a company (e.g. Amazon Web Services), it is also possible for such resources to be owned and controlled by a group of individuals, such as a community of people with the same disease, members of the same family, or a group of friends. Perhaps if individuals incorporate these devices into their bodily identity then there is a *prima facie* argument that they should have greater control over how these devices are managed, maintained, and improved. However, much more work is needed to draw any definitive conclusions.⁸²

One set of concepts which might be useful for thinking about artificial parts and bodily identity are Havi Carel’s concepts of bodily certainty and doubt. Carel defines bodily certainty as “the natural confidence in [one’s] bodily abilities,” while bodily doubt is a doubt in those bodily abilities that can lead to “helplessness, alarm, and distrust in [one’s] body.”⁸³ Carel speaks about illness as being one state which

⁸⁰ Abbe Brown, Shawn H. E. Harmon, Rory O’Connor, Sita Popat and Sarah Whatley, “Body Extension And The Law: Medical Devices, Intellectual Property, Prosthetics And Marginalisation (Again).”

⁸¹ Rachel Martin, “Stephen Hawking Gets A Voice Upgrade.” *Weekend Edition Sunday*, Dec. 7 2014, <https://www.npr.org/2014/12/07/369108538/stephen-hawking-gets-a-voice-tech-upgrade?t=1549637874457> (accessed Feb. 8, 2019).

⁸² For example, Martha Nussbaum includes the ability to have bodily integrity on her list of core human capabilities, see Chapter 4.1.1. If cloud-computing hardware and software is incorporated into one’s bodily identity, then the ability to achieve bodily integrity may require that one has some measure of control over these external devices.

⁸³ Havi Carel, “Bodily Doubt,” *Journal of Consciousness Studies* 20, Issue Nos. 7-8 (2013): 184.

leads to bodily doubt, but one could imagine other (new) sources of doubt which may be unique to networked devices, such as not living in an environment with adequate cyberhealth or relying on shared computing resources that one does not control. In my conversations with doctors, I have found that the mere presence of artificial parts in one's body may lead to bodily doubt. This is even the case when the artificial part is beneficial to the overall physical health of the individual. For example, while an individual's capacity to survive and reproduce may not be harmed by leaving rods in their leg that have been used to fix a fracture, the bodily doubt that accompanied the fracture may persist as long as the foreign objects remain. This mental distress should be taken seriously when determining whether or not the artificial parts should remain in place.

While the philosophies of public health, medicine, and biology do not provide a single answer as to whether or not one should consider artificial parts and their network enabled functions to be a part of one's body in non-diagnostic contexts, these philosophies do provide theoretical tools for thinking about the question, which are absent from security LoAs, like cybersecurity. By incorporating concepts like bodily certainty and doubt, bodily integrity, and definitions of disease into how we think about ICTs, the Public Cyberhealth LoA can help one to craft technology policies and products that are sensitive to individuals' rights and encourage those in the public health and medical fields to think more deeply about how traditional network threats like malware and network fragility can impact health. One context in which this is particularly salient is in the regulation of medical devices. As Richard Clayton, Ross Anderson, and Éireann Leverett have argued, as networked devices become increasingly ubiquitous, “many regulators who previously thought only in terms of safety will have to start thinking of security as well.”⁸⁴ I would add that even beyond security, they must think about robustness and resilience, or in one word—cyberhealth.

⁸⁴ Ross Anderson, Richard Clayton, Éireann Leverett, “Standardisation and Certification Of Safety, Security And Privacy In The ‘Internet Of Things,’” Joint Research Centre, (Luxembourg: Publications Office of the EU, 2018), <https://publications.europa.eu/en/publication-detail/-/publication/80bb1618-16bb-11e8-9253-01aa75ed71a1/language-en>.

3.4.1 Significance for Cyberhealth Policy

Identifying certain cyberhealth issues as diseases may have a number of significant impacts on technology and public health policy. First, technology policymakers and producers must not only think about technology products as commercial products but as constitutive parts of individuals' health state and (possibly) even bodily identity. As an example, let us return to the idea of contact lenses that use cloud computing resources to treat people with face blindness. If the product is not as financially successful as expected, a company may want to shut down the product line and quickly phase-out support for the product. However, for individuals with face blindness, this particular product may have become an integral part of their health status and self-identity. For someone whose face blindness was effectively cured through the use of the product, shuttering the product may be akin to a form of brain damage. As such, for certain types of products, policymakers may want to require companies to support products for a certain number of years, or force companies to transfer maintenance of the product to another entity rather than allowing companies to summarily drop support.

Second, device manufacturers who have largely been focused on safety of their devices should be more thoughtful about how their networked digital devices interact with the broader Internet of Things. As with the previous point, the cost of failure is higher for the individual who has incorporated the device into their bodily identity and sense of bodily certainty. As such, these devices should be held to higher standards of robustness and resiliency than technologies that might just be considered tools. While one might argue that medical devices are already regulated to a greater degree than non-medical devices, as highlighted in Section 3, the regulations regarding the cybersecurity of medical devices are minimal—states typically do not require third party code review or penetration testing.

Third, if we are going to treat digital devices as being part of someone's body, then threats like malware start to look a lot more like traditional health threats (e.g. malaria, the flu) than matters of property destruction (e.g. someone breaking my laptop). If we think of cyber threats in this way, then improving the cyberhealth of the internet is akin to draining malarial swamps—the removal of a hazardous environment. In Chapter 1, I argued that protecting individuals from hazardous environments, such as malarial swamps, is one of the most fundamental

responsibilities of the state and a core aspect of public health policy. Klosko put it in the same category as public goods like national defence and clean water.⁸⁵ If states are justified in draining (or have an obligation to drain) malarial swamps, then it seems reasonable that they might also be justified in mitigating (or have an obligation to mitigate) malware and ensure a sufficient level of cyberhealth more generally. At least this might be the case when networked biotechnologies become more common. Additionally, from the perspective of public health policymakers, just as draining malarial swamps is a core part of public health policy, so may be ensuring a robust internet.

Lastly, categorizing certain cyberhealth issues as pathologies may give one a stronger claim to have those cyberhealth issues addressed than if they were not considered pathologies. I say “may” because a right to treatment often depends on how one defines disease. Lewens has argued persuasively that naturalist theories of disease, like BST, are unable “to serve as the basis for views that hold the health/disease distinction to be an ethically salient one in itself.”⁸⁶ However, Lewens does temper this point by arguing that certain classes of disease such as chronic pain and degenerative diseases may be ethically salient categories, given the degree to which these diseases limit one’s ability to function in the world and the fact that they clearly require medical treatment.⁸⁷ Within Wakefield’s account of disease, it is easier to argue for a right to treatment given that diseases are by definition *harmful* dysfunctions. Given this, one may have a claim that a company should address a cyberhealth vulnerability that could lead to a pathology before developing new discretionary features. Based on my personal experience developing technology products, there is a constant debate over how to allocate resources between fixing bugs, improving security, and developing new features. While the Public Cyberhealth LoA does not suggest one clear solution to how companies and policymakers should balance these various priorities, the discussion in this chapter suggests that the philosophies of biology and public health can add a level of theoretical sophistication to the discussion which is lacking when one only views the problem through the cybersecurity lens described in Chapter 2.

⁸⁵ Klosko, “Presumptive Benefit, Fairness, and Political Obligation.”

⁸⁶ Lewens, *The Biological Foundations of Bioethics*, 177.

⁸⁷ Ibid., 191-192.

3.4.2 An Alternative Approach: Pluralistic Naturalism

It is worth noting that in this section I have made a stronger claim than I really needed to by arguing that we should consider artificial parts and their associated functions as essentially a part of the individual's body. I have done this, in part, to show how cyberhealth issues can be pathologies even within some of the most popular existing definitions of disease. However, Lewens presents an alternate approach to Boorse's naturalism called 'pluralistic naturalism,' which leads to some of the same conclusions. He argues that when talking about health and disease we should consider whether "the overall ability of the organism to survive and reproduce is at a normal level," and not simply the functional efficiency of specific biological parts.⁸⁸ In making this assessment, Lewens argues we should take into account both technology and environmental adaptations as humans are "*niche constructors* par excellence."⁸⁹ For example, he argues that no humans would survive very long without clothes and shelter, yet we do not consider the health of people who use these technologies to be artificial. Applying Lewens' standard, we can say a person in a wheelchair is essentially healthy as long as they live in an environment with ramps.

However, pluralistic naturalism does not include an account of the boundary between an organism and its environment, and thus is neutral as to whether one should consider pacemakers or wheelchairs to be a part of the individual. In some cases, whether or not these devices are considered a part of the individual may be inconsequential. For instance, Lewens would likely reach the same conclusion I have about whether or not Jim (with his pacemaker) is healthy. In other cases, however, it may matter whether or not these devices are considered a part of the individual. In some cases, the difference may be one of degree. For example, one might think companies have some responsibility to support unprofitable devices for a period of time simply because these devices are important to one's health, but if the device is also a part of one's bodily identity then this claim may be stronger. In other cases, whether or not the device is considered a part of the individual may lead one to fundamentally different conclusions. While further research will be needed, one area where this may be the case is the question of who should own and operate shared

⁸⁸ Ibid., 179.

⁸⁹ Ibid.

computing resources, as in the hypothetical face-blindness example discussed previously. If one does not consider these shared devices to be part of the organism, one may be much more comfortable with these devices being owned and operated by private corporations compared to someone who thinks of these devices as part of one's bodily assemblage.

3.5 Conclusion

In this chapter, I argued that poor cyberhealth has the ability to impact health in a number of significant and sometimes surprising ways. In Section 2, I described four contexts in which matters of cyberhealth impact health or public health in straightforward ways: select sectors of critical infrastructure, hospital networks, medical devices, and network access. In each context, I argued that the Public Cyberhealth LoA was better suited to addressing the health risks than the Cybersecurity LoA, given the Cybersecurity LoA's narrow focus on adversary-based threats and treatment of cybersecurity as a private good. Then, in Sections 3 and 4, I argued that when networked devices are coupled to biological systems, we should consider the functioning of the coupled system when determining if someone has a disease. In these cases, when cyberhealth problems, such as malware or buggy code, reduce the functioning of the coupled system, we should consider the reduction in functioning of the artificial devices to be a pathology. In the case of devices with network enable features, this may require one to accept that the cyberhealth of cloud computing resources and even network infrastructure is part of one's health status. Finally, this argument not only suggested we need to rethink how we define disease but raised interesting questions about how medical devices with networked features should be maintained, developed, regulated, and even owned.

It is clear that poor cyberhealth has the potential to significantly impact individuals' health and potentially exacerbate health inequalities. As networked biotechnologies become more sophisticated and commonplace, these effects will likely become more significant. In some cases, the risks are substantial and actual, as in the case of critical infrastructure's vulnerability to natural disasters and sophisticated cyberattacks. Other risks, such as those associated with speculative biotechnologies, are only now becoming visible on the horizon.

Chapter 4: Informational Wellbeing

In Chapter 3, I explored the ways in which a matter of poor cyberhealth (i.e. a lack of network robustness and resiliency) could also be considered a matter of public health. This included rather straightforward vulnerabilities, such as the fragility and insecurity of medical devices, hospital networks, and critical infrastructure, as well as, more unusual cases where poor cyberhealth could itself be thought of as a pathology. Given that digital networks are often important determinants and constituents of health, I argued it was appropriate and useful to think about these networks using the tools of public health policy. The move to such a framework was driven, in part, by a recognition that the cybersecurity framework, with its focus on financial harm, fails to capture the diversity of roles that information networks play in modern lives. However, the impact of these networks not only goes beyond financial losses to health, but beyond health to other aspects of our lives. Therefore, in this chapter, I will widen my focus and explore how the control, use, accessibility and accuracy of digital information impacts personal wellbeing—of which health is but one important part. By wellbeing I mean essentially “what is non-instrumentally or ultimately good for a person.”¹ These impacts, in turn, link back to my argument in Section 1.3 that the state should promote cyberhealth as part of its responsibility to ensure its citizens can live a minimally decent life.

While there are many definitions of wellbeing, information is an even more difficult term to define. As Floridi says, “Information is notoriously a polymorphic phenomenon and a polysemantic concept.”² In this chapter I will use the General Definition of Information (GDI), which defines information as data plus meaning.³ When I speak of ‘digital information,’ I am referring to the information stored in or transmitted by digital technologies like computers, smartphones, servers, the internet,

¹ Roger Crisp, “Well-Being,” *Stanford Encyclopedia of Philosophy*, Sep. 6, 2017, <https://plato.stanford.edu/entries/well-being> (accessed 30 May 2018).

² Luciano Floridi, *The Philosophy of Information*, 81.

³ *Ibid*, 83.

etc. (as opposed to paper records or magnetic tape), and by ‘personal information’ I mean information whose semantic content is *about* an individual (e.g. medical records, one’s address, photos of a person, a social media profile). In this chapter, I will focus my examples on the control, use, accessibility and accuracy of *digital information*, although much of what I will say about digital information will also be true about non-digital information. For instance, having the ability to ensure the accuracy of one’s medical records is valuable regardless of whether or not that record is digital. Focusing on information, as opposed to solely the robustness and resiliency of network infrastructure, is appropriate considering that digital networks are valuable only insofar as they facilitate the movement of information.

Understanding the connection between wellbeing and digital information and its use, control, accessibility and accuracy is central to my overall project for two reasons. First, as was the case with health impacts, if digital information and its use, control, etc. significantly impact wellbeing, then for the purposes of creating technology policy one would want to use a LoA which includes wellbeing as an important variable. While wellbeing impacts are an important variable in the Public Cyberhealth LoA outlined in Chapter 2, wellbeing is almost never included in cybersecurity cost/benefit analyses. Second, understanding the connection between wellbeing and digital information is necessary to operationalize the Public Cyberhealth LoA. In Chapter 2, I defined the goal of the Public Cyberhealth LoA as: *to promote cyberhealth as part of a broader goal of promoting human health and wellbeing*. One cannot create and maintain networks that promote wellbeing, if one does not have a clear idea of how wellbeing is impacted by digital information and its use, control, accessibility, and accuracy.

In this chapter, I will define a theory of informational wellbeing that enables one to identify and measure how various information practices and policies impact personal wellbeing. First, in Section 1, I will introduce one way of talking about wellbeing—the capability approach to wellbeing. This approach defines wellbeing in terms of one’s “ability to do valuable acts or reach valuable states of being”⁴ (e.g. the ability to live a normal lifespan, the ability to be healthy, etc.). I will then outline one

⁴ Amartya Sen, “Capability and Wellbeing,” in *The Quality of Life*, eds. Martha Nussbaum and Amartya Sen (Oxford: Clarendon Press, 1993), 30.

theory of wellbeing which uses the capability approach—Martha Nussbaum’s list of Central Human Functional Capabilities.

In Section 2, I will then demonstrate that certain ‘informational capabilities’ are central to one’s ability to achieve the capabilities on Nussbaum’s list, and thus are central to achieving a high degree of wellbeing according to at least one influential account of wellbeing. The informational capabilities I will focus on are:

- 1) the ability to control access to one’s personal information
- 2) the ability to use one’s personal information
- 3) the ability to ensure the accuracy of one’s personal information
- 4) the ability to live in an environment with an adequate degree of cyberhealth.

In Section 3, I will then formalize this relationship by articulating a theory of informational wellbeing, which defines the relationship between informational capabilities and overall wellbeing. This concept of informational wellbeing is akin to the way we might, as a shorthand, discuss one’s athletic wellbeing or professional wellbeing, i.e. the aspect of one’s overall wellbeing concerned with one’s profession. While I will discuss the theory in depth in Section 3, the essential form is as follows:

An individual has a high degree of informational wellbeing to the extent they have achieved the ‘informational capabilities and functionings’ (e.g. the ability to control access to their personal information, etc.) which are necessary to achieve fundamental human capabilities (e.g. the ability to live a natural life span, etc.).

This theory is primarily intended to guide social scientists as they seek to measure the impacts of digital technologies on wellbeing and to help technology policymakers assess the success of technology policies in wellbeing terms rather than benchmarks like financial impact and the establishment of infrastructure.⁵

⁵ While this theory shares many traits with what Anna Alexandrova calls mid-level theories of wellbeing, its scope is wider than most mid-level theories—more akin to physical wellbeing than child-wellbeing or the wellbeing of mothers. As such, I generally do not refer to this theory as ‘mid-level.’ Having said this, if one localizes the general theory (e.g. the informational wellbeing of displaced persons), the localized version would likely qualify as a mid-level theory. Anna Alexandrova, *The Science of the Philosophy of Wellbeing* (Oxford: Oxford University Press, 2017).

While this theory is central to the Public Cyberhealth LoA for the reasons stated above, one need not fully adopt that LoA to find this theory useful. I believe this detachability is a feature. It may be of particular use in the fields of development studies and sociology as part of efforts to understand the role digital information and ICTs play in people's lives. Additionally, it may be of use to technology policymakers who are not willing to accept the value of the broader public health inspired approach. Having said this, precisely how the theory would be used and the policies it would inspire would likely depend on the LoA one is using.

While this is the first attempt to use the capabilities approach to create a theory of informational wellbeing, the capabilities approach has been selectively applied to the use of ICTs in the past. Nicholas Garnham has explored the application of the capability approach to communications,⁶ Shirin Madon has outlined a capability approach to evaluating e-governance reforms in India,⁷ and Björn-Sören Gigler suggested using 'informational capabilities' to measure the impact of ICTs in the development context.⁸ In each case, the capability approach was determined to be a better approach for assessing the impact of ICTs on people's lives compared to metrics like the establishment of infrastructure.⁹

Gigler's work is closest to my own but differs in two fundamental ways. First, his work is grounded in development economics, specifically the context of rural Bolivia. By contrast, I intend my theory of informational wellbeing to be more widely applicable. Second, Gigler focuses on capabilities that one actively performs, such as

⁶ Nicholas Garnham, "Amartya Sen's 'Capabilities' Approach to the Evaluation of Welfare: Its Application To Communications," in *Beyond Competition: Broadening the Scope of Telecommunication Policy*, eds. Bare Cammaerts and Jean-Claude Burgelman (Brussels: VUB University Press, 2000), 25-36.

⁷ Shirin Madon, "Evaluating The Developmental Impact Of E-Governance Initiatives: An Exploratory Framework," *The Electronic Journal of Information Systems in Developing Countries* 20, no. 5 (2004): 1-13.

⁸ Björn-Sören Gigler, "'Informational Capabilities' - The Missing Link for the Impact of ICT on Development," *The World Bank*, working paper series no. 1, (March 2011), <http://documents.worldbank.org/curated/en/227571468182366091/pdf/882360NWP0Box30series0no10March2011.pdf> (accessed Oct. 24, 2018).

⁹ Ibid.

using the internet to access the price of grain at the market. While these “athletic” capabilities (to use G.A. Cohen’s term)¹⁰ are important in my work, I also will stress the importance of more “passive” informational capabilities, such as simply having the ability to live in an environment with an adequate degree of cyberhealth. While using the internet may be a valuable functioning in certain contexts, one’s wellbeing may be substantially impacted by the use, control, accessibility and accuracy of digital information even if one has never personally used an ICT. For example, many of those without internet access nonetheless depend on bureaucratic services and critical infrastructure that do rely on digital networks and information. My theory is intended to capture these broader effects of cyberhealth on wellbeing in a way in which Gigler’s does not.

4.1 The Capability Approach to Wellbeing

In this section I will introduce Amartya Sen’s capability approach to wellbeing and outline one influential version of this approach—Martha Nussbaum’s list of Central Human Functional Capabilities.¹¹ In Section 2, I will then demonstrate how one’s ability to achieve these fundamental capabilities—and by extension a high degree of wellbeing—depends on one achieving certain ‘informational capabilities,’ such as the ability to control access to one’s personal information.

The capability, or capabilities, approach is a theoretical framework for describing wellbeing in terms of one’s ability to achieve certain ‘valuable functionings.’¹² Within the capability approach, *functionings* are the various states of

¹⁰ G.A. Cohen, “Equality of What? on Welfare, Goods, and Capabilities,” in *Quality of Life*, eds. Martha Nussbaum and Amartya Sen (Oxford: Clarendon Press, 1993), 25.

¹¹ It is necessary to note that Nussbaum speaks of this list in terms of justice as opposed to wellbeing. By this she means the list is not a comprehensive account of what is good for a person, but that these capabilities are the ones that governments have a responsibility to provide their citizens. In this sense, we are both speaking to what might more accurately be described of as advantage. Jonathan Wolff and Avner De-Shalit, *Disadvantage*, (Oxford: Oxford University Press, 2007).; Martha Nussbaum, *Women and Human Development: The Capability Approach*, (Cambridge: Cambridge University Press, 2000).

¹² Amartya Sen, “Capability and Wellbeing.”

being and doing. Examples of functionings include being vaccinated, living in a warm house, or going to school. A *capability*, meanwhile, is one's ability to achieve various combinations of valuable functionings. An example from Nussbaum's list of capabilities, which I will describe presently, is the ability to enjoy recreational activities.¹³ This capability would be measured by assessing one's ability to enjoy various leisure activities such as games, art, theatre, or watching TV, etc. This may, in turn, include assessments of whether one has free time, the necessary financial resources, people to play with, and knowledge of games.

The capability approach is not a complete theory of wellbeing in and of itself, as Sen does not specify which capabilities contribute to wellbeing. As Serena Olsaretti says, "[The capability approach] only identifies a space for individual and social evaluation, a standard of advantage, which can then be used for descriptive purposes."¹⁴ As a way of talking about wellbeing, it is flexible enough to work with a number of different conceptions of what it means to live a good life. While I will focus on one conception in this chapter—Nussbaum's list of Central Human Functional Capabilities—one can still use the capability approach even if one does not agree that Nussbaum's list is canonical.

In addition to being flexible, a second advantage of the capability approach is that it is not overly prescriptive. For instance, within the capability approach, one should assess whether agents have the *ability* to participate politically, not whether they actually vote. As Nussbaum says:

The conception does not aim at directly producing people who function in certain ways. It aims, instead, at producing people who are capable of functioning in these ways, who have both the training and the resources so to function, should they choose. The choice itself is left to them.¹⁵

While ultimately one must specify the relevant valuable functionings and capabilities, in theory the capability approach has a greater respect for individual freedom and

¹³ Nussbaum, *Women and Human Development: The Capability Approach*, 78-79.

¹⁴ Serena Olsaretti, "Endorsement and Freedom in Amartya Sen's Capability Approach," *Economics & Philosophy* 21, no. 1 (2005): 91.

¹⁵ Martha Nussbaum, "Aristotelian Social Democracy," in *Liberalism and the Good*, eds. R. Douglas, G. Mara and H. Richardson, 203–252 (New York, NY: Routledge, 1990).

avoids some of the dangers associated with paternalism compared to other objective list approaches to wellbeing, such as Hurka's perfectionism.¹⁶ However, in practice one must be careful not to overly rely on the measurement of functionings (compared to capabilities), as doing so undermines Nussbaum's defence of the theory as compatible with liberal norms.¹⁷

For my purpose, two types of capabilities are important: 'fundamental capabilities' and 'informational capabilities.' By fundamental capabilities I mean the capabilities which enable a person to pursue a minimally decent, or "minimally flourishing life"¹⁸ (e.g. the ability to be healthy, the ability to live a normal lifespan). An 'informational capability,' meanwhile, is one's ability to achieve various valuable informational functionings. An example of an informational capability is the 'ability to control who has access to one's personal information,' while valuable informational functionings which contribute to this capability may include living in a country with data regulations that require positive consent for the use of one's information, understanding how one's personal information is used, using strong passwords, and using an email client with end-to-end encryption. The value of specific informational capabilities and functionings will be dependent on the local context. For example, having strong passwords is not a valuable functioning if one does not have personal access to ICTs that require passwords.

4.1.1 Nussbaum's List of Central Human Functional Capabilities

While Sen does not specify the list of relevant capabilities for wellbeing, one such influential list is Nussbaum's Central Human Functional Capabilities. This list is Nussbaum's attempt to identify the capabilities which are required for a life to be "not so impoverished that it is not worthy of the dignity of a human being."¹⁹ The list is not intended to be a comprehensive account of all that is good for a person, but a list of capabilities that states should provide to ensure individuals have the ability to

¹⁶ Thomas Hurka, *Perfectionism*, (Oxford: Oxford University Press, 1993).

¹⁷ Claassen Rutger, "Capability Paternalism," *Economics and Philosophy* 30, no. 1 (2014): 57-73.

¹⁸ Martha Nussbaum, *Creating Capabilities* (Cambridge, MA: The Belknap Press of Harvard University Press, 2011), 33.

¹⁹ Nussbaum, *Women and Human Development: The Capability Approach*, 72.

pursue a “minimally flourishing life.”²⁰ I will use Nussbaum’s list as an example of a set of ‘fundamental capabilities,’ with the understanding that future research may reveal additional fundamental capabilities or suggest alterations to the list below.²¹

Below I will quote at length from Nussbaum’s *Women and Human Development: The Capabilities Approach*, although I have cut down the descriptions of each capability as appropriate:

- 1) **Life.** Being able to live to the end of a human life of normal length...
- 2) **Bodily Health.** Being able to have good health...
- 3) **Bodily Integrity.** Being able to move freely from place to place; having one’s bodily boundaries treated as sovereign...
- 4) **Senses, Imagination, and Thought.** Being able to use the senses, to imagine, think, and reason – and to do these things in a “truly human” way, a way informed and cultivated by an adequate education... Being able to use imagination and thought in connection with experiencing and producing self-expressive works and events of one’s own choice, religious, literary, musical, and so forth. Being able to use one’s mind in ways protected by guarantees of freedom of expression with respect to both political and artistic speech, and freedom of religious exercise. Being able to search for the ultimate meaning of life in one’s own way. Being able to have pleasurable experiences, and to avoid non-necessary pain.
- 5) **Emotions.** Being able to have attachments to things and people outside ourselves...
- 6) **Practical Reason.** Being able to form a conception of the good and to engage in critical reflection about the planning of one’s life...
- 7) **Affiliation.**
 - A. Being able to live with and toward others... to have the capability for both justice and friendship...
 - B. Having the social bases of self-respect and non-humiliation; being able to be treated as a dignified being whose worth is equal to that of others. This entails, at a minimum, protections

²⁰ Martha Nussbaum, *Creating Capabilities*, 32-33.

²¹ Wolff and De-Shalit, *Disadvantage*, 9.

against discrimination on the basis of race, sex, sexual orientation, religion, caste, ethnicity, or national origin....

8) **Other Species.** Being able to live with concern for and in relation to animals, plants, and the world of nature.

9) **Play.** Being able to laugh, to play, to enjoy recreational activities.

10) **Control over One's Environment.**

A. Political. Being able to participate effectively in political choices that govern one's life; having the right of political participation, protections of free speech and association.

B. Material. Being able to hold property (both land and movable goods), not just formally but in terms of real opportunity;...having the right to seek employment on an equal basis with others; having the freedom from unwarranted search and seizure.²²

Nussbaum compiled this list after years of cross-cultural discussion, making it the product of a kind of Rawlsian 'overlapping consensus.'²³ As Nussbaum acknowledges, some of these goods (e.g. health) are 'natural goods.' Whether or not one acquires these natural goods is in part due to luck. For these goods, the role of the state is to try to provide the social basis for the good (e.g. access to healthcare, clean water) rather than the good itself (e.g. health).²⁴

While acknowledging that this list is not without controversy,²⁵ I will use it to demonstrate that informational capabilities are central to at least one influential account of wellbeing. Furthermore, it is worth stating that even if one subscribes to a very different ultimate account of wellbeing, one may still be able to accept that the capabilities on Nussbaum's list are useful for some policy purposes. For example, one may think that wellbeing is ultimately about satisfying one's preferences, but nonetheless accept that having the ability to receive an adequate education and the

²² Nussbaum, *Women and Human Development: The Capability Approach*, 78-80.

²³ John Rawls, *A Theory of Justice*, Revised Ed., (Cambridge, MA: Harvard University Press, 1971, 1999), 340.

²⁴ Nussbaum, *Women and Human Development: The Capability Approach*, 81-82.

²⁵ Richard Arneson, "Perfectionism and Politics," *Ethics* 111 (2000): 37-63.

ability to access healthcare furthers that goal. I will address the concern that Nussbaum's account limits the utility of my theory in Section 4.5.

4.2 Informational Capabilities and Wellbeing

In this section, I will demonstrate how certain 'informational capabilities' play a central role in one's ability to secure the fundamental capabilities listed by Nussbaum.

The 'informational capabilities' I will focus on include:

- 1) the ability to control access to one's personal information
- 2) the ability to use one's personal information
- 3) the ability to ensure the accuracy of one's personal information
- 4) the ability to live in an environment with an adequate level of cyberhealth.

While there are other valuable informational capabilities, these four are adequate to illustrate the importance of informational capabilities (as a class) to one's ability to achieve the capabilities on Nussbaum's list.

In the following subsections, I will present examples which illustrate the importance of each informational capability. For instance, I will look at India's Aadhaar program as an example of having the ability to use one's personal information. While exploring each example, I will highlight the relevant intersections with Nussbaum's list by **bolding** the relevant fundamental capability. As I could write an entire chapter on each capability, these examples should be treated as being illustrative of the importance of the given informational capability and not as a comprehensive treatment of the topic.

4.2.1 The Ability to Control Access to Personal Information

While there are many examples of the importance of *being able to control access to one's personal information*, here I will focus on the inability to control access to personal photos. First, I will explore the example of the actress Jennifer Lawrence, who had a series of nude photos stolen in 2014, and then I will widen the discussion to revenge pornography more generally.

In 2014 hundreds of images of celebrities stored on Apple's iCloud service were stolen and subsequently released online.²⁶ Many of these images depicted the celebrities in various states of undress. While many celebrities were affected, the actor Jennifer Lawrence was particularly vocal about what the experience was like. She described how the theft was not merely a property crime, but a violation of her **bodily integrity**. In describing the theft, Lawrence said, "It's taking somebody's intellectual property but also my body. It was violating on a sexual level."²⁷ As the photos began to appear online, she tried to work on a public statement, but she says, "every single thing that I tried to write made me cry or get angry."²⁸ She described herself as, "Just so afraid." The leak caused her unnecessary emotional pain (**Senses, Imagination Thought**), caused her to lose social bases of self-respect and non-humiliation (**Affiliation**), and plausibly led to employment discrimination (**Control of One's Environment**, i.e. having the right to seek employment on an equal basis with others). Certainly, with more conservative fans, the leaking of the photos harmed Lawrence's reputation—although the scale of this harm is hard to assess. She also acknowledged the important social role of these photos in maintaining her relationship with her partner (**Affiliation**) saying, "I started to write a [public] apology, but I don't have anything to say I'm sorry for. I was in a loving, healthy, great relationship for four years. It was long distance, and either your boyfriend is going to look at porn or he's going to look at you."²⁹ In one particularly revealing

²⁶ Charles Arthur, "Naked Celebrity Hack: Security Experts Focus on iCloud Backup Theory," *The Guardian*, Sep. 1, 2014, <https://www.theguardian.com/technology/2014/sep/01/naked-celebrity-hack-icloud-backup-jennifer-lawrence> (accessed Jan. 24, 2019).

²⁷ Oprah Winfrey, "The Jennifer Lawrence Interview, by Oprah Winfrey," *The Hollywood Reporter*, Dec. 6, 2017, <https://www.hollywoodreporter.com/features/jennifer-lawrence-interview-by-oprah-winfrey-1064576> (accessed April 18, 2018).

²⁸ Sam Kashner, "Both Huntress and Prey," *Vanity Fair*, November 2014, <https://www.vanityfair.com/hollywood/2014/10/jennifer-lawrence-photo-hacking-privacy>.

²⁹ *Ibid.*

quote, Lawrence captures how all these impacts made her question her self-identity and self-worth:

I think, like, a year and a half ago, somebody said something to me about how I was 'a good role model for girls,' and I had to go into the bathroom and sob because I felt like an imposter — I felt like, 'I can't believe somebody still feels that way after what happened.' It's so many different things to process when you've been violated like that.³⁰

In this reaction she both reveals how the experience not only cost her the ‘social bases of self-respect and non-humiliation,’³¹ but also in a sense the ability to define her self. By saying she felt like an ‘imposter,’ she is acknowledging that she had internalized the idea that she was no longer worthy to be a role model for girls despite believing she did not do anything inherently shameful. While Nussbaum does not specifically identify the ability to define one’s self as a fundamental capability, the value of self-determination underlies the entire capability approach’s focus on freedom and is present in many of the fundamental capacities Nussbaum lists, including **Practical Reason; Emotions; Sense, Imagination and Thought**. However, this is a subtle and tricky issue, and I will revisit it in more depth in Section 4.6.2, at which point I will argue that we should consider the ability to define one’s self as a fundamental capability on par with the other capabilities on Nussbaum’s list.

While Lawrence has been particularly vocal about her experience, her experience is not uncommon. Mudasir Kamal and William J. Newman describe the mental effects of revenge pornography more generally:

The distress includes anger, guilt, paranoia, depression, or even suicide. There may also be deterioration in personal relationships and feelings of isolation. The humiliation, powerlessness, and permanence associated with these... crimes leave victims engaged in a lifelong battle to preserve their integrity. Consequently, victims of revenge pornography suffer from similar enduring

³⁰ Erika W. Smith, “Jennifer Lawrence Speaks Out About Reclaiming Her Body After Her Nude Photos Were Published Without Her Consent,” *Bust*, <https://bust.com/feminism/194242-jennifer-lawrence-reclaiming-body-after-nude-photos.html> (accessed April 18, 2018).

³¹ Nussbaum, *Women and Human Development: The Capability Approach*, 78-80.

mental health effects as described by victims of child pornography, such as depression, withdrawal, low self-esteem, and feelings of worthlessness.³² Often an individual's name, address, and social media account are posted alongside the photographs, significantly heightening the mental distress of the primary violation.

The capabilities approach is particularly useful in this context because it emphasizes that it is the loss of the ability to control who has access to one's photos rather than the voluntary sharing of the photos themselves that is the problem. Lawrence not only did not feel shame about taking the photos and sharing them with her boyfriend, but she specifically acknowledged that this activity strengthened their relationship. The capabilities approach values one's freedom by acknowledging that what constitutes a good life is not the same for everyone. While for some people taking nude photos and sharing them with one's romantic partner is not a valuable functioning, for others it is. A more simplistic objective list approach to wellbeing may fail to account for the positive aspects of this activity.

4.2.2 *The Ability to Use One's Personal Information*

The second informational capability I will highlight is *the ability to use one's personal information to achieve a host of secondary goods*. This informational capability is particularly important because bureaucracies have typically used identifying informational artefacts (e.g. IDs, biometrics) as the gateway to everything from securing a home loan, to receiving government benefits, to participating politically. While this informational capability is important in non-digital contexts (e.g. paper passports, drivers' licenses, etc.), I will explore a contemporary digital example—Aadhaar, India's biometric ID program.

The Aadhaar program uses citizens' iris scans, fingerprints, and photographs to generate a twelve-digit ID number and an ID card. The aim of the program is for that ID card and one's biometrics to be the mechanism by which one claims welfare, health services, food rations, pension benefits, registers for certain schools, and votes. In theory at least, having the ability to use one's biometric information is central to

³² Mudasir Kamal and William J. Newman, "Revenge Pornography: Mental Health Implications and Related Legislation," *Journal of the American Academy of Psychiatry and the Law Online* 44, no. 3 (2016): 359-367.

many Indians' ability to live a normal lifespan (**Life**), be healthy (**Health**), participate politically (**Control of One's Environment**), and receive an adequate education (**Senses, Imagination, Thought**). One example of Aadhaar's promise is the experience of Manisha Kamble, a homeless seventeen-year-old living in Mumbai. Speaking of her experience, Manisha said, "In India, you're nothing without Aadhaar."³³ Without a birth certificate or address, Manisha was essentially invisible to the state. After a charity helped her get a card, she was able to register for a school (**Senses, Imagination, Thought**), and when she turns 18 she will be able to use her card to register to vote (**Control of One's Environment**).³⁴

In practice, the program has often failed to achieve its purpose and has been plagued with technical and logistical problems. While Manisha's experience illustrates the promise of Aadhaar, others have lost previously available resources, such as food rations, due to technical problems with the program. In some cases, these technical issues are paradigmatic cyberhealth issues, such as unreliable internet (especially in rural regions), faulty fingerprint readers, and insecure databases.³⁵ In other cases, individuals who lack fingerprints through a lifetime of manual labour, age, or amputation can be turned away from food ration offices or other essential services.³⁶ Jean Dreze, an economist studying Aadhaar, has identified at least a dozen individuals who died of hunger in 2018 after either being unable to enrol in the program or being turned away when their information could not be accessed.³⁷ This is

³³ Lauren Frayer, "India's Biometric ID System Has Led to Starvation For Some Poor, Advocates Say," *NPR*, Oct. 1, 2018, <https://www.npr.org/2018/10/01/652513097/indias-biometric-id-system-has-led-to-starvation-for-some-poor-advocates-say?t=1538831766242>.

³⁴ Ibid.

³⁵ Ibid.

³⁶ Indrani Basu, "AADHAAR: Fading Fingerprints Mean This Ageing Space Scientist Can't Care For His Son," *Huffington Post*, April 19, 2018, https://www.huffingtonpost.in/2018/04/19/an-81-year-old-space-scientist-wants-the-supreme-court-to-save-senior-citizens-from-aadhaar_a_23414358 (accessed April 6, 2019).

³⁷ Frayer, "India's Biometric ID System Has Led To Starvation For Some Poor, Advocates Say."

a particularly stark example of a cyberhealth issue being a literal health issue, which was the focus of Chapter 3. For my purposes, Aadhaar's successes *and* failures highlight the importance of being able to use one's personal information in the digital age.

One instance where the language of informational capabilities is particularly useful is for assessing the impact of the program's poor cybersecurity. As the technology lawyer Mishi Choudhary says, "any compromise of such a database is essentially irreversible for a whole human lifetime: no one can change their genetic data or fingerprints in response to a leak."³⁸ Unfortunately, the system has already been hacked. One investigative journalist was able to buy access to a billion people's information for a mere seven dollars.³⁹ While this seems problematic, it is not always easy to point to immediate financial or reputational harm resulting from such breaches. As a result, one may (and courts often do) conclude that no harm has actually occurred from the breach.⁴⁰ However, using the language of capabilities one can argue that one's wellbeing has been affected by these breaches as one has lost the ability to control access to one's personal information and the ability to use one's biometric information in the future; once biometric data is compromised, it is far less secure as a means of identity in other contexts. These impacts are realized the moment the database is compromised regardless of whether or not one's identity is ever actually stolen.

One may want to argue that the ability to use one's personal information is specifically important in the context of Aadhaar and not necessarily generalizable, but personal information is used to achieve goods and services in many different contexts. Biometrics are now frequently used for identification purposes at border crossings, in refugee camps to register individuals and disburse benefits, and increasingly for

³⁸ Mishi Choudhary, "Viewpoint: The Pitfalls Of India's Biometric ID Scheme," *BBC News*, April 23, 2018, <https://www.bbc.co.uk/news/world-asia-india-43619944>.

³⁹ Frayer, "India's Biometric ID System Has Led To Starvation For Some Poor, Advocates Say."

⁴⁰ Daniel Solove, "'I've Got Nothing to Hide' and Other Misunderstandings of Privacy," *San Diego Law Review* 44 (2007): 768.

everyday tasks like banking.⁴¹ Meanwhile, non-biometric personal information is used even more frequently. One's name, address, date of birth, email address, and the ubiquitous mother's maiden name are a few examples of pieces of personal information that enable one to enrol in schools, vote, bank, own property, travel, work, and receive government benefits. Just because we frequently have the ability to use this information, does not mean that the capability is not important, nor that one will always be able to use that information in the future. Combining massive data breaches with sophisticated AI capable of imitating individual's speech patterns may make identity theft much easier and more common in the future, significantly diminishing our ability to use that information to obtain goods and services.

4.2.3 *The Ability to Ensure the Accuracy of Personal Information*

The third informational capability I will highlight is *the ability to ensure the accuracy of one's personal information*. As should be obvious, all of these informational capabilities are to some degree entwined. The ability to ensure the accuracy of one's information contributes to one's ability to use one's information and is dependent upon (to some degree) the ability to control access to one's personal information. This entwinement is not unique to informational capabilities. If we look at Nussbaum's list, capabilities like the ability to be healthy and the ability to live a normal lifespan potentially impact one's ability to achieve all the rest of the capabilities on her list. A person in a vegetative state, for instance, cannot build relationships, own property, participate politically, etc. As such, some of the examples I will use to discuss this capability will touch on or be relevant to the capabilities already discussed.

There are numerous examples of the importance of the ability to ensure the accuracy of one's personal information ranging from the accuracy of one's digital medical records (**Health, Life, Bodily Integrity**) to the odd case of Constantin Reliu—a Romanian man who was incorrectly declared dead after having lived abroad for several decades. Despite appearing in court *in person*, he was told that his appeal was too late, and he would have to remain officially deceased. Speaking of the impact, Reliu said, "I am officially dead, although I'm alive...I have no income and

⁴¹ Anna Lodinová, "Application Of Biometrics As A Means Of Refugee Registration: Focusing on UNHCR's Strategy," *Development, Environment and Foresight* 2, no. 2 (2016): 91—100.

because I am listed dead, I can't do anything [emphasis mine].”⁴² This inaccuracy and Reliu's inability to correct it prevents him from voting (**Political Control Over One's Environment**), accessing health services (**Life, Health**), owning property or working (**Material Control Over One's Environment**). While Reliu's case may seem like a surreal edge case, less dramatic examples are common, such as individuals having difficulty scrubbing erroneous incidents from their credit history (**Material Control Over One's Environment**) or struggling to recover from identity theft.⁴³

The more a society uses personal information for, the more important it is to be able to ensure its accuracy. China's new social credit system (SCS) is a case in point. This mandatory program assigns points based on what the government considers good behaviour and docks points for what the government considers bad behaviour in four areas: government affairs, judicial affairs, social activities, and commercial behaviours.⁴⁴ Data sources include, but are not limited to, financial records, tax records, social media, and travel information. This data is collected by disparate sources and then shared and integrated into a centralized system.⁴⁵ Examples of bad behaviour include bad driving, buying too many video games, and bribing officials.⁴⁶ People with low scores can be banned from traveling by train and plane, prevented from leaving the country (**Bodily Integrity**), barred from hotels,

⁴² Shaun Walker, “Romanian Court Tells Man He Is Not Alive,” *The Guardian*, March 16, 2018, <https://www.theguardian.com/world/2018/mar/16/romanian-court-tells-man-he-is-not-alive>.

⁴³ Charlene Jennett, Sacha Brostoff, Miguel Malheiros, and M. Angela Sasse, “Adding Insult to Injury: Consumer Experiences Of Being Denied Credit,” *International Journal of Consumer Studies* 36 (2012): 549-555. doi:10.1111/j.1470-6431.2012.01120.x.

⁴⁴ Fan Liang, Vishnupriya Das, Nadiya Kostyuk, Muzammil M. Hussain, “Constructing a Data-Driven Society: China's Social Credit System as a State Surveillance Infrastructure,” *Policy & Internet* 10, no 4 (2018): 415-453.

⁴⁵ Ibid.

⁴⁶ Mara Hvistendahl, “Inside China's Vast New Experiment in Social Ranking,” *Wired*, Dec. 14, 2017, <https://www.wired.com/story/age-of-social-credit/> (accessed April 24, 2018).

have their internet speed throttled, lose out on certain types of jobs (**Material Control of One's Environment**), and possibly be publically shamed (**Affiliation**).⁴⁷

Inaccuracies may occur because of slander, human error, buggy code, or because the data scheme cannot account for the complexities of real life. I will sidestep the question of whether this is a good or bad social program. What is important for my purposes in this chapter is that when such a program is in place, it is important that one has the ability to ensure the information being used is accurate.

This case is similar to the Aadhaar case in a number of ways. In both examples, ensuring the accuracy of one's information is important. In the case of Aadhaar, however, the personal information that is used is more limited in scope and it is more likely the information will be accurate—one's biometric information should stay mostly static over time. By contrast, the SCS is using information from a diversity of sources, that information will need to be standardized as it is aggregated, and then will be run through algorithms to adjust one's credit score—at least for now these algorithms are not publicly available.⁴⁸ Each step could potentially erode the accuracy of one's information. As such, while the ability to ensure the accuracy of one's information is important in many contexts, the SCS highlights the importance of this capability and the ways in which it may be more difficult to achieve in the future as big data-driven systems become more common, automated, and potentially opaque.

While one might want to argue that it is simply better for one's information to actually be accurate than to have the ability to ensure that it is accurate, there are cases where it might be in one's advantage for one's personal information to be inaccurate. For example, someone who has served time in prison may not want their criminal record to show up on their Facebook page, a shorter individual may not mind that their dating profile adds a few inches to their height, or a human rights worker in China may be better off if the SCS says they were playing video games while they were actually meeting with a pro-democracy activist. This is not to say that it is unproblematic for people to make things up or to say that they should have the ability to deliberately falsify personal information. Rather, it is to say that one's wellbeing

⁴⁷ Ibid.

⁴⁸ Fan Liang et al., "Constructing a Data-Driven Society: China's Social Credit System as a State Surveillance Infrastructure."

depends on having the ability to ensure one's information is accurate, not necessarily on it actually being accurate.

4.2.4 The Ability to Live in an Environment with an Adequate Level of Cyberhealth

The final informational capability I will highlight is *the ability to live in an environment with an adequate level of cyberhealth*. As discussed in Chapter 3, this ability is essential to being able to rely on critical infrastructure, such as water treatment plants, emergency services, and chemical plants. Additionally, in the digital age, it is critical to the proper functioning of digital voting machines (**Political Control of One's Environment**), hospital infrastructure (**Health, Life**), government systems like Aadhaar (see 4.2.2), and one's ability to communicate with friends (**Affiliation**). While this ability is less important in societies which do not rely heavily on digital networks, in regions like North America and Europe the reliance of critical infrastructure and bureaucracy on digital networks means the cyberhealth of one's environment impacts one's ability to achieve nearly all of Nussbaum's fundamental capabilities.⁴⁹

These four informational capabilities are merely illustrative of how digital information and its use, control, accessibility, and accuracy can impact one's ability to achieve Nussbaum's fundamental capabilities. Other potentially important informational capabilities include the ability to avoid algorithmic bias, the ability to communicate confidentially, and literacy. In the next section, I will more formally define informational wellbeing and discuss a number of potential applications.

⁴⁹ It is worth noting that this is a capability—like the ability to live in a malaria-free environment—where it is unclear that it is the capability that matters as much as achieving a specific functioning. Would anyone be better off living in a malarial zone or in an environment with poor cyberhealth? This question is also a feature of some of the capabilities on Nussbaum's list, such as health. For further discussion on this matter see: Olsaretti, Serena. "Endorsement and Freedom in Amartya Sen's Capability Approach."

4.3 A Theory of Informational Wellbeing

In the last section, I argued that given the ubiquity of digital technologies and their role in 21st century bureaucracies, one's ability to achieve the fundamental capabilities listed by Nussbaum depends on one's ability to achieve certain informational capabilities. Given this importance, there is value in a concept of 'informational wellbeing' that can act as a shorthand for all of the various ways informational capabilities impact our ability to achieve fundamental capabilities. I will define informational wellbeing as:

An individual has a high degree of informational wellbeing to the extent they have achieved the 'informational capabilities and functionings' (e.g. the ability to control access to their personal information, etc.) which are necessary to achieve fundamental human capabilities (e.g. the ability to live a normal life span, etc.).

This is certainly not the only way one could define a concept of informational wellbeing, but it has a number of features that make it useful for both policymakers and technology producers seeking to promote personal wellbeing.

First, the theory is flexible enough to work in a wide variety of cultural contexts. While I have used Nussbaum's list of fundamental capabilities, one could substitute another list of fundamental capabilities without altering the definition I have presented above. Nussbaum has even acknowledged each entry on her list can be "more concretely specified in accordance with local beliefs and circumstances."⁵⁰ The same can be said for informational capabilities and functionings. In some contexts, internet access may be central to one's ability to use one's information, in other contexts having access to a telephone or the postal service may be sufficient. This theory of informational wellbeing does not assume that just because someone lives in a less technologically advanced society that they necessarily have worse informational wellbeing. In this sense, my concept of cyberhealth can be understood as a subset of a broader notion of information system robustness. My focus on cyberhealth, as opposed to this broader notion, has largely been driven by my interest in the huge growth of digital networks and how they make these capabilities more central to flourishing.

⁵⁰ Martha Nussbaum, *Women and Human Development: The Capabilities Approach*, 77.

Second, the theory is sufficiently non-technical to be understood by policymakers, politicians, technology producers, and citizens. This makes it easier for social scientists to tailor the theory to local contexts and for local stakeholders to participate in the process of specifying valuable capabilities and functionings. This is not to say that technical measures (e.g. encryption strength, bandwidth) should be avoided when measuring one's ability to achieve certain functionings. In some cases, technical assessments may be the best way to measure a given informational capability.

Third, the theory is measurable. Measurability is important for policymakers and technology producers. Without being able to quantify impacts to one's wellbeing, it is difficult to determine proportional responses and to set appropriate funding levels. Furthermore, my theory of informational wellbeing can likely be measured using data available to civilians. Examples include data related to technological literacy, access to and usage of the internet, what technological services people use, and network reliability. If my theory relied upon non-publicly accessible data, its utility as a policy tool would decrease as it would be challenging to verify that assessments of informational wellbeing were defensible and accurate. While intelligence agencies like the NSA or GCHQ may have greater visibility into the security of our personal information, without sophisticated spy tools one can still assess whether the services an individual uses follow security standards, whether these services have had known breaches, whether an individual has strong passwords and two-factor authentication in place, and whether their passwords have shown up in databases of stolen information. Much of this data is already collected for different purposes.

Fourth, my theory takes into account the *perception gap*. By the perception gap I mean that we generally cannot sense changes to the state of our personal information. If one breaks one's arm, one can immediately perceive that one's physical wellbeing has decreased. Similarly, the concept of mental wellbeing, is largely based on the quality of one's immediately perceptible mental experience. By contrast, one could have one's identity stolen, but only realize the harm when applying for a loan, checking one's bank account, or attempting to log into an online account that has been compromised. The 'perception gap' suggests that objective list theories, like the capability approach, are better suited to measuring informational

wellbeing than approaches which rely on one reporting one's affect or satisfaction, such as the Positive and Negative Affect Schedule (PANAS).

One might want to object to the notion that one's wellbeing can be affected despite one not perceiving the change.⁵¹ While this might be true for assessing the experiential quality of one's life,⁵² informational wellbeing is not intended to describe this aspect of wellbeing. Rather it is a standard of advantage—a way to make comparison's between people to guide political action.⁵³ In this context, objective measures are appropriate. These four features are not necessary conditions for a theory of informational wellbeing—a theory would not have to be flexible or measurable or non-technical to count as a plausible account of informational wellbeing. However, they are positive features which make my approach a useful tool for assessing the impact of technology policies and products on people's lives.

Within the context of the Public Cyberhealth LoA, this theory provides value in three primary and interrelated ways:

- 1) It helps one craft technology policies, respond to cyber threats, and design technology products in a manner which improves wellbeing.
- 2) It strengthens the normative justification for states to ensure that certain public goods for cyberhealth are adequately produced.
- 3) It helps one to determine at which point a public good for cyberhealth has been adequately produced.

I will discuss each in turn, although this discussion should only be treated as an introduction to these uses. Further research will be needed to determine precisely how policymakers and producers may use this theory.

⁵¹ For a summary of various objective and subjective accounts of wellbeing see Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 2006), Appendix I.

⁵² T.M. Scanlon, *What We Owe to Each Other*, (Cambridge, MA: The Belknap Press of Harvard University Press, 1998), 112.

⁵³ Wolff and De-Shalit, *Disadvantage*.

4.4 Value of the Theory

4.4.1 Creating Policies, Interventions, and Products Which Improve Wellbeing

The first use of informational wellbeing within the Public Cyberhealth LoA is to assess how technology products and policies impact wellbeing. This is the first step to being able to intentionally design technology policies and products that improve personal wellbeing, which, in turn, is essential to fulfilling the goal of the Public Cyberhealth LoA—to *promote cyberhealth (network robustness and resilience) as part of a broader goal of promoting human health and wellbeing.*

If one was only interested in the impacts of poor cyberhealth and informational use on health, one could use existing (albeit controversial) metrics like QALYs. However, as I have demonstrated that digital information and its use, control, etc. affect our lives in a myriad of ways, a new approach which takes these impacts into account is needed. Using the theory of informational wellbeing laid out in this chapter, one can systematically think through how a given policy, intervention, or product may impact one's ability to achieve fundamental capabilities. Depending on one's purpose, there are a number of ways one can make these assessments. As an example, consider the 'right to erasure'—a key part of the EU's General Data Protection Regulation (GDPR).⁵⁴

As background, the GDPR is a set of data regulations that brings all companies operating within the EU under one set of rules. The regulations seek to give individuals greater control over how their personal information is used and ensure companies are operating on a level playing field.⁵⁵ The right to erasure—Article 17 of the GDPR—requires companies to delete personal information, under

⁵⁴ "Right to Erasure," Information Commissioner's Office, accessed 30 January 2019: <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/right-to-erasure/>.

⁵⁵ European Commission, "2018 Reform Of EU Data Protection Rules," European Commission, https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en#background (accessed March 8, 2019).

certain circumstances, if an individual requests that they do so.⁵⁶ Requests should generally be complied with if the data is no longer needed, if the individual revokes consent for the data to be used, if a person objects to the information being used in a way where they are being profiled, to comply with a legal obligation, and if the information was collected illegally.⁵⁷ Companies may refuse the request if the information is needed for archiving or public health purposes in the public interest, exercising the freedom of expression, establishing a legal defence, and where there is a conflicting compliance obligation.⁵⁸ The ‘right to erasure’ is an updated version of what was previously known as ‘the right to be forgotten.’⁵⁹

The first way one can use the theory of informational wellbeing to assess the impact of technology policies on personal wellbeing is to systematically consider how a given technology policy will impact each fundamental capability on Nussbaum’s list. In the table below, I have provided an example of what that could look like in the case of the right to erasure:

Table 5

Fundamental Capability	Impact of Right to Erasure
Life	Likely minimal impact for most people. Potentially significant for those with certain jobs (e.g. human rights workers, some journalists).
Health	Likely minimal impact for most people. Potentially significant for those with certain jobs (e.g. human rights workers, some journalists).
Bodily Integrity	Empowers individuals to remove personal photos, medical information from the internet.

⁵⁶ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data and repealing Directive 95/46/EC (General Data Protection Regulation), <http://data.europa.eu/eli/reg/2016/679/oj>.

⁵⁷ Ibid.

⁵⁸ Ibid.

⁵⁹ Ibid.

Senses, Imagination, and Thought	None or minimal.
Emotions	None or minimal.
Practical Reason	None or minimal.
Affiliation	Empowers individuals to remove embarrassing information from the internet. Removes barriers to appearing in public without shame.
Other Species	None.
Play	None.
Control Over One's Environment	May contribute to one's ability to speak freely and avoid unreasonable search and seizure.

Another way would be to build a canonical list of informational capabilities that are instrumental to achieving fundamental capabilities, and then assess a policy's impact on those informational capabilities. As an example, in the table below I will assess the right to erasure's impact on the four capabilities discussed in this chapter:

Table 6

Informational Capability	Impact of Right to Erasure on Informational Capabilities
Ability to Control Access to One's Personal Information	Makes it easier for individuals to revoke access to one's personal information by 1) providing mechanism for individuals to delete information and 2) requiring companies to comply with requests.
Ability to Ensure the Accuracy of One's Personal Information	Makes it easier for individuals to delete inaccurate information.
Ability to Use One's Personal Information	Potentially increases ability by preserving the security of private information which can be used for other purposes. Potentially increase ability of the individual to sell their information to other parties as it is less readily available to third parties.
Ability to Live in an Environment with Sufficient Cyberhealth	Minimal impact.

These two tables set out at a schematic level how we might assess policies. Of course, full assessment would require further operationalizing these concepts, and specifying modes of measurement. These tasks are outside the scope of this work, but there are

useful examples in the literature. Two of these include Wolff and De-Shalit's version of the "York Model," which uses both objective and subjective assessments to measure each relevant functioning,⁶⁰ and Sen's work with Mahbub ul Haq creating the Human Development Index (HDI). The first is a more detailed assessment of advantage, while the second is a very high-level index that measures three metrics intended to be indicative of one's general capability—life expectancy, years of schooling, and per capita gross national income.⁶¹

In addition to assessing policies, one can also then use the theory of informational wellbeing to *design* policies, products, infrastructure, and responses to cyber threats which can improve personal wellbeing. If certain informational capabilities are central to one's ability to achieve fundamental capabilities, one can design policies or products with the promotion of those informational capabilities as an explicit goal from conception. For example, if one wanted to create a laptop that improved an individual's wellbeing, one could design the product to have the security and accessibility features that enable one to achieve valuable informational capabilities. Or in the case of malware mitigation efforts, such as the work of the Conficker Working Group discussed in Chapter 1, one could compare how various mitigation strategies would impact one's informational wellbeing by performing the types of simple assessments illustrated above.

As this approach is not how technology policymakers and product designers typically create policies and products today, it represents another context where using an LoA grounded in the philosophy of public health will be helpful, given that public health policies often use more complicated, human-centred metrics to judge the success of various policies and interventions. One example of such a metric is the Quality Adjusted Life Year (QALY). The QALY attempts to combine "the effects of health interventions on mortality and morbidity into a single index," and thus take into effect the quality of one's life in addition to the quantity of an individual's life for

⁶⁰ Wolff and De-Shalit, *Disadvantage*, 110-118.

⁶¹ United Nations Development Programme, "Human Development Index (HDI)," United Nations Development Programme: Human Development Reports, <http://hdr.undp.org/en/content/human-development-index-hdi> (accessed 12 October 2018).

assessing the value of interventions.⁶² Researchers in the field of public health are also currently exploring the feasibility and value of a Well-being Adjusted Life Year, which uses wellbeing, rather than just health, to judge quality of life.⁶³ This is not to imply that QALYs are unproblematic. For example, some argue that QALYs are limited in the health-benefits they capture and that they do not take into account existing social inequalities.⁶⁴ However, I would argue that it is precisely because metrics like QALYs raise complicated practical and theoretical questions that public health experts may be well-suited to thinking through the complexities of operationalizing a theory of informational wellbeing.

4.4.2 Strengthening the Justification for the Production of Public Goods for Cyberhealth

The second way in which my theory of informational wellbeing is useful within the Public Cyberhealth LoA is that it can help states determine which public goods for cyberhealth are most worth producing. Here we can return to the example of the mitigation of malware, one of the public goods for cyberhealth discussed in Chapter 1. While containing malware like Conficker may be valuable for states for numerous reasons (e.g. economic, strategic), the theory of informational wellbeing can help provide a way for states to assess how a given threat might impact individuals' ability to achieve fundamental capabilities. While states may not be justified in protecting or promoting *all* of the fundamental capabilities on Nussbaum's list, many of the fundamental capabilities listed by Nussbaum (e.g. the ability to live a normal lifespan, ability to access healthcare, and the ability to avoid discrimination) do traditionally fall within the purview of the state. Normally, in cases like the Conficker infection the

⁶² Paul Kind, Jennifer Elston Lafata, Karl Matuszewski, and Dennis Raisch, "The Use of QALYs in Clinical and Patient Decision-Making: Issues and Prospects," *Value in Health* 12, supplement 1 (2009): S27-S30.

⁶³ John Brazier and Aki Tsuchiya, "Improving Cross-Sector Comparisons: Going Beyond the Health-Related QALY," *Applied Health Economics and Health Policy* 13, 6 (2015): 557–565.

⁶⁴ Sarah J. Whitehead and Shehzad Ali, "Health Outcomes in Economic Evaluation: the QALY and Utilities," *British Medical Bulletin* 96, no. 1 (2010): 5–21.

problem is measured in terms of metrics like the number of infected machines, which only hint that there may ultimately be human costs. By helping states think through how a million-machine botnet has the potential to impact one's informational and fundamental capabilities, the theory can help states have a clearer sense of the scope of the threat. This, in turn, can help a state to determine its proper role in mitigation and prevention efforts.

States could use similar assessments for determining the value of any public good for cyberhealth, including basic research, standards, and regulatory regimes. If a given public good, such as the creation of tighter standards for critical infrastructure, improves individuals' informational wellbeing (and by extension one's ability to achieve certain fundamental capabilities), then there is (*prima facie*) a stronger justification to promote this good than simply correcting for market failure alone.

4.4.3 Determining Levels of State Support

Third, in addition, to helping states determine which informational capabilities and functionings they may be justified or obligated to expand access to, the theory of informational wellbeing can help states determine the appropriate level of support. Let us assume a state is trying to ensure individuals have access to healthcare, and as part of this responsibility the state is justified in promoting cyberhealth more generally due to its role in the health system (e.g. availability of medical records, security of devices, etc.). The state must then determine what is an adequate degree of cyberhealth. The theory of informational wellbeing provides at least one answer to this question by connecting informational capabilities, such as living in an environment with adequate cyberhealth, to fundamental capabilities, such as the ability to be healthy. In this context, an adequate level of cyberhealth is the level which enables individuals to achieve the fundamental capability of health. While this answer will not be the only factor a state uses to determine the adequate level of cyberhealth, it can help guide funding decisions and network standards.

Thresholds like this are also potentially important for limiting overinvestment or overly restrictive policies. For instance, in this chapter I have described a number of reasons why it may be important to control access to one's personal information, including being able to appear in public without shame and avoid discrimination. While this informational capability may be generally valuable, one should not assume

that *all* personal information is equally precious. It may be that certain kinds of personal information are relatively unimportant to one's ability to achieve fundamental capabilities (e.g. one's height), and therefore they need not be highly protected. When I worked in the technology industry, I frequently had to push back against company lawyers who wanted all forms of personal information, no matter how insignificant, treated as if it were someone's social security number. The theory of informational wellbeing presented in this chapter would have been a valuable tool for determining which pieces of personal information truly deserved a heightened level of protection and which could be used more freely.⁶⁵ It must be noted that one of the benefits of the capabilities approach is that it allows one to make these assessments based on local context. If in a given state one's height is used as a valuable identifying piece of information, then controlling access to that piece of information would be a valuable functioning which would count as part of an assessment of one's capabilities.⁶⁶

4.5 Generalizability

One concern one may have about this theory is that its utility is limited to those subscribing to a eudaimonist account of wellbeing, i.e. a theory which equates wellbeing with some notion of human flourishing. Eudaimonism is most famously associated with Aristotle, although more recent accounts have been developed by

⁶⁵ A number of the ways that personal information is protected (e.g. encryption, anonymisation) can make it more difficult to work with. For example, information may need to be decrypted before it can read by a person, which depending on the number of records can take significant computing resources.

⁶⁶ In regard to the importance of local context, Sen is fond of referencing Adam Smith's discussion of necessities and luxuries. While acknowledging that the Greeks and Roman's got by just fine without linen garments, Smith argues that in 18th century Europe "a creditable day-labourer would be ashamed to appear in publick without a linen shirt, the want of which would be supposed to denote that disgraceful degree of poverty, which, it is presumed, no body can well fall into without extreme bad conduct." Adam Smith, *An Inquiry into the Nature and Causes of the Wealth of Nations, Vol II*, eds. R. H. Campbell and A. S. Skinner (Oxford: Clarendon Press, 1976), 870.

Thomas Hurka (perfectionism),⁶⁷ Stephen Darwall,⁶⁸ and Richard Kraut (developmentalism).⁶⁹ Typically, in eudaimonist theories of wellbeing the conditions of human flourishing are placed on a list (like Nussbaum's), which is why these theories are often referred to as objective list theories. One might argue that hedonists, who argue that wellbeing is the positive balance of pleasure and pain, and subjectivists, who argue that wellbeing is one's ability to satisfy one's preferences, will not buy into the concept of informational wellbeing as a valid account of wellbeing.⁷⁰

While the capabilities approach is an objective list approach to wellbeing, one need not subscribe to an eudaimonist account of wellbeing to use the theory as a way of measuring advantage. As Sen says, "Quite different specific theories of value may be consistent with the capability approach, and share the common feature of selecting value-objects from functionings and capabilities."⁷¹ While Nussbaum's list represents one value system, another list of fundamental capabilities reflecting a different value system could easily be substituted. If one wanted to emphasize pleasurable experiences or desire fulfilment, one could select different fundamental capabilities to reflect these values. One example of the former is Martin Binder's attempt to bridge the gap between the capability approach and research into subjective wellbeing. In his approach, the capabilities which are relevant for wellbeing are 'Subjective Well-being Capabilities'—capabilities that enable "individuals to pursue and achieve happiness."⁷²

⁶⁷ Thomas Hurka, *Perfectionism*.

⁶⁸ Stephen Darwall, *Welfare and Rational Care*, (Princeton, NJ: Princeton University Press, 2002).

⁶⁹ Richard Kraut, *What Is Good and Why: The Ethics of Well-being*, (Cambridge, MA: Harvard University Press, 2007).

⁷⁰ Anna Alexandrova, *The Science of the Philosophy of Wellbeing*, Appendix A.

⁷¹ Sen, "Capability and Well-Being," 49.

⁷² Martin Binder, "Subjective Well-Being Capabilities: Bridging the Gap Between the Capability Approach and Subjective Well-Being Research," *Journal of Happiness Studies* 15, no. 5 (2014): 1197-1217.

Hedonists and subjectivists will not accept that the achievement of a set of capabilities is constitutive of wellbeing, however, they can still accept that the capabilities approach is a useful way of indirectly promoting and measuring wellbeing for policy purposes. For example, consider a subjectivist who believes that wellbeing is equivalent to desire satisfaction. While capabilities like the ability to be healthy and the ability to avoid discrimination might not be constitutive of wellbeing, they nonetheless help enable people to satisfy their preferences. As it may be easier or more politically acceptable to craft policies that promote the ability to be healthy (i.e. access to healthcare) than to craft policies that directly satisfy people's preferences (whatever that might entail), the subjectivist can find value in the capabilities language while denying that capabilities constitute wellbeing. Similarly, for measurement purposes, it may be easier to measure if people have the ability to be healthy or the ability to avoid discrimination than to measure if they are satisfying their preferences.

In the specific case of digital networks, there is an additional reason to believe that the hedonist and subjectivist may find value in the capabilities approach—the 'perception gap.' As a reminder, by the perception gap I mean our inability to immediately perceive the status of our personal information. For example, one is not immediately able to perceive that one's identity is stolen. As such, until one is aware of this fact it will not impact one's affect or one's self-reported preference satisfaction. Yet, it still makes sense to say this individual is worse off than if their information were secure. As such, while admitting that hedonists and subjectivists will likely object to treating Nussbaum's list of capabilities as an account of wellbeing, I believe the theory of informational wellbeing, which I have derived from her theory, should be broadly acceptable in policy contexts as a measure of advantage.

4.6 Informational Wellbeing as a Fundamental Capability

In the last section, I explored what informational wellbeing might mean for hedonists and subjectivists. In this section, I will narrow my focus to those who subscribe to eudaimonist theories similar to Sen and Nussbaum's. By this I mean eudaimonists who believe the achievement of a set of capabilities, such as those described by Nussbaum, constitutes wellbeing. Specifically, I will argue that for these eudaimonists

the ability to achieve an adequate degree of informational wellbeing should be considered a fundamental capability on par with the capabilities on Nussbaum's list.

4.6.1 Health and Informational Wellbeing

The first reason that eudaimonists of the Sen/Nussbaum variety should think of informational wellbeing as a fundamental capability is that informational wellbeing is partly constitutive of health. While people may not accept that all the capabilities on Nussbaum's list are essential, the ability to be healthy is widely (if not universally) considered to be valuable.⁷³ Even Sen, who has resisted making a fixed list of valuable capabilities, frequently mentions the capability to be healthy as valuable.⁷⁴

The argument in this section is an extension of the argument in Section 3.4 regarding the coupling of artificial parts to biological systems.⁷⁵ As a reminder, in Chapter 3, I argued that when artificial parts are coupled to a biological system, then we should consider them a part of the organism, and when a coupled artificial part functions poorly we should consider this reduction in functionality a pathology.⁷⁶ While this may be intuitive in the case of artificial hips, I argued that we should think

⁷³ See Section 3.1.

⁷⁴ Amartya Sen, *Inequality Reexamined*, 40.; Martha Nussbaum, *Creating Capabilities* (Cambridge, MA: The Belknap Press of Harvard University Press, 2011), 19-20.

⁷⁵ While there are competing sets of coupling criteria, I argued for a modified version of the criteria used by Clark and Chalmers in their work on the extended mind. An artificial part may be considered coupled to a biological system if: 1) it is performing a task associated with that biological system (e.g. the pacemaker regulates heart rhythm), 2) it is constantly available, 3) its contribution to the proper functioning of the system is readily and directly provided, and 4) the contribution is automatically endorsed/accepted by the system in question. (Clark and Chalmers, "The Extended Mind.")

⁷⁶ This assumes that the other necessary criteria of a pathology are met. In the case of Boorse, this would be that the reduction in function impacts survivability and one's ability to reproduce. In the case of Wakefield, the reduction of function would have to be considered harmful.

of digital devices with networked capabilities in the same way. Examples of such digital devices include digital pacemakers and defibrillators.

When digital devices which use digital networks are considered a part of one's biological systems, informational wellbeing—in particular the ability to live in an environment with an adequate degree of cyberhealth—is not merely instrumental to health, but partly constitutive of what it means to be healthy. There are many things which are instrumental to health, such as poverty. While being poor or rich may impact one's health, one's wealth is not constitutive of one's health. Likewise, one can argue that many aspects of cyberhealth are instrumental to health. One example may be the cyberhealth of a water treatment facility. If the water treatment facility's network is hacked or fails due to a software bug, individuals may experience poor health outcomes, but the cyberhealth of the facility would not be constitutive of an individual's health. However, as digital pacemakers should be thought of as a part of the organism to which they are coupled, the cyberhealth of the device is not only a determinant, but partly constitutive of the organism's health status. For biotechnologies that require the use of digital networks and cloud computing, then the cyberhealth of those networks and cloud services will also in part constitute what it means to be healthy.

4.6.2 Informational Wellbeing and the Self

While the above argument is rather straightforward, the second reason to consider informational wellbeing a fundamental capability is more complex and will require a little set-up. The second reason to consider informational wellbeing a fundamental capability is that it is central to what I will call the 'capability to be a liberal agent.' While this capability is not listed by Nussbaum, it is central to Sen and Nussbaum's conception of human flourishing.⁷⁷ Specifically, I will argue that having an adequate degree of informational wellbeing partly constitutes one aspect of this capability—the ability to define one's self. First, I will define what I mean by the 'capability to be a liberal agent' in the context of Sen and Nussbaum's eudaimonism. Then, I will

⁷⁷ While acknowledging that the term liberal can mean different things to different theorists, Nussbaum explicitly cites liberal theorists like Kant, Mill, Adam Smith, and T.H. Greene as philosophical influences on the capabilities approach. Nussbaum, *Creating Capabilities*, 123-143.

explore the role information plays in the definition of the self using Daniel Dennett's narrative conception of the self. Finally, I will argue that when the self is conceived of in informational or narrative terms, then informational wellbeing partly constitutes the ability to define one's self and, by extension, the ability to be a liberal agent.

4.6.2.1 *The Capability to be a Liberal Agent in Sen and Nussbaum's Eudaimonism*

Comprehensively exploring and defining the notion of liberal agency in Sen and Nussbaum's eudaimonism is a dissertation of its own, but a comprehensive analysis is not needed for my purpose. While Sen and Nussbaum do not give us a complete conception of the person, we can piece together a sketch of this conception by looking at how they describe the capabilities approach and by considering the kinds of capabilities they think are valuable.

First, let us consider the motivating principles of the capabilities approach as listed by Nussbaum. The first of these principles is that individuals are always to be treated as ends and never means. The capabilities approach is primarily focused on *individual* wellbeing, not collective wellbeing. Nussbaum underscores this point by arguing that the main question to ask when comparing societies is: "What is each person able to do and to be?"⁷⁸ The second principle Nussbaum lists is freedom or choice. This is the primary motivating factor behind focusing on capabilities as opposed to functionings. As Nussbaum says, "It [the capabilities approach]...commits itself to respect for people's powers of self-definition."⁷⁹ One example of the importance of choice in Sen and Nussbaum's eudaimonism is Sen's example of a man who is fasting. In his example, there are two starving men. The first man is choosing to fast, while the second does not have access to food. Sen argues that while both men may have the same level of health, it is necessary to account for the fact that the fasting man has chosen to starve when determining the two men's overall wellbeing.⁸⁰ For Sen and Nussbaum the freedom to choose how one lives has intrinsic value.⁸¹

We also see the importance of freedom to Sen and Nussbaum's conception of human flourishing in our second source—the set of capabilities on Nussbaum's list.

⁷⁸ Nussbaum, *Creating Capabilities*, 18.

⁷⁹ *Ibid.*, 18.

⁸⁰ Amartya Sen, *Inequality Reexamined*, 52-53.

⁸¹ Nussbaum, *Creating Capabilities*, 25.

As a reminder, the capabilities Nussbaum lists represent what a person needs to secure a ‘minimally flourishing life.’⁸² Below are a number of capabilities from Nussbaum’s list that are relevant for the current discussion:

- Being able to use imagination and thought in connection with experiencing and producing self-expressive works and events of one’s own choice, religious, literary, musical, and so forth.
- Being able to use one’s mind in ways protected by guarantees of freedom of expression with respect to both political and artistic speech, and freedom of religious exercise.
- Being able to search for the ultimate meaning of life in one’s own way...
- Being able to form a conception of the good and to engage in critical reflection about the planning of one’s life...⁸³

Each of these capabilities suggest that the freedom of choice Nussbaum values is an expansive freedom of choice. It is not the freedom to simply choose between two options (e.g. the freedom to be a Catholic or a Protestant), but the freedom to decide the kind of life one wants to live broadly construed. This freedom is also central to Sen’s notion of flourishing. In his account of the capabilities approach, he refers to this freedom as “agency freedom”—the “freedom to achieve whatever the person, as a responsible agent, decides he or she should achieve.”⁸⁴

These two sources—the motivating principles of the capabilities approach and the kinds of capabilities Nussbaum selects—suggest that for Sen and Nussbaum flourishing requires that one is able to exercise the freedom to choose one’s own path in life based on one’s own desires. It is this capability that I am calling the ‘capability to be a liberal agent.’ This capability may, in turn, be comprised of multiple aspects. For instance, in order to be the kind of agent that can exercise “agency freedom,” one may need to have a certain degree of psychological robustness and a certain degree of autonomy. I will not attempt to list all of the aspects that comprise the capability to be

⁸² Ibid., 33.

⁸³ Nussbaum, *Women and Human Development: The Capabilities Approach*, 78-80.

⁸⁴ Amartya Sen, “Well-Being, Agency and Freedom: The Dewey Lectures 1984,” *The Journal of Philosophy* 82, No. 4 (1985): 203-204.

a liberal agent, but rather I will focus on one which is relevant to the notion of informational wellbeing—the ability to define one’s self.

While acknowledging that the meaning of the term ‘self’ can be elusive, by defining one’s self I mean the process of individuation—what makes a person one individual and not another. I do not mean this in a biological sense, i.e. the physical boundaries of the organism, but psychologically or narratively.⁸⁵ This is the self Daniel Dennett calls the “owner of record.”⁸⁶ It is the *you* that owns your desires, your values, your path in life. According to narrative conceptions of the self, the self is, as Floridi puts it, a “socio- or auto-biographical artifact”⁸⁷—a web of stories and facts which gravitate around a ‘centre of narrative gravity.’

In the following section, I will provide an account of the self based on Daniel Dennett’s narrative theory of the self. Then, I will argue that when the self is conceived of in narrative terms, having an adequate degree of informational wellbeing partly constitutes the ability to define one’s self. From this claim it follows that informational wellbeing is important to one’s capability to be a liberal agent, which I have just argued is central to Sen and Nussbaum’s conception of flourishing.

The theory of the self I introduce below is similar to Floridi’s informational conception of the self (which is also based on Dennett’s work).⁸⁸ However, while Floridi attempts to show how the self—and nearly everything else in the world—can be entirely conceived of in informational terms, I aim to demonstrate the weaker claim that certain informational capabilities are central to one’s ability to define one’s self. While it may seem like I am making a grand metaphysical claim about the nature of selves, I believe, like Floridi, that we should think of the informational conception

⁸⁵ Daniel Dennett, *Consciousness Explained* (London: Penguin Books, 1993), 418.

⁸⁶ *Ibid.*, 418.

⁸⁷ Luciano Floridi, “The Informational Nature of Personal Identity,” *Minds & Machines* 21 (2011), <https://libsta28.lib.cam.ac.uk:2090/10.1007/s11023-011-9259-6>.

⁸⁸ Informational conceptions of the self need not be based on narrative conceptions of the self. For example, Locke’s conception of the self is grounded in the continuity of consciousness. Floridi argues that this approach to the self is also fundamentally informational in nature as consciousness, thoughts, and memories can also ultimately be reduced to states of information and information processing.; Floridi, *The Ethics of Information*, 211-260.

of the self as an LoA, not as the only way in which to conceive of the self. As a reminder, a LoA is a way of modelling the world for the purpose of answering specific questions. This is the same conceptual tool I used to describe the Cybersecurity and Public Cyberhealth frameworks in Chapter 2. In this light, the informational conception of the self is a way of describing the self using variables like personal information, narratives, and other people's beliefs about a person, while downplaying phenomena like consciousness, the soul, and the continuity of mental states, which feature prominently in other conceptions. By saying that this informational conception of the self is a LoA, I am not saying that it is merely a convenient metaphor. LoAs can be assessed on their coherence and utility, and if found to be useful and coherent in the long run, one should accept, modestly and provisionally, that they accurately describe reality. I will explore this connection between LoAs and truth in greater depth in the Conclusion to this dissertation.

If we are seeking to understand how information and its use, control, accuracy, and accessibility impact how we define ourselves, then it is useful to think of the self in informational terms. However, note that there are good reasons to use an informational conception of the self apart from my account of informational wellbeing. Informational conceptions of the self (in the form of narrative conceptions of the self) have been around for decades, with Paul Ricoeur and Daniel Dennett developing their (surprisingly similar) theories in the 1980s.⁸⁹ Informational theories of the self are appealing because informational artefacts like pictures, stories, journals, medical records, and the like do intuitively seem to play an important role in the process of self-definition. This has become increasingly clear in the digital age with the advent of online identities and digital avatars.⁹⁰

While I will use Dennett and Floridi's theories to illustrate the importance of information to the process of self-definition, it is important to note that my argument is not dependent on their specific approaches. Information plays a central role in all

⁸⁹ Paul Ricoeur, *Oneself as Another (Soi-même Comme un Autre)*, trans. Kathleen Blamey, (Chicago: University of Chicago Press, 1992). Published version of Ricoeur's 1986 Gifford Lectures; Daniel Dennett, "The Origins of Selves," *Cogito*, 3 (1989): 163-73.

⁹⁰ Sherry Turkle, *Life on the Screen: Identity in the Age of the Internet* (New York: Simon & Schuster, 1995).

narrative theories of the self, and thus my central point would not fundamentally change if I replaced Dennett or Floridi's theory with another narrative conception of the self.

4.6.2.2 Dennett's Narrative Conception of the Self

Dennett grounds his concept of the self in the narratives that consciousness produces when making sense of information. In this view, consciousness works like an algorithm automatically generating narratives when presented with related pieces of information.⁹¹ These narratives gravitate around what Dennett calls a 'centre of narrative gravity'—which like the centre of gravity of a hoop is both real and yet abstract. As Dennett puts it:

A self, according to my theory, is... an abstraction defined by the myriads of attributions and interpretations (including self-attributions and self-interpretations) that have composed the biography of the living body whose Center of Narrative Gravity it is.⁹²

It is important to clarify that while Dennett uses terms like narrative and biography in the quote above, the "the myriads of attributions and interpretations" should be thought of as informational building blocks of narrative, rather than narratives as such. For example, an individual's blood pressure is not a narrative on its own, but when it is added to other information about a person—one's personal and family medical history, place of birth, diet, etc.—it contributes to a certain narrative about one's life.

Within Dennett's theory, all manner of information can contribute to how one's self is defined: your memories, informational artefacts (emails, text messages, letters, photographs), your physical characteristics, and the things others say about you. It is important to note that as an abstraction, the centre of narrative gravity is inseparable from the 'myriad of attributions and interpretations' which define it. This leads Dennett to occasionally refer to the self as an "organization of information."⁹³

While Dennett focuses on one's centre of narrative gravity from one's own perspective, I believe it makes sense to identify two types of centres of narrative

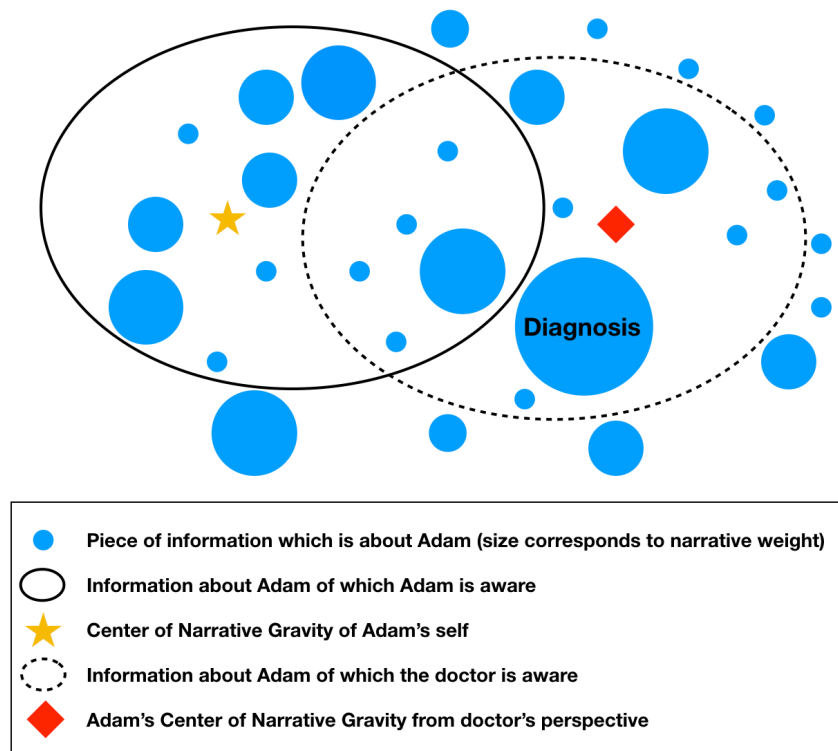
⁹¹ Dennett, "The Origins of Selves."

⁹² Daniel Dennett, *Consciousness Explained*, 426-427.

⁹³ Ibid., 430.

gravity. The first is one's centre of narrative gravity from one's own perspective, while the second is one's centre of narrative gravity from the perspective of others. Within this view, everyone who interacts with you may have a slightly different sense of your self. For the capability to be a liberal agent, the centre of gravity from one's own perspective is of primary importance. However, this third-person perspective is important as other people's conception of a person can influence how one thinks of oneself. As Dennett says, "Parents, friends, and even enemies may all contribute to the image of 'what it means to be me.'"⁹⁴

To illustrate this point, let us say a doctor knows that an individual called Adam is HIV positive, but the doctor has yet to share that information with Adam. Without this information Adam has a certain conception of his self as a completely healthy middle-aged person who will live to a ripe old age. On the one hand, it is clear that until he becomes aware of his diagnosis, this information does not change Adam's sense of self. But on the other hand, from the perspective of the doctor, Adam's centre of narrative gravity has changed. From the doctor's perspective, Adam's diagnosis may in fact be one of Adam's defining features. The figure below illustrates this example:



⁹⁴ Daniel Dennett, "The Origins of Selves," *Cogito*, 3 (1989):

In this illustration, the blue dots represent all known information about Adam. Some of that information is known to Adam, and within that cloud of information there is a centre of narrative gravity which represents Adam's self. There is then another set of information which the doctor knows about Adam. This set has its own centre of narrative gravity and represents Adam's self from the perspective of the doctor. While one can say that Adam's sense of self is not directly impacted by information of which he is not aware, this information may still influence his centre of narrative gravity insofar as those that are aware of it may treat him differently. Additionally, while Adam is not currently aware of the diagnosis, he could very easily become aware of it, at which point it would dramatically shift his centre of narrative gravity. As such, while information one is not aware of may not *yet* influence one's sense of self, one may still have an interest in how that information is treated if it will either impact how others treat one, or if one is likely to become aware of that piece of information in the future.

We can also use this diagram to think about the role of misinformation, by which I mean information which is not true. Let us say that Adam becomes aware of his HIV diagnosis, but that the diagnosis is incorrect. Until that error is corrected, that incorrect diagnosis will shift both the doctor's concept of who Adam is and Adam's own sense of self. Even after the diagnosis is revealed to be incorrect, Adam will always have the experience of having lived with a HIV diagnosis. That experience will indelibly shift his own sense of self and likely his centre of narrative gravity from a third person perspective as well. Some lies are in fact stickier than the truth. For instance, some parents will tell their children that they (the child) are worthless, until the child accepts this as fact. This piece of misinformation may have a long-lasting or permanent impact on that child's sense of self even after the child recognizes intellectually that this belief is untrue. As such, it seems reasonable to conclude that both information and misinformation can contribute to one's centre of narrative gravity from a first- and third-person perspective.

To summarize up to this point, I first argued that the capability to be a liberal agent, i.e. the ability to choose one's own path in life, is fundamental to Sen and Nussbaum's conception of human flourishing. I then argued that this capability was comprised of a number of aspects, including potentially some measure of autonomy and psychological robustness. I argued that one of these aspects is the ability to define one's self. I then defined what I meant by "self" using Dennett's narrative theory of

the self. This LoA describes selves as a ‘centre of narrative gravity’ within an organization of personal information. Finally, I argued that misinformation can also influence one’s centre of narrative gravity from both a first- and third- person perspective. I will now argue that when the self is conceived of in informational terms, informational wellbeing in part constitutes the ability to define one’s self, and, therefore, partly constitutes the capability to be a liberal agent in the sense that is fundamental to Sen and Nussbaum’s concept of human flourishing.

4.6.2.3 *Informational Wellbeing and the Ability of Self-definition*

In the first half of the chapter, I defined informational wellbeing as one’s ability to achieve the informational capabilities and functionings which enable one to achieve fundamental capabilities. In this section, I will demonstrate that a number of the informational capabilities discussed in the first half are not merely instrumental to the ability to define one’s self, but partly constitutive of this ability. These informational capabilities include the ability to control access to one’s personal information, the ability to ensure the accuracy of one’s personal information, and the ability to use certain kinds of ICTs. Note that by the ability to define one’s self I do not mean that one has complete autonomy to compose one’s narrative self. As Floridi says, “Most of our selves, understood as narratives, are written by other authors, what is left to the each of us to contribute must be carefully protected and fostered.”⁹⁵ Rather, this ability consists of having some reasonable ability to determine who one is as an individual apart from external influences. While what counts as reasonable is up for interpretation, I believe the discussion that follows will provide a clearer sense of what I have in mind.

The first informational capability which enables one to define one’s self is *the ability to control who has access to one’s personal information*. As Dennett said, how one’s parents, friends, and even enemies think about one influences how one defines one’s self.⁹⁶ While it is unreasonable to think one should have the ability to control all information about oneself, being able to limit access to certain types of personal

⁹⁵ Luciano Floridi, “On Human Dignity as a Foundation for the Right to Privacy,” *Philosophy & Technology*, (2016) 29: 307. <https://doi.org/10.1007/s13347-016-0220-8>.

⁹⁶ Daniel Dennett, “The Origins of Selves.”

information ensures others' opinions about oneself do not play an outsized role in one's self-definition. I believe the Lawrence case discussed in Section 4.2.1 is useful for illustrating this impact. As a reminder, Lawrence took nude photos of herself to share with her boyfriend as a way of strengthening their relationship. When these photos were stolen and leaked online, her sense of self changed as a result of the public shaming. Despite not feeling like the photos were something to be ashamed of, she began sobbing and had to run away when someone called her a role model for girls. As she said, "I can't believe somebody still feels that way after what happened."⁹⁷ The public shaming that accompanied the information theft dramatically shifted Lawrence's centre of narrative gravity such that she no longer saw herself as someone worthy of being a role model. Lawrence seems to have ultimately been able to reclaim the ability to define herself. Several years after the hack, she accepted a role which required her to do a nude scene. She described this choice as "taking something back"⁹⁸ and said of the experience, "I walked off that set *feeling like a different person* [emphasis mine]."⁹⁹ However, other victims of non-consensual pornography, according to Kamal and Newman, find themselves engaged in a "lifelong battle to preserve their integrity."¹⁰⁰

The second informational capability which enables one to define one's self is *the ability to ensure the accuracy of one's personal information*. To illustrate the value of the capability, let us stay with the case of Lawrence. In addition to having her nude photos leaked online, she is also the frequent target of fake pornography, including sophisticated machine learning generated videos referred to as deepfakes.¹⁰¹

⁹⁷ Smith, "Jennifer Lawrence Speaks Out About Reclaiming Her Body After Her Nude Photos Were Published Without Her Consent."

⁹⁸ Oprah Winfrey, "The Jennifer Lawrence Interview, by Oprah Winfrey."

⁹⁹ Smith, "Jennifer Lawrence Speaks Out About Reclaiming Her Body After Her Nude Photos Were Published Without Her Consent."

¹⁰⁰ Mudasir Kamal and William J. Newman, "Revenge Pornography: Mental Health Implications and Related Legislation."

¹⁰¹ Jon Sharman, "Pornhub and Twitter Ban Ai-Generated 'Deepfakes' Videos that Put Female Celebrities' Faces on Adult Actresses' Bodies," *The Independent*, Feb. 7, 2018, <https://www.independent.co.uk/life-style/gadgets-and-tech/pornhub-twitter->

While today these videos are easy enough to expose as fake, it is easy to imagine that in the future they may be virtually indistinguishable from genuine videos and impact one's self in similar ways as the leaking of real images or videos described above. Additionally, note that while I have focused on Lawrence in this section, this technology is also used to create fake non-consensual pornography of non-famous people and can be used in non-sexual contexts to create false narratives about a person (e.g. make it seem like they were somewhere they were not). In an age when many people take dozens of images of themselves a day and roughly 1 in 5 American adults between the ages of 25-34 have sent nude photos of themselves another person,¹⁰² ensuring the accuracy of one's personal information may not be as simple as searching one's memory for the truth or publishing a denial.

While in the previous two paragraphs I have used examples of other people's opinions about a person influencing that person's self, these informational capabilities are also valuable for influencing how non-human agents define an individual. Retail companies, financial institutions, search engines, political campaigns, social media websites, and, in some cases, courts all use personal information to develop profiles of who a given person is—essentially your self from the perspective of that agent. These profiles are then used to shape the advertisements, search results, and news stories one sees, and may determine whether a person receives a loan or is granted bail. In each of these cases, it is as if these companies and institutions are essentially communicating something back to you about who you are. It is as if they are saying: you are a person who would buy X product, you are a person who would like X news story, you are a person who cannot be trusted to pay back a loan. While in most cases the most significant harm will have nothing to do with one's ability to define one's self (e.g. being denied bail restricts one's freedom), we should not underestimate how in aggregate these machine generated versions of our selves can limit our ability to define ourselves.

deepfakes-ban-ai-celebrity-faces-porn-actress-bodies-emma-watson-jennifer-lawrence-a8199131.html (accessed Feb. 6, 2019).

¹⁰² Amanda Lenhart and Maeve Duggan, "Main Report: Couples, the Internet, and Social Media," *Pew Research Center*, Feb. 11, 2014, <http://www.pewinternet.org/2014/02/11/main-report-30/> (accessed Feb. 6 2019).

One particularly dramatic example, described by Ysabel Gerrard, is how social media sites will recommend pro-eating disorder content to individuals based on who is in their social network. Gerrard describes the experience saying, “[pro-eating disorder content] became almost inescapable once I was embedded in these spaces.”¹⁰³ This content would show up in her social media feeds, daily emails, and recommendations. Note that in this case the harm someone potentially faces is not just physical (developing an eating disorder), but to some degree the ability to define what one values and who one wants to be in life. Zeynap Tufekci has discussed this phenomenon in the context of YouTube’s recommendation engine leading to political radicalization.¹⁰⁴ I do not mean to imply that this type of profiling is always harmful or unjustified, but rather that as it becomes more commonplace, the informational capabilities discussed in this chapter become increasingly important to being able to define one’s self and choose one’s own path in life.

The third and final informational capability I will highlight is *the ability to use certain kinds of ICTs*. Floridi describes how certain ICTs influence one’s self, saying:

Obviously, any technology, the primary goal of which is to manage memories, is going to have an immense influence on how individuals develop and shape their own personal identities. It is not just a matter of mere quantity; the quality, availability, accessibility, and replaying of...personal memories may deeply affect who we think we are and may become.¹⁰⁵

One example of this kind of ICT would be the photo sharing site Flickr which, as a cost cutting measure, deleted millions of personal photos in March 2019.¹⁰⁶ The website, which at one point had close to 90 million users, was one of the most popular

¹⁰³ Ysabel Gerrard, “Beyond the Hashtag: Circumventing Content Moderation on Social Media,” *New Media & Society* 20, no. 12 (2008): 4505.

¹⁰⁴ Zeynap Tufekci, “YouTube, the Great Radicalizer,” *The New York Times*, March 10, 2018, <https://www.nytimes.com/2018/03/10/opinion/sunday/youtube-politics-radical.html> (accessed 14 April, 2019).

¹⁰⁵ Floridi, *The Ethics of Information*, 223.

¹⁰⁶ Kaitlyn Tiffany, “Flickr Will Soon Start Deleting Photos — and Massive Chunks of Internet History,” *Vox*, Feb. 6, 2019, <https://www.vox.com/the-goods/2019/2/6/18214046/flickr-free-storage-ends-digital-photo-archive-history> (accessed Feb. 7 2019).

ways people could store their personal photos as the company offered a terabyte of free storage. Within the narrative conception of the self, these photos are not merely property but data points which influence a person's centre of narrative gravity. While Flickr is just one site, in 2012 Webshots shut down deleting 690 million photos,¹⁰⁷ and there is no guarantee the massive social media companies of today will not follow suit. As the technology journalist Katie Notopoulos said at the time of Webshots demise, "It's not a stretch to imagine a day when all our words and images hosted on these services are removed as the companies collapse or morph. Friendster is now a video gaming service, MySpace is music streaming."¹⁰⁸ The articles describing Flickr's deletion of personal photos tend to lament the loss of internet history. While this may be a valuable observation, part of the value of the narrative conception of the self is that it highlights that this information loss can also profoundly impact individuals' self from both a first- and third-person perspective. As Floridi says, "one's informational sphere and one's personal identity are co-referential, or two sides of the same coin."¹⁰⁹ Additionally, it is worth noting that while in the case of Flickr this information is being purposefully deleted, this type of information can also be deleted or become inaccessible due to any number of threats to cyberhealth (human error, natural disasters, buggy code, malware). For instance, the same month that Flickr began its purge, MySpace accidentally deleted twelve years worth of music, images, and videos during a botched server migration.¹¹⁰

When one's self is conceived of as an organization of information, having access to technologies like social media, blogging platforms, and the like is a valuable functioning which contributes to one's ability to define one's self. While people in the

¹⁰⁷ Ibid.

¹⁰⁸ Katie Notopoulos, "Your Internet Photos Are Already Starting To Die," *BuzzFeed.News*, Oct. 3, 2012, <https://www.buzzfeednews.com/article/katienotopoulos/your-internet-photos-are-already-starting-to-die> (accessed Feb. 7, 2019).

¹⁰⁹ Floridi, *The Ethics of Information*, 244.

¹¹⁰ Kaitlyn Tiffany, "Myspace, Which Still Exists, Accidentally Deleted 12 Years' Worth Of Music," *Vox*, March 18, 2019, <https://www.vox.com/the-goods/2019/3/18/18271088/myspace-music-deleted-internet-archive-flickr-tumblr> (accessed April 11, 2019).

past were able to define their selves without these technologies, they become increasingly central to one's ability to define one's self as a greater share of personal information migrates from journals and photo albums to social media sites, from people's heads to cloud storage.

I will readily admit that much of the personal information which makes up one's narrative self is not the kind of digital information I have generally been discussing in this dissertation. While memories can be posted to blogs or recorded in digital photos, they can also simply exist in people's minds. While one may record one's emotional reactions, they also may just come and go without leaving a digital trace. Having said this, the ratio of digital information to non-digital information in one's 'cloud' of personal information is only growing. Additionally, this digital information is the only information to which non-human agents often have access. As such, while having a high degree of informational wellbeing is not synonymous with the ability to define one's self, it is becoming increasingly central to this ability.

At the start of this section, I argued that when networked biotechnologies are considered a part of a person, then informational wellbeing in part constitutes what it means to be healthy. In the second half, I made a similar argument—when the self is conceived of as an organization of information, then informational wellbeing in part constitutes what it means to be able to define one's self and, by extension, one's capability to be a liberal agent in the sense valued by Sen and Nussbaum. As both the capability to be healthy and the capability to be a liberal agent are central to Sen and Nussbaum's conception of human flourishing, they should also accept that 'having the ability to achieve a sufficient level of informational wellbeing' is a fundamental human capability on par with the other capabilities listed by Nussbaum.

4.7 Conclusion

In this chapter I defined a theory of informational wellbeing that can help one to identify how digital information and its use, control, accessibility, and accuracy impact an individual's wellbeing. Basing my theory on Sen and Nussbaum's capabilities approach to wellbeing, I defined informational wellbeing as:

An individual has a high degree of informational wellbeing to the extent they have achieved the 'informational capabilities and functionings' (e.g. the ability to control access to their personal information, etc.) which are

necessary to achieve fundamental human capabilities (e.g. the ability to live a normal life span, etc.).

This theory not only connects informational capabilities to an established theory of wellbeing, but it is flexible, measurable, and non-technical.

I argued that this theory is valuable in at least three primary ways: 1) it helps one to create technology policies, cyber responses, and digital products which improve wellbeing; 2) it strengthens the normative justification for states to produce the public goods for cyberhealth discussed in Chapter 1; and 3) it helps one to determine when a given public good for cyberhealth has been adequately provisioned.

In the second half of the chapter, I argued that those subscribing to Sen and Nussbaum's version of eudaimonism should consider the ability to achieve an adequate degree of informational wellbeing to be a fundamental capability on par with the other capabilities listed by Nussbaum. First, I argued that informational wellbeing in part constitutes what it means to be healthy. Then, I argued that it partly constitutes the ability to define one's self and, by extension, the capability to be a liberal agent.

Finally, note that this theory is particularly useful for understanding why the promotion of cyberhealth is a worthwhile goal at all. As mentioned in the introduction to this chapter, digital networks are valuable because of the information they contain and transmit. By connecting this information and its use, control, accuracy, and accessibility to wellbeing, I have not only provided a way to more fully understand the value of—and justification for—the public goods for cyberhealth discussed in Chapter 1, but also the specific tools like ethical review boards and the Intervention Ladder discussed in Chapter 2. While promoting health (the focus of Chapter 3) may be a worthwhile reason to promote cyberhealth on its own, this chapter helped to put that capability in context by illustrating the complicated and interconnected ways that cyberhealth and technology policy can impact one's life more broadly.

Conclusion

This work was motivated by two primary concerns. The first concern was that despite the ubiquity of digital technologies, we often struggle to articulate how these technologies impact our wellbeing—data breaches rarely leave dead bodies. The second concern was that despite digital networks playing an ever larger role in everyday life, the dominant cybersecurity paradigm seemed to serve business interests at the expense of society on the whole. In some cases, this cybersecurity paradigm even seemed to make things worse by encouraging an increasingly dangerous cyber arms race.

In this dissertation, I presented an alternative framework for conceptualizing the digital landscape called Public Cyberhealth. This framework, or level of abstraction (LoA), was inspired by the philosophy of public health and differs from the dominant cybersecurity approach in four primary ways. First, while the Cybersecurity LoA is focused on malicious attacks, the Public Cyberhealth LoA aims to promote network robustness and resilience more generally. While taking malicious attacks seriously (see Conficker case in Section 1.2), it also highlighted and suggested ways to address non-malicious points of failure, such as buggy code, natural disasters, and human error, as seen in the formal articulation of the LoA in Section 2.3. Second, while the Cybersecurity LoA is largely focused on business and military interests, the Public Cyberhealth LoA captures the myriad of ways in which network threats and interventions can impact individuals' health, wellbeing, and rights. Third, while the Cybersecurity LoA generally limits the role of the state to the protection of government information and the investigation of cybercrimes (see Section 2.2.2.1), the Public Cyberhealth LoA uses the philosophy of public health to establish the normative justification for—and ethical limits on—state interventions in cyberspace. This can most clearly be seen in my use of the LoA to define spheres of public and private cyberhealth in Section 2.4.1. And fourth, while the goal of the Cybersecurity LoA is to prevent unauthorized access to information and networks, the Public Cyberhealth LoA promotes cyberhealth as part of a broader goal of promoting health and wellbeing.

Throughout this dissertation, I have argued that viewing the digital landscape through this public health inspired lens would profoundly change how one thinks

about technology policy, product design, and potentially even the concepts of health and wellbeing. In Chapter 1, using the example of the Conficker worm, I demonstrated how this public health inspired approach could help policymakers understand the value of public goods for cyberhealth and the obligations that states, corporations, and individuals have to contribute to said goods. In Chapter 2, I illustrated how the philosophy of public health could help define spheres of public and private cyberhealth, and how public health tools like ethical review boards and the Intervention Ladder could help ensure that cyberhealth interventions are justifiable and proportional. In Chapter 3, I explored how various health related concepts like health inequality, bodily integrity, disease, and bodily certainty and doubt can help policymakers think in more sophisticated and nuanced ways about how poor cyberhealth impacts health. And in Chapter 4, I defined a theory of informational wellbeing, which enables policymakers to better assess the impact of technology policies and products on personal wellbeing. Together these various elements form a cohesive way of thinking about the digital landscape and can help policymakers craft more consistent technology policies that positively impact wellbeing and protect individuals' rights.

While the Public Cyberhealth LoA is useful for policymakers and product designers today, there are good reasons to believe that it will only become more useful in the future as digital technologies become increasingly sophisticated and ubiquitous. Programs like India's Aadhaar program and China's social credit system (see Sections 4.2.2 and 4.2.3) are new and continuing to evolve, but they suggest that in the future bureaucracies will increasingly rely on biometrics and AI systems to manage access to goods like health, education, and the ability to travel. While these programs can be plagued by traditional cybersecurity threats like malware, they can significantly impact one's health and wellbeing in a myriad of ways that have nothing to do with traditional matters of cybersecurity. Another example is the growing sophistication of medical devices. In Chapter 3, I argued that in some cases cyberhealth problems should be thought of as pathologies. In that chapter, I largely focused on digital pacemakers because there are relatively few internal digital devices that meet the coupling criteria and use digital networks. However, this class of device will likely become increasingly common as the technology becomes more reliable and AI systems develop new ways of using real-time device data for predicting health outcomes and modifying treatment on the fly.

Future Research

In this dissertation, I have presented an overview of the Public Cyberhealth LoA. While I have been able to describe its most fundamental features and a number of practical uses, many of the topics I have discussed are good candidates for further research. In many cases, this future research will need to be conducted in collaboration with experts from other fields, including economics, computer science, public health policy, technology policy, and development studies.

First, in regard to public goods, I believe there is value in a more thorough comparison between public goods for public health and public goods for cyberhealth. In Chapter 1, I demonstrated there is a broad similarity between these two classes of public goods, but more work can be done to understand the similarities and differences of specific goods. For instance, while Conficker was a kind of weakest-link problem (like infectious disease), not all kinds of malware follow this pattern. Additionally, further work must be conducted to determine precisely what obligations states, corporations, and individuals have to contribute to various public goods for cyberhealth—specifically we need a better sense of the likelihood and costs of various network threats and a clearer sense of the benefits of various interventions. This empirical work can help one determine which specific public goods for cyberhealth meet Klosko's criteria described in Section 1.3.2. While computer scientists and social scientists are better suited to this kind of empirical work, the Public Cyberhealth LoA can help ensure that this empirical research takes into account health and wellbeing impacts alongside financial costs and benefits.

The second area of future research is determining how best to operationalize the theory of informational wellbeing introduced in Chapter 4. I believe this theory might be the most useful contribution of this dissertation, but there is still much work to be done for the theory to be put to use. The next step to operationalize this theory is to develop a process for identifying the specific informational capabilities and functionings which constitute informational wellbeing in a given context.¹ Researchers from the field of development studies may be particularly helpful in this research, as the capabilities approach has most commonly been deployed in the development context.

¹ Alexandrova, *A Philosophy for the Science of Well-Being*.

The third area of future research is further fleshing out the variables, observables, and behaviours that comprise the Public Cyberhealth LoA. For the purposes of this dissertation, it would often have been counterproductive to describe examples in the formalized LoA language introduced in Chapter 2. Having said this, for specific practitioners, such as policymakers, technology producers, lawyers, social scientists, and IT professionals, there is value in having a more detailed LoA to guide their work. For example, if a product manager is designing a product using the Public Cyberhealth LoA, it would be very useful to understand specifically what threats should be accounted for, what health risks should be assessed, and what individual rights should be considered. A detailed LoA is even more critical for researchers who are looking to build models of the digital landscape in order to estimate the cost of various threats and the benefits of various interventions. As with operationalizing my theory of informational wellbeing, this work will best be accomplished by working in collaboration with practitioners and subject matter experts, including computer scientists, public health experts, and policymakers.

Lastly, there is substantial work to be done to figure out how best to incorporate the Public Cyberhealth LoA into the policymaking process. One of the main aspects of this work is determining how the Public Cyberhealth LoA can incorporate, or at least co-exist with, existing cybersecurity institutions and policies. A second, related aspect is determining how to weight informational wellbeing alongside the economic and strategic interests that I have largely ignored in this work. While public health policymakers have developed mechanisms for balancing health metrics (e.g. QALYs) and economic costs, there is less consensus about how to determine the economic value of wellbeing (let alone informational wellbeing).

LoAs, Coherence, and Truth

While the areas of research mentioned above will be the focus of research in the near future, there is also the longer-term project of assessing the coherence and utility of the Public Cyberhealth LoA over time. This work is not simply important for deciding if and how the LoA should be used as a policy tool, but for determining if the LoA is adequately describing the digital landscape. In this sense I follow William James in thinking that, “‘The true’, to put it very briefly, is only the expedient in the way of our

thinking... Expedient in almost any fashion; and expedient in the long run and on the whole, of course.”²

Putting to the side the larger debate about metaphysical realism, when one is speaking of complex human/object systems, it is reasonable to accept that one generally lacks direct experiential access to the system as a whole—one can hold an apple in one’s hands, but not the complicated relationship of humans and digital networks. To help us to understand these complicated systems, we can use LoAs to build models, make predictions, and answer specific questions. Floridi argues that LoAs are essentially interfaces which mediate between the world and epistemic agents, and in the case of complex systems these interfaces are all we have to determine how the world works. As such, he argues that the method of LoA supports a kind of ‘liminal realism,’ i.e. a realism which falls somewhere between an internal realism (rooted in conceptual schemes) and a strong metaphysical realism in which one can describe the world as is. While LoAs do not serve a mimetic function, through the creation of models one can generate reliable knowledge about the world.³

While our experience is mediated, one need not fall into a relativistic trap. As Hasok Chang argues in his work on ‘pragmatic realism’:

If our use of a theory has led to successful outcomes and not as a result of any strange accident or coincidence as far as we can see, then we can and should say, modestly and provisionally, that the relevant statements made in this theory are ‘true.’⁴

As such, I am ‘modestly and provisionally’ suggesting that if the Public Cyberhealth LoA is found to be useful and coherent ‘in the long run and on the whole,’ it should be accepted as an accurate (and not merely convenient) way of modelling the digital landscape. That is to say informational wellbeing really is part of what is ultimately good for a person, cyberhealth is a constitutive part of a person’s health status, and our selves are truly informational in nature.

² William James, *Pragmatism: A New Name for some Old Ways of Thinking*, Cambridge, MA: Harvard University Press, 1975, 106.

³ While the world may not directly knowable, it is ‘epistemically interactable.’ Floridi, *Philosophy of Information*, 370.

⁴ Hasok Chang, *Pragmatic Realism*, 118.

In thinking about assessing our two competing LoAs, it is important to recognize that Chang's 'pragmatic realism' and Floridi's 'liminal realism' encourages one to be more accepting of alternate approaches. As Chang argues:

In the absence of what else we might operationally mean by 'real', and with the recognition that this concept of reality is not something we can do without, we should have the courage to admit that a lot of different kinds of things are real, even if the concepts pointing to them belong to mutually incommensurable systems of practice.⁵

In the short term, we should have the open-mindedness to evaluate our competing LoAs without falling back on the lazy idea that the Cybersecurity LoA is the right frame simply because it is the dominate paradigm. If both approaches have utility and coherence, then we should accept both conceptions of the digital landscape as real. And, in the long run, if the Public Cyberhealth LoA is deemed more useful and coherent than the Cybersecurity LoA, we should be prepared to abandon that old paradigm as not only unhelpful, but also as essentially untrue.

⁵ Chang, *Pragmatic Realism*, 119

Bibliography

- “Fact on the Fall-Out.” *The Washington Post*, December 16, 1954: 20.
- “Fact Sheet: U.S.-EU Cyber Cooperation.” The White House, 26 March 2014.
<https://obamawhitehouse.archives.gov/the-press-office/2014/03/26/fact-sheet-us-eu-cyber-cooperation> (accessed May 18, 2017).
- “Frequently Asked Questions About the International Health Regulations,” World Health Organization, 2009, <http://www.who.int/ihr/about/FAQ2009.pdf>.
- “FY 2018 Budget in Brief.” United States Department of Homeland Security.
<https://www.dhs.gov/sites/default/files/publications/DHS%20FY18%20BIB%20Final.pdf>.
- “Intervention Ladder Informs Lords Behaviour Change Report.” Nuffield Council on Bioethics, July 19, 2011. <http://nuffieldbioethics.org/news/2011/intervention-ladder-informs-lords-behaviour-change-report> (accessed Dec. 7, 2018).
- “National Vaccine Injury Compensation Program.” Health Resources & Services Administration. October 2018. <https://www.hrsa.gov/vaccine-compensation/index.html> (accessed December 13, 2018).
- “Smallpox.” Centers for Disease Control and Prevention, July 12, 2017.
<https://www.cdc.gov/smallpox/index.html> (accessed Dec. 30, 2018).
- Acheson, Donald. *Independent Inquiry into Inequalities in Health Report*. The Stationary Office, 1998.
https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/265503/ih.pdf.
- Alexandrova, Anna. *The Science of the Philosophy of Wellbeing*. Oxford: Oxford University Press, 2017.
- Amazon.com. “Summary of the Amazon S3 Service Disruption.” Amazon.com.
<https://aws.amazon.com/message/41926/> (accessed Feb. 15, 2019).
- Anand, Sudhir. “The Concern for Equity in Health.” In *Public Health, Ethics, and Equity*, edited by Sudhir Anand, Fabienne Peter and Amartya Sen, 15-20. Oxford: Oxford University Press, 2004.
- Anderson, Ross; Richard Clayton; Éireann Leverett. “Standardisation and Certification Of Safety, Security And Privacy In The ‘Internet Of Things.’” Joint Research Centre. Luxembourg: Publications Office of the EU, 2018.

<https://publications.europa.eu/en/publication-detail/-/publication/80bb1618-16bb-11e8-9253-01aa75ed71a1/language-en>.

Arneson, Richard. "Perfectionism and Politics." *Ethics* 111 (2000): 37–63.

Arthur, Charles. "Naked Celebrity Hack: Security Experts Focus on iCloud Backup Theory." *The Guardian*, Sep. 1, 2014.

<https://www.theguardian.com/technology/2014/sep/01/naked-celebrity-hack-icloud-backup-jennifer-lawrence> (accessed Jan. 24, 2019).

Asada, Yukiko. *Health Inequality*. Toronto: University of Toronto Press, 2007.

Asllani, Arben, Charles Stephen White, and Lawrence Etkin. "Viewing Cybersecurity as a Public Good: The Role of Governments, Businesses, And Individuals." *Journal of Legal, Ethical and Regulatory Issues* 16, no. 1 (2013): 7-14.

Aylward, Bruce, Arnab Acharya, Sarah England, Mary Agocs, and Jennifer Linkins. "Polio Eradication." In *Global Public Goods for Health: Health Economic and Public Health Perspectives*, edited by Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager, 33-53. Oxford: Oxford University Press, 2003.

Babb, Jr., Stanley E. "Accuracy of Planetary Theories, Particularly for Mars." *Isis* 68, no. 3 (1977): 426-434.

Basu, Indrani. "AADHAAR: Fading Fingerprints Mean This Ageing Space Scientist Can't Care For His Son." *Huffington Post*, April 19, 2018.

https://www.huffingtonpost.in/2018/04/19/an-81-year-old-space-scientist-wants-the-supreme-court-to-save-senior-citizens-from-aadhaar_a_23414358 (accessed April 6, 2019).

Batchelder, Dennis et al. *Microsoft Security Intelligence Report*, volume 18. (2015).

<http://www.microsoft.com/security/sir/archive/default.aspx>.

Bauer, Johannes and Michel Van Eeten. "Cybersecurity: Stakeholder Incentives, Externalities, and Policy Options." *Telecommunications Policy* 33, no. 10 (2009): 706-19.

Berger, Joseph. "A Dam, Small and Unsung, Is Caught Up in an Iranian Hacking Case." *The New York Times*, 25 March 2016.

https://www.nytimes.com/2016/03/26/nyregion/rye-brook-dam-caught-in-computer-hacking-case.html?_r=0 (accessed Nov. 17, 2017).

- Binder, Martin. "Subjective Well-Being Capabilities: Bridging the Gap Between the Capability Approach and Subjective Well-Being Research," *Journal of Happiness Studies* 15, no. 5 (2014): 1197-1217.
<https://doi.org/10.1007/s10902-013-9471-6>.
- Boorse, Christopher. "A Rebuttal on Health." In *What Is Disease?*, edited by James M. Humber and Robert F. Almeder, 3-134. Totowa, NJ: Humana Press, 1997.
- Boorse, Christopher. "On the Distinction Between Disease and Illness." *Philosophy and Public Affairs* 5, no. 1 (1975): 49-68.
- Braveman, Paula. "Health Disparities and Health Equity: Concepts and Measurement." *Annual Review of Public Health* 27 (2006): 167-194.
<https://doi.org/10.1146/annurev.publhealth.27.021405.102103>.
- Brazier, John and Aki Tsuchiya. "Improving Cross-Sector Comparisons: Going Beyond the Health-Related QALY." *Applied Health Economics and Health Policy* 13, 6 (2015): 557-565.
- Broad, William J., John Markoff and David E. Sanger. "Israeli Test on Worm Called Crucial in Iran Nuclear Delay." *The New York Times*, Jan. 15, 2011.
<http://www.nytimes.com/2011/01/16/world/middleeast/16stuxnet.html>
 (accessed Nov. 17, 2017).
- Brown, Abbe, Shawn H. E. Harmon, Rory O'Connor, Sita Popat, and Sarah Whatley. "Body Extension And The Law: Medical Devices, Intellectual Property, Prosthetics And Marginalisation (Again)," *Law, Innovation and Technology* 10, no. 2 (2018). <https://doi-org.ezp.lib.cam.ac.uk/10.1080/17579961.2018.1526853>.
- Buchanan, Allen, Dan Brock, Norman Daniels, and Daniel Wikler. "Introduction." In *From Chance to Choice*, edited by Allen Buchanan, Dan Brock, Norman Daniels, and Daniel Wikler, 1-26. Cambridge: Cambridge University Press, 2000.
- Buchanan, James and Gordon Tullock. *The Calculus of Consent: Logical Foundations of Constitutional Democracy*. Ann Arbor: University of Michigan Press, 1962.
- Buchanan, James. *The Demand and Supply of Public Goods*. Chicago: Rand McNally & Company, 1968.
- Buchanan, James. *The Limits of Liberty: Between Anarchy and Leviathan*. Chicago: University of Chicago Press, 1975.

- Burri, Haran and David Senouf. "Remote Monitoring and Follow-Up of Pacemakers and Implantable Cardioverter Defibrillators." *Europace* 11, no. 6 (2009): 701–709.
- Carel, Havi. "Bodily Doubt," *Journal of Consciousness Studies* 20, Issue Nos. 7-8 (2013): 178-197.
- Centers for Disease Control and Prevention. "Public Health Ethics." Centers for Disease Control and Prevention, 10 May 2015.
<https://www.cdc.gov/od/science/integrity/phethics/> (accessed May 16, 2017).
- Chang, Hasok. "Pragmatic Realism." *Humanities Journal of Valparaíso*, no. 8 (2016): 107-122.
- Charney, Scott. "Collective Defense: Applying the Public-Health Model to the Internet." *Security & Privacy IEEE* 10, no. 2 (2012): 54-59.
- Chen, Lincoln, Tim G. Evans, and Richard A. Cash. "Health as a Global Public Good." In *Global Public Goods*, edited by Inge Kaul, Isabelle Grunberg, and Marc Stern, 284-305. New York; Oxford: Oxford University Press, 1999.
- Chiel, Ethan. "Here Are the Sites You Can't Access Because Someone Took the Internet Down." *Splinter*, Oct. 21, 2016. <https://splinternews.com/here-are-the-sites-you-cant-access-because-someone-took-1793863079> (accessed Feb. 15, 2019).
- Choucri, Nazli, Stuart Madnick, and Priscilla Koepke. "Institutions for Cyber Security: International Responses and Data Sharing Initiatives." Working Paper Cybersecurity Interdisciplinary Systems Laboratory. MIT, October 2016. <http://web.mit.edu/smadnick/www/wp/2016-10.pdf> (accessed May 29, 2017).
- Choudhary, Mishi. "Viewpoint: The Pitfalls Of India's Biometric ID Scheme." *BBC News*, April 23, 2018. <https://www.bbc.co.uk/news/world-asia-india-43619944>.
- Clark, Andy and David J. Chalmers. "The Extended Mind." *Analysis* 58, no. 1 (1998): 7–19.
- Clark, Andy and David Chalmers. "The Extended Mind." *InterAction* 8, no. 1 (2016): 48-64, <https://search.proquest.com/docview/1808003977?accountid=9851> (accessed March 28, 2019).

- Coggon, John. *What Makes Health Public?: A Critical Evaluation of Moral, Legal, and Political Claims in Public Health*. Cambridge: Cambridge University Press, 2012.
- Cohen, G. A. "Equality of What? on Welfare, Goods, and Capabilities." In *Quality of Life*, edited by Martha Nussbaum and Amartya Sen, 54-61. Oxford: Clarendon Press, 1993.
- Comptroller and Auditor General. "Investigation: WannaCry Cyber Attack and the NHS." National Audit Office, HC 414 Session 2017–2019, April 25, 2018. <https://www.nao.org.uk/wp-content/uploads/2017/10/Investigation-WannaCry-cyber-attack-and-the-NHS.pdf> (accessed Dec. 20 2018).
- Cooper, Rachel. "Disease," *Studies in History and Philosophy of Biological and Biomedical Sciences* 33, no. 2 (2002): 263-282.
- Crisp, Roger. "Well-Being." *Stanford Encyclopedia of Philosophy*. Sep. 6, 2017. <https://plato.stanford.edu/entries/well-being> (accessed 30 May 2018).
- Cummings, K. Michael and Robert N. Proctor. "The Changing Public Image of Smoking in the United States: 1964–2014." *Cancer, Epidemiology, Biomarkers & Prevention* 23, no. 1 (2014): 32-36.
- Darwall, Stephen. *Welfare and Rational Care*. Princeton, NJ: Princeton University Press, 2002.
- Dawson, Angus. "Herd Protection as a Public Good: Vaccination and our Obligations to Others." In *Ethics, Prevention, and Public Health*, edited by Angus Dawson and Marcel Verweij, 160-178. Oxford: Oxford University Press, 2007.
- Dennett, Daniel. "The Origins of Selves." *Cogito* 3 (1989): 163-73.
- Dennett, Daniel. *Consciousness Explained*. London: Penguin Books, 1993.
- Department of Health and Social Care. *Our Healthier Nation: A Contract for Health*. Cm 3852, Feb. 9, 1998. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/265721/title.pdf.
- Department of Homeland Security. "Critical Infrastructure Sectors." Department of Homeland Security, 11 July 2017. <https://www.dhs.gov/critical-infrastructure-sectors> (accessed Jan. 5 2018).
- Dercon, Stefan. "Risk, Poverty, and Vulnerability in Africa." *Journal of African Economies* 14, no. 4 (2005): 483-488.
- Dewey, John. *Logic: The Theory of Inquiry*. New York: Henry Holt & Co., 1938.

- Director General. "Global Strategy for the Prevention and Control of Noncommunicable Diseases." *World Health Organization*. A53/14, March 22, 2000. http://apps.who.int/gb/archive/pdf_files/WHA53/ea14.pdf (accessed December 13, 2018).
- Edwards, Charlie. *National Security for the Twenty-first Century*. London: Demos, 2007.
- Etherington, Darrell. "Amazon AWS S3 Outage Is Breaking Things For A Lot Of Websites And Apps." *TechCrunch*, Feb. 28, 2017. <https://techcrunch.com/2017/02/28/amazon-aws-s3-outage-is-breaking-things-for-a-lot-of-websites-and-apps/> (accessed April 2, 2017).
- European Commission. "2018 Reform Of EU Data Protection Rules." European Commission. https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en#background (accessed March 8, 2019).
- FBI's Internet Crime Complaint Center. "2017 Internet Crime Report." FBI's Internet Crime Complaint Center, May 7, 2017. https://pdf.ic3.gov/2017_IC3Report.pdf (accessed March 12, 2019).
- Fine, Paul, Ken Eames, and David L. Heymann. "Herd Immunity." *Clinical Infectious Diseases* 52, no. 7, April 1, 2011: 911–916. <https://doi.org/10.1093/cid/cir007>.
- Finkle, Jim. "J&J Warns Diabetic Patients: Insulin Pump Vulnerable to Hacking." *Reuters*, Oct. 4, 2016. <https://www.reuters.com/article/us-johnson-johnson-cyber-insulin-pumps-e/jj-warns-diabetic-patients-insulin-pump-vulnerable-to-hacking-idUSKCN12411L> (accessed Dec. 3, 2017).
- Floridi, Luciano. "The Informational Nature of Personal Identity." *Minds & Machines* 21 (2011). <https://libsta28.lib.cam.ac.uk:2090/10.1007/s11023-011-9259-6>.
- Floridi, Luciano. *The Ethics of Information*. Oxford: Oxford University Press, 2013.
- Floridi, Luciano. *The Philosophy of Information*. Oxford: Oxford University Press, 2011.
- Freyer, Lauren. "India's Biometric ID System Has Led To Starvation For Some Poor, Advocates Say." *NPR*, Oct. 1, 2018. <https://www.npr.org/2018/10/01/652513097/indias-biometric-id-system-has-led-to-starvation-for-some-poor-advocates-say?t=1538831766242>.
- Gakidou, Emmanuela; Christopher Murray; and Julio Frenk. "Defining and Measuring Health Inequality: An Approach Based on the Distribution of

- Health Expectancy.” *Bulletin of the World Health Organization* 78 (2000): 42-54.
- Garnham, Nicholas. “Amartya Sen's 'Capabilities' Approach to the Evaluation of Welfare: Its Application To Communications.” In *Beyond Competition: Broadening the Scope of Telecommunication Policy*, eds. Bare Cammaerts and Jean-Claude Burgelman, 25-36. Brussels: VUB University Press, 2000.
- Gerrard, Ysabel. “Beyond the Hashtag: Circumventing Content Moderation on Social Media.” *New Media & Society* 20, no. 12 (2008): 4492-4511.
- Gibbs, Samuel and Lois Beckett. “Dark Web Marketplaces AlphaBay and Hansa Shut Down.” *The Guardian*, July 20, 2017.
<https://www.theguardian.com/technology/2017/jul/20/dark-web-marketplaces-alphabay-hansa-shut-down> (accessed Nov. 26, 2018).
- Gigler, Björn-Sören. “‘Informational Capabilities’- The Missing Link for the Impact of ICT on Development.” *The World Bank*, working paper series no. 1, (March 2011).
<http://documents.worldbank.org/curated/en/227571468182366091/pdf/882360NWP0Box30series0no10March2011.pdf> (accessed Oct. 24, 2018).
- Goel, Vindu. “Facebook Tinkers With Users’ Emotions in News Feed Experiment, Stirring Outcry.” *New York Times*, June 29, 2014.
<https://www.nytimes.com/2014/06/30/technology/facebook-tinkers-with-users-emotions-in-news-feed-experiment-stirring-outcry.html> (accessed March 4, 2019).
- Greenberg, Andy. “Russian Hacker False Flags Work—Even After They're Exposed.” *Wired*, February 27, 2018. <https://www.wired.com/story/russia-false-flag-hacks/> (accessed March 22, 2019).
- Head, John G. *Public Goods and Public Welfare*. Durham, N.C.: Duke University Press, 1975.
- Health Resources and Services Administration. “National Vaccine Injury Compensation Program.” Health Resources and Services Administration, October 2018. <https://www.hrsa.gov/vaccine-compensation/index.html> (accessed Feb. 26, 2019).
- Hern, Alex. “Hacking Risk Leads to Recall of 500,000 Pacemakers Due to Patient Death Fears.” *The Guardian*, Aug. 31, 2017.

- <https://www.theguardian.com/technology/2017/aug/31/hacking-risk-recall-pacemakers-patient-death-fears-fda-firmware-update> (accessed Dec. 3, 2017).
- Humboldt, Wilhelm von. *The Limits of State Action*. Cambridge Studies in the History and Theory of Politics. Cambridge: Cambridge University Press, 1969.
doi:10.1017/CBO9781316036372.
- Hurka, Thomas. *Perfectionism*. Oxford: Oxford University Press, 1993.
- Husak, Douglas N. "Drugs, Crime and Public Health: A Lesson From Criminology." In *Criminal Law, Philosophy and Public Health Practice*, edited by A. M. Viens, John Coggon, and Anthony Kessel, 42-61. Cambridge: Cambridge University Press, 2013. <https://doi.org/10.1017/CBO9781139137065.003>.
- Hvistendahl, Mara. "Inside China's Vast New Experiment in Social Ranking." *Wired*, Dec. 14, 2017. <https://www.wired.com/story/age-of-social-credit/> (accessed April 24, 2018).
- IEEE Standards Association. "P7010 - Wellbeing Metrics Standard for Ethical Artificial Intelligence and Autonomous Systems." IEEE Standards Association. <https://standards.ieee.org/project/7010.html> (accessed Feb. 17, 2019).
- International Telecommunications Union. "ICT Facts and Figures 2017." *International Telecommunications Union*, July 2017.
<https://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2017.pdf> (accessed Jan. 13, 2019).
- Jackson, Amy Berman. United States District Court for the District of Columbia. Memorandum Opinion. Misc. Action no. 15-1394 (ABJ), MDL docket no. 2664. https://ecf.dcd.uscourts.gov/cgi-bin/show_public_doc?2015mc1394-117.
- James, William. *Pragmatism: A New Name for some Old Ways of Thinking*. Cambridge, MA: Harvard University Press, 1975.
- Jennings, Bruce. "Public Health and Civic Republicanism." In *Ethics, Prevention, and Public Health*, edited by Angus Dawson and Marcel Verweij, 30-58. Oxford : New York: Clarendon Press ; Oxford University Press, 2007.
- Jennett, Charlene, Sacha Brostoff, Miguel Malheiros, and M. Angela Sasse. "Adding Insult To Injury: Consumer Experiences Of Being Denied Credit."

- International Journal of Consumer Studies* 36 (2012): 549-555.
doi:10.1111/j.1470-6431.2012.01120.x.
- John, Stephen. "Why 'Health' Is Not a Central Category for Public Health Policy." *Journal of Applied Philosophy* 26, no. 2 (2009): 129-143.
- Kallberg, Jan and Rosemary A. Burk. "Cyberdefense as Environmental Protection—The Broader Potential Impact of Failed Defensive Counter Cyber Operations." In *Conflict and Cooperation in Cyberspace*, edited by Panayotis Yannakogeorgos and Adam Lowther, [Page Numbers]. Boca Raton, London, Paris: Taylor and Francis (2014).
- Kamal, Mudasir and William J. Newman. "Revenge Pornography: Mental Health Implications and Related Legislation." *Journal of the American Academy of Psychiatry and the Law Online* 44, no. 3 (2016): 359-367.
- Kaplan, Fred. *Dark Territory*. New York: Simon and Schuster Paperbacks, 2016.
- Kashner, Sam. "Both Huntress and Prey." *Vanity Fair*, November 2014.
<https://www.vanityfair.com/hollywood/2014/10/jennifer-lawrence-photo-hacking-privacy>.
- Kaul, Inge, Isabelle Grunberg, and Marc A. Stern. "Defining Global Public Goods." In *Global Public Goods: International Cooperation in the 21st Century*, edited by Inge Kaul, Isabelle Grunberg, Marc A. Stern, 2-19. New York ; Oxford: Oxford University Press, 1999.
- Kaul, Inge, Isabelle Grunberg, and Marc A. Stern. Introduction to *Global Public Goods: International Cooperation in the 21st Century*, edited by Inge Kaul, Isabelle Grunberg, Marc A. Stern. New York ; Oxford: Oxford University Press, 1999.
- Kaul, Inge, Isabelle Grunberg, and Marc Stern, eds. *Global Public Goods*. New York; Oxford: Oxford University Press, 1999.
- Kephart, Jeffrey O., Steve R. White, and David M. Chess. "Computers and Epidemiology." *Spectrum IEEE* 30, no. 5 (1993): 20-26.
- Kim, Jim Yong, Aaron Shakow, Arachu Castro, Chris Vanderwarker, and Paul Farmer. "Tuberculosis Control." In *Global Public Goods for Health: Health Economic and Public Health Perspectives*, edited by Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager, 54-72. Oxford: Oxford University Press, 2003.

- Kind, Paul, Jennifer Elston Lafata, Karl Matuszewski, and Dennis Raisch. "The Use of QALYs in Clinical and Patient Decision-Making: Issues and Prospects." *Value in Health* 12, supplement 1 (2009): S27-S30.
- Kirk, Susan. "How Children and Young People Construct and Negotiate Living with Medical Technology." *Social Science & Medicine* 71, no. 10 (2010): 1796-1803. <https://doi.org/10.1016/j.socscimed.2010.07.044>.
- Klosko, George. "Presumptive Benefit, Fairness, and Political Obligation." *Philosophy & Public Affairs* 16, no. 3 (1987): 241-259.
- Klosko, George. *The Principle of Fairness and Political Obligation*. Lanham: Rowman & Littlefield Publishers, Inc., 1992.
- Kramer, Andrew. "Russia, This Time the Victim of a Cyberattack, Voices Outrage." *New York Times*, May 14, 2017, <https://www.nytimes.com/2017/05/14/world/europe/russia-cyberattack-wannacry-ransomware.html> (accessed May 28, 2017).
- Kraut, Richard. *What Is Good and Why: The Ethics of Well-being*. Cambridge, MA: Harvard University Press, 2007.
- Krebs, John. "The Importance of Public Health Ethics." *Bulletin of the World Health Organization* 86, no. 8 (2008): 577-656.
- Krekel, Bryan. "Capability of the People's Republic of China to Conduct Cyber Warfare and Computer Network Exploitation." The US-China Economic and Security Review Commission. 2009. <http://www.thing.net/~rdom/ucsd/CyberwarUSChina2009.pdf>.
- Kunreuther, Howard, and Geoffrey Heal. "Interdependent Security." *Journal of Risk and Uncertainty* 26, no. 2 (2003): 231-49.
- Ledyard, John O. "Market Failure." In *The New Palgrave Dictionary of Economics*. Edited by Palgrave Macmillan. London: Palgrave Macmillan, 2008. https://doi.org/10.1057/978-1-349-95121-5_1052-2.
- Le Grand, Julian. *Equity and Choice: An Essay in Economics and Applied Philosophy*. London: HarperCollins Academic, 1991.
- Lenhart, Amanda and Maeve Duggan. "Main Report: Couples, the Internet, and Social Media." *Pew Research Center*, Feb. 11, 2014. <http://www.pewinternet.org/2014/02/11/main-report-30/> (accessed Feb. 6 2019).

- Lewens, Tim. *The Biological Foundations of Bioethics*. Oxford: Oxford University Press, 2015.
- Liao, Matthew. "Do You Have a Moral Duty to Leave Facebook?." *The New York Times*, November 24, 2018.
<https://www.nytimes.com/2018/11/24/opinion/sunday/facebook-immoral.html>
 (accessed Jan. 10, 2019).
- Liang, Fan, Vishnupriya Das, Nadiya Kostyuk, Muzammil M. Hussain. "Constructing a Data - Driven Society: China's Social Credit System as a State Surveillance Infrastructure." *Policy & Internet* 10, no 4 (2018): 415-453.
- Lloyd, Alun and Robert May. "How Viruses Spread Among Computers and People." *Science* 292 (2001): 1316-1317.
- Lodinová, Anna. "Application Of Biometrics As A Means Of Refugee Registration: Focusing on UNHCR's Strategy." *Development, Environment and Foresight* 2, no. 2 (2016): 91—100.
- Lohr, Steve. "Digital Divide Is Wider Than We Think, Study Says." *New York Times*, Dec. 4, 2018. <https://www.nytimes.com/2018/12/04/technology/digital-divide-us-fcc-microsoft.html> (accessed Dec. 5, 2018).
- Lovelace Jr., Berkeley and Antonio José Vielma. "Friday's Third Cyberattack on Dyn 'Has Been Resolved,' Company Says." *CNBC*, Oct. 21, 2016.
<https://www.cnn.com/2016/10/21/major-websites-across-east-coast-knocked-out-in-apparent-ddos-attack.html> (accessed Feb. 15, 2019).
- Luciano Floridi. *The Philosophy of Information*. Oxford: Oxford University Press, 2011.
- Madon, Shirin. "Evaluating The Developmental Impact Of E-Governance Initiatives: An Exploratory Framework." *The Electronic Journal of Information Systems in Developing Countries* 20, no. 5 (2004): 1-13.
- Martin, Rachel. "Stephen Hawking Gets A Voice Upgrade." *Weekend Edition Sunday*, Dec. 7 2014. <https://www.npr.org/2014/12/07/369108538/stephen-hawking-gets-a-voice-tech-upgrade?t=1549637874457> (accessed Feb. 8, 2019).
- Mayo Clinic Staff. "Bradycardia." *Mayoclinic.org*, Aug. 23, 2017.
<https://www.mayoclinic.org/diseases-conditions/bradycardia/symptoms-causes/syc-20355474> (accessed Dec. 6, 2017).

- Mayo Clinic Staff. "Conjoined Twins." Mayoclinic.org, March 7, 2018.
<https://www.mayoclinic.org/diseases-conditions/conjoined-twins/symptoms-causes/syc-20353910> (accessed March 6, 2019).
- McGuinness, Damien. "How a Cyber Attack Transformed Estonia." *BBC News*, April 27, 2017. <https://www.bbc.co.uk/news/39655415> (accessed Feb. 28, 2019)
- McMichael, Anthony, Colin Butler, and Michael Ahern. "Global Environment." In *Global Public Goods for Health: Health Economic and Public Health Perspectives*, edited by Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager, 94-118. Oxford: Oxford University Press, 2003.
- McMillan, Robert. "Sharks Want To Bite Google's Undersea Cables." *Wired*, Aug. 15, 2014. <https://www.wired.com/2014/08/shark-cable/> (accessed Dec. 19, 2018).
- Meade, James E. *The Theory of Economic Externalities: The Control of Environmental Pollution and Similar Social Costs*. Geneva: Sijthoff-Leiden, 1973.
- Mill, John Stuart. *On Liberty*, in *Utilitarianism and Other Writings*. 1859. Edited by Mary Warnock. Glasgow: Collins, 2003.
- Mill, John Stuart. *On Liberty: in Focus*. Edited by John Gray and G. W. Smith, London and New York: Routledge, 1991.
- Milman, Oliver. "Six weeks after Hurricane Maria, Puerto Ricans Still Waiting for Help from Fema." *The Guardian*, November 9, 2017.
<https://www.theguardian.com/world/2017/nov/09/six-weeks-after-hurricane-maria-puerto-ricans-still-waiting-for-help-from-fema> (accessed December 14, 2018).
- Moeler, Michael. *Minimal Morality: A Multilevel Social Contract Theory*. Oxford: Oxford University Press, 2018.
- Mulligan, Deirdre K. and Fred B. Schneider. "Doctrine for Cybersecurity." *Dædalus* 140, no. 4 (2011): 70-92.
- Newman, Lily Hay. "Medical Devices Are The Next Security Nightmare." *Wired*, Mar. 2, 2017. <https://www.wired.com/2017/03/medical-devices-next-security-nightmare> (accessed Nov. 20 2017).
- Nordhaus, William. "Paul Samuelson and Global Public Goods." *Samuelsonian Economics and the Twenty-First Century*, edited by Michael Szenberg, Lall

- Ramrattan, and Aron Gottesman, 88-98. Oxford: Oxford University Press, 2006. doi:10.1093/acprof:oso/9780199298839.003.0006.
- Notopoulos, Katie. "Your Internet Photos Are Already Starting To Die." *BuzzFeed.News*, Oct. 3, 2012.
<https://www.buzzfeednews.com/article/katienotopoulos/your-internet-photos-are-already-starting-to-die> (accessed Feb. 7, 2019).
- Nozick, Robert. *Anarchy, State, and Utopia*. Oxford: Basil Blackwell, 1974.
- Nuffield Council on Bioethics. "Public Health: Ethical Issues." London: Nuffield Council on Bioethics, 2007.
- Nussbaum, Martha. "Aristotelian Social Democracy." In *Liberalism and the Good*, edited by R. Douglas, G. Mara and H. Richardson, 203–252. New York, NY: Routledge, 1990.
- Nussbaum, Martha. *Creating Capabilities*. Cambridge, MA: The Belknap Press of Harvard University Press, 2011.
- Nussbaum, Martha. *Women and Human Development: The Capability Approach*. Cambridge: Cambridge University Press, 2000.
- Olsaretti, Serena. "Endorsement and Freedom in Amartya Sen's Capability Approach." *Economics & Philosophy* 21, no. 1 (2005): 89-108.
- Olsen, Hilgunn. "Open Drug Scenes and Police Strategies in Oslo, Norway." *Journal of Scandinavian Studies in Criminology and Crime Prevention* (2017): 141-156.
- O'Neill, Patrick Howell. "Conficker Worm Still Spreading Despite Being Nearly 10 Years Old." *Cyberscoop*, Dec. 8, 2017.
<https://www.cyberscoop.com/conficker-trend-micro-2017/> (accessed Feb 27, 2019).
- Orend, Brian. "Fog in the Fifth Dimension: The Ethics of Cyber-War." In *The Ethics of Information Warfare*, edited by Luciano Floridi and Mariarosaria Taddeo, 3-24. Cham: Springer, 2014.
- Oughton, Edward J., Mike Hapgood, Gemma Richardson, Ciaran Beggan, Alan Thomson, Mark Gibbs, Catherine Burnett, Trevor Gaunt, Markos Trichas, Rabia Dada, and Richard Horne. "A Risk Assessment Framework for the Socioeconomic Impacts of Electricity Transmission Infrastructure Failure Due to Space Weather: An Application to the United Kingdom." *Risk Analysis*, November 8, 2018. [https://doi: 10.1111/risa.13229](https://doi:10.1111/risa.13229).

- Parfit, Derek. *Reasons and Persons*. Oxford: Oxford University Press, 2006.
- Parfitt, Tom. "Georgian Woman Cuts Off Web Access to Whole of Armenia." *The Guardian*, April 6, 2011.
<https://www.theguardian.com/world/2011/apr/06/georgian-woman-cuts-web-access> (accessed Nov. 6, 2018).
- Park, Jihun, Joohee Kim, So-Yun Kim, Woon Hyung Cheong, Jiuk Jang, Young-Geun Park, Kyungmin Na, Yun-Tae Kim, Jun Hyuk Heo, Chang Young Lee, Jung Heon Lee, Franklin Bien, and Jang-Ung Park. "Smart Contact Lenses With Integrations Of Wireless Circuits, Glucose Sensors, And Displays." *Science Advances* 4, no. 1 (2018): DOI: 10.1126/sciadv.aap9841.
- Pew Research Center. "Internet/Broadband Fact Sheet." Pew Research Center, Feb. 5, 2018. <http://www.pewinternet.org/fact-sheet/internet-broadband> (accessed May 23, 2018).
- Piscitello, Dave. "Conficker Summary and Review." ICANN, 2010.
<https://www.icann.org/en/system/files/files/conficker-summary-review-07may10-en.pdf>.
- Porter, Dorothy. *Health, Civilization and the State: A History of Public Health from Ancient to Modern Times*. London: Routledge, 1997.
- Porter, Roy. *The Scientific Revolution in National Context*. Cambridge, UK: Cambridge University Press, 1992.
- Poushter, Jacob. "Smartphone Ownership and Internet Usage Continues to Climb in Emerging Economies." *Pew Research Center*, Feb. 22, 2016.
<http://www.pewglobal.org/2016/02/22/smartphone-ownership-and-internet-usage-continues-to-climb-in-emerging-economies> (accessed Feb. 1 2019).
- Public Health England. "Complete Routine Immunisation Schedule." GOV.UK, May 7, 2014. <https://www.gov.uk/government/publications/the-complete-routine-immunisation-schedule> (accessed Feb. 28, 2019).
- Rawls, John. *A Theory of Justice*. Revised Ed. Cambridge, MA: Harvard University Press, 1971, 1999.
- Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC.
<http://data.europa.eu/eli/reg/2017/745/oj>.

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).

<http://data.europa.eu/eli/reg/2016/679/oj>.

Ricoeur, Paul. *Oneself as Another (Soi-même Comme un Autre)*. Translated by Kathleen Blamey. Chicago: University of Chicago Press, 1992 (1990).

Rowe, Brent, Michael Halpern, and Tony Lentz. "Is a Public Health Framework the Cure for Cyber Security?." *Cross-Talk* 25, no. 6 (2012): 30-38.

Rutger, Claassen. "Capability Paternalism." *Economics and Philosophy* 30, no. 1 (2014): 57-73. doi:10.1017/S0266267114000042.

Samuels, Kate; Mark B. McClellan; Mohit Kaushal; Kavita Patel; and Margaret Darling. "Closing the Rural Health Connectivity Gap: How Broadband Funding Can Improve Care." *USC-Brookings Schaeffer On Health Policy*, April 1, 2015. <https://www.brookings.edu/blog/usc-brookings-schaeffer-on-health-policy/2015/04/01/closing-the-rural-health-connectivity-gap-how-broadband-funding-can-improve-care> (accessed May 26, 2018).

Scanlon, T. M. *What We Owe to Each Other*. Cambridge, MA: The Belknap Press of Harvard University Press, 1998.

Sedenberg, Elaine and Deirdre Mulligan. "Public Health as a Model for Cybersecurity Information Sharing." *Berkeley Technology Law Journal* 30, no. 3 (2015): 1687-1739.

Selby, Nick. "Local Police Don't Go After Most Cybercriminals. We Need Better Training." *Washington Post*, April 21, 2017, <https://www.washingtonpost.com/posteverything/wp/2017/04/21/local-police-dont-go-after-most-cybercriminals-we-need-better-training> (accessed Jan. 19, 2019).

Sen, Amartya. "Well-Being, Agency and Freedom: The Dewey Lectures 1984." *The Journal of Philosophy* 82, No. 4 (1985): 169-221.

Sen, Amartya. "Capability and Wellbeing." In *The Quality of Life*, eds. Martha Nussbaum and Amartya Sen, 30-53. Oxford: Clarendon Press, 1993.

Sen, Amartya. "Dialogue Capabilities, Lists, And Public Reason: Continuing The Conversation." *Feminist Economics* 10, no. 3, November 2004: 77-80.

Sen, Amartya. *Commodities and Capabilities*. Amsterdam: North-Holland, 1985.

- Sen, Amartya. *Inequality Reexamined*. Oxford: Oxford University Press, 1995.
- Shakarian, Paulo, Jana Shakarian, Andrew Reuf. *Introduction to Cyber-Warfare*. Waltham: Syngress, 2013.
- Sharman, Jon. "Pornhub and Twitter Ban Ai-Generated 'Deepfakes' Videos that Put Female Celebrities' Faces on Adult Actresses' Bodies." *The Independent*, Feb. 7, 2018. <https://www.independent.co.uk/life-style/gadgets-and-tech/pornhub-twitter-deepfakes-ban-ai-celebrity-faces-porn-actress-bodies-emma-watson-jennifer-lawrence-a8199131.html> (accessed Feb. 6, 2019).
- Sherman, Amy. "Fact-checking the Death Toll Estimates from Hurricane Maria in Puerto Rico." *Politifact*, June 5, 2018. <https://www.politifact.com/truth-o-meter/article/2018/jun/05/fact-checking-death-toll-estimates-hurricane-maria/> (accessed Jan. 6, 2019).
- Simpson, Thomas W. "The Wrong in Cyberattacks." In *The Ethics of Information Warfare*, edited by Luciano Floridi and Mariarosaria Taddeo, 141-154. Heidelberg: Springer, 2014.
- Singer, Peter and Allan Friedman. *Cybersecurity and Cyberwar*. Oxford; New York: Oxford University Press, 2014.
- Smith, Adam. *An Inquiry into the Nature and Causes of the Wealth of Nations, Vol II*. Edited by R. H. Campbell and A. S. Skinner. Oxford: Clarendon Press, 1976.
- Smith, Erika W. "Jennifer Lawrence Speaks Out About Reclaiming Her Body After Her Nude Photos Were Published Without Her Consent." *Bust*. <https://bust.com/feminism/194242-jennifer-lawrence-reclaiming-body-after-nude-photos.html> (accessed April 18, 2018).
- Smith, Richard, and Joanna Coast. "Antimicrobial Drug Resistance." In *Global Public Goods for Health: Health Economic and Public Health Perspectives*, edited by Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager, 73-93. Oxford: Oxford University Press, 2003.
- Smith, Richard, Robert Beaglehole, David Woodward, and Nick Drager, eds. *Global Public Goods for Health: Health Economic and Public Health Perspectives*. Oxford: Oxford University Press, 2003.
- Solove, Daniel. "'I've Got Nothing to Hide' and Other Misunderstandings of Privacy." *San Diego Law Review* 44 (2007): 745-772.

- Spar, Deborah. "The Public Face of Cyberspace." In *Global Public Goods*, edited by Inge Kaul, Isabelle Grunberg, and Marc Stern, 344-363. New York; Oxford: Oxford University Press, 1999.
- Summers, Judith. *Soho: A History of London's Most Colourful Neighborhood*. London: Bloomsbury, 1989.
- The Rendon Group. "Conficker Working Group: Lessons Learned." *The Rendon Group*. 2010.
http://www.confickerworkinggroup.org/wiki/uploads/Conficker_Working_Group_Lessons_Learned_17_June_2010_final.pdf (accessed Feb. 20, 2019).
- Thieme, Nick. "After Hurricane Maria, Puerto Rico's Internet Problems Go from Bad to Worse." *PBS*, Nov. 23, 2018.
<https://www.pbs.org/wgbh/nova/article/puerto-rico-hurricane-maria-internet/> (accessed Nov. 6, 2018).
- Thorsheim, Peter. *Inventing Pollution*. Athens, OH: Ohio University Press, 2006.
- Tiffany, Kaitlyn. "Flickr Will Soon Start Deleting Photos — and Massive Chunks of Internet History." *Vox*, Feb. 6, 2019. <https://www.vox.com/the-goods/2019/2/6/18214046/flickr-free-storage-ends-digital-photo-archive-history> (accessed Feb. 7 2019).
- Tiffany, Kaitlyn. "Myspace, Which Still Exists, Accidentally Deleted 12 Years' Worth Of Music." *Vox*, March 18, 2019. <https://www.vox.com/the-goods/2019/3/18/18271088/myspace-music-deleted-internet-archive-flickr-tumblr> (accessed April 11, 2019).
- Trim, Peter and David Upton. *Cyber Security Culture: Countering Cyber Threats Through Organizational Learning and Training*. Farnam: Gower Publishing, 2013.
- Tufekci, Zeynap. "YouTube, the Great Radicalizer." *The New York Times*, March 10, 2018. <https://www.nytimes.com/2018/03/10/opinion/sunday/youtube-politics-radical.html> (accessed 14 April, 2019).
- Turkle, Sherry. *Life on the Screen: Identity in the Age of the Internet*. New York: Simon & Schuster, 1995.
- U.S. Congress. House. Active Cyber Defense Certainty Act. HR 4036. 115th Cong., 1st sess. Introduced in House October 12, 2017.
<https://www.congress.gov/115/bills/hr4036/BILLS-115hr4036ih.pdf>.

- United Nations Development Programme. "Human Development Index (HDI)." United Nations Development Programme: Human Development Reports. <http://hdr.undp.org/en/content/human-development-index-hdi> (accessed 12 October 2018).
- United Nations. *Universal Declaration of Human Rights*. (1948). <http://www.un.org/en/universal-declaration-human-rights/>.
- United States Department of Homeland Security. "Enabling Distributed Security in Cyberspace." *United States Department of Homeland Security*, 2011. <https://www.dhs.gov/xlibrary/assets/nppd-cyber-ecosystem-white-paper-03-23-2011.pdf>.
- US-CERT. "Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors." US-CERT, March 16, 2018. <https://www.us-cert.gov/ncas/alerts/TA18-074A> (accessed May 25, 2018).
- U.S. Food & Drug Administration. "Cybersecurity." U.S. Food & Drug Administration. <https://www.fda.gov/medicaldevices/digitalhealth/ucm373213.htm> (accessed March 4, 2019).
- Ventola, Lee. "Immunization in the United States: Recommendations, Barriers, and Measures to Improve Compliance." *Pharmacy and Therapeutics* 41, no. 7 (2016): 426–436.
- Wakefield, Jerome C. "The Concept of Mental Disorder: Diagnostic Implications of The Harmful Dysfunction Analysis." *World Psychiatry* 6, no. 3 (2007): 149–156.
- Walker, Shaun. "Romanian Court Tells Man He Is Not Alive." *The Guardian*, March 16, 2018. <https://www.theguardian.com/world/2018/mar/16/romanian-court-tells-man-he-is-not-alive>.
- Wang, Youfa, May A. Beydoun, Lan Liang, Benjamin Caballero, Shiriki K. Kumanyika. "Will All Americans Become Overweight or Obese? Estimating the Progression and Cost of the US Obesity Epidemic." *Obesity* 16, no. 10 (2008): 2323–2330.
- Weise, Elizabeth. "Massive Amazon Cloud Service Outage Disrupts Sites." *USA TODAY*, Feb. 28, 2017. <https://www.usatoday.com/story/tech/news/2017/02/28/amazons-cloud-service-goes-down-sites-scramble/98530914/> (accessed April 15, 2017).

- Whitehead, Sarah J. and Shehzad Ali. "Health Outcomes in Economic Evaluation: the QALY and Utilities." *British Medical Bulletin* 96, no. 1 (2010): 5–21.
- Winfrey, Oprah. "The Jennifer Lawrence Interview, by Oprah Winfrey." *The Hollywood Reporter*, Dec. 6, 2017.
<https://www.hollywoodreporter.com/features/jennifer-lawrence-interview-by-oprah-winfrey-1064576> (accessed April 18, 2018).
- Wolff, Jonathan and Avner De-Shalit. *Disadvantage*. Oxford: Oxford University Press, 2007.
- Wolff, Jonathan. *The Human Right to Health*. New York: W.W. Norton & Co., 2012.
- Wolff, Josephine. *You'll See This Message When It Is Too Late: The Legal and Economic Aftermath of Cybersecurity Breaches*. Cambridge, MA: MIT Press, 2018.
- Woodward, David, and Richard Smith. "Global Public Goods and Health: Concepts and Issues." In *Global Public Goods for Health: Health Economic and Public Health Perspectives*, edited by Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager, 3-32. Oxford: Oxford University Press, 2003.
- Woodward, David, and Richard Smith. "Global Public Goods and Health: Concepts and Issues." In *Global Public Goods for Health: Health Economic and Public Health Perspectives*, edited by Richard Smith, Robert Beaglehole, David Woodward, and Nick Drager, 3-32. Oxford: Oxford University Press, 2003.
- Woolf, Nicky. "Massive Cyber-Attack Grinds Liberia's Internet to a Halt." *The Guardian*, Nov. 3, 2016.
<https://www.theguardian.com/technology/2016/nov/03/cyberattack-internet-liberia-ddos-hack-botnet> (accessed Dec. 29, 2018).
- World Health Organization. "International Standards on Drug Use and Prevention." 2nd edition. *World Health Organization*. 2018.
https://www.unodc.org/documents/prevention/standards_180412.pdf (accessed Dec. 13, 2018).
- World Health Organization. "Management of Substance Abuse: Terminology & Classification." World Health Organization.
https://www.who.int/substance_abuse/terminology/en/ (accessed Mar. 1, 2019).
- World Health Organization. "Programmes and Projects." World Health Organization.
<https://www.who.int/entity/en/> (accessed March 11, 2019).

- World Health Organization. "Research Ethics Review Committee." World Health Organization. <http://www.who.int/ethics/review-committee/en/> (accessed May 16, 2017).
- World Health Organization. "World Health Assembly." World Health Organization. <https://www.who.int/mediacentre/events/governance/wha/en/> (accessed Feb. 27, 2019).
- World Health Organization. *Basic Documents*. 48th edition. World Health Organization, 2014.
- World Health Organization and USAID, *WHO Standards for Prosthetics and Orthotics*, World Health Organization (Geneva: World Health Organization, 2017).
- Yannakogeorgos, Panayotis and Adam B. Lowther. *Conflict and Cooperation in Cyberspace: The Challenge of National Security*. Edited by Panayotis A. Yannakogeorgos and Adam B. Lowther. Boca Raton: Taylor & Francis, 2014.
- Zacher, Mark. "Global Epidemiological Surveillance: International Cooperation to Monitor Infectious Disease." In *Global Public Goods*, edited by Inge Kaul, Isabelle Grunberg, and Marc Stern, 266-283. New York; Oxford: Oxford University Press, 1999.
- Zetter, Kim. "It's Insanely Easy to Hack Hospital Equipment." *Wired*, April 25, 2015. <https://www.wired.com/2014/04/hospital-equipment-vulnerable> (accessed Dec. 3, 2017).
- Zhang, Meng, Guohua Song, and Lansun Chen. "A State Feedback Impulse Model For Computer Worm Control." *Nonlinear Dynamics* 85 (2016). <https://doi.org/10.1007/s11071-016-2779-0>.