JOURNAL OF CYBERSECURITY

Journal of Cybersecurity, 3(3), 2017, 173–183
doi: 10.1093/cybsec/tyx005
Advance Access Publication Date: 20 May 2017
Research paper

Research paper

Coercion in cybersecurity: What public health models reveal

Steven Weber*

Department of Political Science, School of Information, University of California Berkeley, 203 B South Hall, Berkeley, CA 94720, USA

*Corresponding author: E-mail: steve_weber@berkeley.edu

Received 9 May 2016; revised 15 March 2017; accepted 17 April 2017

Abstract

Insights from public health theory and practice have been put forward as elements of doctrine to inform theory and policy frameworks for cybersecurity. Analogies between public health and cybersecurity are superficially appealing but fail on closer examination in two distinct ways: the "publicness" of the goods in question, and the readiness of the relevant actors and institutions to exert and accept coercive authority. This article assesses the analogy in depth, starting with a review of foundational arguments from public goods theory. I demonstrate how policy choices not technological ground truths have configured many cybersecurity "goods" or goals as public goods. I then assess the public goods provision problem in context by examining the history of important public health challenges and responses. I argue that this framing presents difficult choices regarding the use of coercive power to supply public goods. These are choices that public health officials have largely settled, but that internet society and the technology community have not because the requirements are counter-cultural to basic mindsets in those communities. Pushing past cultural resistance around the idea of "coercion in the interests of security" does not fully determine any specific cybersecurity policy outcome, but it does force a more straightforward assessment of what tradeoffs are at stake. The level of coercion that public policy will have to grapple with for cybersecurity goals is higher than generally understood.

Key words: public health; public goods; authority.

In the search for a cybersecurity "doctrine," both scholars and practitioners have turned to metaphorical stories about "public health" as a potential model for thinking about the problem. Is this an intellectually defensible conceptual move? More importantly, will it lead to better policy and decision making?

This article unpacks and answers these questions in five steps. Step 1 explains why the search for a high-level cybersecurity doctrine is both understandable and somewhat treacherous. Steps 2 and 3 review the foundational arguments of public goods theory to demonstrate how the way in which the internet has become organized configures many cybersecurity "goods" or goals as public goods. Step 4 assesses the public goods provision problem in the context of the history of important public health challenges. I argue that this particular framing presents difficult choices regarding the use of coercive power to supply public goods, choices that public health officials have largely settled but that internet society has not wanted to face head-on. Step 5 asks if there is a way to soften the terms of that dilemma. In the end there really is not, and that has important consequences for the level of coercion that public policy may have to grapple with to achieve some high priority cybersecurity goals.

This conclusion is unsettling because it challenges common foundational notions of what the internet is and should be, notions grounded in a soft libertarianism dating back to the early days of modern personal and networked computing. I argue that underpinning many of today's cybersecurity policy fights about encryption, liability, and the like is a "first principles" debate about coercion and coercive authority that democratic societies in particular would prefer to avoid. This is made trickier by the deep—and now validated in some respects (for example, by the 2013 release of confidential National Security Agency documents by Edward Snowden

and other disclosures)—distrust of governments' cyber-intentions, honesty, and basic trustworthiness.

But that's no excuse for avoiding the debate. The public health metaphor should be used instead to press to the forefront some inconvenient choices and gut-wrenching tradeoffs around coercive authority that cybersecurity policy debates need to grapple with, sooner rather than later. Pushing through the cultural resistance around the idea of "coercion in the interests of security" does not determine any particular outcome, but it does force a more straightforward assessment of what tradeoffs are at stake.

Searching for doctrine through metaphor

Doctrines are high-level conceptual frames that specify goals, means, and core logics that connect goals and means in the service of a desired outcome. Doctrines are obviously useful tools to structure thought, debate, politics, and ultimately decision-making, particularly when decisions seem to require difficult trade-offs. The iconic example of doctrine in the nuclear security world—Mutually Assured Destruction or MAD—tied a clear goal (deterrence) to a clear means (maintenance of second-strike capability) in the service of a clearly desired outcome (the prevention of nuclear war) (Friedberg, 1980).

If nuclear doctrine—developed over decades with vast investments of intellectual and political resources—could be mined for a starting point template of cybersecurity doctrine, then that would certainly help make sense of this incredibly complex and seemingly unfamiliar territory. But the world of cybersecurity is of course much messier than the nuclear world, with vastly more and different kinds of actors, massively more distributed technologies, and many other differences that complicate any simplistic mapping of nuclear doctrine to this space. A small sample of known challenges and issues includes malware, denial of service attacks, espionage, phishing, state and non-state actors, data breaches, encryption, and so on.

Thus even at the simplest level of comparison, the exercise gets into trouble. There is not a shared consensual outcome (like preventing a nuclear war) in cybersecurity—rather, relevant actors have meaningfully different goals. A prominent example of this is that the US Government and most Americans would put "protection of free speech" under the umbrella of a principal cybersecurity outcome, while the Chinese (and some other governments) would put "stability of the ruling regime" first. There is not a clearly specified and precise goal (like deterrence of any and all nuclear attacks). Are we trying to avoid all cyber-attacks on other states' critical infrastructures and if not, under what conditions might such attacks make sense or be considered legitimate?[1] And there's nearly no agreement on acceptable means (like the maintenance of second-strike capability,) even among actors who seem to share a point of view on goals and outcomes. Should zero-day exploits be stockpiled in military arsenals? Should private network providers be permitted, incentivized, or even required to inspect packet traffic for unusual signatures? These are some of the reasons why efforts to extend doctrinal logics from the nuclear era to cybersecurity issues have not yielded much practical payoff (Cyber Analogies, 2014; Cirenza, 2016; Lawson, 2012).

Of course, a simple extension of doctrine from any existing domain was never really likely to work. In today's internet context—still the early days of one of history's most dynamic and far-reaching technological developments—it would be too much to hope for a simple, coherent, and compelling doctrinal statement that could unify these debates.[2] Instead, what we have is a search for partial models and metaphors that can draw on familiar conceptual schemes and relevant experiences for at least some guiding principles that could inform the search for doctrine (a much softer and less ambitious but still useful objective).

In that context, researchers and commentators have frequently turned to stories and arguments from 'public health' for what might be a more appropriate metaphor. This move has been made by economists (Rowe et al., 2012), lawyers and computer scientists (Charney, 2012; Mulligan and Schneider, 2011), security analysts (Rice et al., 2010; Betz and Stevens, 2013; Moore et al., 2008), journalists (Martijn, 2014), and others. It is a theme frequently encountered in public talks and conference presentations. These arguments share a basic common thread: The problems that "cybersecurity" presents to modern societies mimic some of the key challenges that the field of public health has been dealing with for centuries. The most important similarities lie in the "public good' characteristics of many public health and cybersecurity objectives (more on this point below) that may account for why both are hard for large societies to achieve. But in the last century in particular, the science of public health and—equally important—the practices of public health authorities have advanced to a place where in many if not most cases solution concepts are understood (even if they are not always implemented successfully). Some of those solution concepts can be mapped on to the cybersecurity domain. Doing that, can provide specific policy prescriptions for particular kind of security threats. More ambitious articulations suggest a broader doctrinal justification and umbrella logic for how to think about security in a networked environment.[3]

It is easy to understand the attraction to metaphors from public health since there are similarities on the surface in language, and some familiar stories of success in public health that make the metaphor seductive. With the obvious parallelisms lurking in people's minds—computers get attacked by *viruses*, storage systems suffer from *infections*, attack vectors *spread and transmit* and so on—it was almost inevitable that "public health" would rise to the fore as a candidate metaphor for cybersecurity doctrine. But inevitable does not mean accurate, justified, or pragmatically useful.

I am highly skeptical of how this metaphor is typically used, for three reasons. First, there is an analytic distinction that needs to be cleared up, about the extent to which public health outcomes (and some of the specific cybersecurity parallels) are actually public goods (the word "public" means different things in these settings). Second, where public goods are in fact at stake, I believe the public health metaphor glosses over the key role of government in providing public goods. Third, I argue most importantly that it sidesteps or simply

1 George Perkovich and Ariel Levite at the Carnegie Endowment for International Peace are in the process (2016) of carrying out a large-scale study documenting the wide range of beliefs on this subject across different relevant governments.

2 Note that the development of a coherent and (mostly) consensual body of nuclear doctrine took several decades—and was never fully stable. The nuclear era began in the mid-1940s. The doctrinal combination of MAD, deterrence, and strategic nuclear arms control evolved over the course of the 1960s and wasn't institutionalized (internationally) until 1973 and SALT 1. In the 1980s, it was called into question by the Reagan Administration's interest in "nuclear war fighting" and strategic defense.

3 It is the nature of the networked environment, of course, that prima facie makes the public health analogy more interesting and potentially revealing, than simply mapping cybersecurity onto "health care" per se. Richard Clayton (2011) "Might Governments Clean-up Malware?" *Communications and Strategies* 81:87–104 is useful on this point.

neglects just how important coercion and coercive authority have been in major public health achievements. But I have more in mind than simply to offer a conceptual or linguistic critique here.

As problematic as the public health metaphor may be, it can still advance the debate significantly by forcing attention to these hard, uncomfortable, and very important questions: Is cybersecurity a public good in any meaningful sense and if so, why? To the extent that it is a public good, what is the role of coercion in provisioning it? Do we think cybersecurity is important enough to merit that level of coercion, applied to internet society? And would it take a major cybersecurity crisis event to make this level of coercive power legitimate for governments? These are the questions I take up in the rest of this article.

A caveat is necessary to bridge levels of analysis problem, that sometimes causes confusion as it differentiates policy debates from technical and economic ones. It's obvious that to talk about "cybersecurity" as if it were a single thing is, for almost all purposes, over-simplified and over-aggregated. At the policy level, there is no stable consensus on what falls within the term. This is only set to get worse, as more and more "things" become connected to digital networks as part of the Internet of Things, and cybersecurity thus begins to converge in some respects with simply "security". At the same time, researchers and professionals know perfectly well, for example, that the challenges of spear phishing are technically, strategically, and economically distinct from the challenges of encryption, and so on. A parallel tension exists in the public health world, where (for example) the management of communicable disease is conceptually distinct in critical ways from the management of non-communicable disease. But we still use the term "public health" as an umbrella concept to start; governments still have public health departments responsible for action; and there are still public health policy frameworks that structure debates and decision making.

One response to this levels-of-analysis mismatch would be to reject any efforts that seek a general cybersecurity framework or doctrine even as a starting point. That might be intellectually defensible in the abstract. But it is not a pragmatic response, since the policy community and other decision makers as well as scholarly writings that are meant to speak to policy debates continue to use the term and seek some level of generalizable insights in that context. In this article, I use the term (for the moment) in a softer and more heuristic fashion. The question I pose next is intentionally circumspect: does cybersecurity *on the whole* have important characteristics of "publicness" that would invoke public goods theory sufficiently, that this theory should inform the debate over cybersecurity ends, means, and goals? Grappling with that question first, at that level, sets up a later analysis of particular public health analogies that can be mapped more directly against particular cybersecurity challenges (in Section 4).

Public Health and Cybersecurity: Public Goods?

This is the classic chart presenting the criteria which formally define a public good:

	Excludable	Non-Excludable
Rival	Private goods	Common-Pool Resources
Non-Rival	Club goods	**Public Goods**

Rival versus non-rival is a measure of the degree to which my consumption of a good diminishes your ability to consume the same good. Excludable versus non-excludable is a measure of the degree to which it is possible to limit or restrict access to the good—for example, to those who have paid for it or shared in the burden of providing it. The distinctions are not always hard and fast but they are nonetheless meaningful. Seats at a N.Y. Giants football game in New Meadowlands Stadium are certainly a private good. The televised broadcast of that game, however, can be either a public good (if it's over the air on a broadcast network like NBC) or a club good (if it's on a paid-for cable channel like ESPN).

This second distinction is the critical one because it shows that while "nature" often determines rival-ness, nature often does *not* determine excludability, which can be structured in different ways depending on social, legal, and technical conventions. In the case of the Giants game excludability on TV is a function of networks' business models. A broadcast New York Giants game can be either a public good or a club good, and the difference is a product of human choice.

This analytic distinction can be made confusing by language, particularly in discussions about public health. The following question is an important prerequisite to using public health metaphors to inform debates about cybersecurity policy: are *public health goals and outcomes* actually *public goods* in the theoretical sense? The problem is that analytically, many public health goods are actually *not* public goods—at least not in some "natural" sense that transcends human choice.

Consider the common example of the flu vaccine. A population that has had a sufficient proportion of its members vaccinated against flu acquires a degree of herd immunity which would *appear* to be a public good, since the absence of disease is non-rival and is often portrayed as non-excludable (in the sense that the non-vaccinated individuals within the population benefit just as much as those who did vaccinate) (A classic reference is Anderson and May, 1985). But that's too simple an abstraction and takes too much for granted.

Non-excludability (like football on over-the-air TV) is a choice made by people and enacted by institutions. Nature does not decide this; people do. The government of a city or a country could very well choose to structure the problem differently and exclude from the benefits of herd immunity those who do not contribute. To do that, the government could pass and enforce a law that prohibits the unvaccinated from entering public spaces, and quarantines or even deports them. It could demand proof of vaccination at entry points like airport immigration booths. Herd immunity then becomes a club good—non-rival, but very much excludable, just like a cable TV broadcast.[4]

Now consider the high level analogy to cybersecurity. Some simple cybersecurity goods—for example, email attachments free of malware—are non-rival in the most basic sense that my consumption of "clean" email does not diminish your ability to consume clean email. But are these kinds of goods naturally non-excludable? They certainly don't have to be, any more than herd immunity is. Consider a software patch that fixes a security loophole—clearly non-rival (it's just a piece of code) but it could quite easily be

4 Would some number of people "get through" anyway by cheating or spoofing a vaccination certificate, or other means of proof? Of course they would. But that doesn't change the argument, any more than the observation that some people find ways to 'steal' a cable TV broadcast for which they have not paid.

excludable, for example if it were designed to function only on licensed copies of the latest release of an operating system or something similar.

Excludability can be taken further in simple thought experiments. Imagine a private network where only authenticated and authorized people willing to fully identify themselves to the network in particularly demanding ways can send attachments. Or imagine an ISP that offers a premium network product for a high price that much more aggressively tests and scans attachments for viruses and malware prior to delivery.[5] Now the consumption of clean email has been recast into a club good. Nature did n't change; what happened is that different business models were adopted by commercial players in the network.

If it feels like a struggle to imagine these things, recognize that the struggle is with assumptions about regulatory regimes and business models, and preferences for network characteristics like openness, net neutrality, and interconnection that constrain those business models. It is not with nature or technology. Without prejudging what is the better choice, the point here is simply that treating cybersecurity goods as public goods, if we do so, is in fact a policy choice. The choice to treat cybersecurity goods as club goods (or in some cases, even private goods) is a choice equally available.

Getting to public goods

The distinction between public goods and club goods and the importance of human and political choice in determining that distinction might seem slightly arcane, but it matters tremendously because it frames up the question of how these goods are going to be provided (or not).

The major insight of public goods theory is not good news: it is the tendency toward market failure and fewer public goods than people actually want. Put simply, private actors operating within free markets generally will not provide an optimal level of public goods. They won't do so because they can't capture the positive externalities that the public good creates. Users, because they can't be excluded, will free-ride on public goods and the private costs of providing them will then exceed the private, internalized benefits. The market fails and as the jargon goes, public goods are "underprovided" (The roots of the argument are Samuelson, 1954. A classic reference is Olson, 1965).

There are exceptions to this baseline expectation of market failure. The simplest exception is a form of "hegemony"—when one player in the market is large and powerful enough that providing the public good essentially by itself (and allowing others to free-ride) represents a dominant or at least plausible strategy for that player. Olsen acknowledged another possibility in his original work by allowing for public goods provision by a "privileged" or what he called "K" group whose capacity depends on the ability of K group members to efficiently monitor each other's contributions. Russell Hardin developed the argument further in part by exploring the conditions under which $K = 1$. In international trade theory (where open trade is conceptualized as a public good) the $K = 1$ argument is known as "hegemonic stability theory". And Elinor Ostrom famously explored empirical examples of collective goods provision in communities managing what she called "common pool resources" (mostly environmental goods). Her explanation for these cases is laid out in a set of design principles which rely on high levels of social capital including wide participation in rule making, low-cost means of conflict resolution, effective internal community monitoring, and more.[6]

But from the perspective of a global internet system trying to deal with security, these exceptions to the market failure problem are exceptions that essentially prove the rule. They depend on restrictive conditions that do not apply to many key problems on the internet, and so the "tendency" toward market failure as above remains the baseline expectation.

This baseline is why public goods theory typically turns to non-market solution concepts, including ideologies, social norms, subsidies, and—most importantly and reliably in practice—institutions like governments with coercive authority. Getting past public goods under-provision means calling into play some mix of these concepts above and beyond a market solution.

Governments, for example, typically see nuclear deterrence as a public good and want higher levels than the market would provide—so they use some of their coercive power to extract taxes to pay for nuclear weapons on an involuntary (i.e., not market) basis. They also invoke ideologies (like nationalism) and social norms (like a sense of obligation and fairness). No modern government would rely solely on non-coercive tools to solve the public goods provision problem around a core national security objective.

The same goes for herd immunity and vaccination programs: ideologies and social norms play a part in getting people to contribute to the public good. Public health authorities in early autumn typically launch large-scale publicity campaigns about the dangers of flu as well as the safety and efficacy of the flu vaccine. Many of the ads carry a message of normative responsibility for others—you should get a flu shot not only to keep yourself healthy but to reduce the risks to all the "innocent" people with whom you will come into contact. Supplementing this is usually some attention to what behavioral economics calls channel factors, simply making it easier for people in their day-to-day lives to get vaccinated. (The practical importance of channel factors is explored in Bertrand et al., 2004). Public health authorities also invest heavily in information collection and monitoring programs—how prevalent is the flu in any particular city or state, what are the vaccination rates achieved, and so on. Do these normative, behavioral, and information provisioning moves get all the way to where we need to go?

Under normal circumstances in most countries, it's basically good enough. But when the threat of pandemic-scale outbreak becomes serious, public health authorities (as well as other government agencies) can, under a wide variety of circumstances, wield coercive power. Schools can reject admission to students that cannot prove a vaccination record. Most modern public health authorities are empowered to order mandatory reporting, mandatory screening and symptom surveillance, and sometimes even mandatory treatment [In the USA, most of these authorities (though not all) reside with state or local public health authorities. A good summary of authorities is State Quarantine and Isolation Statutes, 2016]. The option remains (more likely to be deployed in some countries than in others) to

5 ISPs have on occasion experimented with related, preliminary schemes: for example, Qwest in 2007 and Comcast in 2009 developed automated processes for detecting botnet traffic and notifying customers that their computers were infected.

6 Olson, M. (1965) *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge: Harvard University Press; Russell, H. (1982) *Collective Action*. Baltimore: Johns Hopkins University Press; Kindleberger, C. (1973) *The World in Depression 1929-1938*. Berkeley: University of California Press; Ostrom, E. (1990) *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

restructure the overall situation so that a public good becomes a club good, with coercive authority providing the means to do that as well. For example, an immigration guard with coercive authority can present the visitor with a stark choice: either take a vaccine or exit the country.

What does this tell us about the cybersecurity public good landscape? To generalize from the prior examples, compensating for under-provision of public goods in a market failure situation is usually a matter of combining at least three approaches: persuasion through things like education, ideologies, and social norms; monitoring through transparency and information gathering; and government-backed coercion of behavior. It's seductive to imagine a high-level mapping of many of these approaches and capacities to cybersecurity. At the "low" end of the activism spectrum, public education campaigns to encourage behaviors like rapid installation of software patches and cautious treatment of emails with web-links are essentially no-regret moves (the major point of controversy here is cost-efficiency). Going further, some governments might consider actively subsidizing security-promoting behaviors and/or firms that offer effective security products to consumers.

But when the mapping begins to touch on higher levels of activism that include coercive authority and coercive measures in the cybersecurity realm, practice becomes much more complicated and controversial.

Consider for example the basic question of whether governments should legislate mandatory reporting (to government authorities and/or to the public) of major and even minor security intrusions—and if so, on what terms should these requirements cover firms, government agencies, and even individuals? There is deep disagreement over the quantity of private information, including information about sources and methods for detection of intrusion, that could be revealed, above and beyond debates about efficacy in incentivizing future security investments.

These questions are controversial, the subject of profound and sometimes bitter debate. They are deeply tangled in matters of secrecy, reciprocity, and legal liability. They have also become (in the USA at least) intensely politicized following the Snowden revelations. The foundation of basic trust—a shared sense that the players in this game desire a common goal even as they argue (sometimes bitterly) over the best ways to achieve it and the distribution of costs and benefits on the way to that goal—may be lost, at least for a time.

But even if that foundation of trust can be built or rebuilt, beneath these controversies lie even deeper and more profound disagreements over the (assumed) mechanisms of change that are embedded in coercive authority. Put differently, what level of coercion, exerted at what point in the process, and against whom, is most likely to work? Consider a simple example: Is it better for cybersecurity outcomes overall for researchers to give Microsoft 30-day private notice of a zero-day exploit with the promise of public disclosure after 30 days; or would it be better to keep the secret; or share it privately as a research finding with no promise/threat of public disclosure? A conceptually related question is should governments devise specific regulations to require the development and installation of patches, with concrete penalties for non-compliance on the part of software developers and perhaps even users. This obviously is an argument in part over how one believes Microsoft and the users of its software can best be incentivized to respond, and in part a function of beliefs about the balance of risk at the simplest level between actors that are loosely labeled "good guys" and "bad guys". If we substitute Google or IBM or Apple for Microsoft in that zero-day exploit scenario, does the assumed balance of risk change? More concretely, who in practice makes those judgments?

A substantial body of research addresses these kinds of questions about security flaw disclosures through formal security economics models; empirical estimates of various costs to firms, users, and governments; and other methods.[7] This work elucidates some of the large number of variables and key strategic interaction dynamics at play that would bear on a rational public policy maker's attempt to fashion an optimal reporting regime. For example, how large should sanctions be for non-compliance, vis-a-vis the costs of a reporting regime and the efficiency with which information about breaches can be shared and used; and how likely must it be that non-compliant firms are "caught" by audits? These are only some of the relevant variables that the models reveal but are extremely hard to measure in practice. And while this research has certainly made progress in clarifying some of the most important considerations, it remains clear to most people who work in the policy world around these issues that cybersecurity professionals and hackers as well have very strong and sometimes nearly theological views on questions like these, that have not been reconciled within the models.

And so, foundational questions remain about what are plausible conditions for higher levels of coercion in cybersecurity policy. The more questions you pose in this fashion, the more controversial coercive authority seems to become. Under what conditions should a government agency be empowered to force disconnection from the network of an ISP, a firm, or an individual? Or require mandatory treatments—firewall installation, security audits? When are active and intrusive surveillance means acceptable or necessary? Just how much coercive authority can governments justifiably wield against exploits and attacks that are shown to be possible in theory but not yet demonstrated in practice?

These questions are certainly not new. But even to pose questions like these about the internet may seem out-of-bounds to some. How, for example, can the free and open internet be maintained if governments can mandate the use of one or another kind of software code, or permit (or even promote) private actors to engage in "active defense"? Yet, if we choose to structure our thinking about cybersecurity around public goods theory, these are precisely the questions that the theory demands be asked, if the problem of under-provision is going to be managed more effectively.

Theory only goes so far. To get a broader sense of how tough the choices around coercion really might be in practice, it's instructive to go beyond the theory of public goods and into the historical record of what has actually worked in major public health issues. The record suggests some further uncomfortable hypotheses about the level of coercive authority that will likely be needed to get an adequate level of public goods provided in analogically similar security problems.

7 Ross Anderson maintains a repository of relevant papers (including but not limited to the economic and strategic interaction dynamics around notification) at http://www.cl.cam.ac.uk/~rja14/econsec.html. Anderson, "Why Information Security is Hard: An Economic Perspective", in *ACSAC '01 Proceedings of the 17th Annual Computer Security Applications Conference* provides a classic introductory overview. Two examples of contrasting models of security flaw disclosures are in Eric Rescorla (2005) 'Is Finding Security Holes a Good Idea', *IEEE Security and Privacy*, 3: 14–19; and Ashish Arora, Rahul Telling, and Hao Xu (2008) 'Optimal Policy for Software Vulnerability Disclosure,' *Management Science* 54:642–56.

Coercion in the service of public health

This section reviews some key examples from the modern history of public health science and policy, in order to assess how challenges of public goods provision have been managed in that realm, and infer what the implications of those insights might be for cybersecurity policies.[8] Consider quarantine to start, almost certainly the oldest "active" public health measure known to humans. Quarantine has been practiced for centuries, long before the germ theory of disease, on what medical professionals today would call an "empirical basis" [Shorthand for observations of efficacy that are not always accompanied by a fully scientific, casual explanation. Empirical quarantine goes back at least as far as 14th Century Croatia and Venice, when ships traveling from ports where plague had been observed, were required by authorities to sit at anchor for 30–40 days before landing (e.g. see Mackowiak and Sehdev, 2002). And quarantine was often coercive—perhaps even more so because of the lack of a scientific basis that might have aided in persuasion. After Pasteur and Koch, compulsory quarantine gained a scientific rationale and was used regularly to prevent the spread of disease, particularly in the dense urban populations that were growing at the time (A history and justification is Huber, 1926). The point is that restricting the free movement of people and goods is effective against many infectious diseases but generally requires persistent enforcement and significant coercive power. It was not unusual (and still is not) for police forces (and occasionally, military forces) to be deployed in quarantine situations.

The deeper understanding of the problem that came with the germ theory of disease did not make it much easier to persuade mass populations against free-riding on others—much as greater awareness of software vulnerabilities by individuals has not translated into consequential behavior change around the negative externalities of inadequate security practices.

Consider second the development of sanitation and hygiene, almost certainly the oldest "passive" public health measures known to humans. As the germ theory of disease made clear why these measures worked, urban leaders developed sewage, water treatment, and food safety programs that also were not voluntary but imposed by municipal and building codes (Cook. 2001). The rationale for coercion in this setting was laid out boldly by Herman Biggs, General Medical Officer of the New York City Department of Health at the end of the 19th century. In 1894, he convinced his colleagues on the NYC Health Board that they should consider tuberculosis a "notifiable" disease and keep records of its spread. It took four more years for notification to be made obligatory and several additional years before the city forced some resistant physicians to comply. Chastened by the experience, Biggs later called for public health officials to be granted broader general authorities including those for mandatory vaccination, enforced quarantines, surveillance, and mandatory reporting of certain infections (with peoples' names attached) to public health registries (Biggs, 1897; Duffy 1992).

All would become part of the standard tool-kit of public health officials, even in the face of predictable resistance to this expansion of government authority. Forms of opposition that sound familiar to contemporary public health debates would have sounded just as familiar in the early 1900s: anti-vaccination movements that question the underlying science, the safety of vaccines, or the legitimate right of governments to coerce individuals to vaccinate; concerns about anti-immigrant or racial biases in the context of quarantine; resistance to notification programs that require naming the infected persons (Duffy, 1992).

But on the whole, public health authorities won more battles than they lost around issues of coercive authority. They did so, and with the general support of the courts, because the visible benefits of reduced morbidity and mortality from infectious disease along with a general sense of trust in the public health authority overcame concerns about coercion. It was, of course, a different time with broadly different and more accepting attitudes toward science, authority, and government power deployed in the service of public goods. It was also a different time with regard to general levels of trust in public sector institutions. The analogy to cybersecurity debates today is clearly complicated by a landscape that has shifted along each of these dimensions, as well as the salient fact that system penetrations, data breaches, and even state-sponsored cyber-attacks have not yet led to the kinds of vivid, lethal, and spreading harms of the most serious infectious diseases. But the history of sanitation and hygiene measures suggests, as with quarantine, that voluntary measures will not sustain the kind of public goods contributions that are needed.

In any case, *de facto* grants of legitimate coercive authority in the public health realm came under pressure toward the end of the 1960s from a number of different directions.

The Warren Court's "due process revolution" along with the further development of privacy jurisprudence began to shift the balance against coercion (Gostin, 2000). The HIV/AIDS epidemic of the 1980s and 90s both illustrated and magnified this shift. The activism and politics of the HIV epidemic rendered controversial a number of coercive government actions that in the past would have been much less controversial for a disease of similar severity and (at the time) morbidity, like mandatory screening and even quarantine. Fearful of discrimination, the gay community built a distinctive level of political cohesion and took actions like screening, prevention, and care into their own hands. They were highly organized in part because HIV/AIDS threatened the very existence of the gay community (e.g. see Cohen and Elder, 1989; Wright, 2013).

Because the disease struck a relatively well-defined community with such force, and with the overhang of discrimination and stigma, this outcome is consistent with the concept of a "privileged group" in public goods theory, which is more likely to be capable of solving the problem of public goods under-provision without external coercion (Olson, 1965). Olson's model does not provide a precise set of parameters that define a privileged group; the effective size of the group (small vs. large) is in part a function of collective identity, peer pressure, "direct contact", and other possible variables that affect individual incentives and monitoring costs]. The point is that this community likely avoided government coercion because of its ability to self-organize.

Consider a thought experiment where the conditions that support privileged group provision of public goods do not exist: an alternative world in which the pathophysiology of HIV is more like a widespread computer virus. The simple way to see this is to imagine

8 From a methodological standpoint, precisely paired comparisons between specific public health problems and specific cybersecurity issues would of course be optimal. But creating such paired comparisons rests on a heavy set of ceteris paribus assumptions, which I have argued in previous sections are not quite justified. This makes precisely paired comparisons premature. Instead, I offer here a rougher set of analogical insights that should be treated as hypotheses, not conclusions.

that the HIV virus underwent a mutation so that it was quite easily transmitted and infected people indiscriminately. Now, imagine two cities that follow different policy paths. City One relies on education and voluntary behavior change to stop the spread of disease. City Two enacts emergency coercive public health measures: mandatory reporting with full identification of all new cases; quarantine of infected individuals and required social distancing for populations at risk, and the like. Which city would have been labeled a public health success? The question is whether voluntary measures would have survived as the preferred solution—or perhaps whether voluntary measures would have survived at all.

Extending the HIV/AIDS example to a thought experiment like this is a means to put pressure on hopeful assumptions about the generalizability and extensibility of the HIV "success" case to public goods in security. It suggests that the conditions under which privileged groups form are restrictive in practice and much less likely to be present in most cybersecurity challenges. Collective identity, peer pressure, "direct contact", and other variables that increase individual incentives and reduce monitoring costs are not characteristic of such challenges as resistance to phishing schemes or even simple patching of identified software vulnerabilities.

A more generalizable "success" case relevant to cybersecurity issues is automobile safety. In fact, the decline in motor vehicle accident deaths inside the USA should probably be seen as one of the greatest public health achievements of the modern era. The last 100 years have seen a 7× increase in number of drivers, a 12× increase in motor vehicles, and a 10× increase in miles traveled; but about a 90% decrease in deaths per mile traveled (Detailed recent data at FARS Encyclopedia, 2016; historical data at National Safety Council, 1998). As a base level analogy, the potential "threat environment" on roads has multiplied in density and intensity, but the roads have become a much safer place at the same time. How did this happen?

Technology plays a major role in this case (and will continue to do so as autonomous vehicles come on line), but technological change by itself did not overwhelm or solve the road safety issue. It was systematic safety regulation by governments, which accelerated in the 1960s, that prompted the deployment and use of many of the relevant technologies as well as their integration with driver training and behavioral change. It is notable that The US National Highway Safety Bureau, created by the National Highway Safety Act in 1966, hired a public health doctor as its first director. It then moved to set new and stringent regulations for the design of head rests, seat belts, and barriers between lanes. At the same time, states began to implement tougher laws and more stringent enforcement against driving under the influence of alcohol (Committee on Injury Prevention and Control, Institute of Medicine, 1999; Graham, 1993; Transportation Research Board, 1990).

It's not surprising to most Americans that there was significant ideological opposition to the public health application of coercive authority in automobiles—for example, in making seat belt use mandatory (e.g. see Morelock et al., 1985;Weissert and Weissert, 2012) An important ideological parallel is that cars in America were and are symbols of freedom—for earlier generations, the "Chevy Impala on the Interstate" had an emotional valence much like the "free and open internet" does today, and contemporary language about the "right" to drive a car has resonance much like language about the "right" to connect to the Internet. But the numbers tell an important story about coercion overcoming ideology in the case of auto safety. In 1981, 3 years before the first state seat belt law, only 11% of people regularly used seat belts. In 2016, that number was close to 87%. And there is evidence for a direct causal relationship with law: according to the National Highway Traffic Safety Administration's 2013 survey, seat belt use in jurisdictions with stronger seatbelt enforcement laws continue to exhibit generally higher use rates than those with weaker laws (Rivara et al., 1999; Pickrell and Liu, 2014).

When it comes to seat belts, coercion works. And of course this is true not just of seat belts—many coercive restrictions make driving a meaningfully safer activity. Indeed, this is such a background assumption with regard to motor vehicles that it would be peculiar or even outlandish to hear in 2016 a serious argument that cars shouldn't have to be inspected and licensed; or that anyone should be able to drive at whatever level of blood alcohol they desire; or text while driving.

In many respects, these kinds of legitimate coercive authorities have come to be seen not so much as an intrusion on rights, but instead as a manageable balance that we've stumbled on, in a situation where even higher levels of coercion were available and under serious consideration. For example, seat belt laws require that drivers wear them, but the automatic seat belt that some cars had built into them in the 80s and 90s were rejected (Williams et al., 1988, 1992). It's accurate that automatic seat belts were also made less useful and in some situations even detrimental by the development of new air bag technology; but public opinion had turned against their use before that]. Government regulations require that all cars be designed so that you cannot start the ignition without your foot on the brake, but there's no general requirement to have the ignition wired to a breathalyzer. One of the lessons of this case is the simple but instructive contrast among levels of coercive authority. For example, when three factor authentication systems are seen as a realistic requirement, two-factor authentication systems are perceived as a lesser requirement burden than they would be in a stark contrast to simple passwords.

The story of public health interventions around smoking is another useful example. The proliferation of cheap mass-market cigarettes (a new technology, much like automobiles and the Internet in that respect) created a mass public health threat, driving up lung cancer rates in the USA by almost 15 times in the 20th century (American Cancer Society, 1999). This number references only lung cancer and does not include other detrimental health consequences and negative externalities tied to cigarettes]. Public health efforts to reduce tobacco use ramped up substantially after the US Surgeon General in 1964 explicitly linked cigarettes to lung cancer (US Public Health Service, 1964). The level of coercion around antismoking has risen gradually since then, spanning restrictions on speech (advertising bans) to punitive taxes to full bans on smoking in an increasing number of locations including airplanes, restaurants, bars, and in some cases the entire physical space of college campuses, as at the University of California Berkeley.

These measures evolved in large part because less coercive measures—public education campaigns about the dangers of tobacco being most notable—failed to have adequate impact [(Centers for Disease Control and Prevention, 1996; U.S. Department of Health and Human Services, 1994). A broad and useful survey is National Center for Chronic Disease Prevention and Health Promotion (US) Office on Smoking and Health, 2012]. In 2014, the Center for Disease Control estimated that 16.8% of Americans continued to smoke (Office on Smoking and Health, National Center for Chronic Disease Prevention and Health Promotion, CDC, 2015). Yet more coercive restrictions, including extensions to alternative technologies like vaping, are on the horizon (Electronic Cigarettes, 2016). The Family Smoking Prevention and Control Act of 2009 extended the FDA's purview over tobacco products and set in motion a

several year process toward greater restriction—in 2011, the FDA announced it would treat e-cigarettes as a tobacco product; and in 2014, the first set of proposed regulations with serious restrictions was released. The debate continues but most observers believe that restrictions will rise and the only question is how quickly and how aggressively]. The recent uptick in tobacco use among Millennials suggests an additional insight, that the effectiveness of a particular level of coercion degrades over time and has to be re-invigorated with tighter restrictions to maintain the same level of efficacy, at least when dealing with an addictive substance like nicotine (Office on Smoking and Health, National Center for Chronic Disease Prevention and Health Promotion, CDC, 2015). Also, see reports of an IPSOS survey with comparable results at What Millennials Don't Want Anyone to Know, 2015.

If habituation to restrictive measures in addictive circumstances does require progressively higher levels of coercion to maintain a stable safety level, what does that suggest about individuals' security behavior on the Internet? We know, for example, that users habituate quickly to voluntary security warnings and protocols for example in Web browsers (see the experiment and summary of others' findings in Felt et al., 2015). We know that there are addictive characteristics (and sometimes, frank addiction) to some internet behaviors that could also contribute to users ignoring security warnings (Young 2004).

What's notable—and somewhat ironic—is that the obvious "market" approaches to these dilemmas, even in the case of smoking, continue to encounter substantial resistance and are often seen as less desirable than coercion when the question is framed as a choice. It would be relatively straightforward in principle to have tobacco users bear the full costs of the risk they assume by charging them insurance rates that fully internalize tobacco-related disease. It would be more complicated—though not impossible—to develop a workable formula to internalize the externality costs of some cyber-negligent behaviors. But, it seems at least in the tobacco example that in practice most people prefer to have the societies they live in try to manage the dilemma through coercion. This is partially but not only because there are negative externalities to non-smoking individuals associated with second-hand exposure that make market-based approaches complex and imperfect. And it is certainly not a matter of concern about discrimination: It would be almost laughable today to argue that smokers are a protected minority and cannot be discriminated against (simply look at the scarlet-letter style embarrassments that smokers encounter in places like smoking "lounges" at airports). Still, many Americans apparently prefer coercive regulation on tobacco to fully differentiated insurance rates, even though the latter would be consistent with ideologies of voluntary choice. There is no a priori reason why the logic would or should be reversed for cybersecurity behaviors that have similar characteristics.

Consider as a last example the public health experience with obesity, a de-facto epidemic of chronic degenerative disease with morbidity and mortality consequences comparable to or possibly worse than smoking (a comprehensive review and comparison in National Center for Health Statistics, 2015). It is difficult to identify in the USA a successful, validated public health intervention that is able to achieve sustained weight loss among a significant proportion of an increasingly obese population (review in Overweight and Obesity Statistics, 2016; Yang and Colditz, 2015).

There is certainly no shortage of non-coercive intervention on these issues. Communication and education around obesity and related health problems is now so prevalent in American life that you can hardly avoid it on a daily basis. But that has had little demonstrable impact on outcomes over time.

Part of the problem is behavioral complexity: public health education campaigns on obesity tend to rely on a simple model of energy balance (telling people to reduce caloric intake and increase physical activity or put simply, eat less, exercise more). This is easy to communicate and easy to understand. But, the energy balance model and the simple behavioral message associated with it is out of step with contemporary developments in metabolic science, which have moved toward much more complicated homeostatic feedback models to explain obesity and weight loss (for a representative study, see Weinsier et al., 2000). In other words, the relevant public health question is not as simple as "how do we get people to eat less and exercise more?" It is a much more complicated and ambiguous question, like "what personal, genetic, environmental, and other factors cause dis-regulation of energy balance?" That question does not come with simple behavioral consequences that are broadly applicable to a diverse population and easily communicated. And, while obesity is not a sentient, determined adversary (like a sophisticated hacker), who knowingly adapts her strategy to counter your defensive moves, the homeostatic feedback model does have some related characteristics as the body's energy system adapts to try to maintain set points in the face of energy balance-related behavioral changes.

The relevant analogies are obvious. An admonishment to "patch your software" or "don't click on links you don't recognize" is in some sense the equivalent of "eat less"—a good idea, but far too simple and not enough to counter a complex process and certainly not a sentient adversary. Mass education campaigns may have had some impact persuading people that obesity and un-patched software are both bad things, but they have had very little success advising and convincing people how to act instead. Layering on additional behavioral modification messages on the conviction that people do not consistently adhere to "eat less, exercise more" or that their determination to stick with a diet wanes over time has not helped. Nor have the confusing, inconsistent, and constantly changing messages about what kinds of calories are "good" and what kinds are "bad" (is it most important to avoid sugar in drinks? trans-fats? hydrogenated oils? simple carbohydrates?) (Fairchild, 2013). A classic political justification of this form of paternalism is Rousseau, 1762. Unfortunately, none of this bodes well for the kind of complicated messages that voluntary and educational interventions around cybersecurity behaviors would have to transmit to mass populations and regularly update as the threat adapts and morphs strategically over time.

One could almost be forgiven for wishing that some liberal government would experiment with bringing substantially more coercive power to bear on the problem of obesity—not only to collect evidence on how well such an approach would work, but also on how the public would react over time to its costs and benefits as they would manifest and be experienced *in practice*. New York City's 2012 ban on supersize soda containers was rejected by the New York State Court of Appeals as an overreach by the City's Board of Health in 2014, but would it have been rejected by the citizens if given a chance? (The court ruling is at New York Statewide Coalition of Hispanic Chambers of Commerce et al., 2016)

To see how we may confront a similar quandary when it comes to coercive measures for internet security, consider a thought experiment where then-Mayor Michael Bloomberg had taken on cybersecurity, instead of soda and obesity, as an issue critical to the future well-being and competitiveness of New York City. Imagine then, for

example, that any business located within New York City was required by law to use two-factor authentication for customers logging on to its website. Bloomberg would likely have encountered different sources of resistance, some ideological. ISPs might have cried foul. Some advocacy groups might have seen the beginnings of a slippery slope toward over-regulation. Other groups might have argued that these requirements discriminate against people who can't afford or otherwise do not have easy access to a second factor authentication device, like a mobile phone. All of these objections have some merit and Bloomberg might very well have had this initiative blocked by the courts, as happened with the supersize soda ban. But, would he have necessarily been wrong to try and test the proposition that the security benefits exceed the costs—or perhaps more importantly that the public might have seen it that way and backed a more coercive approach?

The important point is that that these kinds of questions and difficult, even gut-wrenching, trade-offs around coercive authority will become more common in cybersecurity issues going forward. It is certainly happening in the public health world. In 2015, California passed new laws that further restrict exemptions and require immunization to fight the return of measles and whooping cough, as well as drug resistant tuberculosis (California statue at SB-277 Public health: vaccinations, 2016. Recent data on immunization rates in California children entering kindergarten at 2015-16 Kindergarten Immunization Assessment, 2016]. In contrast, voluntary and informational efforts to reduce the unnecessary use of antibiotics (a clear public bad) have mainly failed (McCullough et al., 2015). Another example is motorcycle helmet laws—when laws that require the use of head protection are repealed, it may be a victory for anti-coercion advocates, but it is certainly a loss for the safety of motorcyclists, as severe injuries and deaths rise apace (Striker et al., 2016).

This isn't a failure of education or information provision—motorcycle riders know well the risks of head injury (acknowledging that the population of motorcyclists includes people with a wide variety of risk propensities, a significant subjective probability of massive head injury is an outcome no rational motorcyclist would accept). It isn't a monitoring problem—there are technologies that could easily be deployed to prevent the riding of a motorcycle without a helmet. It isn't an externality problem because (despite arguments by some motorcyclists to the contrary) a motorcycle accident does create obvious negative externalities for the cars, drivers, and any pedestrians that happen to be present when the bike goes down, and, of course, for the broader society that will later have to subsidize emergency services and, possibly, long-term care for the victims of the collision. It is first and foremost a failure to exert an appropriate level of coercive authority that is demonstrably needed to insure the provision of an important public good.

Ideology and action

Public health metaphors have been put forward as a way to try to create a high-level doctrinal logic that would confront the increasing threat of security breakdowns placing much of what we value about the Internet at risk. I have argued in this paper that the logical coherence of these metaphors deteriorate under close scrutiny. The way in which that deterioration occurs—conceptually and empirically—points to a greater role for coercive authority in providing cybersecurity public goods.

This is not to say that we are facing a pitched battle between stark hierarchies and fully decentralized cooperation, either as a matter of ideology or practice. The internet is of course a mixed landscape. Internet governance in the broadest sense (not just with regard to security) depends on a complex blend of hierarchical elements (which are not always pulling in the same direction, as is the case with national regulators), collaborative professional networks (for example, among network administrators), state-sponsored or enabled institutions that foster cooperation among private actors (such as FS-ISAC, the Financial Services Information Sharing and Analysis Center) and market contracts, among other elements. This mix has varied over time and among issues, and will continue to do so. The central claim of this article is simply that in a world where many cybersecurity problems and goals have been configured as public goods, applying public health analogies to cybersecurity policy options suggests that the mix for this particular aspect of governance will tilt more heavily toward coercion.

The point is that if internet society wants to treat many cybersecurity issues as public goods, then what public health metaphors demonstrate is that it may be necessary to accept a higher level of coercion on behavior than internet society is accustomed to. That will seem to many a counter-cultural argument that will cause profound friction, particularly with regard to the "free and open" cultural trope that has emerged to supplant techno-libertarian ideologies of an earlier era. It may also run up against resistance from contemporary "innovation" discourse, which generally portrays innovation as something desirable that individuals and private firms achieve, and that coercive authority tends to quash (Barlow, 1996; Bradner 1999; Markoff, 2005; The Tao of IETF, 2012).

The friction and resistance is real but needs to be confronted openly. Contemporary Internet society is not a space that takes comfortably to discussions around the uses of coercive authority, even when it is said to be in the best interests of the network or the users of the network overall. But avoiding those discussions and the choices they present will lead to public good underprovision and continued, possibly accelerated corrosion in security and trust.

Another way to put the point is this: a metaphor is not a doctrine and it is certainly not a theory. But to the degree that theory and doctrine development for cybersecurity want to draw from the theory and experience of public health, the case for coercion will be significantly stronger than most of the proponents of this kind of thinking would prefer.

Acknowledgements

The author acknowledges the assistance and useful criticisms of colleagues including Naazneen Barma, Jesse Goldhammer, Richmond Wong, Betsy Cooper, Frank Smith, and Michael Dalby; two anonymous reviewers and the area editor of the Journal of Cybersecurity.

This work was supported by the Hewlett Foundation via UC Berkeley Center for Long Term Cybersecurity.

References

American Cancer Society. (1999) *Cancer Facts and Figures - 1999*. Atlanta, GA: American Cancer Society.

Analogies, C. (2014) In: Goldman, E, Arquilla, J (eds). Monterey, CA: Naval Postgraduate School, 2014, 132.

Anderson, R.M. and May, R.M. (1985) 'Vaccination and Herd Immunity to Infectious Diseases', *Nature*, **318**:323–29.

Barlow, J.P. (1996) 'A Declaration of the Independence of Cyberspace', *Electronic Frontier Foundation, EFF*. https://www.eff.org/cyberspace-independence

Bertrand, M., Mullainathan, S., and Shafir, E. (2004) 'A Behavioral-Economics View of Poverty', *The American Economic Review*, **94**:419–23.

Betz, D.J. and Stevens, T. (2013) *Analogical Reasoning and Cyber Security, Security Dialogue* **44**:147–64.

Biggs, H.M. (1897) 'Address in Public Medicine. Preventive Medicine in the City of New York', *The British Medical Journey* **2**:629–38.

Bradner, S. (1999) 'The Internet Engineering Task Force'. In: DiBona, C, Ockman, S, Stone, M (eds), *Open Sources: Voices from the Open Source Revolution*. USA: O'Reilly Media, 47–52.

Centers for Disease Control and Prevention, CDC. (1997) 'State-specific Prevalence of Cigarette Smoking Among Adults, and Children's and Adolescents' Exposure to Environmental Tobacco Smoke – United States, 1996', *Morbidity Mortal Weekly Report*, **46**:1038–43.

Charney, S. (2012) 'Collective Defense: Applying the Public-Health Model to the Internet', *IEEE Security & Privacy*, **10**:54–59.

Cirenza, P. (2016) The flawed analogy between nuclear and cyber deterrence. *Bulletin of the Atomic Scientists*. February 2016. http://thebulletin.org/flawed-analogy-between-nuclear-and-cyber-deterrence9179 (22 April 2016, date last accessed).

Cohen, I. and Elder, A. (1989) 'Major Cities and Disease Crises: A Comparative Perspective', *Social Science History*, **13**:25–63.

Committee on Injury Prevention and Control, Institute of Medicine. (1999) *Reducing the Burden of Injury: Advancing Prevention and Treatment*. Washington, DC: The National Academy Press.

Cook, G.C. (2001) 'Construction of London's Victorian Sewers: The Vital Role of Joseph Bazalgette', *Postgraduate Medical Journal* **77**:802–4.

Duffy, J. (1992) *The Sanitarians: A History of American Public Health*. Urbana: University of Illinois Press.

Electronic Cigarettes (e-Cigarettes). 2016. http://www.fda.gov/NewsEvents/PublicHealthFocus/ucm172906.htm (23 April 2016, date last accessed).

Fairchild, A.L. (2013) 'Half Empty or Half Full? New York's Soda Rule in Historical Perspective', *New England Journal of Medicine* **368**:1765–67.

Fatality Analysis Reporting System (FARS) Encyclopedia. 2016. http://www-fars.nhtsa.dot.gov/Main/index.aspx (23 April 2016, date last accessed).

Felt AP, Ainslie A, Reeder RW, *et al*. (2015) Improving SSL Warnings: Comprehension and Adherence. In: *Proceedings of the Conference on Human Factors and Computing Systems*. Association for Computing Machinery, New York City, NY, USA.

Friedberg, A. (1980) 'A History of U.S. Strategic "'doctrine"'- 1945 to 1980',. *Journal of Strategic Studies*, 1980; **3**:37–71.

Gostin, L.O. (2000) *Public Health Law: Power, Duty, Restraint*. Berkeley: University of California Press.

Graham, J.D. (1993) 'Injuries From Traffic Crashes: Meeting the Challenge', *Annual Review of Public Health*, **14**:515–43.

Huber, E.G. (1926) 'The Control of Communicable Diseases Prevalent in Massachusetts', *The Boston Medical Surgery Journal* **195**:122–27.

2015-16 Kindergarten Immunization Assessment. California Department of Public Health, Immunization Branch. (2016) http://www.cdph.ca.gov/programs/immunize/Documents/2015-16_CA_KindergartenSummaryReport.pdf (23 April 2016, date last accessed).

Lawson, S. (2012) 'Putting the "War" in Cyberwar: Metaphor, Analogy, and Cybersecurity Discourse in the United States', *First Monday*, 2012. ISSN 13960466. http://firstmonday.org/ojs/index.php/fm/article/view/3848/3270 (2 May 2017, date last accessed)

Mackowiak, P.A. and Sehdev, P.S. (2002) 'The Origin of Quarantine', *Clinical Infectious Diseases*, **35**:1071–72.

Markoff, J. (2005) *What the Dormouse Said: How the 60s Counterculture Shaped the Personal Computer*. New York: Viking Adult.

Martijn, M. (2014) Here's Why Public Wifi is a Public Health Hazard. *Matter*: Medium, October 2014.https://medium.com/matter/heres-why-public-wifi-is-a-public-health-hazard-dd5b8dcb55e6#.r61idh89e (20 December 2016, date last accessed).

McCullough, A.R., Parekh, S., and Rathbone, J., *et al*. (2015) 'A Systematic Review of the Public's Knowledge and Beliefs About Antibiotic Resistance', *The Journal of Antimicrobial Chemotherapy* **71**:27–33.

Moore, J.H., Parrott, L.K., and Karas, T.H. (2008) *Metaphors for Cyber Security*. Alberquerque, New Mexico: Sandia National Laboratories.

Morelock, S., Hingson, R.W., and Smith, R.A., *et al*. (1985) 'Mandatory Seatbelt Law Support and Opposition in New England: A Survey', *Public Health Reports (1974-)* **100**:357–63.

Mulligan, D.K. and Schneider, F.B. (2011) 'Doctrine for Cybersecurity', *Daedalus* **140**:70–92.

National Center for Chronic Disease Prevention and Health Promotion (US) Office on Smoking and Health. (2012) *Preventing Tobacco Use Among Youth and Young Adults: A Report of the Surgeon General*. Atlanta, GA: Centers for Disease Control and Prevention, CDC.

National Center for Health Statistics. (2015) *Health, United States, 2014: With Special Feature on Adults Aged 55–64*. Hyattsville, MD: U.S. Department of Health and Human Services (DHHS Pub No. 2015-1232).

National Safety Council. (1998) *Accident Facts, 1998 Edition*. Itasca, Illinois: National Safety Council.

New York Statewide Coalition of Hispanic Chambers of Commerce, *et al*. vs. The New York City Department of Health and Mental Hygiene, *et al*. New York State Unified Court System, 2014. https://www.nycourts.gov/ctapps/Decisions/2014/Jun14/134opn14-Decision.pdf (23 April 2016, date last accessed).

Office on Smoking and Health, National Center for Chronic Disease Prevention and Health Promotion, CDC. (2015) 'Current Cigarette Smoking Among Adults — United States, 2005–2014', *Morbidity Mortal Weekly Report* **64**:1233–40.

Olson M. (1965) *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge: Harvard University Press.

Overweight and Obesity Statistics. (2016) U.S. Department of Health and Human Services. http://www.niddk.nih.gov/health-information/health-statistics/Pages/overweight-obesity-statistics.aspx (23 April 2016, date last accessed).

Pickrell, T.M. and Liu, C. (2014) *Traffic Safety Facts Research Note: Seat Belt Use in 2013–Overall Results*. Washington, DC: National Highway Traffic Safety Administration (Pub No. DOT-HS-811-875).

Rice, M., Butts, J., Miller, R., *et al*. (2010) 'Applying Public Health Strategies to the Protection of Cyberspace', *International Journal of Critical Infrastructure Protection*, **3**:118–27.

Rivara, F.P., Thompson, D.C., and Cummings, P. (1999) 'Effectiveness of Primary and Secondary Enforced Seat Belt Laws', *American Journal of Preventive Medicine*, **16(1, Supplement 1)**: 30–39.

Rousseau, J.-J. (1762) *The Social Contract*. France: Chez Marc-Michel Rey.

Rowe, B., Halpern, M., and Lentz, T. (2012) 'Is a Public Health Framework the Cure for Cyber Security?', *CrossTalk*, **25**:30–39.

Samuelson, P.A. (1954) 'The Pure Theory of Public Expenditure', *The Review of Economics and Statistics* **36**:387–89.

SB-277 Public health: vaccinations. 113th Cong, 2015, enacted. http://leginfo.legislature.ca.gov/faces/billCompareClient.xhtml?bill_id=201520160SB277 (23 April 2016, date last accessed).

State Quarantine and Isolation Statutes. (2016) http://www.ncsl.org/research/health/state-quarantine-and-isolation-statutes.aspx (23 April 2016, date last accessed).

Striker, R.H., Chapman, A.J., Titus, R.A., *et al*. (2016) 'Repeal of the Michigan Helmet Law: The Evolving Clinical Impact', *The American Journal of Surgery*, **211**:529–33.

The Tao of IETF. (2012) 'A Novice's Guide to the Internet Engineering Task Force' In Hoffman, P. (ed.). IETF Trust 2012, accessible at https://www.ietf.org/tao.html (2 May 2017, date last accessed); the most relevant sections are 4, 6, and Appendix A.

Transportation Research Board. (1990) *Safety Research for a Changing Highway Environment*. Washington, DC: National Research Council, Transportation Research Board (Special Report No. 229).

U.S. Department of Health and Human Services. (1994) *For a Healthy Nation: Returns on Investment in Public Health*. Atlanta, GA: Office of Disease Prevention and Health Promotion and Centers for Disease Control and Prevention.

US Public Health Service. (1964) *Smoking and Health: Report of the Advisory Committee to the Surgeon General of the Public Health Service*. Washington, DC: US Department of Health, Education, and Welfare, Public Health Service (PHS Pub No. 1103).

Weinsier, R.L., Nagy, T.R., Hunter, G.R., *et al*. (2000) 'Do Adaptive Changes in Metabolic Rate Favor Weight Regain in Weight-Reduced Individuals? An Examination of the Set-Point Theory', *The American Journal of Clinical Nutrition*, **72**:1088–94.

Weissert, W.G. and Weissert, C.S. (2012) *Governing Health: The Politics of Health Policy*. Baltimore: Johns Hopkins University Press.

What Millennials Don't Want Anyone to Know. (2015) *CSP Daily News*. http://www.cspnet.com/category-news/tobacco/articles/what-millennials-dont-want-anyone-know (23 April 2016, date last accessed).

Williams, A.F., Wells, J.K., Lund, A.K., *et al*. (1992) Use of seatbelts in cars with automatic belts. *Public Health Report*, **107**:182–88.

Williams, A.F., Wells, J.K., and Lund, A.K. (1988) 'Driver Use of, and Reaction to, Automatic Seat Belts in Ford and GM Cars', *Journal of Public Health Policy* **9**:222–32.

Wright, J. (2013) 'Only Your Calamity: The Beginnings of Activism by and for People With AIDS', *American Journal of Public Health* **103**:1788–98.

Yang, L. and Colditz, G.A. (2015) 'Prevalence of Overweight and Obesity in the United States, 2007-2012', *JAMA Internal Medicine* **175**:1412–13.

Young, K.S. (2004) 'Internet Addiction: A New Clinical Phenomenon and Its Consequences', *American Behavioral Scientist* **48**:402–15.