Doctrine for Cybersecurity

Deirdre K. Mulligan & Fred B. Schneider

Abstract: A succession of doctrines for enhancing cybersecurity has been advocated in the past, including prevention, risk management, and deterrence through accountability. None has proved effective. Proposals that are now being made view cybersecurity as a public good and adopt mechanisms inspired by those used for public health. This essay discusses the failings of previous doctrines and surveys the landscape of cybersecurity through the lens that a new doctrine, public cybersecurity, provides.

Governments, businesses, and individuals are growing increasingly worried about the security of networked computing systems. This concern is justified. Press reports of successful attacks grow ever more frequent: cross-site scripting used to pilfer consumers' passwords, large-scale breaches of corporate customers' personal information, distributed denial-of-service attacks on websites, cyber espionage aimed at classified documents, and attacks on civil critical infrastructures.

Consequently, computer scientists and their funders are investing heavily in technological means for improving cybersecurity. But technological solutions are useless if they are not deployed or if operating practices allow attackers to circumvent them. Policy must create incentives for system developers, operators, and users to act in ways that enhance rather than weaken system security. Moreover, neither technologists nor policy-makers have the luxury of starting with a clean slate. All must labor in the shadows of legacy networks and end systems that are not secure (nor easily made so) and in the context of extant policy that reflects societal values from a time when dependence on networked information systems was minimal.

Enhanced levels of cybersecurity can create tensions over cost, function, convenience, and societal values such as openness, privacy, freedom of expres-

DEIRDRE K. MULLIGAN is an Assistant Professor in the School of Information at the University of California, Berkeley, where she is also a Faculty Director of the Berkeley Center for Law and Technology.

FRED B. SCHNEIDER is the Samuel B. Eckert Professor of Computer Science at Cornell University.

*(*See endnotes for complete contributor biographies.)*

© 2011 by Deirdre K. Mulligan & Fred B. Schneider

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

sion, and innovation. Absent a widely accepted doctrine, evaluation of proposals for improvement is difficult, and debate about their adoption can be neither compelling nor conclusive. The utility of a doctrine is thus determined by the extent to which it offers a framework for resolving these tensions while not imposing, ignoring, or ruling out possible technical or policy solutions.

We thus conclude that a prerequisite for achieving enhanced cybersecurity is articulating a *cybersecurity doctrine*, which specifies *goals* and *means*.

- *Goals* define what system properties will be preserved as well as what policies will be enforced, for whom, at what costs (monetary expenses as well as costs to convenience and compromised societal values), and against what kinds of threats. Goals might be absolute, or they might specify a range of permissible trade-offs. In allowing trade-offs, we acknowledge the political nature of cybersecurity and the need for conversations among those affected when goals are set.
- *Means* might involve technological, educational, and/or regulatory *measures*. We should expect means to include policy that creates incentives – which might range from market-based to coercive – that foster adoption and/or deployment of the measures proposed.

Through incentives provided as part of its means, a cybersecurity doctrine can address barriers to market production of cybersecurity that reflect a lack of will rather than a lack of ability, as others have aptly noted.[1] Incentives can also prompt continued improvement to address the constantly emerging landscape of threats and the new needs that arise as a growing range of applications is being migrated to networked information systems.

Our doctrine of *public cybersecurity*, the subject of this essay, is rooted in the thesis that cybersecurity is a public good. The doctrine focuses on the collective interest rather than on any single individual's or entity's computer, network, or assets. It can be compared with public health, another public good, and we make this comparison below.

We begin by analyzing the limitations of various cybersecurity doctrines that have been proposed. Next, we discuss in detail our new doctrine of public cybersecurity. Then we describe how to support public cybersecurity, starting with approaches to building systems that have fewer vulnerabilities. Subsequent sections explore approaches for managing insecurity: diversity, surveillance, installation of patches, isolation, and the role of intermediaries. Finally, we put this work into a larger perspective and offer some conclusions.

The advent of time-sharing in the 1960s meant that computations on behalf of multiple users were interleaved on a single computer. Each user's computation and data thus had to be protected from misbehavior by programs run by other users. Confronted by a problem born of technology, engineers of early time-sharing systems sought solutions in technology. Therefore, early cybersecurity doctrine focused on developing new technology. Societal values could be and were ignored because the doctrine respected the shared values of the small population that used these early computing systems.

Technological solutions for creating the needed isolation were beyond early capabilities, especially when users could be motivated, capable adversaries bent on disrupting another user's computation or stealing information. Improved technology, however, is not the only way to solve

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

problems that technology has created, and subsequent cybersecurity doctrines focused on policy to leverage those technological solutions that were at hand. But these doctrines, too, were unsuccessful. Even if they had succeeded, they ultimately would have been inadequate because the problem was changing.

Computer systems were becoming pervasive, which had two broad consequences. First, the information technology sector became a significant economic force, raising concerns about the freedom to innovate and success in the marketplace. Second, computer systems increasingly touched the lives of ordinary people. Citizens' records were stored electronically, and information technology allowed workers to be more efficient. Eventually, computer networks and the Web changed how people shopped, communicated, socialized, and engaged in politics. As a result, privacy and other societal values have become crucial considerations in developing cybersecurity doctrine.

Because their effectiveness has been limited, it is instructive to review the three doctrines – prevention, risk management, and deterrence through accountability – that have dominated cybersecurity thinking for the past fifty years. In particular, analyzing the measures each doctrine proposes offers insights into properties that affect whether a cybersecurity doctrine fails or succeeds.

Doctrine of Prevention. The goal of this doctrine is to render systems completely free of vulnerabilities. Absent vulnerabilities, attacks are not possible, so the system is secure. Such *absolute cybersecurity*, though a worthwhile undertaking, is unlikely ever to be achieved.

To secure systems that incorporate humans as users and operators, we would need some way to prevent social engineering attacks and intentional insider malfeasance. Here, prevention requires overcoming the frailty of humans, which is likely to involve more than technology.

If we ignore human involvement in a system, then the problem is different but no less challenging. Software systems today are too large and complicated to be verified using formal logics. Researchers, assisted by computers, have been able to devise formal proofs for small systems[2] (those with fewer than ten thousand lines of code), and software producers regularly employ automated checking for analyzing specifications as well as for relatively simple properties of large bodies of code. For the reason that smaller code bases are more amenable to formal verification, techniques to reduce the size of the code bases for certain key systems are also being explored.[3] But revolutionary advances are needed before formal verification could be used to validate the entire code base that runs a desktop system or server. Thus, this approach to prevention is not a practical solution for the near term.

System testing is the clear alternative to ensure that a system has no vulnerabilities. Tests, however, can reveal only the presence of vulnerabilities – not their absence. Demonstrating the absence of vulnerabilities requires exhaustive testing; the amount of work involved is prohibitive even for small components, much less for large systems.

Formal proofs and testing are performed relative to some expectations about what the system must do and the environments in which it will operate. In other words, the doctrine of prevention establishes the absence of vulnerabilities only for settings where certain assumptions hold. Unfortunately, reasonable assumptions about the environment today might subsequently be invalidated. Attacks evolve in sophistication in response to better defenses. Threats emerge to exploit new opportunities for disruption

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

that are created when cyberspace provides access to new forms of value. Therefore, a system that is deemed to be secure might not remain so for long.

In light of this dynamic, expectations about the environment must be periodically revisited and, if necessary, revised. Thus, the doctrine of prevention involves a recurring expense. That expense is inconsistent with the business model employed by many of today's software providers, which favors reuse and extension of existing hardware and software in order to lower the cost of producing new systems.

The adoption of mandatory standards can be viewed as a way to support the doctrine of prevention because implementing standards increases the chances that what is built and/or deployed will have fewer vulnerabilities. Some standards concern functions an artifact must or must not support; some govern its internal structure; others prescribe the process by which the artifact is constructed or maintained; and still others stipulate qualifications the personnel who are involved in creating the artifact must have. Examples include the Department of Defense Trusted Computer System Evaluation Criteria, or TCSEC[4] (also known as the Orange Book); its successor, the Common Criteria for Information Technology Security Evaluation[5]; security provisions in information privacy laws[6]; the Federal Information Security Management Act[7]; and the Voluntary Voting System Guidelines.[8] Current market activity suggests that such mandates show value in some areas. However, a correlation between the absence of vulnerabilities and compliance with standards has not yet been documented. The stated goal for the doctrine of prevention is unlikely to be achieved through these measures.

Doctrine of Risk Management. Absolute cybersecurity is cost prohibitive, but fortunately, it is unnecessary for most systems. The doctrine of risk management stipulates a more modest goal: that investments in security reduce expected losses from attacks. To adopt this doctrine is to admit that all vulnerabilities are not equal, that one should focus only on vulnerabilities whose exploitation (i) is sufficiently likely to occur based on perceived threats and (ii) could enable expensive (by some cost measure) system compromises. In contrast to the doctrine of prevention, the objective of defending a smaller body of code against a more restricted set of threats is likely within our capabilities. Moreover, maintaining that steady state would require fewer assumptions about the environment to be revisited periodically.

In theory, the doctrine of risk management seems sensible. But a lack of information about vulnerabilities, incidents, and attendant losses makes actual risk calculations difficult. Companies and individuals do not know how to value (i) confidentiality of information, (ii) integrity of information, or (iii) the pain of dealing with recovery from an attack's effects (bad credit ratings, for example). These costs appear to vary tremendously. Also, people have difficulty reasoning about extremely low-probability events. Finally, when costs are borne by third parties, investment incentives materialize only well after a breach occurs, and causation is difficult to discern, much less prove.

Accurate information about threats and attacks may not be publicly available because those with that knowledge fear tarnishing their reputations or compromising their intelligence methods and/or sources. Even were that information accessible, deployment of replacement systems and upgrades would alter the functions that systems perform and the set of relevant vulnerabilities, which, in turn,

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

could lead to new attacks. These differences mean that the past is not a good predictor of the future. As a consequence, actuarial models cannot be constructed, and insurance to transfer risk is impossible to price in a way that ensures profits to the policy underwriter.[9]

Were there a method to analyze a system mechanically and obtain a quantity that indicates just how secure that system is, we could form a basis for assessing what is gained from specific investments in cybersecurity. At present, such cybersecurity metrics do not exist. Neither can investments be justified by quantities derived entirely from empirical observations. The absence of detected system compromises could indicate that investments in defenses worked, that attacks have not been attempted, or that the compromise escaped notice (as in theft of confidential information, for example). Whether prior security investments were well targeted is impossible to know; such ambiguity leaves security professionals to justify investments based solely on non-events.

Risk management approaches are further confounded by externalities that arise from the emergent nature of cybersecurity in networks. Individuals and entities employing these tactics can neither fully reap the benefit of their security investment nor entirely control their vulnerability through investments.[10] For example, a single compromised system anywhere in a network can serve as a launch point for attacks on other systems connected to that network. Thus, local investment in defenses not only provides local benefits but also benefits others; likewise, underinvestment in defenses elsewhere in the network could facilitate local harms. Absent coordination, the sole logical strategy would be to invest in insurance (if it were available); only with insurance can an entity reap the entire benefit of its investment.[11] However, this strategy does nothing to improve security, and as noted above, viable long-term business models for insurance do not exist today.

The outlook for risk management is not entirely bleak. In the policy arena, state security breach notification laws[12] are a form of risk management intervention. Significant costs are incurred to notify individuals and to manage the adverse publicity surrounding reportable breaches. These potential costs act as a proxy for the costs of security failures to customers, forcing companies to internalize previously externalized security failures. The price tag on breaches also means that these laws have created a set of data to use in calculating risk and return on investments. Nonetheless, current laws focus on only a narrow set of breaches and, as a result, might artificially skew investments.

Doctrine of Deterrence through Accountability. This doctrine treats attacks as crimes; it focuses on infrastructure to perform forensics, identify perpetrators, and prosecute them. In theory, attacks are deterred by increasing the chances that perpetrators will be found and convicted.[13] Implementations of this doctrine require strong authentication technologies and surveillance of network activity. Robust forms of user identity would allow us to overcome the loose binding that exists today between individuals and machines.[14]

Absent an effective means for retribution, this doctrine has no teeth, and fails as a result. Moreover, punishment of perpetrators of cyber-attacks is not always feasible in today's global environment. Attribution of actions by machines to individuals is complicated, agreement about illegality is illusive, and cross-border enforcement requires more cooperation than is likely to emerge between nations. Recent attacks against U.S. and other systems suggest that we cannot ignore

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

non–nation-state actors that engage in terrorism and large-scale financial crimes. The very features that make the Internet a profitable environment for criminals – worldwide reach, connectedness, neutral treatment of packets, and weak binding of machines to individuals – make it difficult for law enforcement to identify and catch perpetrators. Other features of the international landscape complicate efforts to bring them to justice.[15]

Conceptual obstacles also limit the effectiveness of the doctrine. First, the doctrine is punitive. Like most criminal law, it is aimed primarily at using punishment to produce both general and specific deterrence. This approach does little to keep networks up and running when they are under siege, nor does it prompt proactive security investments. Second, the doctrine could require individuals to sacrifice privacy and, in the extreme case, abandon the possibility of anonymity and the protections for freedom of speech and association that it affords.

Nevertheless, many attacks are indeed carried out by criminals plying their trade. We are likely to benefit if criminal activity in cyberspace faces the risk of retribution that we employ to deter crime in the physical world. Thus, the doctrine of deterrence through accountability has value. But in cyberspace, unlike in the physical world, terrorists or state actors are difficult to distinguish from common criminals.[16] Deterrence through accountability is not necessarily effective against these transnational threats; other doctrine is also required.

Cybersecurity is non-rivalrous and non-excludable; by definition, therefore, it is a *public good*. It is non-rivalrous because one user's capacity to benefit from the security of a networked system does not diminish the ability of any other user to enjoy the same advantage. It is non-excludable because users of a secure system cannot easily be excluded from the benefits security brings. Measures intended to foster the production of public goods thus constitute a sensible starting point in our search for doctrines that promote cybersecurity.

Economists define a *common good* as one that is rivalrous and non-excludable. The sea, outer space, and air are examples. Insofar as common goods are inherently different from public goods, doctrines for common goods are likely unsuitable for enhancing cybersecurity. Indeed, this irrelevance is apparent in laws for protecting common goods, which typically aim to ensure rights of equal use, and in the mechanisms these laws introduce, which are intended to manage the depletion and inequitable consumption by first-comers or more sophisticated users. The production of cybersecurity bears little relation to these issues.

Public health – the prevention of disease and promotion of good health in populations writ large – is a public good. It is non-rivalrous because having a healthy population implies a lower prevalence of disease, which in turn decreases the chances any member will fall ill. It is non-excludable because no one can limit an individual's ability to profit from the health benefits that living among a healthy population brings.

The essential characteristics of public health law are a focus on the health of the population as a whole and the singular role of governments in that enterprise.[17] To discharge these responsibilities, the law authorizes various agencies to engage in a broad set of activities, including:

- Public education about the causes and effects of disease, as well as methods of prevention. Education empowers individuals to act in ways that optimize their own health, which in turn furthers public health.

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

- Creation and use of methods for the prevention and treatment of specific diseases. Methods could involve (i) providing subsidies to procure care needed by those who could not otherwise afford it[18] or (ii) imposing specific health standards as eligibility requirements for receiving various societal benefits (for example, public education, which we discuss below).
- Identification and management of disease and infected individuals though surveillance, information gathering, and analysis. Methods include mandatory reporting requirements for certain diseases and conditions, mandatory testing or screening for others, symptom surveillance that seeks to identify obscure public health threats in masses of routine records, and mandatory treatment.

The interests of individuals and the public often align; public health law speaks to the points of conflict. It offers frameworks to mediate tensions between the rights of individuals as sovereigns over their physical bodies and the obligation of the state to protect the population as a whole.

For example, public health mandates that children be vaccinated because, in a generally healthy population, such vaccinations cannot be justified based on the benefit to the individual. In fact, the optimal choice for any given child might be to avoid vaccination and thus avoid the risk of side effects. Mandatory vaccination creates *herd immunity*, which benefits the collective by reducing the total number of hosts available to carry a disease, thereby decreasing the risk to individuals who have not been vaccinated. However, if too many individuals act in self-interest and eschew vaccinations, then the herd immunity that gives some protection to the unvaccinated may disappear. This is a "tragedy of the commons" whereby individuals acting rationally leave everyone worse off.

Every state in the United States conditions a child's attendance at school on satisfying some specified regimen of vaccinations. In addition, vaccine manufacturers are indemnified from liability for side effects users might experience. Specifically, the Vaccine Injury Compensation Program (VICP) provides certain payouts from public coffers to children injured as a result of a vaccine. VICP also offers a compensation mechanism outside seeking damages from the vaccine manufacturers, thereby establishing an environment conducive to both the production and willing use of vaccines. While the program cannot fully compensate for negative health consequences from vaccinations, it is an important component of the overall public health strategy.

Public health is a logical outgrowth of disease detection and prevention mechanisms, which transformed societal perception of health from a primarily private concern to a concern of the collective. Ultimately, this development led to the perception of public health as a public good that the government should enable.[19] Health now becomes intertwined with societal values. For example, we impinge on societal values when we introduce mandatory reporting and surveillance systems that (i) alert individuals at specific risk so they can be tested and treated and (ii) allow isolation, quarantine, and even mandatory treatment to be imposed. At the same time, public health interventions aim to minimize their intrusiveness because of the chilling effect that may have on access to health care; for example, we see anonymous HIV testing and needle exchange.

This public health framework (of laws, agencies, and measures) applies equally well to weaponized pathogens. This is not

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

to suggest that motive is irrelevant in considering public health strategies. For example, weaponized pathogens may change more quickly than those that evolve in nature, and certainly, the transmission vectors may differ when pathogens are used as weapons. But the basic tools of public health – public education (to minimize exposure and facilitate early detection), investments to create means for prevention and treatment (antidotes and vaccines), and surveillance and analysis (facilitating isolation and quarantine as defenses) – still apply.

Both public health and cybersecurity aim to achieve a positive state (health or security) in a loosely affiliated but highly interdependent network. The former is a network comprised primarily of people existing in an environment over which they have some limited control; the latter is a network of people, software, and hardware (for communications, storage, and processing). Given that the positive state is ultimately unachievable, both struggle with how to manage in its absence as well as with how to work toward attaining it. Success ultimately depends not only on technical progress but on reaching a political agreement about (i) the relative value of a public good in comparison to other societal values and (ii) the institutions granted authority to resolve conflicts (and the methods they use).

We define a *doctrine of public cybersecurity* to be any cybersecurity doctrine whose goals are (i) to produce cybersecurity and (ii) to manage insecurity[20] that remains, where political agreement balances individual rights and public welfare. There is no single doctrine of public cybersecurity, for the reason that there are different meanings attached to "cybersecurity" and "insecurity."[21] Also, different choices of measures and incentives result in different doctrines of public cybersecurity. Notice, though, that none of the doctrines discussed above has all the elements we require for a doctrine of public cybersecurity.[22]

The analogy to public health inspires cybersecurity measures such as prevention, containment, mitigation, and recovery – that is, strategies that direct resources toward production and preservation of cybersecurity. But modern public health doctrine does not compensate victims of disease; thus, a parallel doctrine of public cybersecurity would not focus on restitution. Indeed, restitution is economically efficient only when attacks are infrequent, and that assumption cannot realistically be made today.

Furthermore, modern public health does not punish victims of disease, but there is some nuance. Using quarantine to limit the spread of disease benefits the collective by depriving an individual of certain freedoms. Such a response could be considered a "harsh consequence," which is one definition of "punishment." By analogy, a doctrine of public cybersecurity could dictate responses that deprive individuals of actions, but only if those responses benefit the collective. Punishments solely for retribution could not be part of a public cybersecurity doctrine (because retribution does not benefit public welfare); however, nothing precludes implementing a doctrine of public cybersecurity alongside a cybersecurity doctrine that incorporates retribution. Finally, the parallel with public health also suggests that prevention be preferred to recovery.

With regard to incentives, ensuring that actors contribute to public cybersecurity requires interventions to overcome positive and negative externalities that lead rational individuals to underinvest, as occurs in public health. When incentives are insufficient to motivate private provisioning, the public interest re-

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

quires making value-ridden choices to interfere with the rights and interests of individuals and organizations. Those choices are embodied in goals that reflect political agreement about how to define the good in question; the socially desirable level that should be maintained, given competing priorities and values; and provisions for determining when the individual's desires yield to the collective's need. For example, an agreement might stipulate that state coercion is permitted only when certain incursions into the rights and interests of individuals are tightly circumscribed.

Public health solutions do not always translate into sensible support for public cybersecurity, but the former often inspires strategies for the latter. The examples we explore below illustrate how doctrines of public cybersecurity can be useful for evaluating current cybersecurity proposals. Our choice of examples should not, however, be seen as an endorsement for any particular proposed set of interventions.

Underutilized approaches (formal methods, testing, and improved software engineering processes and standards, for example), developed in part to serve the doctrine of prevention, are effective in producing cybersecurity (by reducing the number of vulnerabilities present in a system), even if they cannot produce absolute cybersecurity. Thus, existing methods could serve as a means for a doctrine of public cybersecurity, just as disease prevention through vaccination and the monitoring of our food and water supplies fosters public health. The question is: What incentive structures would ensure that these methods are used?

Education could play a key role in defect reduction. Knowledgeable developers are less likely to build systems that have vulnerabilities. They are also better able, and thus more likely, to embrace leading-edge preventions and mitigations. There is, however, no agreement about what should be taught. Reaching such a consensus would require a dialogue among universities and practitioners.

Were there an agreed-upon body of knowledge for cybersecurity practitioners, mandatory certification could ensure that practitioners master that material as a condition for practice. But the details of how certification is handled can be subtle. Possession of a certificate does not by itself compel the use of best practices, and it is easy to imagine certified system-builders who cut corners by choice (out of laziness, for example) or by mandate (because management is trying to reduce costs). Moreover, unless the certification process imposes a continuing education requirement to ensure that certificate holders stay current with new developments, it might impede rather than promote the spread of innovation. Even when continuing education is mandated, old habits die hard; for example, physicians who have been shown new methods that are empirically demonstrated to be superior nevertheless tend to stick with familiar practices.[23]

Utilizing techniques to reduce defects during system development and employing better-educated practitioners will mean that systems become more expensive to produce. Today's software-procurement market does not provide developers with compelling incentives to incur those additional expenses. Moreover, purchasers are unable to predict the costs of a system's vulnerability to attack and, without ways to measure a system's security, cannot rationalize paying higher prices. The doctrine of risk management failed for the same underlying reasons.

Law could force system producers and/or purchasers to make the necessary investments. Software distributors cur-

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

rently disclaim liability beyond the purchase price for damages caused by their products. This practice probably reduces the time and energy that developers devote to eliminating defects, as evidenced by the number of buffer overruns and other exploitable coding errors still being discovered and exploited by attackers. Existing law could, for example, be revised to disallow limits on damages flowing from attacks taking advantage of poor coding practices that lead to buffer overflows and other easily exploited vulnerabilities. Limits on liability could depend on the use of formal methods, type-safe languages, or specific forms of testing (such as fuzz testing[24]). Creation of a class of certified security professionals could also provide the basis for a professional duty-of-care supporting liability for shoddy security.

Furthermore, law could require that software developers adhere to security standards. Alternatively, safe harbor provisions could be created to protect software developers against future findings of liability for those systems built according to specified standards. In fact, the law arguably already mandates that companies follow certain standards regarding personally identifiable information. Through a series of settlement agreements, the Federal Trade Commission established a de facto standard that requires a company collecting and handling the personal information of consumers (i) to establish reasonable security processes and (ii) to mitigate system vulnerabilities that are known in the marketplace and for which mitigations exist. A first step in determining whether law should more broadly mandate the adoption of security standards might be research that identifies connections between security development processes and positive security outcomes.[25]

Monocultures in nature risk extinction from pathogens and are less able to adapt to changing conditions. *Diversity* – of the individuals within each species and by virtue of many species coexisting within an ecosystem – creates a resilient ecosystem. By extension, public health benefits from individuals in a population having different inherent resistance to pathogens and, by virtue of different exposures[26] to diseases, having different immunities.

Although nature abhors monocultures, cyberspace seems to favor them. A collection of identical computing platforms is easier, and hence cheaper, to manage because it demands that users master only one interface and managers make only one set of configuration decisions. In addition, user-training costs are reduced when job transfers do not have the overhead of learning another operating system and suite of applications; in a monoculture, investments in educating system users or managers can be amortized over a larger user base. Finally, a monoculture facilitates networking: interoperability of a few different kinds of systems is far easier to orchestrate than integrating a diverse collection, standards notwithstanding. Mindful of these advantages, both the public and private sectors tend to adopt procurement policies that foster creating computer monocultures.[27]

Methods exist, however, for artificially and automatically creating diversity in software systems without sacrificing the advantages a monoculture provides. These methods involve tools that randomly transform code and/or stored information while preserving its semantics. Once such *artificial diversity* is introduced, internal details of an individual system are no longer predictable. Thus, an attack that depends on knowledge of internal details is more likely, after a small number of instructions, to cause a system crash than to give an attacker con-

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

trol of that system. In many settings, a system crash is preferable to attacker control. Moreover, a platform that crashes in response to an attack cannot then help spread that attack to other platforms. By (implicitly) signaling to system operators that something is wrong, a crash also creates an opportunity for initiating other means to block an attack's spread.

Like the diversity found in nature, artificial diversity is inherently a probabilistic defense. An attack against any one component might not be derailed by the specific random transformations that were made to that component. Also, by converting some attacks into crashes, artificial diversity can adversely affect a system's availability.

Despite these limitations, artificial diversity facilitates public cybersecurity by providing a means to cope with residual vulnerabilities, thereby supplying a way to manage insecurity. Today, artificial diversity is used often in operating systems but less so in applications[28] (even though, increasingly, it is applications that attackers target). However, the various legal approaches (discussed above) for incentivizing defect reduction during development are equally well suited for incentivizing system producers to support artificial diversity. There is no shortage of incentives at hand for encouraging broader adoption of the measure.

Public health relies extensively on *surveillance*. Data collected through a variety of means enable disease containment and mitigation through:

- dissemination of information that facilitates individual actions;
- isolation and quarantine, which limit the interaction of affected individuals with the rest of the population to avoid exposure to infection; and
- mandatory treatment to reduce danger to the public.[29]

Data collection for public health occurs at many levels. At the lowest level is the inclination of individuals to assess their own well-being. Education equips individuals with a basic level of knowledge about health indicators – normal body temperature, pulse, blood pressure, and respiratory rate – as well as with simple precautions to limit infection and the spread of disease (frequent hand washing, for example). Primary care providers collect other data in conjunction with annual check-ups and, when symptoms require further analysis, at hospitals and other more advanced diagnostic facilities. Each successively higher level is concerned with the overall health of a larger population and thus provides a natural venue for constructing and analyzing larger data aggregations.

By minimizing disclosure of information about an individual's health, public health law strives to reduce one potential deterrent to seeking health care: that is, an individual's fear of being shunned because of a publicized health condition. In general, identifying information should flow away from primary health care providers only in instances where aggregation and/or analysis is necessary to identify significant trends. Even in this case, efforts are undertaken to protect individuals' privacy.

In contrast to public health, cybersecurity is not supported today by extensive coordinated surveillance, yet it would be feasible and advantageous to do so. Low-level indicators about the basic "health" of a computer can be made available by running built-in checking software (such as virus scanners and intrusion detection systems). Each of the Internet Service Providers (ISPs) that constitutes the network has an infrastructure that facilitates

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

monitoring of events internal to its network as well as interactions with other networks.

Surveillance of network traffic (including volume, distribution over time, and destinations) could be a powerful potential source of information about certain attacks and vulnerabilities. Denial-of-service attacks, for example, have a clear manifestation and a natural mitigation based on traffic filtering by ISPs. However, the source(s) of such attack packets, the target(s), and the intermediaries are likely to span multiple ISPs, which would have to share data and coordinate for mitigation. Unfortunately, data sharing among ISPs today is inhibited by competition and, in some cases, varied interpretations of privacy law.[30] ISPs thus do not always have the situational awareness that would enable them to suppress packets delivering attacks. Widespread sharing of information, however, can introduce a risk by increasing chances that attackers learn about vulnerabilities for specific sites.

Just as there are privacy issues with collecting data about an individual's health, network traffic surveillance raises privacy concerns. The extent to which collecting packets actually impinges on privacy depends on what information is recorded, how long it is stored, how it is used, and who can access the information. For example, real-time responses to protect networks can be accomplished by authenticating machines, a far less politically fraught solution than proposals for "Internet drivers' licenses" and other tight bindings between machines and individuals.[31]

ISP cooperation and information sharing is less likely to raise privacy concerns than the collection of information by centralized government organizations. Yet given that defense of its citizens is a clear responsibility of government, seeing the packets themselves can be invaluable to a government seeking situational awareness about threats in cyberspace. Unfortunately, packet inspection is also easily abused if a government intends to spy on citizens; critics cite this fear (among others) when discussing the Einstein[32] systems recently deployed by the U.S. government for monitoring connections to the Internet at federal civilian agencies. As with public health, political agreement must weigh the expected benefits of surveillance (backed by sound research and field experience) against the risks it poses to other values.

An understanding of the kinds of vulnerabilities found in systems is a form of situational awareness of potentially great value to system builders. In the absence of mandatory reporting requirements for cybersecurity incidents, diverse public and private reporting mechanisms have evolved. The U.S. Department of Homeland Security's Computer Emergency Readiness Team (US-CERT) and the National Institute of Standards and Technology (NIST) Computer Security Division maintain databases of common vulnerabilities. Many organizations contribute, but these are not the only such databases and none provides more than a partial view. Some vulnerabilities never reach the public databases. For example, a private-sector community of "security researchers" report their findings on system vulnerabilities to middlemen, who offer the information for sale to companies that build and sell anti-malware or intrusion prevention/detection products.[33] Yet the ad hoc Conficker Working Group[34] is an example of a rather successful coordinated private-sector activity involving information sharing about risks.

A *patch* is an update that can be applied to an installed system in order to eliminate one or more previously identified vulnerabilities. Exploitation of an unpatched

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

vulnerability on a computer could target that machine, and an individual's assets contained therein, and therefore be fully internalized as a result. Alternatively, the exploitation could target the machines and systems of others, producing a negative externality.[35] The uncertainty about consequences means that self-interest is not a strong incentive for machine owners to apply patches.

Various policy interventions could raise patch rates. Choosing among them requires additional information about why people and businesses delay or outright fail to apply patches.

- Research might conclude that low patch rates in the consumer market are caused by an underappreciation of the risks. Public education to inform individuals that applying patches improves cybersecurity might dramatically increase patching.
- We might find that individuals lack awareness of vulnerabilities present on their machines. Here, built-in software to check whether all current patches have been applied might suffice for triggering consumers to be more attentive to downloading patches for their machines.
- Feedback about what others do could create new behavioral norms that might lead to better patching practices. Researchers in other areas have found that showing individuals how their behavior compares to others' taps into competitive and/or social consciousness. Simply stating that a significant percentage of others have patched their machines, and are thus doing their part for cybersecurity, might push laggards into applying patches.
- Research might find that individuals or enterprises hesitate to install patches for fear of destabilizing their other software. Greater transparency about the specific configurations and applications software vendors have tested might help individuals overcome their reluctance. The fears of enterprises that depend on homegrown software might be somewhat assuaged by providing test suites to patch developers.[36] As a final safety net, all software could be required to contain mechanisms whereby a patch that has been applied can easily be removed and the system and data restored to the pre-patch state.
- If the impediment to installing security patches is time or expertise, then vendors could mitigate the problem by configuring defaults that automatically download and apply security patches.
- Another reason consumers forgo installing patches could be that they are charged for Internet access in proportion to the amount of bandwidth they use and will incur lower costs by not downloading patches. Here, one solution is to subsidize the bandwidth required for such downloads; another is to introduce tariffs that distinguish between different kinds of traffic.

Those who run pirated software might hesitate to install patches for fear that the installation process would disable the illegal software or detect and report it. Regulations could address this obstacle by prohibiting security-patch installation from implementing functionality in support of license enforcement or any other form of intellectual property protection. Thus, even pirated software could be patched, so that herd immunity can be achieved.

Incentives to apply patches could also have a useful indirect effect. If patch installations are frequent and disruptive, then consumers have reason to prefer products with fewer security vulnerabilities. Consumer demand would then pres-

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

sure software producers to build and deploy more secure products.

Mandates to apply patches raise concerns about subsidizing the installation of patches[37] and compensating injured parties when patches cause harm. Losses from applying mandated patches, particularly where unacknowledged and uncompensated, will breed suspicion and resistance to patching efforts. Thus, it seems advisable to consider backstop measures, analogous to what VICP provides to incentivize the use and production of vaccines and the process used by the Food and Drug Administration to ensure vaccine efficacy.

Geological features, such as mountains and oceans, have proved valuable in protecting individuals and populations. When natural boundaries are absent, we build our own: fences surround buildings and nations, often with guards to control who is allowed to transit the border. Such boundaries protect activities on one side from activities occurring on the other. A boundary may limit travel in one direction or in both directions; it may be entirely impervious or may selectively limit who or what may pass. Neither is a panacea: an impervious boundary could bar the good with the bad; a selective boundary must employ some kind of *filter*, and that filter might block what should not be blocked or let pass what should.

Firewalls, so-called network guards, intrusion detection/prevention systems, and "air gaps" are examples of mechanisms that implement boundaries in networked systems.[38] Data collected through surveillance can serve as the basis for *signatures*, which are then used to define filters, effectively creating dynamic boundaries. Surveillance thus can lead to automatically imposed quarantines. Given that attacks in networks propagate rapidly, automatic response is especially attractive.

Ideally, we would deploy selective boundaries that block only attacks. In practice, though, filters will be far from perfect.

- Filters that inspect packet payloads (known as *deep-packet inspection*) in addition to checking packet headers are ineffective when packet payloads are encrypted or otherwise obfuscated. Encryption is not used extensively in networks today, but that could easily change. Attackers often use encryption to evade detection. Moreover, what is being spread are often malware variants, where each variant is obfuscated by the application of a different random set of semantics-preserving transformations. It is difficult and often impossible to construct a signature that matches all variants by generalizing from a few.
- A filter might be designed either to (i) block packets and protocols corresponding to known attacks or (ii) pass packets from protocols or conveying content that is known to be normal. Filters that implement (i) are fooled by new attacks (in addition to suffering limitations described above) and those that implement (ii) could block previously unseen protocols and kinds of packets, thereby stifling innovation.
- Whether a packet is part of an attack could depend only on sender intent. Consider a large number of request packets being sent to a Web server. Are many people trying to access the same particularly topical content, or is a denial-of-service attack in progress? Sender intent is the sole differentiator.

There is also a human element to consider. Boundaries and filters must be installed, configured, and managed by human operators, and people make mistakes. Moreover, when such a mistake allows

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

unimpeded flow, then the error might be difficult to detect until it is too late.

Network providers are understandably reluctant to publicize details of defenses, because revealing that information could help attackers. Yet we see example defenses in today's commercial networks, which create and reference "black lists" of sites whose communications will be ignored and "white lists" of sites that are known to be trustworthy. Some ISPs create a competitive advantage by offering their customers a service whereby suspicious inbound-traffic spikes directed at the customer's site will automatically prompt upstream filtering to block those suspicious packets. As a result, denial-of-service attacks in such networks are more difficult to undertake. Other ISPs monitor each endpoint, disconnecting a given endpoint if outgoing traffic suggests that the endpoint is compromised.[39]

A boundary may be deployed around a system (be it a single computer or a network) that must be protected from attacks, or around a system that is likely to harbor attackers. Different incentives are effective in each case. One natural scenario for direct government investment exists when security boundaries and national ones overlap. Systems in various countries are subject to different laws, typically reflecting a range of societal values. A government might therefore justify installing a boundary whenever systems subject to its laws are connected to systems located in a jurisdiction that allows system behavior the first considers an attack.

Boundaries are more likely to be accepted and work effectively when initiated by the collective rather than by individuals. First, an individual is unlikely to have the necessary authority to mandate changes to defenses on all the remote systems that could be involved in creating a quarantine. Second, the possibility of free-loading limits the incentives for owners or operators of networks or individual systems to make the investments to support enforced isolation. Finally, an agent of the collective, equipped with a broader view of system vulnerabilities, would define better signatures for filters.

An example of such boundaries is found in recent proposals for deterrence through accountability. Some have suggested that the Internet be partitioned into national or multinational enclaves. Those enclaves that serve the population whose network security is of concern (i) run protocols that enable packet-sender tracing and (ii) do not carry traffic from enclaves where packet-sender tracing is not supported or cannot be trusted. The ability to trace attack packets back to an individual machine enables support for accountability in those enclaves that serve the populations the boundary is intended to protect.

Boundaries with sufficiently powerful filters have the potential to intrude on societal values. One concern arises when the defining filters not only block packets that contain attacks but can be configured to block other kinds of packets. Such a filter could be used to prevent data from leaving an enclave, which makes it well suited for protecting confidential information against theft. But content filters also permit government censorship, as illustrated by the firewalls China has installed to protect that nation's computing systems from receiving information in violation of local laws regarding allowed speech. Deployed in the reverse direction, a content filter could block someone from sharing information with others, thereby stifling debate.

So there are trade-offs, with social values and potential benefits for the collective requiring constraints on activities by individuals and businesses. Moreover, no criteria for deciding where a system should be segregated will be infallible.

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

The result is a complex risk-management decision procedure that society must prescribe, with imperfect information and unknowable consequences.

The public health system leverages health professionals and other institutions to influence individuals' behavior. For example, health professionals educate individuals about the benefits of vaccinations, schools demand conformance with vaccination schedules, and airports screen passengers for symptoms during some infectious disease outbreaks. Intermediaries clearly play an important role in public health strategies.

Intermediaries also have an important part in fostering cybersecurity. For example, many network operators, such as employers and universities, require that all machines on their networks run virus detectors or malware detectors (with up-to-date signature files). These intermediaries could require that all machines are up-to-date on security patches. Similarly, some ISPs have chosen to notify subscribers when a computer appears to be infected.[40] At least one ISP restricts Web surfing until the infected machine is cleaned up, while another ISP reportedly quarantines any compromised machine until it is clean.[41]

ISPs are well positioned to facilitate patching and, by monitoring traffic, to enforce isolation of machines harboring certain malware. Yet they currently have little incentive to engage in such practices because they would then incur the bulk of the security costs, but any costs from infected machines would be more widely dispersed. Moreover, an ISP that disables or limits a machine's access to the Internet will likely bear the burden of assisting that customer as she attempts the necessary repairs. Analysis[42] suggests that the cost incurred by an ISP in fielding a customer's tech-support call approaches the ISP's annual revenue from that customer. Making this sort of monitoring and clean-up a mandatory obligation for ISPs would not only force action but would also prevent consumers from contracting with ISPs that enforce weaker security requirements.

More daunting are the potential costs an ISP might incur from making an incorrect decision to disconnect a customer.[43] To limit spam email, for example, an ISP might block all bulk sending of email. But missives sent by a political organization might then be blocked, resulting in unwanted attention from advocacy groups and the press.[44] While the law is evolving to provide ISPs that take steps to protect security with immunity from suits brought by providers of malware, those users who experience losses after installing required patches or system upgrades, or who suffer because of isolation, might also file legal complaints. In sum, the costs of ISP intervention present a formidable barrier to such action; nonetheless, the law could remove these disincentives.

In a recent proposal,[45] legal scholars Doug Lichtman and Eric Posner argue that expanding ISPs' liability "for violations of cyber-security" would improve cybersecurity because (i) individual attackers are often either beyond the reach of the law or are judgment-proof and (ii) ISPs "can detect, deter, or otherwise influence the bad acts in question." Similar to the way cardholder purchases are monitored by credit card companies, ISPs could detect cybersecurity violations by building profiles of their users and looking for traffic anomalies. But as Lichtman and Posner openly admit, anomaly detection with usable levels of fidelity has eluded cybersecurity researchers for decades. Thus, implementing the proposal is not feasible at present.

Still, a policy holding ISPs liable for the damage caused by infected machines running on their networks might encourage

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

more diligence in monitoring and fixing their subscribers' machines. The details are subtle and depend on the standard for liability – whether strict, knowledge-based, or otherwise defined. To complicate matters, the policy could have an undesirable outcome: the ISP could undertake less monitoring as a way to avoid its duty to intervene.

Alternatively, governments could provide indirect or direct subsidies to foster cybersecurity-preserving activities by ISPs. For example, creating a centralized service for hosting patches or subsidizing bandwidth to all endpoints could ensure that cost or delay to download a patch would not become an impediment to installing that patch.

Given the decentralized and private provisioning of network resources in the United States and many other countries, understanding the role of intermediaries in driving cybersecurity is essential. As in other areas, such as copyright, the challenge is to establish policies that incentivize desirable behavior while minimizing impact on other values.

Computer scientists have discussed a biological basis for cybersecurity for at least two decades. The thrust of that research is to understand whether computer networks can benefit from implementing defenses similar to those that protect living things. Developers have explored intrusion detection systems that mimic pathogen detection in the human immune system[46] and software defenses based on artificial diversity.[47] A recent Department of Homeland Security white paper[48] describes how a human immune system's response mechanisms might serve as the blueprint for software that defends individual computers and networks against cyber-attacks. Much research remains to be done, however, before those ideas are reduced to running code.

In contrast to the biological metaphor, which focuses on technical measures for blocking cyber-attacks, the analogy between public health and cybersecurity is primarily concerned with new policy and new institutions. Proposals for a Cyber-CDC, for example, have attracted considerable interest.[49] Inspired by the existing Centers for Disease Control and Prevention, the Cyber-CDC is envisioned as a government institution that organizes public- and private-sector strategies to enhance cybersecurity. It would also undertake data collection about threats and attacks, analyze and disseminate that information (perhaps in partnership with the private sector), serve as a repository for technical remedies, and educate the public about best practices, defenses, and remedies.

An IBM white paper[50] broadens the analogy. Borrowing not only from public health but also from public safety, the paper recommends establishing a Cyber Federal Emergency Management Agency and devising a Cyber National Response Framework. Independently, Microsoft's Corporate Vice President for Trustworthy Computing, Scott Charney, has advocated measuring "device health," with device "health certificates" serving as a basis for authorizing device access to network resources.[51] Rather than focusing on institutions, cybersecurity expert Jeffrey Hunker looks to public health as a model for behavioral norms.[52] Individuals would be expected to satisfy certain norms, and government institutions would focus on supporting those norms.

None of the aforementioned work includes a compelling argument for why the analogy to public health is a suitable starting point for a cybersecurity doctrine. Public health informs people's behaviors (seemingly an obvious route to enhanced cybersecurity), but so does religion (which nobody is advocating as a

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

cybersecurity solution). In formulating our doctrine of public cybersecurity, we use economic theory to justify the shared status of public health and cybersecurity as public goods because economics explains the externalities and incentives that arise in cybersecurity. Viewing cybersecurity as a public good is not new,[53] but we do appear to be the first to employ insights from economic theory to justify the public health model for cybersecurity.

Our public cybersecurity doctrine goes beyond prior work that explores cybersecurity counterparts for institutions and policies that have served public health well. Public cybersecurity is obtained by identifying cybersecurity counterparts to the *goals* of public health – not the *institutions* of public health. First, public health law provides a powerful framework for balancing collective versus individual interests. Second, just as managing disease is an important goal of public health, managing insecurity is an important goal of public cybersecurity. The siren call for the production of "secure" systems and networks must be – and, with public cybersecurity, is – augmented with a mandate to manage the inevitable insecurity that comes from the constant vulnerabilities and adversaries that networked systems face.

The goals of public cybersecurity focus on the collective. Individual high-consequence systems, such as those that control critical infrastructures, are not singled out. Why not focus on the seemingly smaller problem of making only the high-consequence systems secure? For the same reason that the public health system does not focus on keeping only "important" people healthy, isolation is not a realistic proposition for cybersecurity. Public health teaches that it is easier to keep specific individuals healthy when everyone is healthy. The same is true with cybersecurity. If we foster the production of cybersecurity generally, building our networks' capacity to manage insecurity, we will be better able to ensure that our high-consequence systems are secure.

Cybersecurity, like security in so many other contexts, involves trade-offs with other values.[54] Conflicts will arise between public cybersecurity and the interests of specific individuals, entities, and society at large. A cybersecurity doctrine is obliged to provide principles and processes to negotiate and resolve these conflicts. Public health already offers such guidelines to benefit the public good.

First, the state intervenes most drastically when an individual's health decision might directly impact the health of others. The state is generally unable to coerce an individual's decision when the health of only the individual is implicated. Substitute "health" for "security," and we have sensible guidelines for public cybersecurity.

Second, public health guidance, applied to managing the externalities associated with public cybersecurity, suggests the following:

- The state's obligations and abilities to shape and override private choices should turn on the extent to which they have a direct impact on the security of the broader public rather than the security of an individual or entity.
- To facilitate better decisions by individuals, the state should provide information or gentle interventions that influence the perception of risk but that do not supplant the decision-making.
- Where security choices of the individual will impact the security of others, the state should use a wider array of tools to alter behavior.
- Even where state action is permissible, impact on other societal values must be considered in choosing among solutions.

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

- Whenever possible, the state should opt for minimal interventions implemented in a decentralized manner, so as to limit the negative impact they may have on willingness to participate.

Inadequate cybersecurity is the obstacle to success in the information age. Though the problem resides in technologies, the solution involves policies. It requires intervention in the private choices of individuals, hard trade-offs, and political agreements that could span nations. We believe that a doctrine of public cybersecurity can be the basis for those policies. Our doctrine of public cybersecurity establishes a framework for state incentives and coercion that we believe is rational, defensible, and legitimate. It directs the focus of cybersecurity away from the individual and toward the collective. It advocates building systems with fewer vulnerabilities while acknowledging that systems cannot be rid of all vulnerabilities and must therefore be resilient in the face of attacks. If adopted, public cybersecurity will reorient public policy and discourse toward the proper goals of encouraging collective action to produce the public good of cybersecurity and managing the insecurity that remains.

ENDNOTES

* Contributor Biographies: DEIRDRE K. MULLIGAN is an Assistant Professor in the School of Information at the University of California, Berkeley, where she is also a Faculty Director of the Berkeley Center for Law and Technology. She is the Policy Lead for the National Science Foundation's TRUST Science and Technology Center, Chair of the Board of the Center for Democracy and Technology, and Cochair of Microsoft's Trustworthy Computing Academic Advisory Board. Her recent publications include "Privacy on the Books and on the Ground" (with Kenneth A. Bamberger), *Stanford Law Review* (2011); and "Catalyzing Privacy: New Governance, Information Practices, and the Business Organization" (with Kenneth A. Bamberger), *Law & Policy* (2011).

FRED B. SCHNEIDER is the Samuel B. Eckert Professor of Computer Science at Cornell University. He also serves as the Chief Scientist for the National Science Foundation's TRUST Science and Technology Center and is Cochair of Microsoft's Trustworthy Computing Academic Advisory Board. He is a Fellow of the Association for Computing Machinery, the American Association for the Advancement of Science, and IEEE, and is a member of the U.S. National Academy of Engineering and its Norwegian counterpart (NTV). He was awarded a D.Sc. (*honoris causa*) by the University of Newcastle-upon-Tyne in 2003.

Acknowledgments: We benefited from comments on early drafts of this paper and discussions with Marjory Blumenthal, Aaron Burstein, Scott Charney, John Chuang, David Clark, Craig Fields, R. Kelly Garrett, Jens Grossklags, Joshua Gruenspecht, Joseph Lorenzo Hall, Carl Landwehr, Susan Landau, Steve Lipner, Greg Morrisett, Helen Nissenbaum, Shari Pfleeger, Audrey Plonk, Ashkan Soltani, participants at the 2009 Workshop on the Economics of Securing the Information Infrastructure, and attendees at several planning meetings for this issue of *Dædalus*.

This essay is supported in part by the Air Force Office of Scientific Research (AFOSR) grant F9550-06-0019; National Science Foundation grants 0430161, 0964409, CNS-0524745 (ACCURATE), and CCF-0424422 (TRUST); Office of Naval Research grants N00014-01-1-0968 and N00014-09-1-0652; and a grant from Microsoft. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of these organizations or the U.S. government.

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

[1] Ross Anderson and Tyler Moore, "The Economics of Information Security," *Science* 314 (5799) (October 27, 2006): 610 – 613.

[2] Gerwin Klien et al., "seL4: Formal Verification of an OS Kernel," *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles* (SOSP '09), Big Sky, Montana, October 11 – 14, 2009 (New York: Association for Computing Machinery, 2009), 207 – 220.

[3] For example, "prerendering" can reduce the code required to generate the user interface in a voting system, thereby simplifying the vote-entry software and making it more amenable to verification; see Ka-Ping Yee, "Building Reliable Voting Machine Software," Ph.D. dissertation, Department of Computer Science, University of California, Berkeley, Fall 2007.

[4] Department of Defense Computer Security Center, *Department of Defense Trusted Computer System Evaluation Criteria*, CSC-STD-001-83 (Fort George G. Meade, Md.: Department of Defense, August 1983).

[5] *Common Criteria for Information Technology Security Evaluation*, International Organization for Standardization (ISO) Standard 15408, August 1999, http://www.niap-ccevs.org/cc-scheme.

[6] Health Insurance Portability and Accountability Act of 1996 (HIPAA), Public Law No. 104-191, 110 Stat. 1936 (1996) (regulating the use and disclosure of "Protected Health Information"); Gramm-Leach-Bliley Act (GLBA), Title V, Public Law No. 106-102, 113 Stat. 1338 (1999) (codified at 15 U.S.C. sec. 6801 – 6827 [2006]), 15 U.S.C. sec. 6801, sec. 6805 (empowering various agencies to promulgate data-security regulations for financial institutions); Children's Online Privacy Protection Act, 15 U.S.C. sec. 6501 et seq. (prohibiting the collection of personally identifiable information from young children without their parents' consent).

[7] Federal Information Security Management Act of 2002 (FISMA), 44 U.S.C. sec. 3541 et seq.

[8] *Draft Voluntary Voting System Guidelines: Version 1.1* (Washington, D.C.: The U.S. Election Assistance Commission, May 27, 2009), http://www.eac.gov/testing_and_certification/voluntary_voting_system_guidelines.aspx.

[9] Rainer Böhme and Galina Schwartz, "Modeling Cyber-Insurance: Towards a Unifying Framework," working paper presented at the Workshop on Economics of Information Security, Harvard University, Cambridge, Massachusetts, June 2010, http://weis2010.econinfosec.org/papers/session5/weis2010_boehme.pdf.

[10] Anderson and Moore, "The Economics of Information Security."

[11] Jens Grossklags, Nicolas Christin, and John Chuang, "A Game-Theoretic Analysis of Information Security Games," *Proceedings of the 17th International World Wide Web Conference* (WWW2008), Beijing, China, April 21 – 25, 2008.

[12] Security breach notification statutes that require companies to notify individuals when certain personal data have been accessed or disclosed without authorization are in place in forty-six states, the District of Columbia, Puerto Rico, and the Virgin Islands. The first was California's Notice of Security Breach Law, California Civil Code, sec. 1798.29 (2002).

[13] Butler W. Lampson, "Computer Security in the Real World," *Computer* 37 (6) (June 2004): 37 – 46.

[14] South Korea currently requires Internet users to attach their real names and resident identification numbers when they post messages on the Internet. Websites that allow posting must collect and confirm names and resident IDs with a government server; see Se Jung Park, Yon Soo Lim, Steven Sams, Sang Me Nam, and Han Woo Park, "Networked Politics on Cyworld: The Text and Sentiment of Korean Political Profiles," *Social Science Computer Review* (September 21, 2010).

[15] Jason Franklin, Vern Paxson, Adrian Perrig, and Stefan Savage, "An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants," *Proceedings of the 14th ACM Conference on Computer and Communications Security* (CCS '07), Alexandria, Virginia, October 29 – November 2, 2007 (New York: Association for Computing Machinery, 2007), 375 – 388.

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

16 William J. Lynn III, "Defending a New Domain: The Pentagon's Cyberstrategy," *Foreign Affairs*, September/October 2010.

17 The mission of public health was defined by an influential Institute of Medicine committee as "fulfilling society's interest in assuring conditions in which people can be healthy"; Institute of Medicine, *The Future of Public Health* (Washington, D.C.: National Academies Press, 1988), 7. A rich description of the legal framework is set out in Lawrence O. Gostin, *Public Health Law: Power, Duty, Restraint* (Berkeley: University of California Press, 2001).

18 For example, after several outbreaks of measles among primarily unvaccinated children, a federal law was passed to provide free vaccines to certain groups of children and funds to states for supporting efforts to enhance vaccination levels.

19 Institute of Medicine, *The Future of Public Health*, 3.

20 Systems that employ technical means to enable continued operation in the face of attacks are sometimes called *intrusion tolerant*. A sampling of specific techniques for achieving intrusion tolerance is discussed in Jaynarayan Lala, ed., *Foundations of Intrusion Tolerant Systems* (Los Alamitos, Calif.: IEEE Computer Society, 2003). In this essay, the term *managing insecurity* is intended to denote something broader, admitting nontechnical means as well as intrusion tolerance techniques.

21 Helen Nissenbaum, "Where Computer Security Meets National Security," *Ethics and Information Technology* 7 (2) (June 2005): 61–73.

22 The doctrine of prevention is not concerned with managing insecurity; the doctrine of risk management and doctrine of deterrence through accountability are not concerned with producing cybersecurity. None concern trade-offs of individual rights for public welfare.

23 Deborah G. Mayo and Rachelle D. Hollander, *Acceptable Evidence: Science and Values in Risk Management* (New York: Oxford University Press, 1994).

24 In fuzz testing, a system is exposed to random inputs of unexpected kinds. This form of testing reveals inadequacies in the input validation routines of a system. Several classes of attacks are blocked by implementing stringent input validation.

25 For a case study comparison of four vulnerability reduction techniques, see Koen Buyens, Bart De Win, and Wouter Joosen, "Empirical and Statistical Analysis of Risk Analysis-Driven Techniques for Threat Management," *Proceedings of the Second International Conference on Availability, Reliability and Security* (ARES '07), Vienna, Austria, April 10–13, 2007 (Washington, D.C.: IEEE Computer Society, 2007), 1034–1041. For a theoretical comparison of two high-profile development processes (Microsoft SDL and the Open Web Application Security Project's Comprehensive, Lightweight Application Security Process [CLASP]), see Johan Grégoire, Koen Buyens, Bart De Win, Riccardo Scandariato, and Wouter Joosen, "On the Secure Software Engineering Process: CLASP and SDL Compared," *Proceedings of the Third International Workshop on Software Engineering for Secure Systems* (SESS '07), Minneapolis, Minnesota, May 2007 (Washington, D.C.: IEEE Computer Society, 2007). An argument in favor of evidence-based practices appears in Daniel Jackson, Martyn Thomas, and Lynette Millett, eds., *Software for Dependable Systems: Sufficient Evidence?* (Washington, D.C.: The National Academies Press, 2007).

26 Vaccination works by causing exposure to a relatively benign form of the disease against which protection is being sought.

27 A notable example is "Implementation of Commonly Accepted Security Configurations for Windows Operating Systems," policy memorandum M-07-11 (Washington, D.C.: U.S. Office of Management and Budget, March 22, 2007), http://www.cio.gov/documents/FDCC_memo.pdf. The memorandum lists the few versions of Windows that certain civilian federal agencies are permitted to use.

28 David Ladd et al., *The SDL Progress Report* (Microsoft Corporation, 2011), 23–24, http://www.microsoft.com/downloads/en/details.aspx?FamilyID=918179A7-61C9-487A-A2E2-8DA73FB9EADE&displaylang=en.

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

29 The general rule allows individuals to refuse treatments. Some states mandate treatments for communicable diseases, such as tuberculosis, that pose a danger to the public.

30 Michel J.G. van Eeten and Johannes M. Bauer, "Economics of Malware: Security Decisions, Incentives and Externalities," OECD STI Working Paper 2008/1 (Organisation for Economic Co-operation and Development Directorate for Science, Technology and Industry, May 2008), http://www.oecd.org/dataoecd/53/17/40722462.pdf.

31 See David D. Clark and Susan Landau, "Untangling Attribution," *Harvard National Security Journal* 2 (2) (2011), http://harvardnsj.com/2011/03/untangling-attribution-2.

32 "Einstein Intrusion Detection System: Questions That Should Be Addressed" (Washington, D.C.: Center For Democracy and Technology, July 28, 2009), http://cdt.org/security/20090728_einstein_rpt.pdf.

33 The risk that such vulnerabilities might be sold or disclosed to irresponsible or hostile parties is cause for concern, but this market structure flourishes in the absence of alternatives.

34 See http://www.confickerworkinggroup.org.

35 For example, unpatched machines can be co-opted by attackers into a botnet. Such collections of remotely controlled machines are used today for a variety of illicit activities, including generation of spam email and distributed denial-of-service attacks.

36 However, some enterprises regard their software and data as proprietary. They might not be comfortable providing test suites to patch developers.

37 Richard Clayton, "Might Governments Clean Up Malware?" *Proceedings of the Ninth Annual Workshop on Economics and Information Security* (WEIS10), Cambridge, Massachusetts, June 7–8, 2010.

38 The term *air gap* originally referred to isolation caused when no wires are connected to a given component. With the advent of wireless networking, the term's meaning has broadened to denote isolation caused when the Laws of Physics ensure no signal can reach the component.

39 *Malicious Software (Malware): A Security Threat to the Internet Economy*, Ministerial Background Report DSTI/ICCP/REG(2007)5/FINAL, Organisation for Economic Co-operation and Development, June 2008, http://www.oecd.org/dataoecd/53/34/40724457.pdf.

40 In October 2009, Comcast announced a trial of an automated service that "warn[s] broadband customers of possible virus infections." The Comcast Constant Guard issues a pop-up notice directing customers to resources to rid their machine of infection; see Elinor Mills, "Comcast Pop-ups Alert Customers to PC Infections," CNET News, October 8, 2009. A similar service, the Qwest Customer Internet Protection Program, displayed a Web page warning to customers with options for removing the detected infection for free; customers were obliged to do so before they were allowed to return to surfing the Web. Similarly, an older, now discontinued SBC service quarantined computers until they were cleaned.

41 Ibid.

42 Clayton, "Might Governments Clean Up Malware?" 5 n.2.

43 For a discussion of some of the issues raised by an ISP's decision to cut off broadband access due to infection, see George Ou, "Comcast Heading the Right Direction on Cybersecurity," Digital Society, October 9, 2009, http://www.digitalsociety.org/2009/10/comcast-heading-th-right-direction-on-cybersecurity.

44 See http://www.eff.org/wp/noncommercial-email-lists-collateral-damage-fight-against-spam.

45 See generally, Doug Lichtman and Eric Posner, "Holding Internet Service Providers Accountable," *Supreme Court Economic Review* 14 (2006): 221, 233–234.

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021

46 Stephanie Forrest, Alan S. Perelson, Lawrence Allen, and Rajesh Cherukuri, "Self-Nonself Discrimination in a Computer," *Proceedings of the 1994 IEEE Symposium on Research in Security and Privacy*, Oakland, California, May 16 – 18, 1994.

47 Stephanie Forrest, Anil Somayaji, and David H. Ackley, "Building Diverse Computer Systems," *Proceedings of the Sixth Workshop on Hot Topics in Operating Systems*, Cape Cod, Massachusetts, 1997.

48 *Enabling Distributed Security in Cyberspace: Building a Healthy and Resilient Cyber Ecosystem with Automated Collective Action* (Washington, D.C.: Department of Homeland Security, March 23, 2011), http://www.dhs.gov/xlibrary/assets/nppd-cyber-ecosystem-white-paper-03-23-2011.pdf.

49 *National Cyber Leap Year Summit 2009 Co-chairs Report*, September 16, 2009, http://www.nitrd.gov/NCLYSummit.aspx.

50 Daniel B. Prieto and Steven Bucci, "Meeting the Cybersecurity Challenge: Empowering Stakeholders and Ensuring Coordination," IBM U.S. Federal White Paper (Somers, N.Y.: IBM Corporation, February 2010).

51 Scott Charney, "Collective Defense: Applying Public Health Models to the Internet," white paper (Redmond, Wash.: Microsoft Corporation, 2010), http://www.microsoft.com/security/internethealth.

52 Jeffrey Hunker, "U.S. International Policy for Cybersecurity: Five Issues That Won't Go Away," *Journal of National Security Law & Policy* 4 (2010): 197 – 216.

53 Benjamin Powell, "Is Cybersecurity a Public Good? Evidence from the Financial Services Industry," Working Paper #57 (Washington, D.C.: The Independent Institute, March 14, 2005); Bruce H. Kobayashi, "An Economic Analysis of the Private and Social Costs of the Provision of Cybersecurity and Other Public Security Goals," Working Paper #26 (Arlington, Va.: George Mason University School of Law, 2005), http://law.bepress.com/gmulwps/gmule/art26; Brent Rowe and Michael Gallagher, *Private Sector Cyber Security Investment Strategies: An Empirical Analysis* (Research Triangle Park, N.C.: RTI International, March 2006).

54 Nissenbaum, "Where Computer Security Meets National Security."

Downloaded from http://www.mitpressjournals.org/doi/pdf/10.1162/DAED_a_00116 by guest on 04 December 2021