international risk
governance center

OPINION PIECE

PUBLIC CYBERSECURITY AND RATIONALIZING INFORMATION SHARING

AN OPINION PIECE FOR IRGC
Fred B. Schneider, Computer Science Department, Cornell University
Elaine M. Sedenberg, School of Information, U.C. Berkeley
Deirdre K. Mulligan, School of Information, U.C. Berkeley

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

international risk
governance council

This paper was supported in part by AFOSR grants F9550-06-0019 and FA9550-11-1-0137, National Science Foundation grants 0430161, 0964409, and CCF-0424422 (TRUST), ONR grants N00014-01-1-0968 and N00014-09-1-0652, and grants from Microsoft.

Authors:
Fred B. Schneider, Computer Science Department, Gates Hall, Cornell University, Ithaca, NY, 14853, USA.
fbs@cs.cornell.edu
Elaine M. Sedenberg, School of Information, U.C. Berkeley, Berkeley, CA, 94720-4600, USA.
elaine@ischool.berkeley.edu
Deirdre K. Mulligan, School of Information, U.C. Berkeley, Berkeley, CA, 94720-4600, USA.
dkm@ischool.berkeley.edu

This paper should be cited as:
Schneider, F., Sedenberg, E., Mulligan, D.: Public Cybersecurity and Rationalizing Information Sharing, Opinion Piece for the International Risk Governance Center (IRGC). Lausanne: IRGC.

Available from www.irgc.org

FOREWORD

Much work in cybersecurity is focused on the problems of making specific systems more secure and well behaved. Doing that is clearly very important. However, just as treating sick patients one at a time is not sufficient to stop the spread of an epidemic, patching or improving the design of individual software systems one at a time is almost certainly not sufficient to produce wide-spread cybersecurity.

Medical doctors can make important contributions, but preventing, slowing and stopping epidemics is the domain of public health—the creation of an entire environment designed to promote sanitary practice and provide protection at the level of a community.

Drawing this parallel, in 2011, Deirdre K. Mulligan, then an Assistant Professor in the School of Information at the University of California, Berkeley and Fred B. Schneider, the Samuel B. Eckert Professor and Chairman of the Department of Computer Science at Cornell University, argued that achieving widespread cybersecurity required a similar general community approach. They elaborated this argument in a paper titled "Doctrine for Cybersecurity" in *Dædalus*, the journal of the American Academy of Arts and Science.

Building on the metaphor of public health, Mulligan and Schneider argued that individual strategies (standards, adherence to good practice in software engineering, formal methods, red/green machines, filters and firewalls, etc.) are all valid, but, as with health, there are aspects of cybersecurity that are a "public good". They went on to outline a range of strategies, some technical but others behavioral, educational or legal, that they believed necessary to create an effective "doctrine for cybersecurity."

In order to make these ideas more widely available to the risk governance community IRGC invited Fred Schneider, Elaine Sedenberg and Deirdre Mulligan to summarize and elaborate their ideas in this short opinion piece titled *Public Cybersecurity and Rationalizing Information Sharing*. We hope readers will find their ideas as useful and stimulating as we have at IRGC.

M. Granger Morgan
Chair, Scientific and Technical Council, IRGC Foundation

Contents

PREFACE BY IRGC

Cyber-technology and associated cybersecurity are today central to our economic and social lives. Secure systems contribute to creating a sense of confidence that the technologies and processes aimed at improving performance and welfare will not endanger data privacy, confidentiality, integrity or availability—areas of value both to people and businesses.

There are conflicting views about what cybersecurity entails and how information sharing can promote it. Questions also include how to protect individual interests and other social values, and how to prioritize what is important in information sharing arrangements.

The lack of a clear framework under which information sharing proposals can be evaluated, increases the controversy around them. For example, the recently enacted U.S. Cybersecurity Act of 2015 contains many cybersecurity information sharing provisions and mandates, yet it fails to connect them to specific cybersecurity goals. There is an absence of any clear objectives. The new sharing initiatives have been met with skepticism as to their utility and, additionally, objections have been raised based on the newly created risks to privacy. The new European Network and Information Security (NIS) Directive[1] that establishes a compulsory incident reporting scheme for operators of essential services to collect evidence about cyberattacks and other breaches has also been greeted with concern because some industry players fear that it will force them into releasing confidential information. In both contexts, advocates of information sharing are hampered by a failure to tie information sharing generally and specifically to the advancement of particular cybersecurity goals.

In the midst of the debate about the trade-offs that governments, businesses, and individuals ought to make, the idea that cybersecurity should be understood as a public good suggests both the need and the approach to clarifying cybersecurity goals, as well as a way to ground conversations about information sharing and other policies.

Public goods include fresh air, national security, and public health. Public goods are non-excludable and non-rivalrous. They are to be maintained or developed at a societal level. The concept is an economic concept, not a value judgment. Applied to public health, it justifies the allocation of public resources to prevention and ongoing efforts to manage disease, and the data collection and other intrusions on individual interests necessary to the benefit of all. In the field of cybersecurity, it suggests that there is a collective responsibility to develop cybersecurity and manage cyber-insecurity, and that doing so requires a shared perspective on what cybersecurity entails—who and what should be protected—and when and under what circumstances its pursuit can cause harm or prejudice to individuals or other national priorities, such as innovation. This framing provides a way to evaluate information sharing proposals.

If, like public health, cybersecurity is considered a public good, which has to be protected and developed with established high-level principles and criteria, then it may be that some of the trade-offs mentioned above would be easier to resolve, at least at a national policy level. IRGC is not convinced this is true, but we pose the question. **Establishing cybersecurity as a public good would create the overarching policy principle to define goals and means, to bring cohesion to sectoral and specific, purpose-led policies and programs**. It would also suggest that it is important to pursue both international collaboration in harmonizing technical choices and institutional and regulatory measures.

With this objective in mind IRGC invited Prof. Fred Schneider, Elaine Sedenberg and Prof. Deirdre K. Mulligan to write an opinion piece for publication by IRGC. We proffer our sincere thanks to them for their valuable contribution.

IRGC opinion pieces

IRGC opinion pieces are authored papers that reflect the opinion of recognized scientists on governance issues. They are sometimes controversial because of the topics they discuss. In providing such opinion pieces IRGC's intention is to trigger a discussion between scientists and policymakers on possible scientific or governance approaches that might contribute to solving a current risk and governance problem.

IRGC's initiative on cyber risk governance

During 2015 IRGC conducted two workshops to explore issues of cybersecurity. The first, organized by Professor Granger Morgan and held in Washington DC on May 28-29, compared methods for assessing terrorism risk with those for cybersecurity and explored the potential for each field to learn from the other. Our deliberations in this workshop were informed by the public health model and related ideas advanced by Mulligan and Schneider.

The second workshop, on cybersecurity risk governance, was held at the Swiss Re Centre for Global Dialogue, Zurich, Switzerland on October 29-30. It focused in greater detail on the private sector and discussed the changing cyber threat landscape, various techniques for improving cybersecurity, and methods for dealing with the residual risk while focusing on quantification and transfer to insurance.

1. Introduction

Achieving any specific level of cybersecurity inevitably entails making compromises with regard to cost, function, and convenience, as well as trade-offs between societal values, such as openness, privacy, freedom of expression, and innovation. In defining regulations and incentives, decisions have to be made about how to balance these trade-offs while optimizing security outcomes. To further complicate matters, neither technologists nor policymakers have the luxury of starting with a clean slate. Instead they work within the shadows of legacy networks and end systems that are neither secure, nor easily made so. Moreover, current security postures often reflect societal values from a time when dependence on networked information systems was minimal.

A *cybersecurity doctrine* prescribes a set of goals, a basis for making trade-offs among these goals, and various means to achieve the goals. Its utility is determined, in part, by the extent to which it offers a framework for achieving goals without imposing, ignoring, or ruling out possible technical or policy solutions. And the value of cybersecurity doctrines per se is measured by the extent to which they bring clarity to policy questions and proposed incentives.

The Doctrine of Public Cybersecurity[2] has as its goals the production of security and the management of its absence. The doctrine derives from the observation that cybersecurity is non-rivalrous and non-excludable and, thus, satisfies the definition of a public good. Cybersecurity is non-rivalrous, since one user benefiting from the security of a networked system does not diminish the ability of any other user to benefit from the security of that system; it is non-excludable, because users of a secure system cannot be excluded easily from benefits security brings.

Notice that the Doctrine of Public Cybersecurity targets the collective rather than any single individual's or entity's computer, network, or assets. Also, it steers policy makers away from deterrence-oriented strategies ("doctrines of accountability") reflected in current criminal law, doing so because deterrence does little to encourage investments in the production of cybersecurity or in managing its absence.

This paper briefly explores how information sharing fits into the Doctrine of Public Cybersecurity and how laws and policies around these activities can be tailored to promote security with limited intrusions on privacy and autonomy.[3] Some in the U.S. and elsewhere have argued that information sharing is an attractive means for supporting cybersecurity; others worry that compromises to societal values (such as privacy) seem inevitable. This paper revisits the Doctrine of Public Cybersecurity in order to shed light on debates about information sharing, by exploring its potential utility, and considering policies to mitigate its impact on other societal values.[4] Its goal is not to advocate for the creation of specific institutions (government or otherwise), but rather to explore the potential utility of information sharing to promote public cybersecurity.

2. In Analogy with Public Health

Public health—the prevention of disease and promotion of good health in populations writ large—is a public good. It is non-rivalrous since having the population healthy implies a lower prevalence of disease which, in turn, decreases the chances that any member can fall ill. And it is non-excludable, because nobody can limit an individual's ability to profit from the health benefits that living among a healthy population brings. Public health law focuses on the health of the population as a whole and the singular responsibility of government in that enterprise.

Public health and cybersecurity are both thus public goods that aim to achieve a positive state (health or security) in a loosely affiliated but highly interdependent network. With one, it is a network of people existing in an environment over which they have some limited control; with the other, the network comprises people, software, and hardware (for communications, storage, and processing). And because the sought-after positive state is ultimately unachievable, public health and public cybersecurity must struggle with how to manage its absence as well as with how to prompt its production. Success ultimately depends not only on technical progress but also on reaching a political agreement about (i) the relative value of some public good in comparison to other societal values and (ii) the institutions' granted authority to resolve conflicts (and the methods they use).

Just as with public health, ensuring that actors contribute to public cybersecurity requires interventions to overcome positive and negative externalities that lead rational individuals to underinvest. When incentives are insufficient to prompt private provisioning, the public interest requires making value-ridden choices to interfere with the rights and interests of individuals and organizations. Those choices would be embodied in goals that reflect political agreement about the good in question, the socially desirable level, 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.

The analogy with public health inspires cybersecurity measures like prevention, containment, mitigation, and recovery—strategies that direct resources toward production and preservation of cybersecurity[5]. Modern public health doctrine does not compensate victims of disease so, by analogy, a doctrine of public cybersecurity would not focus on restitution. Indeed, restitution is economically efficient only if you assume attacks are infrequent, and that assumption is not realistic today. Quarantine, in response to disease, benefits the collective by limiting the spread of disease. It does so by depriving an individual of certain freedoms. By analogy, a doctrine of public cybersecurity would dictate responses that deprive individuals of actions, but only if those responses benefit the collective. Punishments solely for retribution would not be part of a public cybersecurity doctrine (since retribution does not benefit public welfare). Finally, the parallel with public health also suggests that prevention should be preferred to recovery.

Some express reservations about making an analogy between cybersecurity and public health. These reservations stem from the role that an intelligent, sentient adversary plays in undermining cybersecurity (which is largely absent from public health, though possible through bioterrorism, and at a smaller scale through intentional transmission of communicable diseases which, while rarely seen, has occurred at the individual and nation-state level). Pathogens do evolve biologically and adapt to environmental changes or take advantage of changing social structures (e.g., rapid spreading through urbanization, or growing antibiotic resistance from prescription overuse). Although motivated by survival rather than by a desire to maximize damage, the evolution of pathogens nevertheless embodies the same type of an arms race we see with cybersecurity and development of malware, which, once created or improved by a human, spreads in a rapid, non-sentient fashion. So, motives differ and intelligence in the narrow sense is lacking, but in both domains the public good is subject to a constantly changing set of new exploits from adversaries.

Moreover, focusing only on the sentience of an adversary misses the point: preventative techniques are effective, regardless of motive. This can be seen in public health examples such as vaccination and condom use. While intentionally spreading a disease is relatively uncommon, individuals knowingly have exposed others to HIV. Condoms nevertheless are an effective preventative measure against such hostile acts and vaccines increase herd immunity, lowering the risk of infection. Similarly within cybersecurity, automated software platforms for launching cyberattacks can be purchased on the black market to intentionally spread and infect targets. However, similar to preventative techniques like condoms, patching vulnerabilities or limiting a user's downloads are effective against these attack engines, regardless of the capabilities (or lack of sentience) in the automation. In short, adversarial considerations are simply less relevant when dealing with prevention and risk management orientations—in contrast to deterrence-oriented strategies, which are focused on intent—because harms manifest, and protections work, regardless of intent.

Still, caution is advised when invoking the analogy between health and cybersecurity. We have no qualms about using the analogy for inspiration, but we are reluctant about advocating adoption of a means or strategy from public health until it has been evaluated anew relative to public cybersecurity's stated goals: producing cybersecurity or managing its absence.

3. Monitoring and Information Sharing to Support Public Health

Public health decidedly benefits from collecting and sharing information about the health of a population and the spread of epidemics. This collection and sharing aids in:

- Determination of the origin or current sources of a given disease outbreak. This, in turn, enables treatment of contagious individuals as well as supporting

other restorative actions (e.g., quarantine, vaccination, education, or disposal of contaminated sources) that impede further spread of the disease.
- Sharing facilities and expertise needed for identifying the cause, analyzing conditions favorable for propagation, and developing remediations for disease outbreaks. Such investments are invariably better amortized over a larger population and/or a broader collection of locales. Moreover, the diversity intrinsic to larger regions can be helpful in understanding underlying mechanisms and possible means of control.
- Assessment of the scope of an outbreak which enables predictions that then can inform selecting responses well matched to the urgency of a problem. Also, over the long term, information about the scope of outbreaks fuels research, informs policy decisions, and helps in formulating educational efforts that further reinforce prevention and response measures.

However, information sharing of public health data occurs in the context of complex commitments to other values—particularly individual privacy and maximal participation in the health care systems—that are, at times, in tension with public health information needs. If not carefully considered, information sharing activities will be undermined by individuals and/or their healthcare providers who feel compelled to take evasive measures in order to limit the collection of sensitive information.

The following guidance currently being observed for public health information sharing reduces such problems and encourages participation:

- Resist the temptation to collect exhaustive information and only seek information (or information at a level of detail) that is necessary for implementing effective policies and programs.
- Provide communities with information they need to understand and, when relevant, make decisions about participating in programs that involve collection and use of public health information—including specifications regarding the purpose and use of data collection and assurances of confidentiality.
- Make information held by public health institutions available in a timely manner, consistent with relevant mandates, and resource constraints.
- Protect the confidentiality of information that can bring harm to an individual, community, or organization, and limit disclosures to instances where there is high likelihood of significant harm to the individual or others.

Of course, the elements that make up these guidelines are sometimes in conflict. Seeking and making information accessible in order to facilitate community decision-making can erode individual privacy, for example. It also can harm the wellbeing of individuals or communities that will suffer economic losses if a contagious disease or genetic condition becomes associated with a particular group by location or ethnicity.

Public health surveillance does not require overt patient consent for collecting and sharing incident data (thereby violating one widely accepted tenet of privacy) within the system, since doing so would add an administrative burden to healthcare professionals and potentially slow a reporting process where timeliness is crucial for slowing the spread of a new communicable disease. Rather, individuals and medical professionals are bound by social

contract and a duty to inform the state when an individual's health implicates the wellbeing of others. The collective need to know about a contagious disease is directly at odds with a patient's individual right to privacy. Collectors and holders of an individual's data are required to employ policies, practices, and mechanisms that will protect the confidentiality of an individual's health information to mitigate the privacy lost from the surveillance.

Reporting, minimization, and decentralization are common elements in the public health data collection landscape. Legal frameworks, institutional policies and practices, and technical approaches to data sharing reflect preferences for keeping identifiable and granular data in the hands of the initial collector rather than pooling it. Adherence to these principles erects practical barriers to the misuses or repurposing of public health data at scale; now, multiple systems must be compromised or multiple entities convinced if a shift in use is to occur. And when breaches or shifts in use do occur, the limited nature of the data often reduces the potential for harm.

At times, though, public health goals do require sharing identifiable information in ways that allow officials to link this data to other datasets or to identify persons with a specific disease or health condition. In almost all cases, this identifiable data is only maintained at the level where the intervention occurred, which is usually the state or local level.[6] In limited cases, such as a rare disease outbreak or certain high-risk disease surveillance programs, identifiable data may be shared with other jurisdictions or reported to federal agencies in order to enable public health activities. For example, within the HIV/AIDS surveillance system, experts[7] support the routine sharing of some data with identifiers in order to resolve duplicate case counts across states and territories, so data quality may be assured at a national level.[8] But when identifiable data must be transferred, means are employed to limit risk—encryption, replacing identifiers, etc. Earliest feasible de-identification is especially important.

4. Monitoring and Information Sharing to Promote Cybersecurity

Information sharing has figured prominently in recent policy proposals, and it is a core feature of a new U.S. law to improve cybersecurity. These provisions are motivated by a belief that information currently unavailable to relevant parties is necessary for certain cybersecurity-promoting activities. However, the activities that policymakers want to facilitate are often left unclear. Some activities that would be enabled are consistent with the Doctrine of Public Cybersecurity; facilitating patch development and widespread dissemination is an example. Other information sharing, such as helping law enforcement to prosecute bad actors, would not be consistent in that it promotes deterrence rather than prevention or risk management. Shared information, however, could reveal private communications, associational interests, the physical whereabouts and movements of individuals, other personal details, and it might also disclose confidential information about companies' networks, policies, and proprietary interests. In short, information sharing is often in tension with privacy, and with other values.

Today's cybersecurity environment boasts a wide range of information sharing activities. Some, like industry-specific Information Sharing and Analysis Centers (ISACs) and the United States Computer Emergency Readiness Team (US CERT), are long-standing and supported by the government to promote sharing between trusted communities, industry-specific partners, and, sometimes, even the public. Others have arisen independently in response to specific threats, and they are largely the outcome of private decisions by security practitioners and their employers. Some are aimed at improving specific products; others focus on sharing best practices or on identifying and managing attacks. We briefly examine below some existing efforts, to highlight their relationships to public cybersecurity goals.

Sharing Information about Vulnerabilities. There is a rich market (white hat and black hat) for information about vulnerabilities and exploits. Vulnerability reward programs (VRPs), also known as "bug bounties", incentivize the reporting of information to organizations (often software vendors) so that patches can be developed. These programs are designed to promote disclosure to those in the position to develop patches, because discovered—but unreported—vulnerabilities may be sold on the black market as zero-day exploits (an exploitable software vulnerability unknown to the vendor). The effectiveness of these programs is debated; vulnerabilities often command a higher price on the black market, so the best ones might get sold to the wrong party. In addition, some argue that commercialization of vulnerability information limits the availability of data for security researchers, and thus it is an unwise course.

Vulnerability reporting that leads to the development and installation of patches serves a robust preventative function. However, the need for coordination and the lack of a uniform policy regarding public release of information about vulnerabilities can detract from its utility. Moreover, acting to patch vulnerabilities comes with trade-offs for the affected company and its customers. A vulnerability made public—even where accompanied by a patch—facilitates reverse engineering by attackers seeking to create exploits against unpatched systems. There are also reasons an end-user might delay applying a patch. First, installation of a patch takes time and interrupts on-going operations. Second, after a patch has been applied, the resulting system might not exhibit identical behavior, which could disrupt operations.

Sharing Information about Best Practices. Regulatory models that formally adopt or refer to industry-generated security standards indirectly encourage information sharing about security best practices. Incident response organizations designed to coordinate action or facilitate a response to a security compromise also advise about recommended security practices. US-CERT centers focus on disseminating relevant threats and vulnerability information to targeted parties. They also publish recommended best practices, so that these can be used widely.

Sharing Information about Threats and Risks. In 1998, U.S. Presidential Decision Directive 63 (PDD-63) identified distinct industries and called for the private sector in each to set up Information Sharing and Analysis Centers (ISACs) to mitigate risk and promote effective responses to adverse events, including cyberattacks. Organizing around industry sectors facilitated

more-specific information exchanges about vulnerabilities, threats, and isolated incidents. Exchanging security information does have risks, such as loss of competitive advantage, market share, and stock market value from negative publicity. Still, members benefit from industry-specific information exchanges that assist in prevention efforts, vulnerability identification, and risk management. And, information sharing about emerging and existing threats, particularly before they have been exploited within an industry, can bolster prevention-related activities. Furthermore, after an exploit has been fielded, information sharing assists in coordinating action to manage the resulting insecurity of vulnerable systems.

Sharing Information to Manage and Respond to Vulnerabilities and Threats. Successful coordination for mobilizing the response to an imminent threat requires sharing information. In the past, the mechanism often has been an ad hoc working group of researchers and practitioners. For instance, individuals from Microsoft, ICANN[9], domain registry operators, anti-virus vendors, and academic security researchers spontaneously formed the Conficker Working Group[10] to contain and blunt the effectiveness of an aggressive worm that threatened the Internet in 2008. In an incident in 2008 that highlighted a design vulnerability in the Border Gateway Protocol (BGP), engineers at Google had to coordinate and collaborate quickly with technical personnel who administered the Internet infrastructure across the world to respond and correct a censored link that effectively blocked worldwide access to YouTube. This rerouting from the Pakistan Internet Exchange (PIE) was not intended to spread beyond national borders, but quickly undermined key Internet infrastructure technologies and required rapid international information sharing to correct the mistake and keep the Internet functioning as expected.[11]

Although such information sharing groups have succeeded, there is widespread agreement that government involvement could further facilitate the process to ensure that an organized response occurs and assist with resources (financial, information, or logistical support). Researchers have noted other weaknesses, too, in the cybersecurity information-sharing landscape, including the difficulty of obtaining data in a timely and consistent format, organizational and policy challenges associated with disseminating vulnerability disclosures, and inattention to the privacy risks associated with sharing relevant data.

Sharing Information about Breaches of Certain Personal Information. Nearly all U.S. states have adopted security breach notification (SBN) laws that require companies—and often government agencies as well—to inform individuals, and sometimes relevant state agencies, when certain unencrypted personal information has been accessed by unauthorized parties. In theory, these notifications could assist individuals in avoiding entities with lax security practices, and they provide more comprehensive information about risks to personal information to aid all institutions in shaping their security investments and assist policymakers in identifying areas of market failure. In practice, the laws are not fully effective in any of these areas: the shared information is insufficient for individuals to make judgments about the relative security practices of entities, and the information is often not detailed enough and not pooled to facilitate analysis and learning by effected entities or policymakers. However, even in its imperfect state, the SBN laws have motivated entities

to take steps—such as encryption and improved data management—that improve the security of certain personal information by preventing breaches and facilitating quick responses where they occur. Despite the security gains, entities are reluctant to share information about breaches in the absence of a legal obligation, because doing so can lead to regulatory actions, law suits, reputational damage, and market devaluation. Even where law suits are unsuccessful—as they have often been in the U.S. due to difficulties proving harms considered cognizable for standing or relief—the fear of liability and the cost of defending against suits make entities reluctant to share information about breaches absent a legal mandate to do so.

5. The Public Health Analogy Revisited for Cybersecurity Information

The public goods nature of cybersecurity suggests that proposals for monitoring and sharing cybersecurity information could benefit from existing policies and protocols for monitoring health and sharing that information. This, because in both public health and cybersecurity, information sharing has the potential to compromise privacy and reveal other confidential information about a principal (a person or an institution) that, if divulged, could have economic or other undesirable consequences for that principal.

Yet there also are some important differences. First, disseminating information about vulnerabilities can aid attackers, whereas disseminating information about a disease is unlikely to further an epidemic. In fact, because cybersecurity information often can be weaponized, international sharing will require protocols that have not been needed for sharing health data. Second, health information has considerably narrower scope than the transactional data that information systems handle. This makes it harder to predict the privacy risks from divulging transactional data and, thus, more attention is required for need-based collection, data minimization, de-identification, and confidentiality protection. Third, the timescale of cyber incidents warrants attention in information sharing initiatives. The extent to which malicious behavior can alter the timing of a vulnerabilities' impact—from hoarding zero-day attacks, to delayed exploitation of personal financial information—warrant attention, as does the rapid timescale at which information must be shared to be maximally useful in prevention or mitigation efforts. Finally, there is a matter of trust in institutions. At least in the U.S., governments are not trusted to manage information about people's day-to-day activities. This problem might best be handled if data is collected and stored in a way that no one entity has a full view of the data. Bringing analysis tools to the data rather than data to the government for analysis may offer a more viable solution.

From what public health and cybersecurity have in common, certain principles currently employed for sharing health information would have obvious applications for cybersecurity.

- Decisions to collect and share information should be tied to the production of cybersecurity and/or the management of its absence.

- Limitations should be prescribed to ensure that collected data cannot be used against data subjects for adverse purposes unrelated to the production of cybersecurity and/or the management of its absence.[12]
- As much cybersecurity data as possible should be made open and accessible for public use. Data that cannot be made open should be made accessible through limited data access mechanisms and special use agreements.
- Coordination must also be considered at multiple levels, including globally. Just as public health officials work together through various organizational and governmental entities to share information—despite concerns about potential weaponization and other misuse—there is a need to consider what public cybersecurity information sharing activities may be appropriate at an international scale.[13]

While government plays a central role in the public health area, it is unclear whether that is appropriate for cybersecurity, given the distribution of expertise, spectrum of industries involved, and diversity of data implicated (best practices, vulnerabilities, threat indicators, malware, incident reporting, etc.). There is an obvious role for the governments to play in coordination and agreement on standards for data to be reported. It also is well suited to help orchestrate how stakeholders partition the financial burdens of monitoring and information sharing. Finally, it must lead conversations about cybersecurity goals, otherwise information sharing, and other means, are likely to advance private rather than public needs, and the protections for other values and rights in these efforts.

The need for coordination is complicated by various factors though. First, being a relative newcomer, cybersecurity lacks a hierarchy of agencies at the various levels of government, so it will be difficult to distribute responsibilities for surveillance, aggregation, and information sharing. More importantly, cybersecurity incidents lack clear geographic distinctions, and much useful data is in private (not public) institutions. So multiple industry sectors and state and federal agencies will need to coordinate.

6. Final remarks

The Doctrine of Public Cybersecurity is focused upon society rather than on any one individual. Yet it is vital that data collection and research activities respect the individual. There doubtless will be trade-offs between the rights and autonomy of individuals, versus benefits to the collective. Within public health, monitoring and information sharing has advanced specific goals and outcomes. The same is likely to hold for information about cybersecurity. But, making that case requires articulating those goals, their connection to specific kinds of information sharing activities, and ensuring appropriate protections for privacy and other competing values are built in.

Endnotes

[1] ec.europa.eu/digital-agenda/en/news/network-and-information-security-directive-co-legislators-agree-first-eu-wide-legislation

[2] Mulligan, Deirdre K., and Schneider Fred B., 2011. "Doctrine for Cybersecurity." Daedalus 140 (4): 70–92. www.cs.cornell.edu/fbs/publications/publicCYbersecDaed.pdf

[3] Information sharing activities as a way of promoting security have public good characteristics, but also must contain some restrictions and limitations in order to balance other societal values. Rosenzweig, Paul, "Cybersecurity and Public goods: the Public/Private 'Partnership'" Emerging Threats in National Security and Law, (2011). Examples given in this paper are biased to the U.S. because the authors are most familiar with events there—not because cybersecurity problems and policy challenges aren't also occurring in Europe.

[4] Sedenberg, Elaine M. and Mulligan, Deirdre K., "Public Health as a model for cybersecurity information sharing" Forthcoming in Berkeley Technology Law Journal 30:3 (2016).

[5] Rowe et al. provide a taxonomy of public health-derived interventions and maps them to specific cybersecurity threats. By further breaking down threats into communicable, non-communicable, risky behaviors, and coordinated activities, the paper presents a taxonomy of system-level interventions and prevention activities that reflect the principles discussed here. Rowe, Brent, Halpern, Michael and Lentz, Tony, "Is a Public Health Framework the Cure for Cyber Security?" CrossTalk, (2012).

[6] Bernstein, Amy B. and Haring Sweeney, Marie, Public Health Surveillance Data: Legal, Policy, Ethical, Regulatory, and Practical Issues, 61 CDC MORBIDITY & MORTALITY WKLY. REP. (SUPP.) 30, 33 (July 27, 2012), www.cdc.gov/mmwr/pdf/other/su6103.pdf

[7] In this case, "experts" refers to the Council of State and Territorial Epidemiologists (CSTE) which is an association of public health professionals that makes recommendations regarding public health information systems.

[8] Fairchild, Amy L. et al., Public Goods, Private Data: HIV and the History, Ethics, and Uses of Identifiable Public Health Information, 122 PUB. HEALTH REPS. (SUPP.) 7 (2007), www.ncbi.nlm.nih.gov/pmc/articles/PMC1804110/pdf/phr122S10007.pdf

[9] ICANN, an organization commonly referred to only in acronym form, stands for the Internet Corporation for Assigned Names and Numbers (www.icann.org).

[10] "Conficker Working Group" confickerworkinggroup.org/wiki (last modified 2011)

[11] Matthew, Ashwin Jacob, "Where in the World is the Internet? Locating Political Power in Internet Infrastructure." Dissertation. (2014) escholarship.org/uc/item/13m8k8ns#page-1

[12] Note the public health system encourages participation by reducing the possibility that information collected for public health will be used to the detriment of individuals.

[13] Given the public good nature of cybersecurity, there are already some existing programs within the U.S. to administer cybersecurity aid to developing natures to promote improvement to their networks and systems and thus aiming to decrease overall global risk. Similar international attention should be given within information sharing proposals. Hansel, Mischa, "Cyber Security Governance and the Theory of Public Goods." E-International Relations, (2013).

ACKNOWLEDGEMENTS

This opinion piece was written for IRGC by Fred B. Schneider (Cornell University), Elaine M. Sedenberg (U.C. Berkeley) and Deirdre K. Mulligan (U.C. Berkeley). Comments and suggestions from Granger Morgan (CMU) and Marie-Valentine Florin (IRGC) are gratefully acknowledged.

External contributors provided valuable input through an external peer-review, and we wish to thank in particular Maya Bundt (Swiss Re), Thomas Haeberlen (BIS) and James Larus (EPFL) who provided comments and suggestions for improvement.

Disclaimer

The views and policy prescriptions contained in this document are those of the authors, and are not a consensus judgment by IRGC, its reviewers, or its sponsors.

About IRGC

The **International Risk Governance Council** (IRGC) is an independent non-profit foundation whose purpose is to help improve the understanding and governance of systemic risks that have impacts on human health and safety, the environment, the economy and society at large. IRGC's mission includes developing risk governance concepts and providing risk governance policy advice to decision-makers in the private and public sectors on key emerging or neglected issues. IRGC was established in 2003 on the initiative of the Swiss government and works with partners in Asia, the US and Europe.
For more information, please visit www.irgc.org.

The **International Risk Governance Center** at the Ecole Polytechnique Fédérale de Lausanne (EPFL) was established in 2016 to organise IRGC activities.

Members of the IRGC Foundation Board

Philippe Gillet (Chairman), Vice-President and Provost, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland; **Charles Kleiber** (Vice-chairman), Former State Secretary for Education and Research, Switzerland; **John Drzik**, President, Global Risk and Specialties, Marsh, Inc.; Chairman, Marsh & McLennan Companies Global Risk Center; **Christian Mumenthaler**, CEO, Reinsurance and Member of the Executive Committee, Swiss Re, Switzerland; **Daniele Tonella**, CEO, AXA Technology Services, France; **Margareta Wahlström**, former Assistant Secretary-General, Special Representative of the Secretary-General for Disaster Reduction (UNISDR), Switzerland; **Wang Weizhong**, former Vice-minister, Ministry of Science and Technology, People's Republic of China

Members of the IRGC Scientific and Technical Council

Prof. M. Granger Morgan (Chairman), University and Lord Chair Professor, Department of Engineering and Public Policy, Carnegie Mellon University, USA; **Dr. V.S. Arunachalam**, Founder Chairman, Center for Study of Science, Technology and Policy (CSTEP), India; **Prof. Wändi Bruine de Bruin**, Professor of Behavioural Decision Making, Leeds University Business School, UK; **Prof. Luis Abdón Cifuentes**, Associate Professor of Industrial Engineering, Pontificia Universidad Católica, Chile; **Dr. Gérard Escher**, Senior Advisor to the President, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland; **Dr. John D. Graham**, Dean, Indiana University School of Public and Environmental Affairs, USA; **Prof. Manuel Heitor**, Professor, Instituto Superior Técnico, Technical University of Lisbon, Portugal; **Prof. Janet Hering**, Director, Swiss Federal Institute of Aquatic Science and Technology (EAWAG), Switzerland; **Prof. Kenneth Oye**, Associate Professor of Political Science and Engineering Systems, Massachusetts Institute of Technology (MIT), USA; **Prof. Arthur Petersen**, Professor of Science, Technology and Public Policy, University College London (UCL), UK; Prof. Ortwin Renn, Dean of the Economic and Social Science Department, University of Stuttgart, Germany; **Prof. Jonathan B. Wiener**, William R. and Thomas L. Perkins Professor of Law, Duke Law School, USA; **Prof. Xue Lan**, Professor and Dean, School of Public Policy and Management, Tsinghua University, People's Republic of China

international risk
governance center

International Risk Governance Center

Ecole Polytechnique Fédérale de Lausanne
CM 1 517
Station 10
1015 Lausanne
Switzerland

Tel +41 21 693 82 90
Fax +41 21 693 82 95

irgc@epfl.ch
irgc.epfl.ch
irgc.org

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

international risk
governance council