



Cybersecurity Dilemmas: Technology, Policy, and Incentives: Summary of Discussions at the 2014 Raymond and Beverly Sackler U.S.-U.K. Scientific Forum

DETAILS

32 pages | 8.5 x 11 | | ISBN 978-0-309-38005-8

AUTHORS

National Academy of Sciences; Royal Society

BUY THIS BOOK

FIND RELATED TITLES

Visit the National Academies Press at NAP.edu and login or register to get:

- Access to free PDF downloads of thousands of scientific reports
- 10% off the price of print titles
- Email or social media notifications of new titles related to your interests
- Special offers and discounts



Distribution, posting, or copying of this PDF is strictly prohibited without written permission of the National Academies Press. (Request Permission) Unless otherwise indicated, all materials in this PDF are copyrighted by the National Academy of Sciences.

Cybersecurity Dilemmas

TECHNOLOGY, POLICY, AND INCENTIVES

SUMMARY OF DISCUSSIONS AT THE
2014 Raymond and Beverly Sackler
U.S.-U.K. Scientific Forum



NATIONAL ACADEMY
OF SCIENCES

THE
ROYAL
SOCIETY

FOREWORD

The introduction of the Internet and the World Wide Web has revolutionized the ways we work, socialize, shop, and access information. More and more aspects of our lives are being transferred online to create a world that is steadily more reliant on digital technology. A single global digital infrastructure has been created as a platform that must meet the diverse demands of different countries and sectors. As a result, cybersecurity is a growing concern for individuals, public and private organizations, and nations alike. More and more data are being shared and stored online, creating massive pools of personal information that are vulnerable to attack and exploitation by criminal and state actors.

The truly international nature of digital infrastructure creates a medium in which criminals can act maliciously, crossing borders with ease. As a result there are important international dimensions of cybersecurity and an increased need for communication and coordination between governments and companies, not just at a national level but also on a global scale. Cybersecurity is no longer solely the interest of cryptographers and software developers; it affects all of our lives, personal and professional.

The U.S. National Academy of Sciences and the Royal Society share a mission: to promote the use of science to benefit society and inform critical policy debates. This summary of the discussions that took place on the subject of cybersecurity on December 8 and 9, 2014, in Washington, D.C., will serve as a reference for decision makers, educators, and others seeking an overview of the cybersecurity dilemmas facing the world.

Since 2008, the Raymond and Beverly Sackler U.S.-U.K. scientific forums have sparked new excitement and enthusiasm for the exchange of ideas among thought leaders from the United States and United Kingdom on topics of worldwide scientific concern. This most recent forum, on cybersecurity, demonstrates how much remains to be achieved through collaboration and discussion between the two nations.

As presidents of the National Academy of Sciences and the Royal Society, we are pleased to introduce the latest piece of work supported by the Sacklers' inspired generosity.

Dr. Ralph Cicerone

President, National Academy of Sciences

Sir Paul Nurse

President, Royal Society

Contents

SUMMARY 3

1 SECURITY IN CYBERSPACE 5

- Trade-offs in Cyberspace 6
- Costs and Benefits 7
- Law Enforcement and the Internet 8
- Technological Innovations 8
- Setting Priorities 9
- Retroactive Security 10
- Regulation and Deterrence 10
- Cybersecurity Research 11

2 SAFEGUARDING PRIVACY 12

- Big Data 13
- Notice and Consent 14
- Other Approaches to Privacy Protection 15
- Balancing Competing Demands for Privacy 16
- Controlling the Use of Data 17
- Controlling the Collection of Data 18
- Norms for Data Use 19

3 INTERNATIONAL RELATIONS AND NATIONAL SECURITY 20

- Law Enforcement and Intelligence 21
- Network Effects in Surveillance 22
- Private Companies in an International Arena 23
- Arms Control as a Model for Cybersecurity 25
- Making Progress on International Issues 26

4 ACCELERATING PROGRESS IN CYBERSECURITY 29

- Gaining and Maintaining Trust 29
- Strengthening the Workforce 30
- Exerting Leadership 31
- Preparing for an Uncertain Future 32

ACKNOWLEDGMENTS

The following individuals served on the steering committee of distinguished researchers:

Ross Anderson FRS, University of Cambridge,
Eric Grosse, Google, Inc.,
Andrew Hopper FRS, University of Cambridge,¹
Butler Lampson NAS, Microsoft Corporation,
Susan Landau, Worcester Polytechnic Institute,
John McCanny FRS, Queen's University Belfast,
William Press NAS, University of Texas at Austin,¹
Angela Sasse, University College London, and
Fred Schneider, Cornell University.

This summary of the forum is drawn from the presentations and discussions of participants at the meeting. It was reviewed in draft form by Steven Bellovin, Richard Clayton, and Kieron O'Hara. The reviewers provided comments and suggestions but were not asked to endorse the views in the document, nor did they see the final draft before its release.

Oversight of the review process was provided by NAS Council member Stephen Fienberg. The summary was prepared by consultant writer Steve Olson and with staff assistance from Lynette Millett, Alice Jamieson, and Jon Eisenberg.

Sincere thanks to the Raymond and Beverly Sackler U.S.-U.K. Scientific Forum for support of this activity.

THE NATIONAL ACADEMY OF SCIENCES (NAS) was established to advise the United States on scientific and technical issues when President Lincoln signed a Congressional charter in 1863. The National Academies of Sciences, Engineering, and Medicine have issued numerous reports on topics related to cybersecurity, privacy in the information age, and societal implications of information technology.

THE ROYAL SOCIETY is a self-governing Fellowship of many of the world's most distinguished scientists. Its members are drawn from all areas of science, engineering, and medicine. It is the national academy of science in the U.K. The Society's fundamental purpose, reflected in its founding Charters of the 1660s, is to recognize, promote, and support excellence in science, and to encourage the development and use of science for the benefit of humanity.

¹Profs. Hopper and Press served as co-chairs of the steering committee and of the forum.



NATIONAL ACADEMY
OF SCIENCES

THE
ROYAL
SOCIETY

Summary

The Raymond and Beverly Sackler U.S.-U.K. Scientific Forum “Cybersecurity Dilemmas: Technology, Policy, and Incentives” was held on December 8 and 9, 2014, at the Washington, D.C., headquarters of the National Academies of Sciences, Engineering, and Medicine. With support from the Computer Science and Telecommunications Board (CSTB) of the Academies, the forum was organized by a steering committee of distinguished researchers from the United States and the United Kingdom.

The forum brought together approximately 60 participants from academia, government, industry, philanthropy, and nongovernmental organizations. Participants included former senior government officials from the United States and the United Kingdom as well as individuals from both countries who have been critical of the policies of their respective governments. The forum was held under the Chatham House Rule, which specifies that the ideas expressed at a meeting may not be attributed to any particular individual or institution and that the list of attendees may not be circulated beyond those who participated. The intention was to create a setting where participants could speak frankly as individuals, even about issues that affect their own organizations or countries. The two-day meeting included presentations and discussions on such topics as cybersecurity and international relations, privacy, rational cybersecurity, and accelerating progress in cybersecurity.

This summary of the forum is drawn from the comments made by participants at the meeting but does not reflect a consensus of those present or of the sponsoring organizations. However, the observations and proposed actions in this document provide an overview of key issues in cybersecurity from a group of people working at the forefront of the field.

Cybersecurity can be seen as demanding a trade-off between functionality and security: users demand flexibility and complexity in the systems they use, but this demand significantly increases the difficulty of ensuring the security of the system. Although perfect cybersecurity is not possible, there are many opportunities to improve systems and better protect their users.

A major concern for individuals is how they can protect their privacy in a world where data about them are increasingly collected, stored, and used for a variety of purposes. Different stakeholders have conflicting interests in the balance between privacy and data collection. Although some service providers are primarily interested in collecting as much data as possible, even if it is not immediately useful, individual customers value their privacy and autonomy. Customers' stored data may be anonymized, but such data can be stitched back together to create a detailed profile of an individual with relative ease. If data collection and storage are not carefully controlled, they can introduce new opportunities for criminals to gain access to them for malicious purposes.

In our interconnected world, cyberspace is a key topic that transcends borders and should influence (as well as be influenced by) international relations. As such, both national and international laws will need careful evaluation to help ensure the conviction of cybercriminals, support companies that work internationally, and protect national security. To meet the growing demand for protecting national security, international law and norms could be strengthened to reduce the risk of international cyberattacks. In addition, there is a growing need for future leaders in both the private and public sectors understand and articulate the implications of cybersecurity risks for their own organizations and for the wider economic and social system.



1 Security in Cyberspace

Individuals, businesses, governments, and society at large have tied their future to information technologies, and activities carried out in cyberspace have become integral to daily life. Yet these activities—many of them drivers of economic development—are under constant attack from vandals, criminals, terrorists, hostile states, and other malevolent actors. In addition, a variety of legitimate actors, including businesses and governments, have an interest in collecting, analyzing, and storing information from and about individuals and organizations, potentially creating security and privacy risks.

Cybersecurity encompasses all the activities designed to protect work being carried out in cyberspace from the hostile actions of adversaries. Cybersecurity is made extremely difficult by the incredible complexity and scale of cyberspace. The challenges to achieving cybersecurity constantly change as technologies advance, new applications of information technologies emerge, and societal norms evolve.

On December 8 and 9, 2014, the Raymond and Beverly Sackler U.S.-U.K. Scientific Forum “Cybersecurity Dilemmas: Technology, Policy, and Incentives” examined a broad range of issues associated with cybersecurity. Organized by the National Academy of Sciences and the Royal Society, the forum brought together about 60 invited participants in Washington, D.C., for a day and a half of presentations and discussions on such topics as cybersecurity and international relations, privacy, rational cybersecurity, and accelerating progress in cybersecurity.

This summary of the forum is drawn from the comments made by participants at the meeting and does not reflect a consensus of those present or of the sponsoring organizations. Rather, it explores some of the more prominent dilemmas surrounding cybersecurity, identified in italicized boldface text, as well as issues related to those dilemmas.

Trade-offs in Cyberspace



A dilemma at the heart of cybersecurity is that people want conflicting things from computer and communications technologies.

They want a technology to have the most modern and powerful features, be convenient to use, offer anonymity in certain circumstances, and be secure. But these attributes have competing requirements. For example, simpler systems are fairly easily made secure, but over time people demand more functionality, and the greater complexity that results

makes systems less secure. Similarly, although a complex system can be better protected by isolating it or by sanitizing all input, doing so makes the system less useful and less valuable to its users.

Users of computing and communications technologies understandably focus on getting the job done. If a security solution gets in the way, these people will find ways around it—for example, by remotely connecting an unsecured laptop so they can work at home or demanding that a particular information technology work in almost any circumstance or setting.

Because of these conflicting desires, many abstract cybersecurity goals are not realistic. Security is often a relatively low priority for the individuals using information systems. Indeed, unless users see a clear advantage in the security being provided, they generally are unwilling to tolerate systems that are slower or more expensive or less capable simply because they are more secure.

It is hard to estimate the total cost of cybersecurity breaches. Security experts understandably tend to focus on the worst things that could happen to systems, and users and cybersecurity vendors likewise often claim very large estimates of the damage resulting from breaches. Individual users, on the other hand, tend to think more about what has happened to them, to people they know, or to people they recognize as being similar to themselves. Moreover, the people being harmed by security lapses or measures may not be the same people who decide which security approaches and methods to use.

The harms that result from cybersecurity breaches can go well beyond economic costs. They include embarrassment and disruption, such as private pictures being distributed or data about corporate salaries being released. Economic and other harms to particular individuals or companies (or nations) can be significant. The harms from breaches may be small for any one individual but large in the aggregate. An ongoing challenge in cybersecurity is to understand the costs of breaches as compared to the costs (sometimes in the form of inefficiencies) of additional security measures.

An ongoing challenge in cybersecurity is to understand the costs of breaches as compared to the costs of additional security measures.

Costs and Benefits



Another prominent dilemma involves misaligned incentives. One way to think about the cybersecurity problem is to see it purely in terms of attackers and targets. A target has something that an attacker wants, and an attacker uses information technologies to try to get it. This view diminishes the importance of the context, including the value of the asset and the cost of the attack.

But the values of assets differ. Some targets are the equivalent of nuclear launch codes, which need extremely high-assurance protections. Others are online newspaper subscriptions, which need lower assurance protections. Also, attacks generally have costs for the attacker, both in terms of the resources required to mount the attack and potential costs if an attack is detected and punished. Unless the expected return from an attack is greater than the cost of the attack, the attack will be uneconomical.

These trade-offs require that decisions be made about the effort devoted to protecting assets. For example, what needs to be done to protect high-assurance assets, and what can be neglected in protecting low-assurance assets? Treating low-assurance assets

as valuable assets—as is done, for example, when complex password rules are applied to low-assurance assets, or when people exaggerate the costs of cybercrime—leads to the irrational use of resources. In addition, sometimes it may be easier and cheaper to disrupt criminal activity down the line rather than to thwart it in advance by introducing rigorous security measures. For instance, it could be made harder to use stolen data, or the markets where criminal goods and services are exchanged could be disrupted. Further, the complexity that makes security hard also makes it difficult for individuals to be successful cybervillains on their own. Even if one person could, say, steal data, that person would

still need a network of other specialists (such as malware designers, fake website designers, and money launderers) to carry out a criminal enterprise that exploits the data. If these networks can be disrupted, then the potential payoff of cybercrime can be limited.

Today, no good answer exists to the question of how rigorously people should protect their Internet accounts, or how much money should be spent on improving computer security. Even a simple question such as whether to mask passwords as people type them in (that is, replacing the symbols with a bullet, asterisk, or some other character) is difficult to answer, because the threat of shoulder surfing, where people steal someone else's password by reading it as they type it in, could be replaced by other threats when passwords are masked. In fact, many claims about which practices are most effective in computer security are difficult to refute, both because there is no relevant evidence available and because gathering such evidence, if it existed, would be difficult.



Law Enforcement and the Internet

Law enforcement activities frequently engage with information and activities in cyberspace, whether testing an alibi or attempting to uncover terrorist plots. In addition, access to communications data has become an important investigative tool for the police. For example, the large majority of serious crime cases prosecuted in the United Kingdom are said to rely on such access, in part because the data are relatively easy to obtain, whereas in the past such cases were prosecuted by other means.

Nonetheless, online criminal activities can run far ahead of the capabilities of law enforcement. Highly sophisticated gangs are using computer and communications technology to steal, smuggle, blackmail, sell drugs, and conduct other criminal activities on a large scale. Software to facilitate criminal acts can be purchased from hacking specialists, so those who benefit from a crime no longer need to be cyberexperts themselves. The most serious criminals then can base themselves in jurisdictions that do not have established mechanisms for assisting other countries with law enforcement cases. At the same time, an understanding of criminal motives and structures can aid law enforcement efforts. Criminal coalitions will need to generate specific trust-promoting structures and systems, which (given they are to be used by criminals) is a nontrivial problem. In these situations, ***the dilemma is that advancing information technologies facilitates cybercrime at the same time as it helps the efforts of law enforcement to prevent and solve such crime.***



Technological Innovations

While technology cannot provide perfect security, new technology could provide greater security than exists today. For example, tamperproof audit trails and logs could cover all uses of data and enhance deterrence. Robust identity systems could be applied to people, programs, and machines. Technology development could change the balance and nature of trade-offs, and careful analysis of problems could yield improvements.

Stronger security technologies and procedures have been developed, but evidence of their efficacy and cost effectiveness is still lacking. High-assurance systems are possible, but they would likely be less functional and agile. Fundamental challenges include deciding which parts of the computing world need which levels of protection, determining how much added security will cost, and agreeing on how those costs should be distributed.

Setting Priorities


Providing cybersecurity typically means that trade-offs have to be made among the desired attributes of systems. Setting priorities can guide these trade-offs. One option is to limit the aspirations of systems by not trying to make everything secure. For example, the designers or users of a system could decide what really needs to be protected and what is not as important, just as people decide which assets to put in a bank vault and which to protect less securely. In this way, rational security policies would protect people only as much as they need to be protected.

As an example of setting priorities, the computing world could be divided into sectors that are more safe and accountable and sectors that are less safe and accountable. The sectors that call for more security might require centralized management and ways to control the input to systems, since they would still have vulnerabilities. One way to implement such an approach would be to identify a sector that handles only fully authenticated interactions, with accountability achieved by allowing interactions only with parties that are fully identified. Identity could be established through a combination of the person, the machine, and the program, and full audit logs could further enhance accountability.

One challenge with this approach would be moving information from a less secure zone to a more secure zone. Information could be sanitized—for example, by taking out all the scripting language and other executable code. That would reduce functionality, but there would have to be reasonable assurance that the more secure zone had not been compromised.

Another way to establish priorities would be to make it harder to target important assets. Most accounts contain relatively low-value assets, and attackers cannot target everyone with the most sophisticated possible attacks, since such attacks are expensive. Furthermore, it can be costly for an attacker to figure out which accounts contain more valuable assets. Systems also could be designed to enable their users to remain obscure on the Internet—for example, by dividing their information among multiple unlinked accounts, which would make it harder to identify valuable targets.

Approaches like these would have to overcome difficulties. For example, weakly protected accounts can become more valuable over time as people use them more and for more things. Today, for instance, basic e-mail accounts should be viewed as extremely sensitive, since they are often used to reset passwords for a wide variety of services. Yet the security on the accounts may not be upgraded in line with their increase in value over time, rendering their users more vulnerable.



Weakly protected accounts can become more valuable over time as people use them for more things.

Retroactive Security

Given that perfect security is not possible in cyberspace, one possibility is to move toward retroactive security measures rather than try to prevent all the bad things that could happen. For example, in the financial system, the fundamental basis for security is that almost any transaction can be undone. Preventive measures would still exist, but the emphasis would be on reacting to security issues after the fact rather than on trying to anticipate all possible threats. In this way, the focus would be on actual problems rather than hypothetical worst-case problems, as is often the case with physical security systems. Using this approach, actions that cannot be undone would have to be handled much more carefully. A challenge for retroactive security and the setting of priorities as described above is that the question of which things need high security is highly context-dependent. An individual's "mother's maiden name" is not a secret or a security concern until it is used to answer a financial institution's security question. Similarly, whether a transaction can be undone or not is context- and time-dependent (for example, does the institution involved still exist?). It can be a challenge to know in advance whether something is sensitive or whether it can be undone.

Regulation and Deterrence

The government could enhance its response to cybersecurity threats in a number of ways. It could increase its oversight and regulation of computer and communications technologies. It could use its convening power to encourage companies and institutions to comport with best practices. It could mandate a "safety culture" approach (similar to that seen in aviation) to cybersecurity and privacy not only in government agencies but also in the private sector. It could insist that companies provide security and privacy mechanisms in their products. Tort law could be interpreted in new ways or amended to provide increased penalties for cybersecurity breaches.

One aspect of regulation is deterrence through the threat of some kind of punishment. However, the people conducting cyberattacks are usually difficult or impossible to find and punish. Denial is easy, proof is hard, and prompt attribution is particularly difficult. Furthermore, cyberattacks and cyberexploitation are usually indistinguishable until an explicit attack is executed. Cybersecurity can be violated, for example, by the placement of a capability that allows access for some future unknown purposes.

Deterrence also runs the risk of sweeping up unintentional as well as deliberate attempts to contravene security. People who face a choice between getting their work done and observing unrealistic security guidelines make rational choices, so they need rational security systems.

Although it would be difficult to require accountability throughout a communications network, the nodes (i.e., servers or end-user devices) of a network could be made accountable for their cybersecurity provisions. For example, they could be locked out of a network or strongly isolated if they were found to be insufficiently secure. Administrators would need to detect which nodes are vulnerable or acting maliciously and be able to punish or isolate them. This authority could be delegated to a professional third party, with the responsibility decentralized rather than concentrated in a single location.



A dilemma inherent in regulation is that it tends to be a blunt instrument—slow, behind changing technologies and threats, and prone to unintended consequences, such as inhibiting innovation.

Economic incentives can be a more efficient intervention. For example, the U.S. Securities and Exchange Commission and the U.K. Financial Conduct Authority could adopt new rules requiring that data breaches or noncompliance with best security practices be reported to investors in quarterly reports. If companies are seen as acting in ways that harm their customers, they will not keep their customers' business. As is often said in the technology industry, a competitor is just a click away for consumers.

Cybersecurity Research

Additional research could yield substantial progress on many of the questions that still surround cybersecurity. For example, more study of how people apply—or circumvent—security systems would be useful for designing more rational systems. Metrics for levels of security and values of assets could enable good-enough security rather than absolute security. Is it possible to reduce the maximum harm that attackers can do while increasing the level of assurance that can be provided to potential targets given a particular level of available resources? Can an optimal cybersecurity model be envisioned along with pathways to move toward such a model? Can a machine learning system identify patterns of bad behavior in past activities and use those patterns to detect ongoing bad behavior—a goal of many intrusion detection systems today? How quickly can such approaches adapt when adversaries can use similar technologies to understand what patterns of behavior they need to change to remain undetected? Both foundational and more applied research could yield long-term progress.



Safeguarding Privacy

Cybersecurity tools and techniques are one of the foundations for trust that information will be protected, such as that trade secrets will be safeguarded or that personal information will be kept confidential. As people conduct more of their daily lives online, opportunities to acquire and misuse financial, medical, sexual, and other forms of personal information are multiplying. Furthermore, the continued development and spread of computer and communications technologies are creating new ways for companies, governments, and criminals to gain access to information that people would rather keep to themselves. And once data have been generated and exist somewhere, disclosure of those data creates the potential for harm. A particular challenge is that even if disclosure of some data is not likely to cause harm, aggregation of those data with other data may be harmful. Researchers have explored potential technical solutions to some aspects of this problem, such as differential privacy, but these work at best in limited circumstances, and the general challenge persists.

Individuals have many preferences about their privacy, and those preferences are not fixed. They are dynamic, informed by context, shaped by relationships with other people and institutions, and constantly under negotiation. Sometimes these preferences coalesce socially into expectations, norms, or conventions that are associated with particular contexts. At the same time, governments, communities, social networks, and businesses have legitimate interests in acquiring, analyzing, and using data about individuals. These interests may be commercial, governmental, or social, but they all create a desire or a need for personal information.

Big Data

The advent of the era of “big data” is further complicating the protection of privacy. Today, data are being transacted, computed, observed, and sensed, and data from many different sources can be combined. Individuals do business with companies, live in communities, associate with each other in societies, and are overseen by governments. Data are used for health care, law enforcement, intelligence, politics, education, and virtually every business. All of these data can be stored indefinitely, replicated, and combined in unlimited ways. For example, software now exists that can analyze a person’s social media posts, connect them with other data about that person available online, and construct a surprisingly detailed and accurate profile of that individual.

In the modern world, an individual’s physical, mental, and emotional state is constantly being quantified based on the data he or she generates. In some cases, people are aware that they are generating data and may give permission for these data to be used in certain ways. But these data can be used for multiple purposes, some of which

people want and others of which they do not want. As examples, data can be used to identify suspects in a crime, approve loans, sense early Alzheimer’s disease, detect a person’s learning style, infer sexual orientation or political affiliation, estimate income, identify a network of friends or acquaintances, recognize where a person is through public cameras, or detect when a person is home. Furthermore, data that can be used beneficially are the same data that can be misused. For example, data generated by playing a game online could be used to identify health problems among older people, or they could be used to calculate reaction times and discriminate against older employees.

Big data can reveal people’s activities at a continuous and intimate level and can be used in ways that make many people uncomfortable. For example, someone may enter

a query into a search website and, the next day, encounter targeted advertisements for an associated product. But often the only way to acquire a service or product is to divulge the information demanded by the provider of that service or product. People may choose to use specialized ad-free search engines or browse the Web using privacy-enhancing technologies to limit the amount of targeted advertising they receive. However, people using these approaches may experience a lower quality or utility of service or even no service at all.



Data that can
be used
beneficially are
the same data
that can be
misused.

Notice and Consent

One way people may control the collection and use of their data is through the procedure known as notice and consent. It is a contractual agreement that assumes and respects the free exchange of information and services. It serves notice that an institution wants personal data, describes what the institution will do with the data, and explains what an individual will receive in exchange for the use of his or her data. The individual replies to this notice with a yes or a no (for instance, by clicking a button on a webpage). Yes gives consent and enables access to the service; no denies consent

Given the complexity of the digital world, most people would be hard pressed to manage every aspect of their privacy.

and, generally, the individual's access to the service. In this way, individuals manage their privacy by trading it against incentives offered in the marketplace. Notice and consent makes no moral claim about whether privacy is good or not. It is simply an exchange agreement.

As originally developed in the 1970s, notice and consent was a simple and easy-to-understand system designed to respect individual autonomy and the desire to derive value from data. It worked well at a time when data collection was much less pervasive than it is today and did not include the collection of extremely fine-grained bits of data (such as the timing and targets of swipes on a smartphone). Today, notice and consent, as currently used, has serious flaws. First, for consent to be useful, it has to be informed. But to cover all contingencies, consent notices have become long, dense, difficult to read—and usually remain unread. If

an individual is not informed, that person's autonomy is largely an illusion. Furthermore, people cannot make informed decisions every time the use of a technology demands personal data, especially as technology becomes more embedded in everyday activities. The individual user is being asked to assess one of the psychologically more difficult trade-offs: that between an immediate and predictable good and a long-term and unspecified risk. Moreover, cumulative effects are also hard to assess. An individual piece of information may be harmless, but when many such pieces are aggregated, the aggregate may reveal sensitive information.

Notice and consent typically demands a yes or no answer, but someone may want their data used for some purposes and not for others. Also, preferences, technology, and the use of data can change over time, but notice and consent makes no provision for such change. Given the complexity of the digital world, most people would be hard pressed to manage every aspect of their privacy.



Even if people were given a set of options rather than a binary consent option for the use of their data, they generally cannot be told exactly how their data will be used in the future. Companies may do their best to lay out the risks of providing personal information, but they may not be able to anticipate all such risks. For example, a company may discover a use for data that was not apparent when the data were collected.

Asymmetric access to and use of information means that the users of a technology generally do not know much about what is done with their data. Many users also do not care much about the effects of disclosure in the distant future. Firms that depend on mining private data do not go out of their way to publicize their use of the information and consequent threats to privacy.

Notice and consent does not, moreover, necessarily preclude transfer of data to third parties. As a result, information granted for one purpose may be transferred to someone else who uses it for another purpose. The existence of privacy policies does not necessarily safeguard privacy; such policies could specify, for instance, that all of a person's data will be indiscriminately sold.

The provisions of a privacy policy may apply only to personally identifying information and not to information that has been "anonymized" by removing personal identifiers from that information. However, anonymized information often can be re-identified by cross-referencing it with other data sources.

Finally, much of the information being gathered about individuals today is not subject to notice and consent. It is gathered through administrative records, transactions, and other activities of daily life, and what can be inferred by combining such data may be more harmful than any individual piece of data.

Other Approaches to Privacy Protection

As discussed earlier, better cybersecurity protections and stronger accountability can help to ease the dilemmas associated with privacy. However, they cannot completely solve problems with privacy, because like notice and consent they place an undue burden on the user. Third-party privacy services could place the task in the hands of experts, but if such services then had to be purchased by individuals, inequities would be inevitable.

One alternative to notice and consent that is used more commonly in the European Union than in the United States is the concept of legitimate interests. It calls for balancing the interests of the data controller against the interests of the data subject. Under the framework outlined in the U.K.'s Data Protection Act, data controllers receive guidance about how to identify and protect these interests. In the United States, the Federal Trade Commission Act has an unfairness provision that might be used to implement a similar framework. Such a step would be consistent with the responsible use of data and could provide the basis for a universal approach to privacy protections.

One limitation of the legitimate interests approach is that it does not offer guidance on yet-to-be-invented uses of data. Also, how such a concept would be implemented remains uncertain. It could complement notice and consent, but other approaches are needed.

An analogue to the "right to be forgotten" approach has potential to deal with some of the problems of notice and consent. If such a right were part of a consent regime, it could be interpreted conservatively as a person's right to revisit consent over time and withdraw data that have been supplied. Review of consent could also be triggered by any change in a privacy policy, enabling an individual to renegotiate an existing contract. Such an approach would respect the life cycle of information, the importance and social value of which can change as the context changes.

Balancing Competing Demands for Privacy




A pressing dilemma in the era of big data is that different stakeholders have conflicting interests in the balance between privacy and data collection.

Even in a simple abstract model with just one data holder and two data subjects who exchange only cash and data, there are many scenarios in which the resulting flows of cash and data will not necessarily benefit everyone. In more complex situations, different definitions of optimality are similarly liable to lead to mixed distributions of both benefits and costs.

Within neoclassical economic theory, there are contrasting arguments for and against increased privacy protections. One argument is that privacy creates economic inefficiencies and therefore reduces economic welfare. Another argument is that stakeholders in the marketplace tend to overinvest in data collection and use, which is also inefficient and creates the risk that data will be inadvertently released that in the first place were never needed. Similarly, recent empirical research on privacy shows that either the protection of data or the collection of data can have beneficial and negative

consequences. For example, in the United States, states that legislate stricter privacy for medical data have been shown to experience lower adoption of new health technologies, in particular electronic medical records. But other results show that states with more



protections on health information are more likely to see creative and innovative approaches because the innovators have a better sense of what can and cannot be done and are less subject to regulatory uncertainty.

Similarly, the data industry can be viewed in different ways. If it allows a better match between consumers and merchants by enabling them to find each other with minimal costs, then consumers, merchants, and the data industry can all win. But if the data industry is an oligopoly with only a few gatekeepers who control the relationship or contracts between consumers and merchants, the data industry will have the upper hand with both merchants and consumers. In this case, the lack of competition can reduce choice, and resources can be transferred from consumers and merchants to the data industry rather than creating a

bigger economic pie for everyone. The outcome remains an open question.

At the root of many of these discussions is the question of who owns the data. Can an optimal balance between privacy on the one hand and data collection and use on the other be identified or maintained? An even more relevant question may be whether the interests of different stakeholders can be balanced.

The more control consumers have over their data, the more risks they are likely to take with those data, in the same way that adding safety features to cars, such as anti-lock braking systems, may lead drivers to drive faster because they feel secure. Moreover, transparency and control are necessary but not sufficient conditions for privacy protection. In the absence of other protections, there may instead be “responsibilization,” whereby end users are forced to take responsibility for something over which they actually have little control.

Controlling the Use of Data

Given the problems with current privacy regimes in this era of big data, rather than specifying that particular methods be used to protect privacy, government could regulate uses of data that pose risks. These risks could involve financial losses, physical injury, unlawful discrimination, identity theft, loss of confidentiality, and social or economic disadvantage. Under such a system, some uses of some data would be regulated or forbidden,

even if the data were gathered through notice and consent. Data could continue to be used for beneficial purposes, while harmful uses would be avoided because they would be illegal. Regulations could be applied to what might be termed “personally impactful inferences”—the combinations of existing data that represent potentially harmful use. Controls over the use of data also could apply to profiling activity.

Decisions about how data would be used and how such use would be controlled could emerge from individuals, communities, businesses, government, and society at large. These decisions could take the form of legislation, regulation, or informal standards, although different entities would have to negotiate who makes the decision, and the approach would need to be scalable so as to be widely applicable. If such a regime were to be attempted and as people gained familiarity with it, potential harms and benefits would become more apparent, so controls over use could change over time and vary from place to place.

Controlling the Collection of Data

An alternative or complement to controlling the use of data would be to control the collection of data. Disincentives to the bulk collection of data can be put in place. Entities that ask for too much data or permission to do too much with data can be identified and dissuaded from their actions—for example, by bringing those actions to the attention of potential users. The principle of purpose limitation in the European Union’s data protection directive, under which businesses can retain data only for as long as they need them, could be strengthened so that businesses do not retain data just in case a future use should arise.

Users could be given more granular control over the data they generate. For example, they could have more control over the generation of data by technologies such as cell phones. However, other information is also being gathered, such as by municipal cameras that record license plates. Furthermore, computers connected to the Internet typically send out voluminous quantities of data that can be hard to hide, and exceptional efforts to turn on privacy controls can make a user even more visible to those who are looking for such actions. Indeed, people have little control over the generation of “microdata” from everyday activities even though such data can be combined in revealing ways.

Norms for Data Use

A widely accepted set of norms for the use of data could help to protect privacy. For example, the following norms, similar to the framework provided by the Fair Information Practice Principles,² could be promoted and implemented:

- The use of data should benefit users or protect others. Benefits may be hard to pinpoint, but discussion among people representing multiple perspectives can often arrive at conclusions. At the least, the entities collecting the data could be required to explain to people how they or others are benefiting—if, say, such data collection is helping to stop fraud.
- Data should be kept secure. Security is essential to safeguard the uses of data and protect privacy.
- Users should be able to inspect, export, delete, and edit data they have provided. If people are able to review the data they have provided, they can see whether the information is accurate or they can decide to delete it. Allowing the data to be edited can be more of a challenge, since people may misrepresent themselves or their past activities or not understand the context in which the data were gathered and for what purposes. In some cases, moreover, deletion would be inappropriate—such as with financial data that need to be retained for accounting and legal purposes.

²The Fair Information Practice Principles are rooted in a 1973 report from the U.S. Department of Health, Education and Welfare, *Records, Computers, and the Rights of Citizens*.

3

International Relations and National Security

Cyberattacks can come from anywhere in the world. The relevant technology and expertise to conduct them across borders are widespread and exist in both the public and private sectors. Moreover, just as information technologies can be used to conduct crimes, they can be used as weapons to instigate or escalate conflicts and crises. Threats exist in such areas as cybersecurity attacks, electronic warfare, information operations, and psychological operations, with malevolent actors ranging from criminals and terrorists to entire nations. A cyberattack could escalate to the point where one of the parties views it as an act of war. The more apocalyptic scenarios consider what offensive cyber actions can do to highly developed states with critical infrastructures that depend on Internet capabilities.

Conventional weapons require huge investments, whereas small groups with much more modest resources can develop and deploy cyberweapons. Many actors see a cyberattack as an instrument of asymmetric warfare against the United States and the United Kingdom and their allies. They may not be able to compete on the basis of military hardware, but they can compete in cyberspace. For all these reasons, cybersecurity has critical international dimensions.

Governments face a trade-off between on the one hand using new exploits to gain access to the plans and actions of adversaries and, on the other, exposing and fixing the same exploits to increase the security of communications. The public wants transparency, but the public and private sectors must deal with the use of information technologies for national security threats. The private sector wants to protect privacy to maintain the trust of consumers but is subject to demands for information from governments.

Law Enforcement and Intelligence

Law enforcement generally does not have the capability to deal with the high level of criminal activity that is occurring on networks, which is why law enforcement agencies in some places have increasingly turned to intelligence agencies for help. It would be expensive to provide law enforcement with the capabilities already present in intelligence agencies, and duplicating a capability that already exists is inefficient.

An issue here is that intelligence and law enforcement have traditionally had different goals: law enforcement typically has reacted to crimes, while intelligence agencies typically have sought to prevent threats from being realized. Now law enforcement is being asked to prevent crimes as well, which is one reason it has called on the services of the national security community, especially for dealing with foreign threats inside their countries.

In the United Kingdom, legislation passed in 1985 (the Interception of Communications Act), 1994 (the Intelligence Services Act), and 2000 (the Regulation of Investigatory Powers Act) provided for using information technologies to tackle terrorism and serious crime. Though rarely used in the past, these provisions are now used frequently. Government funds national intelligence agencies to protect national security, including the protection of armed forces operating overseas, countering proliferation, and uncovering state-sponsored cyberattacks. These agencies have developed sophisticated means of electronic espionage, and law enforcement is keenly interested in using these same tools to attack crime.

The United Kingdom decided more than 20 years ago, well in advance of its European partners, to impose the same basic regime for limiting intrusive investigative activity on its intelligence activities as on law enforcement. Laws regarding intelligence aim at preventing intelligence from being used for political purposes or commercial advantage. Not all countries can be expected to adopt such a model, but it suggests norms that could increasingly be adopted. Intelligence agencies could be regulated by publicly accessible laws, not by secret laws or presidential directives. Intrusive methods could be authorized by a warranting process. Principles of proportionality and necessity could be written into law and imposed as legal requirements. Intelligence activity could be independently overseen, particularly when it supports law enforcement, by an independent court to assess claims of abuse and award redress if powers have been abused.

The existing regime of mutual legal assistance treaties may require modernization to tackle cybercrime and terrorism on an international scale. Acquiring data through these treaties can take many months, which is too long to prevent many crimes or deal with a national security emergency. There are increasing jurisdictional disputes as more countries pass laws entitling their police and intelligence services to

seize data held in other countries while forbidding foreign agencies to do the same. Minimum standards for warrants, transparency, and jurisdiction could be implemented through a new international agreement.

Network Effects in Surveillance

Technology companies tend to view influence and profits in terms of networks. They try to develop and establish operating systems, social networks, software platforms, and other products in the expectation that other people will add value to those products. This has implications for cybersecurity, in that the emphasis is on rapidly increasing the number of people who use a platform, not on securing it. For example, if there are a lot of users, developers will create apps for them, and if there are a lot of apps, users will find the platform more appealing. In markets ranging from mainframes to personal computers to routers to social networks, security has tended to be added, if at all, only in the later stages of market competition.

Network effects can be seen in the intelligence world as well. As intelligence increasingly acts more like an information industry, network effects related to where most of the information accumulates and who has access to it will come into play. Network effects can influence the actions of intelligence and law enforcement agencies. For example, network effects can entangle countries with other states that use, or provide, the same platforms. Low marginal costs and technical lock-in can make it very expensive for governments or other entities to build independent networks, even if they perceive a strategic advantage in doing so.

No matter their political inclinations, most policy makers have given little thought to network effects, even though these effects could have a powerful influence on the distribution of power in the future. For example, network

effects could convey power from the leading countries to an association of developed democracies, in the same way that network effects have drawn countries outside the European Union into the association.

The economic models used in information technology (IT) and in government have traditionally been quite different. Applying lessons learned about network effects in the IT industry to international security and surveillance could prove fruitful and might illuminate strategic policy questions about surveillance, information sharing, and international affairs.

Most policy makers have given little thought to network effects, even though these effects could have a powerful influence.

Private Companies in an International Arena



Private companies that operate in multiple countries often find themselves facing dilemmas in responding to requests from governments for the data they hold. These companies have to abide by the laws of the countries in which they are based, and these laws typically take one of three forms; they

- prohibit the disclosure of information;
- require the disclosure of information;
- are agnostic as to whether information has to be released.

Two main statutes affect the disclosure of information in the United States. The first is the Stored Communications Act, which prohibits communications companies from sharing or disclosing data except in certain situations. This law does not cover responding to foreign requests in most situations. The second is the Pen Register Act, which is part of the Electronic Communications Privacy Act. It prohibits companies from disclosing data that move across networks unless certain exceptions apply. In the United Kingdom, the main statute that covers the protection of personal information is the Data Protection Act, which implements the European Union's Data Protection Directive. It prohibits the transfer of personal data to any country outside the European Economic Area, unless that country can ensure an adequate level of personal data protection.

As companies receive more and more requests from foreign countries, they have developed policies to try to address these requests. Many of the largest companies have published transparency reports that describe the legal processes associated with the release of information. These processes are very similar, although there are some differences from company to company.

In general, if the foreign country requesting the information respects the rule of law, has a good legal system and a good human rights record, and the request complies with the local law of the jurisdiction in question, then a company is much more likely to disclose the data. However, requests are considered on a case-by-case basis, which is a resource-intensive process. Sometimes companies have no choice but to curtail or eliminate their operations within a given country because of the legal demands or restrictions they face in that country.

The revelation that the U.S. government has conducted large-scale surveillance of entities outside the United States has led some countries to consider enacting laws that would impede such actions. Other countries also have sought to enhance their own surveillance authorities, as a way to protect their own citizens.



A proliferation of such laws would further increase the difficulties companies face in deciding how to respond to data requests. A country where a data subject is located may have a law that prohibits the release of data. Another country without such a law may be interested in those data and request them. Companies try to navigate their way around conflicting sovereign interests, but the situation is difficult and is likely to become more so. Current mechanisms would need to be improved or new ones found to satisfy each country's sovereign interests.

Issues like these have arisen in other contexts, so precedents and models do exist for making decisions. For example, treaties are the classic way for countries to deal with disagreements. In the context of information, the most important treaties are mutual legal assistance treaties. In some cases, countries can take advantage of mechanisms unilaterally. For example, a country could say that it is permissible for companies within its jurisdiction to cooperate with requests from other jurisdictions in particular situations. In such a case, domestic law can facilitate information sharing without going through difficult treaty negotiations.

In a joint investigation, law enforcement in two countries may be interested in the same criminal act, in which case an agency in the first country can get information from data providers in that country and share it with authorities in the second country. Sharing of information among law enforcement agencies also can happen informally without opening a joint investigation. Other options are available for international data sharing, creating several choices for a given situation.

Other countries have been considering whether they should require the use of local service providers instead of nonlocal providers in the possibly naïve hope of blocking efforts by the U.S. government to gain access to data. Similarly, many countries are defining Internet sovereignty in terms of control and censorship, which could affect hardware, software, and conventional practices in those countries.

However, such laws are likely to increase costs, and they will not eliminate all security issues and may introduce new ones. They also will not necessarily advance the economic and social interests of those countries, since they erect what is essentially a tariff barrier, making it more expensive to offer digital services in that country while facilitating censorship and social control.

Companies will continue to struggle with the competing demands from different nation-states, but network effects will press against the desire to establish separate, closed Internets. The existing multistakeholder governance system for the Internet can help resolve some but not all jurisdictional problems.

Arms Control as a Model for Cybersecurity

Protection from hostile cyber actions falls into four broad categories:

- Cyberdefense – protecting important IT assets.
- Cyberdeterrence – dissuading adversaries from launching hostile operations.
- Cyberpreemption and damage limitation – reducing the capability of the forces that an adversary might use.
- Cyber arms control – can entail workable agreements with potential adversaries to reduce the likelihood of hostile cyber operations and reducing damage should hostile operations occur.

Arms control agreements can have varying scope. They can be universal, such as the Geneva conventions that prohibit attacks on certain kinds of targets. They can be multilateral or bilateral, such as the agreements among NATO members or between the United States and Russia. Or they can be unilateral, where one country takes action for such purposes as reassuring others about its true purposes. Arms control agreements also can have varying mechanisms. Treaties, memoranda of understanding, and coordinated unilateral policies can all control the actions of signatories to the agreement.

One application of an arms control framework to cybersecurity might involve limitations on acquiring offensive capabilities. However, verification, a key element in arms control, may be very difficult. The operational capability of such a limit would depend on research and development, not on delivering manufactured systems. Moreover, seeing activities in cyberspace is hard unless they are conducted on a large scale. Cyberoperations depend on deception. Behavior does not always reveal intent, and intent is important in cyberspace, as elsewhere. Understanding intent depends on deeper knowledge, which would if revealed enable the adversary to anticipate actions and mount more effective defenses. Finally, the instrumentation needed to gather data would likely be extensive, highly intrusive, and easy to evade.

Another application of an arms control framework could be limiting the use of cyberattacks, for example, on national financial systems or power grids. Such limits may require cooperative measures, such as electronic identification of prohibited targets, analogous to the time-honored painting of a red cross on a hospital or ambulance. Such arrangements may not ensure compliance, but they could create or reinforce international or national norms regarding the acceptability of such behavior and be enforceable through reciprocal threat. They could also help to inhibit overt threats or to clarify redlines in an escalation ladder.

Cyberdeterrence has major legal and policy implications. It can work at the legal, policy, or operational level. For example, deterrence could involve defining a line past which a response is swift, sure, and damaging. One problem, however, is that redlines are constantly moving as the issues and technologies evolve, thus increasing the need for dialogue. The importance of these issues further emphasizes the importance of simulations and exercises.

The most likely application of an arms control framework would be through confidence-building measures. Examples from traditional arms control include notification of activities that might be observed but misinterpreted, means for communication during times of tension, agreed conventions for behavior, and non-interference with gathering data for verification of compliance.

Even small steps could yield progress. Development of a common vocabulary and conceptual structure could enhance mutual understanding. The desire to curb activities that countries generally agree are illegal could foster international cooperation. And communicating during crises, differentiating espionage from attack, cooperating against third-party provocateurs, or declaring cyber ceasefires could prevent inadvertent escalation.

Making Progress on International Issues

The international dimensions of cybersecurity will have a profound impact on the future of IT. The freedom, governance, and stewardship of the Internet are in play. Issues such as cyber sovereignty, censorship, and net neutrality are all highly salient.

National cyber strategies for peacetime, conflict, crisis, and warfare could be strengthened. Procedures to engage with adversaries could be compared and correlated within a country and perhaps internationally, as through the formation of cyber alliances or confidence-building measures. Cyber architectures, technologies, designs, and innovations in such areas as the cloud, big data, encryption, and identity management could be tracked and their impacts on international relations assessed. Cyber-related command-and-control systems, battle management, and situational awareness could all receive much greater attention. Gaming, exercises, simulations, and other forms of assessment could enhance preparation.

Non-state actors are wild cards for managing stability, because they can instigate or escalate crises. The legal notion that states are responsible for the actions of their citizens is often unenforceable in today's world. However, attribution of actions is not necessarily as difficult as many non-state actors assume it is. Non-state actors could be identified in a noncrisis period so that they do not continue to believe that they are acting anonymously.

International law and norms to protect against international cyberattacks could be strengthened. A nation that finds itself under a massive cyberattack should be able to call for and expect international support. Article 28 of the United Nations Declaration of Human Rights, which protects the rights and freedoms set forth in the declaration, applies in the online world as well as the offline world. International humanitarian law also applies in cyberspace. Principles of protecting civilians and avoiding collateral damage apply in cyberspace. If it is a war crime to drop a bomb on a hospital, it is a war crime to disable a hospital with a cyberattack.

The government cannot delegate to the private sector the responsibility to police the Internet. However, companies do have a responsibility to their shareholders and owners to protect their reputations. If a company makes no reasonable attempt to detect illegal activities or cooperate with authorities, its reputation can suffer. This is another reason for dialogue between the public and private sectors.



At an international level, existing and new norms could be established and reinforced. For example, the U.S. President has suggested one new norm—namely, that the defense should prevail in the choice between keeping a vulnerability for future covert use and disclosing it to bolster cyberdefense. The military logic is that the breach of a defense can be much more serious than losing the hypothetical value of a future tool. Similarly, nations could agree that nuclear command-and-control and space systems are off-limits to cyberattacks because such attacks might irrevocably destabilize an already tense situation. Another potential norm is that intelligence agencies will not monitor the communications of heads of state and government of close friends and allies except when there is a compelling national security purpose. However, attempts to set up a blanket no-spying agreement are not likely to succeed.

In law enforcement, a set of norms could define the principles of cooperation in international law enforcement. The main objective of the Budapest Convention on Cybercrime is to create a common policy for protecting society against cybercrime, especially by adopting appropriate legislation and fostering international cooperation. This first international treaty addressing Internet and computer crime had been ratified by 46 states as of June 2015. Other examples of cooperation include the exchange of airline passenger information, the sharing of watch-list data, and mutual legal assistance arrangements. However, data sharing can prove controversial when it conflicts with existing privacy laws.

Another possibility would be a cyber council, such as a standing body within the United Nations or another international organization, where discussions can take place as the issues and technologies evolve. All participating nations would need to buy in so that everyone has a voice and a stake in the process.

Within countries, organizations could be established to build “cyber bridges” between the needs and capabilities of the public and private sectors. Today such efforts are often piecemeal and temporary, but more permanent and substantial entities could be created. For example, institutions could work to bridge responsibilities and capabilities between law enforcement and intelligence agencies.

International cybersecurity activities, including international surveillance, require oversight. The general public cannot be invited into a national security agency, but proxies for the public could safeguard trust. These individuals would need training and guidance to do their jobs well, and they would need the right level of authority, but general principles could be established to guide their oversight.

As pointed out in Chapter 1, the potential of technology to protect bad actors remains a point of contention, as systems that offer extremely strong protection become increasingly available. Yet the use of unusually strong protections also could heighten the surveillance of the people who chose to use them. At the same time, even if stronger protections become more widely used, existing and new technologies that are less secure will continue to yield tremendous amounts of information about potential threats. As more and more information is digitized, it will become available to supplement traditional intelligence and law enforcement methods.

In many cases, laws do not align among countries. This places companies in the uncomfortable position of having to try to comport with irreconcilable laws simultaneously. Companies try to achieve a balance on these issues. Governments could enhance collaboration by providing more protection for or assistance to the private sector with regard to these challenges.

In addition to the usual conflicts between national interests, cooperation among countries in cyberspace is hampered by policy makers’ unfamiliarity with the issues, rapidly changing technologies, and not many precedents. The sociological issues are as important as, if not more important than, the technological issues in international affairs. These sociological issues comprise public policy, planning, organizational structure, legal affairs, governance, and leadership.

Countries have fundamental differences in their approach to such areas as human rights, free speech, and sovereignty. Views on democracy, privacy, intellectual property, and many other legal protections can have a strong influence on cybersecurity. Many kinds and levels of engagement and dialogue will be needed to accommodate different national perspectives, world views, policies, and technologies. However, network effects make it difficult for countries to withdraw from existing networks. One result is likely to be some degree of sociocultural convergence as people use the same tools and exchange information.

4

Accelerating Progress in Cybersecurity

Progress in cybersecurity has been slow, and government, rather than leading by example, has often lagged behind other sectors. The rate of progress could be accelerated, but this will require a sustained effort by multiple stakeholders to understand the current context, make changes, and monitor the consequences of actions taken. Resilience, flexibility, and adaptability may be more useful than heavyweight defenses.

Gaining and Maintaining Trust

Given the importance of information technologies in modern life, government has a responsibility to take extra precautionary steps. Governments could make new efforts to protect information to the proper level, prioritize resources, and achieve both oversight and transparency.³

Trust has a technological dimension. For example, establishment of identity is being advanced in both the United Kingdom, with the Identity Assurance Programme, and the United States, with the National Strategy for Trusted Identities in Cyberspace program. These programs allow private sector firms providing authentication services to

³For example, in 2013 the U.K. government increased funding for the National Cyber Security Programme by £210 million, putting the total for the 5-year program at £860 million. As part of an upgrade in cybersecurity after recent breaches, senior civil servants now have increased responsibility for managing risks. Organizations that supply services to the U.K. government must now comply with a “Cyber Essentials” scheme by adopting a set of technical controls.

federate identity and use the right identity for the right purpose. Large companies with hundreds of millions of users across the world may be able to provide more trustworthy authentication services than the government. They perform billions of authentications per day and may be better placed to spot attacks and block them faster than smaller players, including small nations. The current trend is for people to use authentication services from large firms such as Google, Facebook, or Microsoft rather than government-issued IDs when accessing private-sector services.

The users of IT have a role in maintaining cybersecurity. User education—for instance, in the area of phishing—can strengthen this role, although it is not clear what kinds of education would be most effective or long-lasting. Moreover, in many cases users have little choice about whether and how to participate in certain systems, for they are compelled to share or use data or use certain technologies. Imposing additional, complex responsibilities could be unfair. In any case, studies are needed to determine how education can be most effective in this domain. For example, it could be focused on areas with the lowest marginal costs for users to change behavior and the highest marginal benefits in terms of cybersecurity.

Strengthening the Workforce

A critical boost to cybersecurity could come through developing national talent, including elite individuals and teams. Today, both the public and the private sectors are having trouble finding enough qualified cybersecurity workers. Furthermore, professions such as the law and psychology also need people with cybersecurity backgrounds. Especially important are people who can translate or mediate between those who focus on organizational intent and those with expertise in technology.

Hiring strictures and lower salaries in government are among the factors that impede progress in the public sector, but not in all agencies. For example, the U.S. National Security Agency generally has been able to get the people it needs, in part by identifying and attracting people with strong backgrounds and providing the necessary specialized training in cybersecurity. The signals intelligence agencies in both the United States and the United Kingdom work with colleges, universities, and schools to interest students in science, technology, engineering, and mathematics and demonstrate how these skills might be applied in government. Intelligence agencies have many different kinds of jobs, allowing people to follow multiple career paths.

Exerting Leadership

Cybersecurity could be enhanced if the leaders of organizations pressed for cybersecurity, not just the people within the organization with responsibility for IT and cybersecurity. If leaders had an understanding of and interest in the topic, cybersecurity could be an ongoing concern, not something to be checked off and forgotten. For example, senior decision makers could be running desktop exercises in the boardroom or at the executive management level to test how their organizations would respond in times of a cyber crisis. They could disseminate informed and proactive messages about organizational resilience.

Leaders do not need to be experts in cybersecurity, but they do need to ask how security fits into their organizations. Can security be managed? What risks are being taken? Can security be outsourced to another organization? These kinds of benchmarking questions are being asked by leaders and in boardrooms today, which is a sign of progress.

Stronger leadership could also provide organizations with greater flexibility. Business executives, for example, might argue that they succeed in part by taking and accepting risk and that accepting some cybersecurity risk, rather than focusing on comprehensive cybersecurity protection, is the best approach. Such an approach provides further incentive for shifting focus from compliance to risk management, a direction already outlined in the U.S. National Institute of Standards and Technology (NIST) framework for critical national infrastructure cybersecurity programs. In this way, the need for security could become more widely accepted by leaders even

though they may not understand all the technical details and even though the risk-based approach also has problems.

While some government agencies respond to ongoing assessments of risk itself, they tend more often to be driven by compliance. But compliance-based measures tend to look to the past, not to future threats, and they can lead to a “box-ticking” approach to security. Again, leadership within government and its agencies can encourage thinking in terms of risk and resilience.



Preparing for an Uncertain Future

Cybersecurity is a high-stakes issue that will continue to grow in importance. What happens with IT will affect many aspects of public and private life, so cybersecurity policies need to be considered carefully. At the same time, cyberspace continues to change very rapidly, creating new opportunities for malevolent actors to disrupt the system. It can be hard to change a system that always has to be on and is used by most of the population almost continually, especially with limited funds and time.

The fundamental importance of the Internet to modern life points to the need for a continuing multistakeholder governance model with open standards. The problems people have are different and require different solutions, which calls for a multifaceted approach. Many entities have interests in these decisions, which requires not only that they have a voice in them but that people have a common understanding of cyberspace. This can be difficult, since different perspectives need to be combined to see the larger whole. Also, since many parties will be involved in improving security, the technical infrastructure will need to accommodate a wide range of inputs into the decisions about what is going to be allowed.

Innovative ways of thinking about the problem—for example, a complex systems approach, or biological metaphors for predator–prey relationships, or evolutionary perspectives on privacy policies over time—may bring progress. Technological developments, too, can yield major progress. For example, moving the operations of a government agency or of a business to the cloud could raise cybersecurity concerns, but such a move could also enable the upgrading and rethinking of an entire network.

In both the public and the private sectors, some groups are farther ahead than others in providing cybersecurity. All groups can benefit from becoming more resilient, which can put one in mind of some other relevant “R-words”: respond, retaliate, restore, repair, reconstitute, reroute, reboot, write out, and recover. Groups are now better at recognizing incidents, but many still have not implemented the cycles of improvement and change that can steadily improve strategies, capabilities, and resources. All organizations would benefit from acknowledging that they are vulnerable to cyberattack and cybersecurity failures and that they have issues that need to be addressed.

The challenges that will arise in the future are difficult to anticipate, since most of the important applications of the future almost certainly have not yet been invented. Even a decade ago, important features of the world that exists today could not have been anticipated, and the pace of innovation shows no signs of slowing down. Cybersecurity is a problem that cannot be fixed quickly or easily. Rather, many partial solutions and potential paths forward exist and will need to be implemented, which will require collaboration, collective action, and—most of all—determination.

FOR FURTHER READING

For more detailed discussion of many of the topics addressed in this document, see the following National Research Council reports, published by the National Academies Press, Washington, D.C. (before 2002, National Academy Press):

At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues, 2014

Protecting Individual Privacy in the Struggle Against Terrorists: A Framework for Program Assessment, 2008

Engaging Privacy and Information Technology in a Digital Age, 2007

Toward a Safer and More Secure Cyberspace, 2007

Trust in Cyberspace, 1999

Cryptography's Role in Securing the Information Society, 1996

Computers at Risk: Safe Computing in the Information Age, 1991



NATIONAL ACADEMY
OF SCIENCES

THE
ROYAL
SOCIETY