

CSIS

CENTER FOR STRATEGIC &
INTERNATIONAL STUDIES

FEBRUARY 2016

Managing Risk for the Internet of Things

AUTHOR

James Andrew Lewis

A Report of the CSIS Strategic Technologies Program

Managing Risk for the Internet of Things

Author

James Andrew Lewis

A Report of the CSIS Strategic Technologies Program

February 2016

CSIS | CENTER FOR STRATEGIC &
INTERNATIONAL STUDIES

About CSIS

For over 50 years, the Center for Strategic and International Studies (CSIS) has worked to develop solutions to the world's greatest policy challenges. Today, CSIS scholars are providing strategic insights and bipartisan policy solutions to help decisionmakers chart a course toward a better world.

CSIS is a nonprofit organization headquartered in Washington, D.C. The Center's 220 full-time staff and large network of affiliated scholars conduct research and analysis and develop policy initiatives that look into the future and anticipate change.

Founded at the height of the Cold War by David M. Abshire and Admiral Arleigh Burke, CSIS was dedicated to finding ways to sustain American prominence and prosperity as a force for good in the world. Since 1962, CSIS has become one of the world's preeminent international institutions focused on defense and security; regional stability; and transnational challenges ranging from energy and climate to global health and economic integration.

Thomas J. Pritzker was named chairman of the CSIS Board of Trustees in November 2015. Former U.S. deputy secretary of defense John J. Hamre has served as the Center's president and chief executive officer since 2000.

CSIS does not take specific policy positions; accordingly, all views expressed herein should be understood to be solely those of the author(s).

© 2016 by the Center for Strategic and International Studies. All rights reserved.

Acknowledgments

This report is made possible by the generous support of Symantec Corporation.

Center for Strategic & International Studies
1616 Rhode Island Avenue, NW
Washington, DC 20036
202-887-0200 | www.csis.org

Contents

Executive Summary	IV
Managing Risk for the Internet of Things	1
What Is the Internet of Things?	1
How Do We Assess Risk for the Internet of Things?	3
IoT and Mass Effect	7
The Question of Intent	9
Autonomy and Risk	11
Authentication and Encryption to Manage IoT Risk	12
Managing IoT Risk	15
Dynamic Societies Accept Risk	19
About the Author	21

Executive Summary

The majority of Internet “users” are machines, not people. The devices that make up “the Internet of Things” (IoT) connect to the Internet, take action, and create immense amounts of data. These devices will perform progressively more functions, creating new risks for safety and security, but we need more than anecdotes to assess risk and devise useful policies. An initial conclusion about security and the Internet of Things is that popular portrayals significantly exaggerate and misrepresent risk.

- The Internet of Things will be no more secure than the conventional Internet and may be more vulnerable, since many IoT devices will use simple computers with limited functionality.
- Increased vulnerability, however, does not mean an increased risk. The benefits of IoT outweigh the potential for harm, and one risk usually not considered is that premature or overreaching measures for security or privacy will stifle economic growth and innovation.
- IoT devices allow hackers to produce physical effects. Researchers have demonstrated many vulnerabilities in IoT devices, but the consequences of these vulnerabilities largely qualify as malicious pranks. Only IoT devices that perform sensitive functions or where disruption can produce mass effect will increase risk. This means most IoT devices pose little risk.
- The state of online privacy is so dreadful it is unlikely that IoT will make it worse.
- The same problems that keep us from making cyberspace more secure will slow progress in IoT security: technological uncertainty, limited international cooperation, lack of incentives for improvement, limited regulatory authority, weak online identities, and an Internet business model based on exploitation of personal data.
- We can accelerate risk reduction with the same approaches we use for general cybersecurity: research, liability, international cooperation, and regulation. The White House could repeat its approach to critical infrastructure and task sector-specific agencies to work with companies to improve the security of IoT devices they use or sell.
- Autonomy will be a key determinant for IoT risk. Limiting device autonomy or providing a way to override autonomy reduces risk. IoT standards should require a higher degree of human intervention and control for sensitive functions.
- A secure device connecting to an unsecured network does little to reduce risk. Given the weak state of security on most networks, making IoT more secure

requires better use of encryption, strong authentication, and increased resilience for both devices and networks.

- We can use three metrics—the value of data, the criticality of a function, and scalability of failure—to assess IoT risk. Devices that create valuable data, perform crucial functions, or can produce mass effect need to be held to higher standards. Those that do not can be left to market forces and the courts to correct.
- Risk is dynamic. It decreases as technology matures and as familiarity and experience grow. As we gain experience with IoT, risk will decrease.

Managing Risk for the Internet of Things

James Andrew Lewis

Dire warnings about the perils of the Internet of Things (IoT) are easy to find. These warnings misunderstand the nature of risk and how innovation makes technology safer.

The term the “Internet of Things” was first used in the 1990s to describe networked devices with computing power and Internet addresses. Like so many Internet predictions, the idea of an IoT was premature, but by 2008, machines outnumbered people as Internet “users.” These machines connect wirelessly, take action, and create data. IoT devices will perform progressively more functions and will be more efficient and cheaper than non-networked devices. Market demand will inexorably drive us to use the Internet of Things as companies and consumers use smart machines to improve existing products and services or create new ones.

This new application of digital network technology has many challenges for policy, ranging from spectrum management, privacy, data localization, and employment. It will take years to develop the policy frameworks to safely maximize the benefits of IoT. This paper looks at risk and how we measure it, as a way to guide the development of policy.

All new technologies come with risk. How much risk is another matter. Even if every single IoT device is vulnerable to attack, this does not translate into immense new risk. We must still ask if attackers will exploit every vulnerability (unlikely) and what the consequences of exploitation would be—and these consequences can range from prank to life threatening, but in only a few cases is there real risk to society. What we need to consider is how much risk is increased compared to the cyber risks we face now, and how we can manage and reduce the risk that comes from using new technologies like IoT without overreacting in ways that damage innovation, entrepreneurship, and economic growth.

What Is the Internet of Things?

The first railroads had a few miles of track, with trains that were not much faster than a horse. At its height, more than a century later, a vast rail network covered the United States with complex locomotives moving at speeds that trains or cars today cannot match. When it comes to networks of computing devices, we are still closer to the horse than the 20th Century Limited express train, but like trains, rapid,

incremental improvements are propelling us into a new environment of social and economic transition.

To continue the train analogy, as performance increased, even a journal as august as *Scientific American* stated in the nineteenth century that the human frame would not be able to withstand speeds above 45 miles an hour. Unfamiliarity and innate caution mean the appearance of risk from new technologies is greater than its reality.

Computers are not trains. They appear to think and to make decisions. This is largely an illusion, although an understandable one given the speed at which machines execute their programs. The computing devices we have today are not truly autonomous, but the fear of autonomous devices (think “Skynet,” the self-aware intelligence that challenged humans in the “Terminator” films) lies at the edge of a discussion of the Internet of Things, a fear that the future may see interconnected “thinking” machines challenge or replace their human masters. The fear that machines create new dangers or will replace humans began with the industrial age and remains both powerful and powerfully wrong. The Internet of Things is only the latest tool that expands human performance, the latest phase in the automation of routine activities that dates back to the start of the industrialization, but it has been greeted with similar fears.

The devices that will make up the Internet of Things usually have an IP address, onboard computing power, some kind of sensing device that lets them sense their environment, and most have network connectivity (often wireless). A device can be anything from a consumer product to a giant industrial machine. IoT devices will run software programs that let the device “decide” when to take a specified action or to select among different actions. IoT devices may have vulnerabilities not present in a non-IoT device.

IoT will automate many routine activities, allowing machines to make decisions without human interaction. Autonomous devices will control inventory, authorize commercial transactions over the Internet, and arrange for shipping and delivery without human intervention. Interactions will be rapid and automatic, executed according to a series of preprogrammed rules whose composition and nature may not be accessible to the user.

These changes offer immense economic benefits by increasing productivity and lowering costs. In this way, the effect of IoT will mirror the economic benefits computer technologies created when there was broad adoption by businesses in the 1980s and 1990s. Similarly, after the Internet was first commercialized, economists concluded that “information technology is the key factor in improving productivity performance in the U.S. economy in recent years.”¹ One analyst wrote: “What is exciting about information technology is not its ability to substitute for other capital,

¹ Stephen D. Oliner and Daniel E. Sichel, “The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?,” Federal Reserve Board Occasional Papers, May 2000, <http://www.federalreserve.gov/pubs/feds/2000/200020/200020pap.pdf>.

but its ability to restructure every aspect of business, in the process creating new types of markets and organizations.”² The boost to productivity provided by computers and the Internet has slowed in recent years; IoT offers the possibility to rejuvenate it.

The expansion of networked computing to devices without human operators is the next phase in a sequence that begins with the massive mainframe computers of 60 years ago. Computers became smaller, faster, connected, and embedded in functions ranging from the mundane to the complex. Computing devices have become ubiquitous and mobile. The challenges created by IoT are new as we wrestle with the question of which functions do we turn over to machines to act without human intervention. These challenges are compounded by the new and central role of data in economies (and the disturbing implications this creates for privacy) and by the flawed security environment where no computer today is adequately protected from external manipulation. Economic opportunity and the potential for increased risk provide the context for thinking about IoT policy.

How Do We Assess Risk for the Internet of Things?

People accept and manage risk. Consumers and companies make decisions based on their tolerance for risk and their estimates of both risk and the value provided by the “risky” activity. Perceptions of risk are shaped by knowledge and assumptions about safety: that manufacturers have made safe products, that standards and regulations provide guidance for production and use, and that courts will provide remedies if safety fails.

Worries about IoT and risk reflect a broader change in American society since the attacks of September 11. Instead of the sunny millennialism of the 1990s, we live in a world often portrayed as dystopic. But by most measures—life expectancy, cause of death, economic well-being, or frequency of violent conflict—risk has declined significantly for much of the world’s population.³ It is hard to argue otherwise when the leading cause of death in the United States is obesity.

Reporting in the media shapes public knowledge and attitudes, and this can distort perception of risk.⁴ The seminal work in this regard is Paul Slovic, “Perception of Risk,” published in *Science* in 1987. He wrote:

Whereas technologically sophisticated analysts employ risk assessment to evaluate hazards . . . [for] the majority of citizens’ experience with hazards tends to come from the news media, which rather thoroughly

² Yanis Bakos, “The Productivity Payoff of Computers: A review of *The Computer Revolution: An Economic Perspective* by Daniel E. Sichel,” New York University, July 3, 1998, http://people.stern.nyu.edu/bakos/sichel_review.htm.

³ This statement requires a longer discussion than warranted for this paper, but in balance, the weight of the evidence from entities like the World Health Organization and the various academic institutions that track conflict supports it.

⁴ Paul Slovic, “Perception of Risk,” *Science* 236, no. 4799 (April 17, 1987).

documents mishaps and threats. . . . The dominant perception for most Americans (and one that contrast sharply with the views of most professional risk assessors), is that they face more risk today than in the past and that future risk will be even greater than today's.

It has become a routine practice in cybersecurity for researchers to announce some vulnerability or threat and for this to be picked up by the media. It is free publicity, but this practice can distort our understanding of risk and exaggerate it by taking an individual case out of context. Anecdote replaces analysis. What counts is an assessment of actual consequences. For IoT, while billions of IoT devices are in use, there has not been a single fatality attributed to them. This may change as the use of IoT devices expands and as the functions performed by IoT devices become more sophisticated. For now, the absence of risk should form the background for any approach to IoT.

Cars are a good example of how IoT will reshape risk. Cars crash routinely from operator error. Some crashes are caused by equipment failure or, less frequently, from flaws in design or manufacturing. No one wants to be involved in a car crash, and we expect manufacturers to take steps to reduce risk.

Many researchers have shown that it is possible to “hack” cars and that the vulnerabilities they have discovered could hypothetically be used to cause a crash. The extreme scenarios are that a hacker could take over the braking and accelerator systems, or could turn the car engine off while it was driving at high speed. These examples are titillating, but not conclusive. The first question would be whether the number of crashes caused by hacking was greater than the number of crashes prevented by IoT cars. If IoT prevents more crashes than hackers can cause, there is a net gain for society. Since almost all car crashes involve operator error, a working hypothesis is that semiautonomous cars are likely to reduce operator error and bring down the number of car accidents and the net benefit to society outweighs the risk.

Risk will be reduced as innovation decreases the likelihood or the consequences of hacking a car or other IoT device. This kind of innovation is likely to be evolutionary, with successive car models more secure than the previous ones. The difference between the safety of a car built 20 years ago and one built today is substantial, and we can expect the same evolutionary process for connected cars and IoT devices. We can accelerate this innovation through regulation, but increased regulation also comes with risk that we must take into account in any decision. Regulation changes business decisions and investment. Usually, this is for the better but badly designed or unnecessary regulation can impose a cost on societies that unnecessarily reduces opportunity and growth.

The car example also helps us think about how to manage IoT risk. While cars now contain many small computers, some of which are vulnerable to hacking, there are only a few systems of great concern. Access to these critical functions that control the operation of the car creates risk. Access to an onboard entertainment system (which is likely to be wirelessly connected to the Internet) only creates risk if the entertainment

system also connects to a critical function. Hacking a car only increases risk if critical operations can be manipulated. How a car has been designed to operate in the event of a computer failure also determines risk—cars that can continue to operate with degraded IoT systems are safer.

Semiautonomous cars that automate functions like collision avoidance and lane-keeping will make cars even safer. But at the same time, semiautonomous cars that are connected to the Internet come with increased vulnerability, and these vulnerabilities can translate into some increase in risk. The public policy problem is to ask how we reduce the cost to society of bad outcomes using measures that do not themselves impose costs to innovation or reasonable freedoms. The “tools” for reducing risk are regulation, and product improvements, and lawsuits.

Manufacturers weigh the risk of lawsuits and liability costs, and the damage to the brand (which could be substantial) if a car is shown to be unsafe because it is vulnerable to hacking. Their decision also depends on regulators and their ability to create standards for safe IoT cars. Through a combination of regulatory incentives and market forces (including liability), auto companies have made cars immensely safer. In 1921, there were 24 fatalities in the United States for every million miles driven. By 2013, improvements in design and technology combined with regulation to reduce this to slightly more than one fatality per million miles and the number continues to shrink.⁵

Societies can apply these same tools to IoT. IoT creates three kinds of risk—an IoT device could malfunction; it could be hacked; or our efforts to protect privacy or make IoT devices more secure will create economic harm that outweighs the reduction in risk. Insurance companies calculate risk using actuarial data, historical records that show how often an event is likely to occur and what that event is likely to cost. We do not have actuarial data for most things in cybersecurity, including IoT. This makes the precise prediction of risk difficult, but we can define the factors that shape the risk equation:

- **Vulnerability:** The ability of an attacker to gain access and control of a computing device, manipulating or extracting data or controlling or interrupting services. Most researchers believe that the computing devices used in the Internet of Things will be even more vulnerable than the Internet technologies to which we are accustomed, given the technical limitations of many IoT computing devices. Many of these devices will lack the computing power to perform traditional security functions of familiar desktops and laptops, which makes them easy targets
- **Intent:** Simply because an IoT device is vulnerable does not mean that someone will take advantage of it for malicious purposes. An attacker has to decide to exploit a vulnerability after calculating whether attack will provide political,

⁵ National Highway Traffic Safety Administration, “Fatality Analysis Reporting System (FARS),” <http://www-fars.nhtsa.dot.gov/Main/index.aspx>.

military, economic, or social benefit. Intent can reflect simple malice, crime, espionage, terrorism, warfare—all of the usual motives seen in cybersecurity.

- Consequences: Computing devices are vulnerable and attackers may exploit these vulnerabilities, but the final question is, so what? There is already a high level of violence, crime, and accident in societies, yet societies have a remarkable ability to absorb such things. Most of the vulnerabilities found in IoT devices lead to events that would qualify as pranks. The larger question is whether IoT introduces systemic vulnerabilities that would lead to a loss of life or significant economic harm.

Vulnerability, intent, and consequences let us estimate the probability of a damaging IoT event. Most analyses have focused on IoT vulnerability, which is demonstrably high. This is not the most important variable for predicting risk. To estimate the risk created by the mix of vulnerable devices, malicious actors, and potentially harmful consequences, we need to ask how likely it is that we will see malicious action to exploit vulnerabilities to produce damaging consequences. One of our tasks in assessing risk is to parse the population of IoT devices into those where criticality of function or scalability of attack creates real risk. It is this intersection of critical function and vulnerable devices where risk is greatest.

A starting point is to assume that IoT will be no more secure than any other Internet technology—and in some cases, may even be less secure. The experience of the last 20 years has shown how difficult it is to write secure code. The sophistication (or lack thereof) of the IoT device creates additional vulnerabilities. Many IoT devices will have a limited ability to patch and update their software. They will face difficulties in managing authentication and encryption. Larger, more sophisticated IoT devices will be better able to perform security functions, but these options come with additional cost and complexity that may reduce demand for them at the consumer level. Limitations on device performance will constrain our ability to secure IoT.

Many IoT devices are consumer goods. Scenarios for causing significant damage by hacking consumer IoT devices become increasingly problematic as we look for plausible situations where hacking consumer devices produced anything other than localized and temporary effect. Turning down refrigerators to cause milk to spoil could put additional stress on cows, dairy farmers, and grocery stores. As attacks go, however, this is not very frightening.

To take an extreme case, if hackers were able to seize control of a critical aircraft system, leading to a crash, the effect could be equivalent to a terrorist bombing. This assumes, however, that the aircraft crew could not regain control. A straightforward precaution would be to ensure that the crew had the ability to override IoT systems or to reset the system to some basic operating configuration. Many devices we use now, such as aircraft, already are designed to deal with component failure, and pilot-training programs take failure into account. Similarly, taking control of an elevator would require defeating the three or four mechanical safety systems used by modern elevators.

The repeatability of an IoT attack also determines its psychological impact. Hacks that appear to be repeatable and unstoppable will create fear and uncertainty, similar to the fear and uncertainty that gripped the United States after 9/11 when it was not clear that the suicide attacks were not the opening rounds of a long campaign of attacks. The ability to cause a plane to crash creates terror, but the inability to predict when and how often these incidents will be repeated increase that fear.

Most accounts of IoT vulnerability assume that a single hacking incident can be duplicated on a mass scale, but in most instances, the challenge is not hacking a single car or refrigerator, it is hacking several thousand in situations and circumstances that produce mass effect. The number of variables involved in this kind of mass incident suggests that this kind of IoT hacking is very improbable. We do not want to extrapolate systemic effect from an example where hackers, under ideal conditions, can cause a single device to malfunction, into some larger threat to safety or security. The average level of dissonance and even chaos that modern economies accept as normal is high. IoT hacks would have to exceed this threshold to be noticeable.

Most IoT devices will not perform critical functions, nor will they generate or store critical data. This is particularly true for consumer IoT devices. This means that even if these consumer devices are hacked, the result is most likely to be annoyance. A nation with greater exposure to pranks does not face a surge in risk. It is systemic risk—the ability to create significant disruption by attacking a single critical node (like FedWire, the power grid, or a nuclear power plant) or by simultaneously attacking a large number of targets to produce significant effect. A simple precaution would be to ensure that some critical systems, which are not now linked to the Internet, remain disconnected until we can better assess and control risk.

IoT and Mass Effect

To significantly increase risk, IoT attacks must be scalable. Systemic, mass effect from hacking IoT devices is determined by two conditions: the ability to hack a single device that controls many others (sometimes known as a single point of failure) or the ability to hack many devices simultaneously. Tampering with brakes in order to murder someone is a staple of movie melodrama, but hacking the brakes of a single car is a malicious prank. Hacking many hundreds of cars at the same time might have a mass effect that increases risk.

The threshold for this mass effect is high. For example, an average of 85 people were killed every day in 2013 by car accidents in the United States, with a total of 30,057 fatal crashes that year.⁶ While tragic, undesirable, and expensive, this did not have a crippling effect on society. Hacking one car is a prank or a crime. Similarly, disabling a car engine during highway travel could well be fatal for the occupants (depending on speed and other factors) and would snarl traffic, but would not create a systemic effect.

⁶ Ibid.

Hackable IoT devices could also be used to create gigantic “botnets,” computing devices taken over by hackers and used to generate traffic to bombard the target of the Distributed Denial of Service (DDoS) attack. IoT will increase the population of capturable devices but the DDoS story is one of opponents continually increasing the scale of attacks and defenders finding ways to deflect them. This back and forth between defender and attacker in DDoS will continue as more IoT devices come online. DDoS attacks will improve, but so will defenses. Botnets are not that frightening anymore.

Catastrophic risk is the probability of an incident where a malicious attack on an IoT device would produce mass fatalities or major economic damage.⁷ Catastrophe is an overused word when applied to threats and the United States is fortunate to have experienced only two catastrophic attacks in its history—Pearl Harbor and September 11. Nuclear weapons would produce catastrophic results, but no cyber attack could duplicate the effect of a nuclear weapon’s blast and radiation.

Nor is a “cyber 9/11” a realistic scenario. Terrorists well understand the shock value of violent acts, acts that hacking cannot duplicate. While estimates of the total cost of 9/11 vary, one estimate puts the cost of the attack at \$55 billion in physical destruction and had \$123 billion in economic effect.⁸ The cost of the loss of over 3,000 lives is inestimable. The social and economic “aftershocks” are also considerable and the \$188 billion in physical destruction and economic effect does not include the upsurge in spending on homeland security or the indirect cost of impediments to economic activity. Hacking cars or toasters, even if they are hacked in their thousands, will not produce the same shock or damage as 9/11.

A more realistic comparison of a dangerous IoT hack would be with 2003 Northeast blackout, which cost an estimated \$6 billion (largely through lost production) and may have caused 11 deaths.⁹ It seems possible that a well-planned and executed attack on IoT systems that control large networks and that also lack safeguards could produce a similar effect. If it was possible to simultaneously tamper with thousands of home air conditioning system to set them maximum power during a high usage day could lead to brownouts or blackouts.

Electrical grid vulnerability to cyber attack has been a topic of concern for more than a decade. There has been some progress in making the grid more secure, but it has been uneven. Given the existing vulnerabilities of the electrical grid, IoT introduces a new kind of vulnerability but how much risk is increased depends on how IoT devices are designed and used.

In the United Kingdom, for example, some smart meters include mechanisms that can

⁷ Charles Meade and Roger C. Molander, “Considering the Effects of a Catastrophic Terrorist Attack,” RAND, 2006, http://www.rand.org/pubs/technical_reports/TR391.

⁸ Shan Carter and Amanda Cox, “One 9/11 Tally: \$3.3 Trillion,” *New York Times*, September 8, 2011, http://www.nytimes.com/interactive/2011/09/08/us/sept-11-reckoning/cost-graphic.html?_r=0.

⁹ J. R. Minkel, “The 2003 Northeast Blackout—Five Years Later,” *Scientific American*, August 13, 2008, www.scientificamerican.com/article/2003-blackout-five-years-later/.

disconnect a household from the grid. The power company controls these mechanisms. Each uses the same software, so a hacker might be able to introduce malicious code and disconnect thousand of meters simultaneously, cutting off power to entire communities, and follow this by disabling the Internet connection with the electrical utility to prevent quick repairs. Similarly, the State of California requires smart grid devices (called PCTs—programmable communicating thermostats) to allow thermostat settings to be altered remotely while denying consumers the ability to change them back. Hackers could exploit this feature. These are design flaws, since in these scenarios and others, the capacity for human intervention has been restricted, but these restrictions could allow a hacker to create a blackout that could not be easily fixed by either the utility or the consumer.¹⁰

In assessing the effect of such events we must consider two factors often left out of calculations of cyber security risk. Societies are resilient and responsive—factors that are almost always undervalued in estimating cybersecurity risk. People react and repair when there is damage. Societies can absorb more punishment than is often expected and there are often alternatives or workarounds that mitigate the scope of a failure. Societies are responsive, and a major IoT incident, or repeated small incidents, would lead to changes or improvements to reduce the possibility of it happening again.

An attacker could overcome resilience and responsiveness by unleashing a rapid series of large-scale attacks, but attacking at scale is difficult, beyond the capabilities of almost all hackers. So far, the experimental attacks against IoT devices have affected only individual systems. Given the distributed nature of much IoT technology, a mass-scale hack would require significant resources and planning capabilities. Nation states possess such capabilities, but even this would not guarantee success, if we define success as strategic effect of the IoT equivalent to a major terrorist incident.

Evil genius hackers rob banks. They do not waste their time plotting to wreak havoc on society. Military planners want immediate effect and know that repeated strikes are necessary to overcome opponents. Terrorists want drama and bloodshed. A good rule of thumb is that the more something sounds like the plot of a Hollywood thriller, the less likely it is to occur. This means that barring some massive attack or an attack on a significant system, the increase in risk from IoT is low.

The Question of Intent

In assessing whether IoT creates opportunities for disruption that potential attackers are likely to exploit, the most important factor is not vulnerability, it is intent. Intent is crucial for understanding how vulnerable IoT devices increase risk. IoT

¹⁰ Louise Downing and Jim Polson, “Hackers Find Open Back Door to Power Grid with Renewables,” Bloomberg Business, July 2, 2014, <http://www.bloomberg.com/news/articles/2014-07-01/renewable-energy-s-expansion-exposing-grids-to-hacking>; and Rick Boland and Karen Herter, “Statewide Demand Response Network Update from California,” Energy Central, December 14, 2007, <http://www.energycentral.com/articles/article/1622>.

vulnerabilities increase opportunities for malicious action, but absent intent, this does not mean that all opportunities will be taken.

We can put IoT risk in perspective by looking at the frequency of malicious events in cyberspace. In the last 15 years, there have been thousands of incidents of cyber espionage and cyber crime, a few dozen coercive acts (where states or nonstate groups—often acting proxies for states—disrupted network and data), and perhaps three or four incidents that produced physical damage or destruction. These harmful incidents took place as the Internet expanded from a few hundred million to billions of users. There have been, to date, no IoT incidents, even though IoT devices now number in the billions.

During this period, there have been routine predictions of massively disruptive cyber events. None has occurred. Despite widespread and frequently exploited vulnerabilities, cyber incidents fall into a largely predictable pattern driven by economic interests and international politics. Vulnerability is not a good predictor for attack: in the current situation in which most digital devices are vulnerable, these vulnerabilities are routinely exploited for crime and spying, but very few are exploited to create physical harm.

What we are measuring here is not whether cyber-dependent infrastructure and services in the United States are vulnerable to cyber attack, but how much IoT increases this vulnerability. We should see these IoT hacks as just another variant of cyber attacks. IoT increases the number and type of targets that an opponent can choose to attack using cyber means, but opponent decisionmaking processes on when to use such attacks will not change because of this.

State opponents may be tempted to use IoT vulnerabilities for political coercion. The Sony and Sands Los Vegas hacks were actions by North Korea and Iran to punish individual Americans. State hackers could target individuals for punishment by interfering with their IoT devices, with effects ranging from the comical to the tragic. True cyber attacks by states (e.g., those intended to coerce or damage) are likely to occur only in the event of armed conflict. In the case of armed conflict with the United States, a state might be tempted to interfere with critical infrastructure that uses IoT, but the same constraints that apply to an attack using any kind of weapon—missiles, airplanes, commandos—on the United States apply to cyber attacks, including attacks on IoT devices. Nations will not want to escalate conflict or invite retaliation.

An attacker could reasonably suspect that an attack on domestic infrastructures could escalate into conflict in the United States and might prefer to keep it regionally confined. Interfering with domestic critical infrastructure, such as the power grid, could risk dangerous escalation and a damaging U.S. response. Interfering with some systems, such as the financial system, could have damaging repercussions for the attacker as well as increasing the risk of retaliation. In a clash over the South China Sea, for example, China would likely prefer to keep conflict localized. It is also not clear that attacking IoT systems would provide military advantage in a conflict, particularly if it was local and short. We should reject scenarios that involve a replay

of World War II, for example, a drawn-out global conflict involving mass mobilization, since this kind of extended war is now unlikely given its expense and risk.

To be sure, there are a few scenarios where IoT could be used for coercive or terrorist purposes. Terrorists may consider attacks that exploit IoT vulnerabilities (putting aside for the moment the question of their ability to do so). Instead of bombing subways, they could perhaps cause trains to crash. But the same constraints that have so far limited terrorist use of cyber attack apply to attacks exploiting IoT vulnerabilities. If terrorists could identify and execute a scenario that caused mass disruption, IoT hacks could be attractive to them, but terrorists have psychological needs that lead them to prefer direct action, bloodshed, and political drama. Hacking does not meet these needs to the same extent as physical attacks.

Political activists could disrupt IoT devices to make a statement. Instead of defacing a website, activists might be able to create discomfort. The addition of many computing devices with weak security to the Internet creates an opportunity for increasing the size and number of botnets, the source of many (but not all) of the denial-of-service attacks favored for political action. Internet trolls who currently post anonymous insults online may be tempted to take things a step further and interfere with cars or household appliances. The issue here (as for terrorist and state coercion) is that targets fall into two categories: individual targets attacked for personal or political effect, and widespread systemic attacks that provide strategic effect.

Autonomy and Risk

As we have seen in the smart grid example, deciding on the appropriate level of IoT autonomy is a fundamental question for security. The balance between autonomous operation and human control shapes IoT risk. An easy way to think about this is to ask whether an IoT device replaces a human (such as in a driverless car) or augments a human (as with a smart car that helps drivers by automating functions like braking and collision avoidance). How autonomous IoT devices should be in their operations continues a debate that dates back to the earlier discussions of computers, where some scientists saw computers as augmenting human performance while others saw them as replacing humans.¹¹

Many people have already interacted with an autonomous computing device when they play video games. The computer-generated “opponent” in a video game “senses” your actions and makes decisions on how to react. This is done by a powerful computer chip contained in the game box that is running software and takes action in response to your moves, based on some preprogrammed menu of options. This happens in milliseconds. The result is an illusion that the opponent is thinking and interacting with its environment. The real world is far more complex than the artificial game environment, requiring many more inputs and much more complex

¹¹ John Markoff, *Machines of Loving Grace: The Quest for Common Ground between Humans and Robots* (New York: Ecco, 2015).

programming, but video games show the potential for autonomous devices and create a template for building autonomous systems. It is interesting to note that one of the leading makers of graphics chips for video games, Nvidia, has become an important supplier of chips for mobile computing and self-driving cars.

Deciding what degree of independence IoT devices should be given is a decision about risk. To use the car example, a car is rolling down the road, driving itself and the driver is preoccupied with texting. An emergency requires the driver to suddenly take control. The experience of autopilot systems on airplanes suggests that this abrupt transition from machine to human control creates risk—pilots are relying on a computer to fly the plane and then must suddenly make decisions without adequate awareness of the situation.¹² Additionally, the experience of aircraft autopilots shows that risk increases over time. People become accustomed to self-operating systems and their skills as a driver (or pilot) are degraded by disuse. Not only does the driver need to react quickly, he or she needs to know what to do and have experience doing it.

How much control to transfer to the autonomous system is scenario-dependent. If the likelihood of abrupt, unexpected changes is high, autonomous devices may not respond effectively. In those instances where a human operator has primary control or can override machine control, risk is considerably lessened. With manned aircraft, for example, if a pilot can correct dangerous aircraft actions in a timely fashion, risk is minimized. Risk can be managed by limiting autonomous functions, ensuring that the human operator is engaged, and by asking where human operators can safely be removed from control.

This suggests that until we can be more confident about the reliability and safety of IoT devices, it may be important to design systems to allow for manual control for systems that provide vital services or that could produce mass effect. The objection that giving consumers the power to disengage will reduce the economic benefits of IoT is true, but for the relatively small set of IoT devices that could produce mass effect or critical function, the increased risk justifies this.

There are also long-standing and serious concerns about the risk that autonomous systems will overtake, compete with, and replace humans. Progress toward such systems is contingent on the development of artificial intelligence, where computers think like humans rather than operate from a set program. For now, the security challenges created by IoT will be more prosaic: protecting data and preventing unauthorized access and control.

Authentication and Encryption to Manage IoT Risk

The technological solutions for making IoT more secure involve encryption and strong authentication of identity. Greater use of encryption and improved authentication

¹² Maria Konnikova, “The Hazards of Going on Autopilot,” *New Yorker*, September 4, 2014, <http://www.newyorker.com/science/maria-konnikova/hazards-automation>.

functions would reduce risk to privacy and to security in all Internet applications, but the adoption of both encryption and authentication have so far posed difficult challenges not just for IoT but for all Internet activities.

IoT does not change the most important problem we currently face in data and network protection—data exfiltration leading to the theft of intellectual property, business confidential information, and personal information. Data theft is a smaller problem for IoT. Most IoT devices will not store intellectual property or business confidential data.

IoT devices will create floods of new data on personal behavior. The introduction of IoT technologies comes as transitional moment for privacy. Americans traded away privacy for explosive growth in Internet services. Europeans made a different trade—they kept 1980s privacy and got a 1980s Internet economy in return. American consumers have very little control over their personal data today, much less than before the Internet was commercialized. The business model of the Internet is to extract personal data, match it with more personal data, aggregate it, and then use it for commercial purposes. Consumers accept this in an implicit trade for services—the Internet changed their preferences and behavior. IoT will make this personal data easier to generate, collect, and store.

Companies will collect data from IoT devices to improve products and services and to generate additional revenue. Although IoT will generate new kinds of data in quantity, much of this data will often be of little value (tire pressure, or the average temperature in a refrigerator, for example), but even insignificant data may be useful when aggregated for analytical purposes or correlated with other data. In general, however, most IoT data, even when aggregated and correlated, will create little risk on privacy or security.

Weak authentication is a major problem for cybersecurity and until now, both user behavior and existing technologies have meant that it is not easily fixed. It is easy to spoof identities or illicitly obtain credentials that allow an intruder to take control of a network or device. Authentication of identity establishes whether a command is legitimate or not (generally referred to as authorization). Online authentication is weak because a desire for convenience and reliability among both consumers and companies has limited the use of strong authentication technologies. People want quick access, not some complex procedure. This is why we still use passwords, which in many cases can be “cracked” in seconds. The same preferences for convenience and reliability will apply to IoT devices.

What does not work for people will not work for IoT. In IoT, a device will receive malicious code pretending to be an update or modification and accept it at face value as coming from a legitimate or trusted source. But strong authentication technologies were not designed for the simple computers with limited memory and processing power that many IoT devices will use. The first few generations of IoT devices will continue to rely on outmoded authentication technologies, and thus be vulnerable.

Americans have been able to buy strong encryption products for more than a decade, but few use them; those who do use them have implementation problems, and some encryption products have basic flaws that are easily exploitable. Encryption works better when it is centrally provided, by a service provider, as when Apple or Gmail encrypt your messages rather than your attempting to do it yourself.

Encryption changes plaintext, which anyone can read, into a jumble of letters and symbols. Strong encryption (meaning it is hard or impossible to break) requires additional computing resources (often in limited supply on IoT devices) and a way to manage the cryptographic keys used to encrypt and decrypt traffic to and from the IoT devices (the key can be a phrase or token that converts encrypted messages to plaintext and vice versa). Encryption programs can also be difficult to write. IoT poses special challenges to encryption since many devices will be simple, mobile, and rely on wireless connectivity (which is easier to intercept). If IoT follows the pattern of the human Internet, people will not make a special effort to use encryption and devices will not be designed to be secure.

Some (but not all) IoT functions will require both data and commands to be encrypted for security. The conventional solutions to encryption include public key infrastructures (PKI) and secure transport layers (such as SSL or TLS, often designated by HTTPS). PKI is a method to securely exchange encryption keys, a method for key management that allows strangers to exchange keys for coding and decoding. Large industrial IoT devices may be able to use existing encryption products, but simple devices may require the creation of lightweight encryption that require less memory and processing power.

SSL is a widely deployed encryption technology used to secure websites. SSL is not invulnerable, but it provides an adequate level of security for many consumer and commercial activities. TLS is an improved version of SSL. SSL and TLS can provide for secure authentication, and authentication technologies are likely to improve more rapidly than onboard IoT encryption. This suggests that a good strategy for IoT requirements would be to focus on strengthening authentication and authorization for IoT first. In the next few years, new authentication technologies will become available on the market. They will use various combinations of smart phones, cloud-based data analytics, behavioral patterns, and biometrics to securely identify those accessing IoT devices and seeking to issue commands.

Encrypting data, access, and control functions increases security, but it comes at a price. Encryption requires additional computing resources (often in limited supply on IoT devices). It also requires “key management,” a way to manage the cryptographic keys used to encrypt and decrypt traffic. Key management can be difficult and expensive, especially when done at scale. Onboard encryption may be unattractive to the designers of IoT devices because it adds cost and complexity. This is particularly true for consumer devices, which are more likely to have limited computing power and memory. Industrial applications of IoT (which will be fewer in number than consumer devices but produce more value for the economy) may not face the same

limitations since they will be big machines. Encrypting at the edge of the network (e.g., on the device) is messy and ineffective. IoT encryption will likely require developing less computationally intensive encryption programs, but further research and development will be required to make IoT encryption easily and securely deployable. This is an area where standard-setting processes, perhaps led by the National Institute of Standards and Technology (NIST) if necessary, could develop requirements for IoT devices appropriate to their function and cost.

IoT encryption will likely require developing less computationally intensive encryption programs and designing IoT networks to have some of the encryption tasks performed off-device, but further research and development will be required to make IoT encryption easily and securely deployable. The difficulties of encrypting data and functions will be an impediment to the IoT market and a blanket requirement that all IoT devices use strong encryption would not make sense for all IoT devices and functions and must take into account the value and sensitivity of the data or function.

Managing IoT Risk

IoT is new, but so is the Internet. The Internet was commercialized only 20 years ago. Fewer than 40 million people used it then. Today there are over 3 billion Internet users and more than 9 billion connected devices. One thing we have learned about the Internet in the last 20 years is that it creates both benefit and risk, and that the benefits outweigh the risks. This same lesson applies to securing the Internet of Things.

Much of the discussion of IoT security has focused on device vulnerability and the need to harden IoT devices. There are limits to the value of this strategy. It is more useful to think about IoT security in different ways and look at the threat environment, the degree of autonomy, and the architecture of IoT networks as key elements of any IoT risk management strategy.

Vulnerability is not a good predictor of risk. Simply because devices are vulnerable does not mean they will be hacked or that the hack will produce damaging consequences. To date, IoT risk remains largely hypothetical. This may change as the number and kind of IoT devices increases, but an increase in risk needs to be weighed against increases in safety and efficiency.

With the first Internet, there was little or no questioning of the idea of digitizing business practices to gain the benefits of lower cost and better performance. But public perceptions of risk are changing, driven by a highly publicized recognition of the number and scope of malicious cyber incidents. These changing perceptions create forces that will reshape private decisions and public policy for IoT.

Business Decisions and Government Action to Manage Risk

When it comes to IoT, we may do ourselves a disservice by pretending that IoT is some

common denominator that applies equally across very different industries and products. It is a “portmanteau” term, a simple way to describe a complex issue, but this simplicity is a bad guide for policy. Managing risk requires action at many levels by many different actors, but a few data-driven principles can help guide decisionmaking. The risks created by the use of IoT devices can be managed and reduced, but this will require some combination of research, rules, and incentives.

The same problems that keep us from making cyberspace more secure will also slow progress in IoT security: technological uncertainty, limited international cooperation, lack of incentives for improvement, limited regulatory authority for safety, weak online identities, and an Internet business model based on exploitation of personal data. At the same time, the same approaches we use to make cyberspace more secure can be used to manage and reduce the risks created by the use of IoT devices: research, incentives, and regulation. In this, the White House could repeat its approach to critical infrastructure found in Executive Order 13636 and task sector-specific agencies to work with companies to improve and ensure the security of IoT devices they use or sell.

If governments, companies, and consumers approach IoT the way they approach the original, human Internet, this will involve extending the rules we have now for liability, privacy, and cost into the Internet of Things and then identifying where the existing legal framework is inadequate, leading to either new laws and regulations, or a “common law” approach of letting the courts sort out liability. This “extension” process has important implications for topics like privacy, security, and identity, since some of the existing laws and policies created for the physical world proved to be inadequate for the human Internet and required the development of new laws, regulation, and policy.

Where and how to use an IoT device will be a business decision that will balance better performance against increases in risk and cost. Good policy can help make these business decisions easier to make. Policy and law can clarify how companies and consumers should treat risk and liability and where they should invest.

We can improve decisionmaking on IoT if we use three metrics to assess risk: the value of data, the criticality of a function, and scalability of failure. These help policymakers, regulators, and legislators identify where government intervention is necessary to secure IoT and where such actions are not needed. Turning off your refrigerator or air conditioning is annoying. Turning off a jet engine in-flight would be life threatening. Those IoT devices providing sensitive functions require a higher degree of scrutiny and effort for security. IoT creates risk when the function it performs is critical for life and safety, when the data it generates is truly sensitive, and when the effects of interference are scalable. Devices that do these things will need to be held to higher standards through government action. Those that do not can be left to market forces and court action to correct.

Decisions about autonomy will be a key determinant for the security of IoT devices. If human operators can intervene in an IoT operation, this will decrease risk. This also

decreases benefits, so societies will need to decide where and to what degree device autonomy is acceptable and where to maintain a capability for human intervention. Using our metrics of sensitivity of function and scalability of effect, we can identify which IoT devices require limits or constraints on autonomous operation.

Scalability of failure helps determine risk. To move beyond a prank or felony, a hacker needs to achieve mass effect. This means simultaneously hacking hundreds or thousands of devices—an unlikely prospect—or finding an IoT device that controls many others. These “command” devices need a higher degree of scrutiny and attention to security; others do not. The Department of Homeland Security (DHS) should take steps to identify single points of failure for critical infrastructure and then identify “perfect storm” conditions for IoT, those improbable circumstance where some combination of IoT failures (whether malicious or natural) could produce catastrophic results. The White House could repeat its February 2013 approach to critical infrastructure and task sector-specific agencies to work with companies to improve the security of IoT devices they use or sell. Similarly, investment into research on lightweight encryption and authentication schemes suitable for IoT will reinforce private-sector actions to make software products for IoT security.

IoT devices will become more secure over time through a process of incremental innovation, but we can accelerate this process with research, rules, and incentives, particularly by increasing R&D funding for cost-effective ways to make IoT devices more secure. Some incentives will be the result of litigation. If IoT devices fail, plaintiffs will seek redress and courts will assign liability. Carefully designed regulation that avoids technological prescription can speed improvement, and existing mandates for safety in transportation, health, and consumer products can be updated to cover IoT.

It is in the public interest to reduce the frequency of preventable accidents and to reduce risk, but this does not require perfect IoT security. Preemptive action based on anecdotal evidence or hypothetical situations will slow inventiveness and block improvements needed for better security. We do not know what paths IoT innovation will take or how consumers will use it, so we must leave room for experimentation and serendipity. Nothing we use comes with perfect security. As technologies mature, innovation, regulation, and market forces reduce risk, but risk still exists even in mature technologies and we live with it, assessing that the benefits we gain are greater than the likelihood of something bad happening.

This is very similar to the decisions that faced the first Internet when it was commercialized: do we choose to rapidly deploy less secure products in order to gain their economic advantage, or do we slow deployment (and innovation) to wait for better security. In the case of the Internet, the decision was to accept risk in order to reap the economic benefits of the Internet. This was the right decision and we need to make the same decision for IoT.

The alternative would be to mandate security or privacy requirements for IoT devices. This is easier to say than to do, in part because even if an IoT device is made more

secure, it will still be connecting to largely unsecured networks. IoT security faces the same problems as cybersecurity generally. The extent of software vulnerabilities is such that if someone wants to hack you, and he is persistent, he will likely succeed. Despite this, everyone uses computers and the Internet. This reflects decisions made by managers, investors, and consumers about risk. IoT increases the number of hackable devices but the consequences of this depend very much on criticality of function, scalability, and sensitivity of data.

One area for reducing IoT risk can only be addressed by government action. Governments have unique responsibility for international security. The threat environment can be reshaped if the United States takes actions internationally to create norms for responsible behavior, increases state-to-state cooperation, and develops cooperative strategies to punish those who engage in malicious cyber actions. This is the best measure for reducing systemic risk in cyberspace, since the risk of malicious actions is created by nation-states or by cybercriminals that a nation-state harbors, protects, and refuses to punish. Expanded enforcement and criminal penalties for mischievous IoT actions that nations cooperate in applying no matter where the hacker is located will reduce risk.

The introduction of IoT devices will be gradual and improvements will be incremental. This means that risk will be greatest for the first generation of IoT devices. Fears that we have only a short time to ensure that IoT is adopted in ways that minimize risk are misplaced, and the projections that there will be billions more IoT devices in the next few years creates a misimpression. The average age of a car in the United States, for example, is 12 years. Refrigerators tend to be replaced every 15 years. This means that 10 years from now, more than half of cars and refrigerators will still be “dumb” and not at increased risk. It also means that people who replace their refrigerators 10 years from now will be buying improved and more secure versions of the first “smart” appliances simply because of experience and innovation.

Critical infrastructures generally have even longer “refresh” cycles, particularly for large capital goods. This long refresh cycle makes minimizes any increase in risk. Care needs to be exercised in mandating the rapid and widespread adoption of IoT—in smart grids, for example—where the pace of adoption exceeds the pace of improvement.

Managing Risk to Data Protection and Privacy

IoT security, and cyber security in general, would benefit from a reorientation in our thinking about privacy and cybersecurity. Much of the effort of the last 20 years has focused on security networks from intrusion and on reducing vulnerabilities. Without intending to sound too pessimistic, these are hopeless tasks. Determined intruders will usually succeed in gaining access to a network, particularly if they are well resourced. Software has become so complex, with many products relying on millions of lines of code, that it is impossible to avoid or find all errors. A reorientation of our approach to cybersecurity would seek to secure data, not networks, and ensure the continued

operation of critical functions even in a degraded environment.

Data protection and privacy present similar choices. A one-size-fits-all approach to the security of data generated by an IoT device creates strange outcomes and unnecessary barriers. To use an example from an aircraft engine manufacturer, aircraft engines now report automatically to maintenance centers on their status. Some reports include the name of the mechanic at a maintenance center. Data-protection rules developed for personally identifiable information (PII) do not work for IoT. Simply because IoT data includes PII does not make it valuable or sensitive. Privacy rules developed for online transactions need adjustment for an IoT world.

IoT will greatly expand the amount of data transferred and stored across national borders, creating another set of challenges for privacy. IoT will only complicate the already complex problems of data localization and efforts to restrict data flows across borders, but (as with these rules in general) they will create obstacles for manufacturers who service a global market. Most of this data will have little value and not affect privacy. Rules to limit the use of data from IoT devices must take into account the value of the information. A German car manufacturer may collect data from foreign customers regarding their tire pressure from onboard monitors and resell it, but the harm to privacy is nonexistent. The value of this IoT data, even when aggregated, will be very low. Regulations and agreements will need to reflect this and categorize IoT data by value and privacy sensitivity (e.g., personal health data as opposed to device data).

Data protection creates unavoidable tensions between public policy and business decisions. The primary tension is the temptation for public policy to impose overarching solutions for security or privacy that do not reflect the value or sensitivity of the function or data. Most IoT data will not need strict privacy safeguards. IoT will require a graduated scale of protections and security measures that reflect the actual degree of risk, determined not just by potential vulnerability but also by the value and sensitivity of both data and functions. To argue that all data is of equal value is antiquated and inaccurate.

Data created by IoT devices will need to be parsed to determine where additional protections like encryption are needed (since encryption increases the cost and complexity of devices) and where data does not require special treatment. To argue that all data is of equal value is antiquated and inaccurate. Differentiating between important and unimportant data (individual and aggregated) and defining criticality (such as safety of life or potential economic loss) identifies which IoT systems are of concern. If we use these metrics, we find that IoT increases risk to safety and privacy at unacceptable levels in only a small number of cases.

Dynamic Societies Accept Risk

IoT is sometimes approached as a “do-over,” an opportunity to avoid the “mistakes” made in designing privacy protection and security for the first Internet. The questions

over IoT security and privacy are a continuation the larger debates over Internet security and privacy in general. The issues are the same. What has changed in unhelpful ways is our attitude toward risk. The Internet was commercialized in a more optimistic era, when people were more tolerant of the risks that the new technology might hold. Internet security and privacy were left to the market, to the decisions of individual companies, and to a very light regulatory touch. This approach produced an Internet where crime is rampant but it also produced immense economic value—crime has cost us billions, but in turn we have gained trillions. Everyone should accept this trade. Had there been security requirements or privacy restrictions at the start of the Internet age, it is very likely that the explosive growth we have seen would not have occurred.

Risk tolerance is dynamic and changes over time, correlating with public perceptions of safety and with greater familiarity with the technology. Perceptions of cyber risk are driven by growing awareness of the number and scope of malicious cyber incidents. These perceptions, however, are too imprecise to serve by themselves as a useful guide for policy, nor are perception improved by anecdotal evidence and discussion of hypothetical situations. We do not know the directions that human inventiveness and market forces will take IoT technology. We do know that over time, with experience and innovation, risk is reduced and new technologies become safer. If the first Internet paid too little attention to security and privacy, we do not want now to overcompensate. Being risk averse makes us poorer, not safer. There is risk in every technology we use. Hold IoT captive to our fears and we will sacrifice opportunity.

About the Author

James Andrew Lewis is a senior fellow and director of the Strategic Technologies Program at the Center for Strategic and International Studies (CSIS). Before joining CSIS, he worked at the Departments of State and Commerce as a Foreign Service officer and as a member of the Senior Executive Service. His government experience includes work on Asian politico-military issues, as a negotiator on conventional arms and technology transfers, and on military and intelligence-related technologies. Lewis led the U.S. delegation to the Wassenaar Arrangement Experts Group on advanced civil and military technologies and was the rapporteur for the UN Group of Government Experts on Information Security for their successful 2010, 2013 and 2015 sessions.

Since coming to CSIS, Lewis has authored numerous publications. His recent work focuses on cybersecurity, including “Cybersecurity for the 44th Presidency,” which was commended by President Obama. He is an internationally recognized expert, and his comments appear frequently in the media. Lewis has a close research partnership with the China Institutes of Contemporary International Relations. His current research examines sovereignty on the Internet, cybersecurity norms, warfare, and technological innovation. Lewis received his Ph.D. from the University of Chicago.

CSIS | CENTER FOR STRATEGIC &
INTERNATIONAL STUDIES

1616 Rhode Island Avenue NW
Washington, DC 20036
202-887-0200 | www.csis.org

Cover photo: GarryKillian / Shutterstock.com