

PART THREE

**TOMORROW'S
TERRORISM**

9 | Moore's Outlaws

In the years since 9/11, we had done a lot to meet the challenge posed by the exponential growth of international travel. Our border procedures weren't airtight, of course; they never had been. But we weren't flying blind, making decisions at the border on thirty seconds of chat and intuition, either. We'd come a long way, and while we still had further to go, at least everyone understood how important it was to finish the journey.

The challenge posed by computer security was different. There had been no dramatic meltdown. Most people still scoffed at the idea that the exponential growth of information technology revolution could lead to disaster.

Yet for some of us, losses from the information technology revolution are already greater than the gains.

Just ask Howard Crank's widow.

Howard Crank lived a quiet life that revolved around his modest California duplex. He was seventy-three years old, after all, and he'd had both legs amputated above the knee due to diabetes. His wife's health was not good. He was living on his Air Force veteran's pension. But he could afford a computer, and he loved it. It helped him find old Vietnam buddies and research new charities to add to the three dozen he already supported. He might be halfway to housebound, but the new technology was a godsend. Thanks to Moore's Law and the Internet, the whole world was at his doorstep.¹

The Internet, it appears, is how he discovered that he'd won \$715,000 in a Spanish lottery. Money was tight on an Air Force pension, so this was amazing news. Of course, it turned out that there were transfer taxes for him to pay before the winnings could be sent to him. It grew expensive, but his share of the lottery was also growing—to \$115 million.

Howard Crank's life savings were \$90,000. Bit by bit, he sent it all.

It wasn't enough. He got calls from Spain, explaining the hassles and delays. He mortgaged his home and sent the proceeds. More calls. When he wondered aloud whether he'd ever see the money, the caller asked him to have faith. They prayed together.

A few weeks later, he took out a second loan on the house. He maxed out two credit cards. All the money, perhaps \$300,000 in all, went to Spain. Even that was not enough to break his lottery winnings free. He asked his stepdaughter for \$40,000. He didn't want to explain why.

She thought that was odd. So when he was hospitalized a few weeks later with a broken thigh, she checked his financial records. She found that Howard Crank had ruined himself and his wife in an apparent Internet hustle. The Spanish scam artists disappeared without a trace. Crank died of a heart attack before he could explain how it happened.

"I think he probably knew it was a fraud at the end," his stepdaughter told me. "But he was hoping against hope. He'd sent them so much money already, and they were so convincing. But by the end he'd lost his zest for life. He was so desperate."² Desperate he should have been. He had not just squandered his own assets. A year after his death in early 2010, his widow had lost her home and been forced into bankruptcy by the debts he left behind. "She's had to move in with us. She's starting over again at the age of eighty," her son-in-law told me.³

Howard Crank would never have let a con man into the quiet life he and his wife were living. But the Internet that brought the world to his doorstep brought the world's con men as well. Information technology

empowered Howard Crank to search the world for old buddies. And it empowered fraudsters to search the world for the handful of people who might be ripe for their scam.

For Howard Crank, the exponential growth in information technology turned out to be a disaster. It was great for a while. He loved what the new technology did for him, and how cheaply it performed its miracles. But in the end, nothing he gained by embracing it was worth what he lost.

Will the rising curve of information technology eventually leave the rest of us where it left Howard Crank? You're probably thinking that Howard Crank, sympathetic as his story may be, just wasn't savvy enough. You would never fall for such a scam. And you will never suffer the harm that he did.

Well, don't be so sure. The science fiction writer William Gibson once declared that, "The future is already here. It's just unevenly distributed."⁴ He was thinking of the wonders of new technology, but bad futures are distributed as unevenly as good ones. And Howard Crank may have been in the vanguard of Americans ruined by information technology.

Thanks to information technology, it is now cheap to screen millions of people to find those who were susceptible to the lottery fraud. That same technology will make it cheap to screen the world for people and machines that are susceptible to other forms of fraud as well. You may not fall for the Spanish lottery, but you're probably susceptible to *something*. And even if you aren't, your machines are.

Are you really sure the fraudsters won't find you in the end?

Exponential technologies always seem to serve dessert first. That's why they grow exponentially. Their benefits are immediate and irresistible, so we use them in numbers that double and double again. In the beginning, it seems implausible that they will be misused. Indeed, at the outset, people do use them mostly in good, socially responsible ways. I leave it to the philosophers whether that's because people are

basically good or because it takes time for people to figure out how to be bad using new technologies.

Whatever the reason, information technology certainly followed the same path as commercial jets. It took decades between the time the technology was first democratized and the first really frightening misuse.

Until the late 1980s, the risks of misuse were almost entirely theoretical. Computer viruses had been invented by then, but mainly just to show how they would work. It wasn't until the mid-1980s that "wild" computer viruses began to spread from one PC to another via floppy disks. Then, in 1988, a worm caused much of the Internet to grind to a halt. For the academic and defense users who then dominated the Internet, the worm was a shock. But they relaxed when they found that the worm's author, Robert Morris, wasn't a spy or a criminal. He was a student, and he claimed he'd been testing a concept that got out of control.

In retrospect, what's most notable about the malware of that era is its comparative innocence. It caused damage, sure. But it was either academic or nihilistic in purpose; it demonstrated the capabilities and perhaps the ill will of the author. It wasn't really much of a threat, although the worst examples could destroy stored data.

Most attacks were the digital equivalent of the Plains Indians "counting coup" by striking an enemy with a stylized stick and escaping. Like counting coup, the purpose of early hacking was to gain prestige—more by demonstrating prowess rather than by causing harm. And computer security only needed to be good enough to outfox adolescent malcontents, a task both industry and government felt fully capable of handling.

By the mid-1990s, though, the Internet had become a fully democratized place, and money had replaced showing off as a motive for hackers. Spam was the earliest form of profitable Internet crime. And when network administrators started blocking spam by refusing to accept mail from spammers' machines, hackers found they could compromise other people's computers in bulk, then use those machines

to send the messages. If the senders of unwanted email were widely distributed, spam couldn't be stopped by quarantining a few suspect computers. Hacking wasn't just fun anymore; it could put money in the hacker's pocket.

Once underground networks of compromised machines had been assembled, it turned out that they could be used for other profitable crimes as well. If all of the captured machines could be induced to send meaningless messages to a single Internet site at the same time, the site would be unable to process them. The site would falter and fail. Legitimate users would be locked out.

Such "distributed" denial of service attacks turned into a new-style protection racket. Gambling sites, for example, simply cannot afford to be unavailable in the days and hours before the Final Four basketball tourney. If a site suffers an effective denial of service attack, there is a good chance that it will pay a reasonable "security" fee just to get back online quickly. That wasn't the only use to which criminals could put herds of zombie machines. The machines could be programmed to visit ad-supported websites and mindlessly click ads, earning illegitimate click-through fees for those sites.

But security professionals at large firms still had confidence in their defenses. Denial of service was a concern, sure, but the risk could usually be managed by retaining an ISP with lots of bandwidth and an ability to filter packets quickly. Distributed spam took away one tool for discouraging spam, but there were plenty of other ways to filter unwanted mail. For most users, spam was at worst a nuisance.

But malware continued to grow more sophisticated, and it could use the Internet to spread rapidly. Several viruses in 2000 and 2001 caught large companies unprepared and forced a shutdown of their networks while the viruses were eradicated. Hackers began to find ways to intrude into important financial and military systems.

This was getting serious.

Even so, most security experts thought the plague could be contained. They blamed systems administrators who didn't patch their systems quickly enough. Most of all they blamed Microsoft. The

company had emphasized new features over security, they complained, and in its drive to be first to market it had written sloppy code. Other operating systems were said to be more secure; and many thought that relying on a variety of operating systems was inherently superior to the “monoculture” created by Microsoft.

Stung, Microsoft fought back. Bill Gates himself took on the problem. Gates was famous for his insight into the future of the personal computer. Previous Gates messages had produced profound changes in Microsoft’s strategic direction, most famously when he wrenched Microsoft into the Internet age, focusing the entire company on the challenge posed by Netscape—and leading to Microsoft’s (temporary) victory in the browser wars.

By January 2002, Gates had a new focus. He announced that security was the key to Microsoft’s future. From now on, all of its products would be built with security in their foundation: “When we face a choice between adding features and resolving security issues, we need to choose security. Our products should emphasize security right out of the box, and we must constantly refine and improve that security as threats evolve.”⁵

The email was a call to arms. All of Microsoft’s employees were expected to bring this new focus to their jobs. In the past, a single-minded focus had enabled Microsoft to beat some of the most talented companies in the world. IBM, Lotus, WordPerfect, Ashton-Tate, Digital Research, Sun, Real, Apple, AOL, and Borland—not to mention much feared and rarely bested Japanese electronics makers like NEC and Toshiba—all had tried to stand between Microsoft and its strategic vision of the future. Microsoft had defeated them all.

Now Microsoft was gearing up for battle again. This time, though, it only had to beat a bunch of punk hackers. That should be a piece of cake. Once it was done, a new age of online security would dawn, with Microsoft’s trusted products at the heart of every online transaction.

More than seven years have passed since Microsoft set out to beat a ragged band of hackers. The company has rewritten its operating

system more or less from scratch. And its code is indeed far more secure than in 2002.

But it has not won the war. The second Tuesday of each month still brings a boatload of corrections and patches that the company must make to even its newest and most secure operating systems. By 2009, the ragged band of hackers was looking a lot more sleek and prosperous than before, and Microsoft had suffered its first revenue decline in history.

More important than Microsoft's security failures are its successes—and how little difference they have made. Microsoft has indeed tightened up the operating system. But the structure of the PC world has made that almost irrelevant. The point of the PC is the control it gives to the user—who can decide what applications to run—and to the developers, who can create new applications quickly and easily. At the end of the day, Microsoft must empower users and developers. And so that's what its security approach does. Windows Vista, for example, was famous for nagging users to confirm their dangerous decisions to run new code or open new attachments—so famous that Windows 7 has had to cut back on the nagging, despite the security risks. The one thing Microsoft can't do is forbid users to make dangerous decisions. If Microsoft tried that, it would leave its users angry and looking for a new operating system. The same is true for applications; Microsoft can't require developers to write secure code without discouraging them from writing Windows applications. And if it does that, it loses its main advantage in the market—the overwhelming number of applications that run only on Windows.

So, to the extent that Microsoft has succeeded, it has simply displaced the risk. Online security is still getting worse, but it's getting harder to blame the operating system. Instead of exploiting the operating system, more and more attacks exploit holes in applications. Or they induce the user to do something he shouldn't do.

Or both.

One night in January 2009, at about the same time that Howard Crank was sending thousands of dollars to Spain, Beny Rubinstein was getting ready to turn off his computer and go to bed.

Suddenly he got an instant message from Bryan Rutberg, a friend who worked for a technology company. Rutberg's message got right to the point.

"Look, I really need your help."⁶ Rutberg had taken a quick trip from Seattle to London, where he'd been robbed. He was broke in a foreign land. His Facebook page said the same thing, carrying an update that said, "Bryan NEEDS HELP URGENTLY!!!" Bryan needed a loan to get home. Could Rubinstein help?

Rubinstein could. He wired \$600. That wasn't enough, so he sent another wire transfer—\$1,143 in all.

In fact, Rutberg was still in Seattle. His Facebook account had been hacked, and the hacker was messaging Rutberg's friends, asking them all for quick wire transfers.

Rutberg, meanwhile, was locked out of his own account. He tried to stop the impostor by posting a comment on his own page, using his wife's account. Rutberg's comment was quickly deleted, and his wife was "unfriended." He had lost control of his online identity to a brazen scam artist.

A couple of days later, Facebook closed down the account, but Rubinstein's money is long gone. Neither Rubinstein nor Rutberg is a technological naïf. But both were defeated by the mass customization of online fraud. It's not hard to write programs that will look for weak Facebook passwords, or that will send urgent instant messages to the friends listed in compromised accounts. Only when someone responds to the messages do the scammers need to become personally involved. The marks are all prequalified.

Best of all, it's possible for the scammers to get in and get out in hours, then disappear halfway around the world. Local police are helpless; they "are not investigating this case," said a police spokesman. "It is pretty much at a dead end."

As Microsoft has tightened the operating system, hackers increasingly rely on mass social engineering and insecure applications to open a hole in the victim's defenses. Facebook is, of course, free, and

the company is famous for not having a revenue model to match its massive user base. So it's not surprising that its site still has security problems. But it's the social engineering that made this scam work. Rutberg's friends may not trust strangers who tell them they've won the Spanish lottery, but they do trust him.

In fact, the combination of "authorized malware" and targeted social engineering is so powerful that, despite Microsoft's efforts, it's now easier than ever to compromise computers, and their networks.

No one can say we weren't warned. The United States government told us all that a computer security crisis was brewing. Twice, in fact, and under two different presidents.

President Clinton cautioned in January 1999 that, "We must be ready—ready if our adversaries try to use computers to disable power grids, banking, communications and transportation networks, police, fire, and health services—or military assets."⁸

A year later President Clinton proposed a series of measures to address the security problem.

Two years later, President George W. Bush created a special adviser on cybersecurity who spent a year developing a computer security strategy.

Neither effort made much headway. The public didn't see the problem. The network attacks that alarmed Washington were classified. Officials couldn't talk about them. Meanwhile, privacy and business interests worked overtime to persuade the public that national security concerns were overwrought. The real risk was government monitoring and government regulation, they insisted.

And, by and large, that was the view that prevailed—twice, and under two presidents. Nothing was done about computer security that anyone in the privacy or business lobbies might object to.

In 2009, a third president promised to make computer security a top priority, and shortly after taking office, the Obama administration also produced a security strategy. Once again, though, the strategy lacked punch. It failed to call for any action that could possibly irritate

business or privacy groups. It spoke of cybersecurity only in alarmed generalities, unable to explain why Americans should be concerned enough to suffer even modest inconvenience.

But this time may be different. Thanks to the work of a band of Canadian security researchers, we now have a remarkable—and completely unclassified—insight into just how easily computer hackers can penetrate even carefully secured computer networks.

The young Tibetan girl waited quietly in line at the border. Call her Dechen, though that's not her real name. She had spent two comfortable years studying in Dharamsala, India. Now she was going home. She had made the long trip across Nepal to the border with Tibet. The border crossing made her uneasy, but she told herself the Chinese border guards had no reason to stop her.

Dechen was a follower of the Dalai Lama, and she had spent much of her time in India conducting computer chat sessions with his supporters inside China. But she had been careful. She really had been a student. There was no way the Chinese government could know what else she had done. Or so she hoped.

Dechen stepped forward and presented her identification to the guards. They looked it over with care. Too much care. Something was wrong. Her heart sank. She was under arrest.

She was sent to a detention center. No one would tell her why. Could she have been compromised somehow in Dharamsala? She couldn't understand it.

Finally, after two months in captivity, she was called to the interrogation facilities, where two plainclothes officers immediately began questioning her about her activities on behalf of the Dalai Lama. She denied the charges, clinging to her story.

"I was a student. They cannot know what else I did," she must have told herself.

But over and over, intelligence officers accused her of working for the Dalai Lama's youth group. Over and over, she denied it. She had gone two months without contact with friends or relatives, but if she

held firm, they would have to let her go. She stuck to her story.

Finally, the officers lost patience. They pulled out a thick file. The folder held a full transcript of her online chat sessions. It covered years. They'd known everything, recording it as she typed. They told her the names of all her coworkers. All the attempts at security, all the work of the Dalai Lama's youth group, had been defeated.

Dechen was devastated. But the officers were more interested in sending a message.

—Go back to your village, they said, and tell your coworkers that we know who they are. They are not welcome in China anymore. They can expect the same treatment you got, or worse, if they return.⁹

It was over.

The Office of His Holiness the Dalai Lama is partly a religious, partly a diplomatic mission. The Dalai Lama travels widely and seeks audiences with foreign diplomats and officials to demonstrate support for his faith and for Tibetan autonomy. The Chinese government in turn vehemently opposes autonomy for Tibet and does all it can to discourage official meetings with the Dalai Lama.

The Dalai Lama's travel schedule is thus a matter of high state interest, and the planning of his meetings has an element of cat and mouse about it. The Dalai Lama's office finds that the best way to set up those meetings is first to send an email to the officials the Dalai Lama hopes to meet and then follow up quickly with a telephone call.

But around the early part of 2008, something odd began to happen. The Dalai Lama's office would send an email to a diplomat as usual, proposing a meeting. Then it would call to discuss the details, again as usual. But the diplomat's office would be strangely cool. "We've already heard from the Chinese government," the diplomat's staff would say, "and they've strongly discouraged us from having this meeting."

The Dalai Lama and his office had been using the Internet since the 1990s. His network administrators know the risks, and they've been careful about computer security for years. They'd implemented the

standard defenses against network attacks. They didn't know what had happened. But the evidence of a serious breach was simply too strong.

They called in a team of Western computer security experts. What the experts found was deeply troubling, and not just for the Dalai Lama.

Some of the Dalai Lama's staff participate in Internet forums. They chat with other, like-minded individuals about the Dalai Lama's goals and activities. Sometimes one of their online acquaintances sends them Word or .pdf documents relevant to those activities.

The experts concluded that hackers had monitored these forums and then forged an email from a forum participant to a member of the Dalai Lama's staff. Attached to the email was a document of mutual interest. When the staff member opened the document, he also activated a piece of malware packed with it. While the staff member was reading the document, the malware installed itself in the background.

The malware was cleverly designed; two-thirds of commercial antivirus software programs would have missed it. (Hackers often subscribe to antivirus software so they can test their malware against it at leisure.) Even if one attachment were stopped, it would be a simple matter to retransmit the message using a different bit of malware; the attackers could keep trying until something got through.

Once installed, the malware would "phone home," uploading information about the victim's computer and files to a control server operated by the hackers. Next, the captured computer would download more malware to install on the staff member's machine. This was often a complete administrative program that would allow the attackers to control the staffer's computer, and in some cases the entire network.

The administrative malware took full advantage of the empowerment made possible by today's technology. It featured a graphic interface with dropdown menus offering even an unsophisticated attacker a wide variety of options.

Want to record every keystroke as the user types so you can steal all his passwords? Check one of the options on the menu.

Want to turn on the user's microphone, turning it into a bug so you can listen to the office conversations? Check another box.

Want video straight from the user's desktop camera? That's just another option on the menu.

In the end, the Dalai Lama's office was living a version of Orwell's *1984*. Telescreens in each room spied on the occupants. But in this version of *1984*, Big Brother didn't even have to pay for this spy equipment. It had been purchased and installed by the victims.

Once the hackers had compromised a single computer on the network, it wasn't hard to compromise more. Every time an infected computer sent a document by email, malware could be attached to the file. The recipient couldn't possibly be suspicious; the email and attachment were exactly what he expected to receive from his colleague. He opened the document. The malware installed itself in the background. The cycle began again. It was an entire network of surveillance, dubbed Ghostnet by the security team.¹⁰

Ghostnet has lessons for all of us. You may be sure you wouldn't fall for the Spanish lottery, and perhaps not even for a Facebook call for help, but it's hard to find any comfort in this story.

Do you rely on standard commercial antivirus software to scan attachments? Do you open documents sent by people you've met online? How about documents from prospective customers or clients? Or old friends you recently connected with online? Do you open mail and documents sent to you by coworkers?

Of course, you do. So do I. And that means that most of us are no more able to defend ourselves from this attack than the Dalai Lama was.

If there were any doubts about the scope of such attacks, they were eliminated by what the security team did next.

They took another look at the IP address of the hacker's control server, and asked a simple question.

"Do you think hackers who need a graphic interface to steal secrets are really good at locking down their own computers?" I imagine the Canadian team sharing a mischievous smile as they asked.

Perhaps a veil should be drawn over exactly what they did next. Hacking is illegal in most jurisdictions, even if you're hacking someone who has just hacked you. Using methods they decline to specify, the security team was able to verify that whoever attacked the Dalai Lama's network was indeed much better at breaking into other people's computers than at keeping intruders out of their own.

Finding themselves inside the hackers' control servers, the security team naturally had a look around. They watched as reports came in from the Dalai Lama's computers. But that's not all. Reports were coming in from other computers as well. Hundreds of them.

The hackers who compromised the Dalai Lama's network were collecting data from nearly thirteen hundred other computers. Who else had been targeted by the attackers? That wasn't hard to find out. All the security team had to do was to ask who owned the IP addresses of the compromised computers.

What they found was a Who's Who of Asian organizations that ought to be highly concerned about—and pretty good at—computer security: Indian embassies in the United States, Germany, and the United Kingdom; the foreign ministries of Iran, Indonesia, and the Philippines; the prime minister's office in Laos. All were in thrall to the attackers' servers. Computers in sensitive businesses, from the Asia Development Bank to Vietnam's petroleum company, were also sending the attackers their data.

And, even though this set of attacks does not seem to have been aimed at the United States, Ghostnet was collecting reports from computers that belonged to the Associated Press and the auditing firm of Deloitte & Touche. Oh, and NATO too. In early 2010, Google announced that it and many human rights campaigners using Gmail had been targeted with the same attacks.

No one was safe.

The security team split on the question of whether to assign responsibility for Ghostnet to China. Some said it must be the Chinese government. Others were willing to let the facts speak for themselves. The Chinese government denies everything.

But there's not much comfort for us in the denials. The attacks happened, and they worked. If a government wasn't responsible, then this kind of capability is already in the hands of organized crime. Indeed, with its script-spy graphic interface and unsecured control servers, the whole episode underlines a troubling fact. Thanks to exponential empowerment, today's hackers don't even have to be very good. Empowered by democratizing technology, they can still beat our best defenses.

In fact, something similar to Ghostnet is already being used by organized crime. Most businesses depend on bank clearinghouse accounts or electronic fund transfers to pay their bills. They log on to bank sites using passwords; for larger amounts they may also be asked a set of "challenge questions" seeking information only the businesses know. But corporate officials also open email attachments from business contacts, and once attackers have access to the officials' keystrokes, neither the password nor the challenge questions offer any security. Hackers have stolen more than \$100 million from U.S. businesses using this technique, the FBI reported in October 2009.¹¹

I wasn't in government in 1998 or 2003, when the Clinton and Bush administrations called for new computer security measures. I didn't get the classified briefings that galvanized both presidents. Now I figure I don't have to.

This is scary enough.

But maybe you're not ready to agree. Maybe you're worried that these security alarms are a little too convenient—perhaps just an excuse for the government to spy on Americans and interfere with the economic engine of Silicon Valley. Surely, you think, there are still a few good defenses left.

Well, let's take a look at some of the top reasons that people think computer security risks can be managed successfully.

It's a Microsoft Problem. I know plenty of people who still believe that Microsoft's products are uniquely insecure, and that all we need to do is get Microsoft to clean up its act or take our business elsewhere.

For some, the security of Linux was an article of faith; its source code is open to inspection by anyone, so it is protected from exploit by all those watching eyes. And Apple, which didn't even offer an antivirus program for decades, was protected by, well, by Steve Jobs's sheer animal magnetism.

The last few years have been hard on those illusions. As Apple gained market share, malware authors began writing for its operating system, and they didn't have any trouble finding holes. It turns out that, according to a 2009 talk at the Black Hat security conference, even Apple's keyboards can be hacked to reveal all the user's keystrokes.¹² Apple now recommends that its users run multiple antivirus programs.¹³

And all those eyes on Linux's code? In August of 2009, two Google researchers discovered a bug in the central core of Linux; it would allow an attacker to acquire complete administrative control of any machine to which he had physical access.¹⁴ You might call that a success for open source, except that the bug had been hiding in plain sight for at least eight years.

Why, then, is there so much more malware running on Windows than on Linux? Almost certainly for the same reason that there are more applications of every sort running on Windows than Linux. Like other application developers, malware authors want to reach the largest number of users with one piece of code. And the way to do that is to write your application for Windows.

It's a Password Problem. I used to take a lot of comfort from the fact that I didn't use just passwords for the things I most wanted to keep secure. I used a token. Every thirty seconds it displayed a different security code, known only to me and my home server. Even if a hacker could compromise my machine and record all my keystrokes, he couldn't know what the token was going to say next.

But this is the age of Twitter—and real-time hacking. For at least the last couple of years, criminals have been able to beat these token systems. Now, when the owner of a compromised machine starts typing in his temporary code, the malware phones home immediately. As

the owner types, each digit is sent to the hacker, who simply logs in with him.¹⁵

Really Important Transactions Can Be Confirmed Offline. If you're really worried, you may have locked down your financial accounts, so no money can leave the institution without a call to verify the transaction. In fact, even if you haven't locked everything down, you may get a call. Like the credit card companies, mutual funds and financial institutions have stopped trusting their customers' computers. For risky transactions, they insist on offline, or out-of-band, confirmation.

Out-of-band communication is today's most common fail-safe solution for computer compromises. To restore control of his Facebook account, for example, Bryan Rutberg had to send Facebook a separate, out-of-band message from a separate account.

But using another line of communication won't solve the problem for long. Hackers have already begun to build blocking programs into their malware. The programs prevent users from getting to Web sites that might detect and cure their infections. In the future, these programs may be able to thwart other efforts to cure an attack—diverting emails, for example, or corrupting the user's attempts to log on to hijacked sites.

The banks' offline solution is also at risk. Finding a truly offline method of communication is going to get harder. Businesses and consumers are switching in large numbers to "voice over IP," or VoIP, telephony. They cannot resist the allure of bringing to voice communications the cheap, flexible features of Internet communications. They cannot resist going just a little faster on the bike.

But the switch means that they are also bringing to voice communications all the insecurity that plagues other Internet communications. This raises the prospect of a whole new set of attacks, from "voice spam" and fraudulent telephone calls to the theft of incoming and outgoing phone calls. If an attacker who has compromised your computer's online bank account is also able to appropriate your Internet telephone, then it will be easy for the attacker to answer the phone when the bank calls—and to confirm that you really do want

to transfer your life savings to Spain or Nigeria. At that point, it will be cold comfort that switching to VoIP cut your monthly phone bill from \$40 to \$10 or even to \$0.

The Military Has Solved the Problem With Classified Networks. The government used to have its own illusions about security. Maybe our unclassified networks are compromised, Defense Department officials used to say, but the *classified* networks are still bombproof. They can't be compromised by all this malware floating around the Internet. Because they aren't connected to the Internet. There's an "air gap" between the two.

That assumes, of course, that network security decrees are perfectly enforced—and that the most important secrets are only discussed on classified networks—notions that contradict everything we know about human nature.

But never mind, because the air gap illusion, too, has fallen prey to the exponential empowerment of hackers that we've seen in recent years.

The French navy's Rafale Marine jets train out of Villacoublay air base, in the southwest suburbs of Paris. These fighters are state of the art, packed with stealth and electronic warfare capabilities and capable of landing on carriers. But to do that, they first have to take off. And for two days in January, the jets couldn't take off. They'd been grounded by a hacker.¹⁶

The "Conficker" computer worm had been exploiting vulnerabilities in Windows servers for months. It was the most ambitious computer infection in years. At the time it had infiltrated as many as 15 million machines around the world. One of the ways it spreads is by infecting the USB thumb drives that carry data from one machine to the next. Even classified or isolated networks could be captured if a bad thumb drive was used to transfer data to a machine on a secured network.

That's what grounded the French fighters. Before the navy even knew it was under attack, the worm was coursing through its internal network. Rushing to contain the damage, the navy told its staff not to

turn on their machines, and its systems administrators began quarantining parts of the network. Too late for Villacoubly. Its systems were already hosed.

The Rafale fighter downloads its flight plans, a far more efficient process than paper-based systems. But once the contagion had spread to Villacoubly no flight plans could be downloaded. Until an alternative method of delivering the flight plans could be cobbled together, the Rafales were no more useful than scrap iron.

The French press reported the embarrassment in detail. Perhaps as consolation, it was careful to note that things could have been worse—and were, in Great Britain. There, the press said, twenty-four Royal Air Force bases and three-quarters of the Royal Navy Fleet had succumbed to Conficker.

The British and French navies may have been unintended victims of a worm designed for criminal ends. But after Conficker, no one can believe that an air gap is a security fail-safe.

They're Not Looking for Me. The last of the illusions, or at least the last of mine, is that I'm just not that interesting. Other people have more money. Other people have more valuable secrets. Who's going to come looking for me?

That's the last hope of every herd animal. The predators can't eat everyone. If you lie low and blend in, they won't pick you.

Wrong on two counts, I'm afraid. First, take this test. Add up your savings, car value, house equity, and investments. Is the total over \$65,000? If so, you've got a lot of company on the globe. Probably 10 percent of the world's 6.8 billion people have assets exceeding that amount—say 700 million in all. Being one in 700 million sounds like pretty good herd-animal odds until you realize that, for every person with more than \$65,000, there are nine people with less.

As computers become exponentially cheaper, most of those nine people will be able to get online. Then there will be nine people looking for ways to take money from you. And another nine for your spouse, nine for your neighbor, and nine for each of your business partners. Maybe nine each for every person you know.

So they *can* eat everyone.

There are already Nigerian hip-hop anthems and videos celebrating the rolling-in-money “Yahoozees” who fleece Americans like Howard Crank. The world is already full of scam artists willing to work for less than minimum wage. Most of them know English and have access to the Internet.

The relentless march of empowerment will soon give the Yahoozees of the Third World new tools for finding you. In a way, that’s what a Spanish lottery email does. Most of us delete lottery spam. But if one in ten thousand responds, even with great caution, that person has selected himself for fleecing, and the pitch can then be tailored precisely to his failings. So what if that part of the scam is a bit labor intensive? There are nine people with nothing better to do than sit around trying to get into the mark’s head.

Remember that real-time password-stealing program? Well, the thieves don’t have to go looking for rich people to infect. Instead, they infect everyone, and let the malware find the rich ones. The password-stealing program consumes an infinitesimal part of a modern chip’s processing power to run quietly in the background, watching and waiting until its victim logs on to one of about fifteen hundred predetermined financial sites. Anyone logging in to one of those sites, the authors figure, probably has enough money to be worth cleaning out. So when an infected computer sets itself apart from the crowd by logging on to a financial site, the malware alerts its author, who can now focus on taking money from that computer’s owner.

Moore’s Law has taken a lot of the work out of the hunt. And, thanks to the empowerment of information technology, it will keep making the job exponentially easier, year in and year out.

Until the predators find you, too.

You might think that’s the worst of it.

But it’s not, quite. It’s not just that you could lose your life savings. Your country could lose its next war. And not just the way we’re used to losing—where we get tired of being unpopular in some

Third-World country and go home. I mean *losing* losing: attacked at home and forced to give up cherished principles or loyal allies to save ourselves.

Plenty of countries are enthusiastic about using hackers' tools as weapons of war. At the start of a 2008 shooting war between Georgia and Russia over South Ossetia, for example, numerous Georgian websites were swamped by "denial of service" attacks. Security researchers found evidence that the attacks were coordinated and organized by Russian intelligence agencies. The year before, Estonian government agencies and banks were also crippled by denial of service attacks after the Estonian government moved a World War II memorial that had become a symbol of Soviet colonial rule. Estonia's foreign minister charged that the Russian government was behind the attacks. Russia denied the allegation. NATO, and European investigators were unable to refute their denial.

China has also been accused publicly of audacious computer attacks. German Chancellor Angela Merkel discovered that her office computers had been compromised in an attack blamed on the People's Liberation Army. India, France, and Taiwan have also suffered intrusions and attacks attributed to China. The compromise of the Dalai Lama's network was also widely blamed on China, as were a series of serious attacks on Google and other large U.S. companies in 2010. Like Russia, China has consistently denied all charges.

As I said before, in a strategic sense, the denials don't really matter. If the attacks weren't carried out by Russian and Chinese government agencies, that just means that there are more organizations and countries with effective cyberintelligence and cyberwarfare capabilities than we thought. And, in fact, five or ten years from now, there will be. That's because cyberattacks don't require heavy capital investments, the way nuclear weapons or stealth fighter jets do. Any nation willing to put ten of its best computer experts to work on a cyberintelligence program could probably have one in a year or two. (The Conficker worm that brought down British and French military systems could easily have been written by a single well-trained person.) Many

cyberattacks are simply a matter of individual effort. Put enough smart people on enough targets, and some of them will get through.

And that's why attacks on computer networks pose such a strategic threat to the United States in particular. We are an important intelligence target for practically every nation on earth. And attacking our networks is nearly risk-free; the list of suspects is about as long as the UN membership roster. In fact, there are incentives for them to help each other break into our networks. ("I've seized control of an email server at USDA, but what I really want is USTR's (Office of the U.S. Trade Representative). Want to trade? I could throw in the Commerce secretary's password to balance the deal.")

If you're a foreign government, breaking into U.S. networks is a twofer. You can start by stealing secrets. But if push comes to shove, you can use your access to destroy the same systems you've been exploiting. Corrupt the backup files, then bring the whole system down. Or start randomly changing data and emails until no one can trust anything in the system.

It wouldn't take much to create chaos. The financial crisis of 2008 became a panic when bankers began to disbelieve each other. No one trusted the other guy's books, so they stopped lending, and the world crashed. Could that same mistrust be created by modifying or destroying a few firms' computer accounting and trading records? We probably don't want to find out.

It's no secret how to fight a war against the United States. Slow us down, then cause us pain at home and wait for antiwar sentiment to grow. Cyberattacks are ideal for that strategy. Everything in the country, from flight plans and phone calls to pipelines and traffic lights, is controlled by networks susceptible to attack. A determined, state-sponsored attacker could bring them all down—and blame it on some hacker liberation front so we wouldn't even know whom to bomb.

The Pentagon has heard fifty years of warnings about not fighting land wars in Asia, where hand-to-hand fighting and sheer numbers can overwhelm an American army's technological edge. But now it turns out we've opened an electronic bridge, not just to Asia but to

the rest of the world, and now we're trying to defend ourselves hand to hand against all comers. It's hard to see how that ends well.

So that's the nub of the problem. No law of nature says that the good guys will win in the end, or even that the benefits of a new technology will always outweigh the harm it causes.

The exponential growth of information technology has made the Pentagon far more efficient at fighting wars; it has made our economy far more productive.

So far, it's been very good to us as a nation.

But it was good to Howard Crank, too, for a while.

Somehow, this problem too ended up on DHS's plate: We were supposed to figure out what could be done to improve the country's network security.

It was a snake-bitten assignment. Two presidential cybersecurity strategies—one devised by the Clinton administration and one by the Bush Administration—had already run into the ground before DHS was created.

Perhaps those who created DHS hoped that it could succeed where two presidents had failed. In any event, they gave the new department responsibility for civilian cybersecurity. The National Communications System, which ensures the availability of telecommunications in the event of an emergency, was transferred from Defense. The FBI gave up its National Infrastructure Protection Center, which focused on cybersecurity (and promptly recreated the capability under another name so that it could keep fighting for the turf). The Federal Computer Incident Response Center, which handled computer incident response for civilian agencies, came over from the General Services Agency.

These offices fit well with other DHS missions. Two of its big components—the Secret Service and Immigration and Customs Enforcement—have cybercrime units. And DHS was supposed to protect from physical attack the critical infrastructure on which the economy depends.

In carrying out these duties, DHS could get technical help from the National Security Agency, which was in charge of protecting military and classified networks. But the responsibility for civilian

cybersecurity obligations left DHS on the hot seat. If we couldn't find a way to head off disaster, no one else in government would.

To be candid, for the first few years of the department's existence, we didn't accomplish much on this front. There were lots of reasons for that. Fixing travel and border security was more urgent. Staff turnover was high and expertise thin in our cybersecurity offices. But the real reason we didn't get far was that the same forces arrayed against change in the travel arena were lined up against change in information technology.

Businesses had staked their futures on continued exponential growth in information technology. They didn't want policy changes that might change the slope of that curve even a little. Privacy groups instinctively opposed anything that would give the government more information about, well, about anything. And even when it was supportive, the international community was so slow to change direction that it posed an obstacle to any policy that was less than twenty years old.

It didn't matter how obviously necessary a security measure was. Resistance to any change was strong. A case in point was the effort to install intrusion monitoring on the federal government's own networks.

To succeed, most cyberattacks must do two things. The hackers first have to get malicious code into the network they've targeted. Then they have to get stolen information out. If we can detect either step, we can thwart the attack. So one way to defend our networks is to do a thorough job of monitoring traffic as it goes in and out.

We've known this for a decade. The Clinton administration's cybersecurity strategy, drafted in 1999 and released in early 2000, called for a network of intrusion detection monitors that could inspect packets going into and out of all federal government networks. President Clinton requested funds for intrusion monitoring in his outgoing budget. But civil libertarians quickly launched a campaign against it.

It was an odd battle for them to choose. The point of the monitoring network was to inspect government communications. Even the

most extreme privacy zealot shouldn't be shocked to discover that the government was reading its *own* mail, much less that it was inspecting its mail for malware. By then, government agencies were already screening emails for spam; the intrusion detection network simply extended that concept to other unwanted packets. What's more, since roughly the 1980s, these computers had been displaying warnings to users that government systems are subject to monitoring.

But privacy groups were spoiling for a fight. They portrayed the proposal as the second coming of Big Brother.

"I think this is a very frightening proposal," an ACLU representative told ZDNet News.¹

"We feel the government should spend its resources closing the security holes that exist, rather than to watch people trying to break in," said a counsel for the Center for Democracy and Technology.²

"I think the threats (of network vulnerability) are completely overblown," said the general counsel for the Electronic Privacy Information Center, adding that claims of a security threat is leading to "'a Cold War mentality' that threatens ordinary citizens' privacy."³

In the end, civil liberties resistance was so strong that only the Defense Department was allowed to build an intrusion detection network. For years thereafter, the civilian agencies experienced intrusions that could have been prevented by the intrusion prevention system proposed by President Clinton. But once burned was twice shy. The privacy groups had thoroughly tainted the idea of intrusion prevention on the Hill, and there was real reluctance to revisit the issue. When the Bush administration wrote its cybersecurity strategy, it did not even try to revive the idea.

Finally, though, five years later, the Bush administration decided to force the issue. Mike McConnell, the director of National Intelligence, had been my boss at NSA, and he had spent the years after leaving NSA building a cybersecurity practice at a large consulting firm. A quiet, self-deprecating Southerner with a talent for briefing higher-ups, McConnell was determined to move cybersecurity to the front burner.

He didn't have to work too hard to persuade DHS to take on the challenge. We were alarmed at the ease with which attacks were being launched against civilian agencies. With the backing of President Bush and Mike McConnell, we again proposed an intrusion detection network for civilian agencies. And civil libertarians once again renewed the fight to stop us—as though nothing had changed in ten years. Without the slightest evidence of irony, they again raised privacy objections to the government monitoring its own communications.

We got further than President Clinton did, but not much. Congress appropriated funds for the project, but it had not been fully implemented when Barack Obama was elected president. Spooked by the privacy outcry, the Obama administration postponed full implementation of intrusion monitoring so that it could again examine all of the privacy issues. Pilot projects are underway, but final decisions about how, when, and whether to implement effective intrusion monitoring are still awaiting consensus among the lawyers.

Meanwhile, attacks similar to those that compromised the Dalai Lama's network are continuing. The privacy debate had caused ten years of delay, and it may yet kill an effective intrusion prevention system.

It's remarkable when you think about it. Right now, this minute, agents of an authoritarian government are covertly turning on cameras and microphones in homes and offices all across America, spying on the unsuspecting and the innocent. They're recording our every thought, our every keystroke, as we prepare private documents or visit websites.

And they're able to do that today thanks to the hard work of *privacy* advocates.

How did the privacy community end up facilitating surveillance and espionage on an unprecedented scale? History, mainly, and a lack of imagination.

The men and women who built the computer industry grew up in a very different era from those who pioneered the air travel industry. Air travel enthusiasts first launched commercial flights between the two world wars, when government was big and military risks were

on everyone's mind. The pioneers were children of their age. They foresaw a world in which air travel was used for military and espionage purposes; they understood that unregulated flights could lead to disaster as the skies filled up. To manage those risks, they helped the government fashion a comprehensive regulatory scheme for pilots, airlines, and airplanes.

Computer technology, in contrast, was born in the wake of World War II, at a time when the challenge of totalitarianism was on everyone's mind. The men and women who built the earliest computers were children of a different era. They most feared that their machines would be misused by authoritarian governments. Unlike an earlier generation of technologists, they struggled to limit government's role in their industry. And they succeeded. From electronic intercepts to information processing practices, for the next forty years, laws on information technology were aimed as much at regulating the government as at regulating the industry.

By the time the threat of widespread computer misuse finally arrived, the privacy groups already had a narrative fixed in their mind. They could not imagine any threat to computer users' privacy that could be worse than the one they saw in the United States government. Saying no to the government was their default position.

By the end of the Bush administration, DHS was used to the idea that even the most obvious security measures would be opposed by privacy groups. We still had an obligation to do what we could to head off the building security risks. We also knew intrusion prevention, valuable as it was, wouldn't do that by itself.

We needed a broader strategy. In mid-2008, the Homeland Security Council asked DHS to provide options for a set of long-term strategy questions. The policy office was assigned to pull them together.

We found a lot of tough tactical questions that needed to be answered, but the real problem was our strategic posture. And only two ideas that offered any hope of curing our strategic vulnerabilities—*attribution and regulation.*

Attribution

Here's our strategic security problem in a nutshell: We are attacked every day by an imaginative, highly motivated, and anonymous adversary. We can prevail only if we mount near-perfect defenses. And, since there's no penalty for mounting an attack, the adversary simply tries again and again until something works.

This defensive strategy is, quite simply, too hard. A wholly passive strategy almost never works in the real world.

Take burglary. We certainly spend money on defense. A good lock on your door can keep burglars out of your home. But the lock isn't all that good by itself. We take it for granted that burglars can't sit on our doorsteps day after day, studying our lock and trying new lock picks every evening to see what works. If they could, they'd find a way in sooner or later.

Burglars don't sit on your doorstep because they're afraid of being busted. It's the threat of the police that makes your lock as effective as it is.

Defending networks is the same kind of problem. Security measures are all well and good, but unless we can also identify and deter attackers, defense alone will never do the job.

We have a lot of ways to punish attackers once we identify them. It's identifying them that's hard.

We began by trying to use the tools of law enforcement to identify the attackers. Practically all computer attacks are crimes, after all. They usually violate fraud, extortion, and computer abuse laws. Many attacks would be deterred if the perpetrators faced a realistic risk of arrest and prosecution.

But crossing international boundaries on the Internet is easy. Attackers discovered very early that they could cover their tracks by breaking into lightly guarded computers in several countries and hopping from one to the next before launching an attack on their real target. That way, the police would have to track them back from

country to country before discovering their real location. And doing that would require subpoenas valid in each country.

That wasn't easy. To get one country to enforce another country's subpoena requires patience and lengthy legal analysis. The country that's being asked to enforce the subpoena will only do so if it too views computer attacks as crimes. It has to have the ability to carry out the search very quickly. Otherwise the logs will be overwritten and the evidence gone. Indeed, unless the information can be gathered nearly instantaneously, the attackers will always have the advantage. They can compromise new machines and add new hops to their route faster than the police can serve subpoenas to track them.

This problem has been obvious for more than two decades. The United States began encountering it in the 1980s and, by 1989, it had persuaded the Council of Europe to propose work on an international cybercrime convention to streamline the identification process. Getting that far took great effort. The Justice Department had to explain over and over to less computer-savvy governments why it needed such an agreement.

Not until late 2001 was there actual agreement in principle on a few very basic steps—making computer hacking a crime and naming a contact point to handle subpoena requests quickly. And that simply marked the start of a long, slow, international law-making process. The cybercrime convention didn't come into effect until 2004, when a grand total of three countries ratified it. As of 2009, fifteen countries had fully ratified and acceded to the convention, and twenty-eight more were in various stages of adopting it. As international efforts go, that is a considerable success (although the numbers are inflated by the European Union, which has pressed its twenty-seven members to join, along with EU satellites like Liechtenstein).

And what does the convention do to solve the attribution problem? In essence, the members of the convention have agreed that they will adopt a common set of computer crimes and that they will assist each other in investigating these crimes.

That's it. A good thing, no doubt, but hardly likely to stop the massive attacks we see today. Hackers have compromised hundreds of thousands, sometimes millions, of machines. If they chose to hop from one of those to the next before launching an attack, the authorities would need to serve hundreds of thousands of subpoenas in dozens of countries—and to do it as fast as the hackers could move from one machine to the next. The hackers can move at the speed of light—literally. The governments can move at the speed of paper, courts, and sealing wax. It's no contest.

At best, the convention offers a partial solution to computer crime as it existed in the 1980s. But building a consensus for even its limited measures took more than a decade. And even then, the consensus was distinctly limited in geographic reach. Neither Russia nor China has shown any inclination to adopt the convention. Nor, for that matter, have thoroughly wired countries like South Korea, Brazil, Nigeria, Singapore, and Australia. So even if we still lived in the 1980s, there would still be plenty of places in the world for hackers to hide.

The only alternative to the convention that the international community has found is worse—and in thoroughly predictable ways. Led by Russia, the United Nations has recently been touting the idea of “disarmament talks” for cyberspace.

There are several possible motivations for such a proposal. One possibility is that the Russians genuinely believe that an arms control treaty for cyberspace would be good for all concerned, demilitarizing and taking the fear of disaster out of the networks on which the world relies. Unfortunately, that's not particularly likely. You can't have a real arms control agreement unless you can verify compliance. But as we've seen, a principal feature of computer attacks is the difficulty of attribution. If attacks continued after “disarmament” how would we know that anyone had disarmed?

The Russians' models seem to be the multilateral chemical and biological weapons conventions that were negotiated in Geneva during the Cold War. By the usual standards of the international community these are wildly successful agreements, adopted by more than

150 countries. They proved wildly successful from the Soviet point of view as well, since the United States actually abandoned its chemical and biological weapons after signing the conventions while the Soviets kept theirs in place. Even more remarkably, the United States managed to get a black eye in the process, because it had the temerity in 2001 to tell the international community that the convention was unverifiable, that it could not prevent proliferation of biological weapons, and that there was no point in establishing intrusive inspection regimes that would not work.

From the Russian point of view, replaying this drama has no downside. If an agreement is reached, the United States, with its hyper-compliant legal culture now fully integrated into military planning, will undoubtedly adhere to any ban the new agreement imposes. But countries that want to use the tools of cyberwarfare will be free to do so, relying on the anonymity that cloaks attackers today. If the United States sees that trap and refuses to accept an unenforceable agreement, the international community will replay the drama that accompanied the U.S. refusal to negotiate an unenforceable biological weapons protocol.

Just agreeing to consider the proposal, as the new administration seems to have done, allows Russia to divide us from our allies in Europe—who always seem eager to put new international legal limits on warfare, even if the limits can't actually be enforced.

In the end, then, our inability to solve the problem of attribution and anonymity poses severe threats not just to our pocketbooks but also to our national security and our international standing. We thought that it was foolish to solve the problem with what Harvard law professor Larry Lessig once called “East Coast code”—laws and treaties. Instead, we thought, the answer would prove to be “West Coast code”—software and hardware design. In the long run, we needed an architecture that automatically and reliably identifies every machine and person in the network.

I knew that privacy groups would melt down if anyone proposed to do that for the Internet. Anonymity has become (wrongly in my

view) equated with online privacy. Any effort to cut back online anonymity will be resisted strongly by privacy groups. And they'd be able to find popular support, at least for a time. Practically everyone does something online that he's ashamed of.

At the same time, practically everyone spends large parts of the day on a network where his every action *is* identified and monitored. Most corporate networks have robust attribution and audit capabilities, and the insecurity of the public networks is forcing private networks to study the conduct of their users ever more closely in the hopes of identifying compromised machines before they can cause damage.

In trying to chart a broad network security strategy, I thought we needed more research and incentives to improve audit and attribution capabilities in hardware and software. And we needed architectural and legal innovations to encourage one secure and attributable network to link up securely with another. In the long run, and perhaps in the short run, that sort of organic linking among attributable systems may be the only way to build a network on which identification is rapid and sure.

That doesn't mean the old, anonymous Internet has to disappear. But I suspect we'll have to create a new network that coexists alongside the old one. Users who value security—who want an assurance that their financial assets and their secrets will not be stolen by hackers—will choose the secure alternative, at least most of the time.

The policy office at DHS put that idea forward as an option for consideration by the Homeland Security Council.

Regulation

Cybersecurity regulation had been talked about for years. The Bush administration floated the possibility in 2002. Or, to be more precise, Richard Clarke floated the idea.

Clarke was a flamboyant bureaucratic warrior camouflaged by the dress and haircut of a high school math teacher. A career official with a knack for building empires—and making enemies—he had risen to

take charge of both cybersecurity and terrorism policy in President Clinton's National Security Council. He later became famous briefly for his scathing denunciation of the Bush White House's response to terrorism warnings. But in 2000 he was better known as the man who had sponsored the failed Clinton administration plan to build a monitoring network.

Clarke was held over by the Bush administration, with the same two portfolios he had held under President Clinton—terrorism and cybersecurity. But he never seemed to gain the same support in the new administration as he had in the old one. After the attacks of 9/11, pushed out of the terrorism job, he poured himself into his cybersecurity role, spending much of 2002 drafting a strategy for the new administration.

Always a hard-charger, Clarke had high ambitions for his new effort. He planned a grand event to unveil the strategy in September of 2002. Reportedly, the strategy sidled up toward new mandates for industry, calling on technology companies to contribute to a security research fund and pressing Internet service providers to bundle firewalls and other security technology with their services. But just days before the event, Clarke's wings were publicly clipped. His long and elaborate strategy, with its nods toward imposing regulatory requirements, was rapidly and harshly cut down. Anything that could offend industry, anything that hinted at government mandates, was stripped out. It was finally unveiled, not as a final document, but as a simple draft for further comment.

For Clarke it must have been the final straw. He'd already been pulled off the terrorism account with brutal swiftness after 9/11, and now his year of effort on cybersecurity had ended in a public rejection of his work.

He stayed in the White House just long enough to produce a final strategy document that was as tepid as the draft. Then he quit.

Industry had claimed another scalp in its long campaign to head off federal mandates aimed at improving computer security. The president (though not industry) eventually paid a heavy price for Clarke's

resentment. The one-time security adviser became a harsh Bush critic, in testimony before the 9/11 Commission and in his other writings.

I thought of Clarke's fate as we put together the report for the Homeland Security Committee. Regulation had become an electrified third rail. Especially in a generally business-friendly administration, advocating more regulation was not likely to be career enhancing.

But the status quo clearly wasn't working. Moore's law was working against us. We had to find a way to change incentives, to get information technologists to start building security into the foundation of our networks. It's not that I thought regulation was always going to be the right answer. But I was sure that it had to be on the table. Especially because regulation didn't have to mean classic command-and-control Federal Register rulemaking.

Government doesn't have to issue mandatory rules to influence private sector behavior. It can use a variety of incentives to encourage security. So the policy office laid out a range of approaches, ranging from soft to hard.

Soft regulation

The softest option was to nudge industry toward security measures by offering liability protection in exchange. This is the most comfortable form of regulation for business, because instead of punishing bad behavior it rewards good behavior. This is something we understood at DHS, where we administered the Safety Act⁴. That act provides liability protection to companies that manufacture and sell qualified antiterrorism technology.

The idea behind the act is simple. Some anti-terrorism technologies work well but not perfectly; they reduce risk but don't eliminate it. Unfortunately, after a terrorist incident, the people who have been fully protected by the technology will be grateful, and the people who haven't been fully protected will sue, claiming that the technology was defective, since it didn't protect them from all harm. That's not a recipe for encouraging the deployment of new technology.

So, to keep fear of liability from squelching advances in technology, the Safety Act sets a cap on liability for approved technologies. There are a lot of conditions built into the act. Companies must, for example, carry whatever level of liability insurance DHS considers necessary to compensate people who may be harmed in a terrorist attack. But in return, the threat of open-ended, company-killing liability is taken off the table.

We thought that DHS could use the Safety Act itself to encourage companies to adopt some cybersecurity technologies. The protections of the act aren't limited to physical products; they also cover services and information technology. We thought the act could even be applied to security services and processes, vulnerability assessments, and cybersecurity standards.

But the Safety Act wasn't perfectly adapted to cybersecurity tools. Most hackers are not terrorists. In addition, network security measures work in layers. There is no single magic bullet that provides all security needs. If many security products fail to prevent an attack, and not all of them are covered by the act, sorting out which ones caused the damage could require endless, expensive lawsuits. And, because network threats change so often, products designated under the act would have to be updated frequently. Even with regular updates, the extent to which a particular technology provides protection will likely erode over time as attackers seek ways around the defense. At what point should protection be modified or withdrawn, we wondered, and who will press for that change? Finally, the insurance market for cybersecurity products remains at best a work in progress, so it wasn't clear that adequate coverage was available. For these reasons, we concluded, the Safety Act was probably better as a model of what could be done without regulation than as a tool that could be used immediately to encourage broad cybersecurity measures.

We also noted a second "soft" way to influence business—government purchasing standards. Many critical infrastructure companies do business with the U.S. government. The government has great weight as a buyer of technologies, and it can influence the market for security

by the standards it sets for its purchases. The government cannot, however, dictate terms to suppliers of technology. The government may be the single largest buyer of some technology, but it is far outweighed in the aggregate by private sector purchasers. Further, without new policies, the government wouldn't really act as a "single" buyer. IT procurement is divided among many agencies, and these agencies would fight security standards that raise costs or reduce competition.

We wanted the government to consider a more unified approach to its procurement of information technologies. We thought the government could establish government-wide contract models that incorporated preferred technologies and security practices requirements into federal contracts. In fact, some steps on this road had already been taken. Federal purchases are required by law to meet certain federal information security standards.

We knew, though, that using procurement to enhance commercial IT security is easier said than done. The U.S. government's first efforts to leverage its procurement power for IT security began in the 1970s, when the government established the Trusted Computer Security Evaluation Criteria—the "Orange Book"—and began to evaluate commercial products that were submitted for review. The idea, then as now, was to use federal contracts as an incentive for vendors to incorporate security measures in their products.

The scheme never had as big a security impact as hoped; the commercial market for computers rapidly outpaced the government market, and private purchasers came to perceive their security needs as different from those of the government. Sellers and buyers alike complained that security evaluation slowed adoption of current IT hardware and software.

For all those reasons, the procurement process has not so far turned out to be an effective way to influence network security.

Hard regulation

And what about the "hard" option—just plain regulating? You know, just putting network security requirements into the Federal Register?

We couldn't ignore that option, I thought. In fact, a lot of the most critical industries were already subject to government regulation. These included financial institutions, energy, and telecommunications. And some of these industries were already subject to cybersecurity regulation. Financial institutions, for example, must follow a unified set of cybersecurity rules. But even financial regulators don't require particular security measures. The rules are largely procedural, resembling the instructions on a bottle of shampoo: Institutions must study their vulnerabilities, cure them, assess the effectiveness of the cure, and repeat.

It's hard to write rules that go beyond such procedural steps, because the attackers change tactics faster than regulations can be amended. What's more, the cost of mandatory security would be very high; it would slow innovation and productivity growth severely.

Even so, there's a case for mandating particular security measures for regulated industries. It's the Howard Crank problem all over again. Every year, the exponential growth of information technology makes our lives a little better, our businesses a little more efficient and profitable. And every year it makes us a little more vulnerable to a military strike on our infrastructure that could leave us without power, money, petroleum, or communications for months.

Large parts of the country could find themselves living like post-Katrina New Orleans—but without the National Guard over the horizon. Protecting against that risk isn't part of most companies' balance sheets. It's not hard to see that as the kind of market failure that requires regulation.

But even if there is a market failure, the government still isn't well-equipped to solve it. At a minimum, the regulatory agencies would have to find a way to coordinate and issue standards much faster than they now write regulations. Today, the practical speed limit is eighteen months from new idea to final rule. There's not much point in replacing a predictable market failure with an equally predictable government failure.

And what about all the vulnerable IT networks that are not in the hands of regulated industries? If they are compromised, the harm

goes beyond the users of those networks. The compromised machines can be used to attack others, including government systems. To set standards in that world would certainly require new legislation.

Industry, we knew, wouldn't like any talk about regulation. But they were fighting the last war. New security legislation had in fact already been enacted, though in an odd, and mostly unfortunate, way. Laws have been adopted in all but five states that require companies to disclose any security breaches that lead to the disclosure of sensitive customer data. The more the federal government has dithered over security rules for industry, the more aggressively the states have moved into the opening. Their breach notification laws are becoming *de facto* security regulations for all companies. First, they punish bad security by forcing companies that are compromised to admit that fact, as long as some personal data was accessed. Second, in a crude way, they recognize that good security measures can make notification unnecessary, and that encourages companies to invest in technologies that are so recognized. For example, many state laws recognize that encrypted data may be safe even if the system it is stored on has been compromised. So, naturally, many companies have expanded their use of encryption to avoid embarrassing breach notifications.

The problem with these laws is that they don't necessarily point companies in the direction of real security improvements. Because they only punish companies for breaches that disclose personal data, they have encouraged the companies to lock up or discard certain kinds of customer data—rather than focusing on keeping hackers out of systems that control their most critical functions.

The problem is particularly acute in the area of stolen and lost laptops. Thousands of business laptops are lost or stolen every day. Usually, the thief wants the laptop, not the data. But if there is personal data in the laptop, that data has technically been compromised, thus forcing companies to send embarrassing notices to everyone concerned. After a few such cases, companies begin to divert their security budget to double-locking laptop drives with passwords and encryption. Those measures won't keep Ghostnet out of their networks, but

they get the highest investment priority because of the peculiarities of state law.

By the same token, state laws expressly recognizing encryption of data as a defense have artificially heightened the priority that security offices assign to the deployment of encryption, even though it too does little to block a sophisticated attack. There are many measures other than encryption that may be equally effective at providing a defense in depth, but state legislatures have not been able to draft laws that reward more comprehensive security.

Finally, state laws vary substantially, creating great tension for law-abiding companies, which find they cannot actually comply with all of them. For all those reasons, there is growing support for a federal law that would set a single breach disclosure standard. Such a law could also create incentives for higher cybersecurity standards. In fact, replacing inconsistent state notification laws with a security-minded federal law would be a victory for both security and innovation.

The Report

By the time we finished the report, I realized that we hadn't just touched the third rail, we were tap-dancing on it. By candidly treating the end of online anonymity and the adoption of tough security regulation as options, we were goading some of the noisiest oxen in Washington.

Well, what the hell, I thought. Maybe the time was right for a reconsideration of security regulation, especially after the hodge-podge the states were making of the issue.

I was wrong.

Memories of Dick Clark's fate were too fresh, and by mid-2008 the administration was running out of time. I showed a draft of the report to the front office and sent the Homeland Security Council a copy. Not much later I got a call. The council didn't want to even raise regulation as an option in the interagency discussions. They feared that industry and Congress would kill the little progress that had

been made if regulation were even treated as an option. In fact, they wanted to bury the report. Instead of thinking about the future, they'd focus only on tasks that could be done in the waning months of the Bush administration.

This was disappointing but understandable. Chertoff, who'd been a rock in other disputes, was now focused only on fights he could win and changes he could implement in six months or less. And we had reached that point in an administration where accomplishing even the simplest and most obvious tasks had become nearly impossible. Energy was draining out of the Bush team, and what remained was soon focused on a cascading financial crisis that left no time for next year's threats.

I thought that there might be value in letting the Obama administration consider these issues without explaining that it was reviewing options proposed under President Bush. The new administration might have more leeway to consider the attribution and regulation issues with an open mind.

I was wrong about that, too.

The Obama administration brought a flurry of energy and apparent determination to the problem. As well it should have. Barack Obama and John McCain, after all, had been the first presidential candidates whose campaign networks were systematically penetrated and exploited by foreign intelligence-collectors. And candidate Obama had pledged that cybersecurity would be a top national security priority in his administration. Nevertheless, the new administration's resolution seemed to waver within weeks of the inauguration.

The new administration did produce a cybersecurity strategy only a few months into the term, but White House watchers learned a lot from what it said and how it was edited. The draft was reportedly produced on the schedule set by the president—within sixty days of his request. But it didn't go to him on that schedule. Instead, it went through a new set of edits, as office after office protected itself, its prerogatives, or its constituencies by removing controversial passages.

The result was mostly pabulum—pabulum of a sort that would have been familiar to the Clinton and Bush White Houses, of

course, since they too had blinked when faced with hard choices over cybersecurity.

For example, the strategy paper recognized that improving authentication of people and machines is a key to improving cybersecurity. While much of its attention was focused on just making sure that federal networks can properly identify users, it acknowledged as a goal the creation of a “global, trusted eco-system” that could form the basis of a secure network. But it called for that system to be built by working with “international partners” and by building an ecosystem that is seen to protect “privacy rights and civil liberties.” Hard experience tells us that if building a secure network depends on the full support of the international and privacy communities, it will never happen.

Business too was fully protected from the specter of security regulation in the Obama administration’s strategy document, which mentioned regulation just once—to declare that it would be considered only “as a last resort.”

By the time the editing was done, Washington knew that nothing dramatic would come from the cybersecurity initiative—or the new cybersecurity coordinator job the president had announced with fanfare. Indeed, the position remained unfilled for nearly a year, until Howard Schmidt agreed to take the job in late December 2009.

Three presidents in a row had tried to change course and head off the worst consequences of Moore’s law for our national and personal security.

All three had failed.

None had been able to defy the privacy and business lobbies, inside and outside government, that guarded the status quo.

The city of Dubai leaps straight out of the flat sands and flat seas of the Arabian Peninsula. One minute you're driving through scrubless desert, the next you're cruising an elevated freeway past a phalanx of thirty-story skyscrapers, most built in the last ten years. Today, with a mountain of debt, Dubai has the look of last year's boomtown; the newest skyscrapers lack tenants and construction has nearly ceased. But during its heyday from 2005 to 2009, Dubai's ambition seemed as unbounded as the desert that it sprang from. And part of its plan was to become the great transshipment port of the Middle East—just as Singapore is the great entrepôt of the Far East. By 2006, with several bustling modern ports, it had largely succeeded.

That's when it encountered—and transformed—the Committee on Foreign Investment in the United States, first handing DHS its best tool for combating network security threats and, eventually, taking that tool out of the department's hands.

The success of the port of Dubai was due in no small part to Dubai Port World, a company owned by the royal family. DPW, as it was called, was the principal terminal operator in Dubai. Its success there led it to branch out, purchasing terminals in many other ports.

Running a port is a lot like running a small city. The government usually provides police, fire protection, and perhaps utilities, while the terminal operators carry out the main economic activity—storing goods and moving them from ship to land and back. To do that, the terminal operator leases land in a port and then builds a pier for ships, cranes to unload the ships, a parking lot for the cargo to rest, plus

perhaps a small management office. The operator makes its money lifting containers out of ships and holding them for shippers to pick up. The terminal operator is thus a lot like a store owner in a city—economically vital but responsible mainly for his own property.

Operating a terminal was once a local business, just like a store. But globalization has come to the industry, and the top five operators in the world now handle more than a quarter of all trade. None of the biggest operators is an American company; in fact, even in the United States, four out of five terminals are operated by foreign companies.

So it didn't exactly set off alarms when one of these foreign terminal operators decided to buy another foreign terminal operator.

Soon we would wish that it had.

The buyer was DPW. The seller was P&O—the Peninsular and Oriental Steam Navigation Company, a two-hundred-year-old British firm that also had terminal operations in much of the world, including the United States. P&O leased terminals in six U.S. ports, and DPW would be getting those along with the rest of the company.

DPW asked the Committee on Foreign Investment in the United States, or CFIUS, to approve the transaction. Created by executive order in 1975, CFIUS conducts national security reviews of foreign investments in U.S. companies. As long as we are running fiscal and balance of payments deficits, the United States pretty much has to keep selling many of its assets to foreign buyers. But we also have to fence off some companies and sectors for national security reasons. We would not let an adversary—Iran or North Korea, say—purchase a major defense contractor. The opportunities for espionage and sabotage are too tempting. But defense contractors are not the only companies that create opportunities for espionage and sabotage. We would not want, say, major U.S. telephone companies to fall into the hands of countries that might use the companies to spy on Americans.

At bottom, CFIUS was charged with deciding which transactions posed unacceptable national security risks. The committee has broad but vague powers. In essence, any foreign company buying a

U.S. company has the option of notifying CFIUS of the transaction. If CFIUS doesn't do anything within thirty days, then the transaction can go forward. If CFIUS has questions, it can launch an investigation. In theory, the investigation is completed in forty-five days and a recommendation is made to the president, who has fifteen days to decide whether to block the transaction on national security grounds.

That's the theory. When Congress first set rules for CFIUS in 1988, it imagined a fairly quick ninety-day process with a sharp yes-or-no decision at the end. Congress took pains to avoid delaying investments. In addition to the short decision deadlines, Congress allowed companies to skip the CFIUS process completely.

Why do companies go through the CFIUS process if they don't have to? It's simple; they want certainty. If they notify a transaction to CFIUS and get no objection, then the United States can't overturn the deal later on national security grounds (unless the information supplied by the parties was false or misleading). So if the parties to a transaction have even a tiny concern about whether the deal will raise national security objections, it's a good idea to make a CFIUS filing. Most investors want to find out about national security objections early, when the deal can still be unwound. If the concerns arise later, one party or the other may be hurt badly. It's almost impossible to unscramble the eggs once a deal has been finalized, and the effort to do so would put both companies at risk.

When we were at DHS, we estimated that only about 10 percent of large transactions received CFIUS review. The rest didn't raise even modest national security concerns. Of that 10 percent, the vast majority were approved without comment. Only about 10 percent of submitted cases led to further action by the committee, meaning that the committee devotes almost all of its attention to roughly 1 percent of all the investments made by foreigners in U.S. companies.

But the stakes for that 1 percent can be enormous. Congress gave the president authority to block any foreign acquisition of a U.S. company if the committee found credible evidence of a threat to national security.

Whether to seek CFIUS approval for the Dubai Port World transaction must have been a close question. Neither company was based in the United States, after all. But CFIUS still could exercise some authority over the transaction. Six U.S. terminals would be getting a new foreign owner.

At the same time, asking for approval didn't look like a big risk. No one in CFIUS had ever raised a national security concern about the ownership of terminals in U.S. ports. And the banks that back these transactions are notoriously risk-averse; if there were a CFIUS issue, they'd want to know right away, not after the deal was done. So DPW decided to go for the certainty of a committee approval. Since DHS was the recognized expert in port security and one of the toughest security advocates on the committee, DPW consulted us early.

I had just taken over as head of policy, and CFIUS had just been assigned to my office. I had no staff of my own, but I wasn't a stranger to CFIUS. After leaving NSA, I had many clients with CFIUS concerns, and I had negotiated some of the detailed agreements that the Justice and Defense departments insisted upon when foreign companies acquired large interests in U.S. telecommunications companies. I knew how valuable CFIUS could be in protecting security, and I was pleased that DHS had already established itself as a leader on the committee.

While the Defense Department had long worried about foreign investments involving its contractors and its technology, its main concern was military threats to our security. But on 9/11 al Qaeda had used civilian technology to kill more Americans at home than any foreign military attack had ever done. So from the start, DHS focused on ways in which foreign ownership might expose the home front to unconventional attacks. Because national security was not defined narrowly, DHS had no trouble fitting this approach into the statute, and in 2007, Congress ratified DHS's approach by explicitly including homeland security and critical infrastructure protection in the new definition of national security.

DHS's broad view of national security covered a lot of ground. But our top worry was sabotage and espionage in the information technology sector. We knew that there were some governments that routinely asked their companies to help spy on other countries. And any technology that allowed spying could be used for sabotage. Once a hostile nation has compromised a computer, it is up to that nation whether to exploit the computer or shut it down. That was too big a risk to take, DHS argued. Some companies and some countries just couldn't be trusted. They shouldn't be allowed to control U.S. networks.

The federal government doesn't have authority to set cybersecurity standards generally for the private sector. It doesn't even have authority to exclude from U.S. markets companies and products that are likely to be used for espionage. It can prosecute spies and companies that conspire with them, of course, but only after the damage is done, and a successful prosecution depends on compiling proof beyond a reasonable doubt and, often, on extradition or other cooperation from the government that ordered the spying. It's not much of a weapon.

CFIUS, however, offers real authority to protect telecom and IT security, and DHS moved quickly to ensure it was used for that purpose. When a company or country with a questionable reputation filed to acquire a U.S. IT or telecom company, DHS often asked the intelligence community whether either had engaged in espionage against the United States or others. (This practice was eventually institutionalized for all applicants.) Even if the company or country hadn't actually been caught in the act, DHS would assess whether the transaction increased U.S. vulnerability to the kind of cyberattacks we knew were likely in the long run.

A number of transactions did increase U.S. vulnerability. The telecom industry is globalizing at the same time that it is shifting from big, specialized telephone switches to Internet technologies. The IT industry has been globalized for decades, but opportunities to compromise components and complete products continue to grow, particularly as companies diversify their software as well as their hardware supply chains. DHS paid special attention to foreign investment in

computer security products and services. Just as terrorists hoping to assassinate an official are most likely to succeed if they can gain control of the official's security detail, so attackers hoping to compromise a network are most likely to succeed if they can gain control of the network's security system.

The terminal deal that DPW was proposing, though, had nothing in common with the transactions that most threatened U.S. interests. In telecommunications and information technology transactions, we knew there was a risk that foreign buyers might use their new acquisition as a base for espionage or network attacks. But why would Dubai want to sabotage a U.S. port? And even if it did, how would owning a terminal make that more likely?

The terminals that DPW was buying were just plots of land and warehouses inside six U.S. ports. Their security was overseen by the port authorities, the local governments, and the Coast Guard.

I polled the DHS components responsible for ports and found no concerns about the transaction; they all said that the current owner cooperated fully and voluntarily in all our security programs, and they had no reason to think that DPW would act differently. None of the other CFIUS agencies took even a passing interest in the deal.

Even so, there was one more thing we could do. I knew that companies had entered into "mitigation agreements" with CFIUS agencies in the past. In fact, I'd negotiated them. Could we get one here, I wondered?

Mitigation agreements weren't anything that Congress had created. When Senator J. James Exon and Representative James Florio drafted the Exon-Florio Amendment¹, they expected CFIUS to ask a straightforward question about a foreign takeover: "Will this transaction put national security at risk?" And the answer, they thought, would be binary: either yes or no.

In the world of real transactions, though, it is rarely that simple. Suppose a company we don't fully trust wants to buy a company that sells software. Most of the software is plain vanilla consumer

stuff—spreadsheets, word processing programs, and the like. But one of the products is a centrally managed security service that screens all the packets that flow in and out of the user's computer. We might be able to live with the risk of compromise to the consumer products, but if the security service is ever compromised, every user's machine will be owned by a foreign intelligence agency from the day they install the software. That's too great a risk. We decide to oppose the transaction.

When it learns of our objection, though, the buyer says something that neither Senator Exon nor Representative Florio expected.

"Actually, we aren't interested in the security service. We've been planning to sell it. Can you approve the deal if we spin it off?" the buyer asks.

The sensible thing is to agree. We'll get everything we want without blocking the deal.

But what if there really isn't time to sell the security company before the CFIUS statute requires us to say yes or no to the transaction? We have only thirty days, after all.

"Well, will you approve the deal today if we *promise* to sell the subsidiary as soon as possible?" the buyer asks next.

Again, the sensible thing is to agree, as long as we know the buyer's promise will be kept. But to make sure it is, we need a strict agreement that can be enforced long after the transaction has been approved.

That simple example shows why CFIUS found itself forced to invent what became known as a "mitigation agreement." If the buyer entered into a binding agreement that mitigated any security risk in the transaction, the committee would approve the deal. It was good for everyone. The buyer and seller got what they wanted. And so did CFIUS—in fact, it got what it wanted without saying no to a foreign investment, something that can give a country a bad name in investment circles.

But a mitigation agreement doesn't have to be limited to something as clear-cut as the sale of a subsidiary. Sometimes the buyer wants to keep a subsidiary but has no interest in running it. It may solve a security

concern by promising to leave the current American management in place. Before CFIUS can rely on that promise, though, the company has to put it in writing and agree to be legally bound by it.

DPW had been telling us it had no interest in changing the management or practices of the U.S. terminals. I decided that we would ask DPW to put that in writing. After all, there was at least a small risk that the new owners would want to reduce costs by cutting security. A mitigation agreement would lock them in to their promises.

It turned out to be an easy sell. DPW agreed that it would stay in the voluntary security programs that P&O had joined. For good measure, DPW agreed to an open-book arrangement with DHS, allowing the department to inspect its records and obtain employee security data at will. These were incremental improvements in security, and DPW was willing to provide them in order to smooth the way for the transaction.

DHS did not need to get the approval of CFIUS to negotiate these provisions; we were the only agency on the committee with the slightest interest in this transaction. Once DHS was satisfied, the rest of the committee quickly okayed it.

By mid-January, CFIUS had finished its thirty-day review, DPW and DHS had signed the mitigation agreement, and the deal had cleared. According to the law, DPW was fully protected by the safe harbor provision of CFIUS. The United States could not legally overturn the deal.

A week went by, then another. Although CFIUS approval was in place, DPW was still in a bidding war with another purchaser. Not until February 11, when its rival bowed out, was DPW's victory announced in the business press. The contest was over.

Or, rather, it would have been but for a small company in Miami. Eller & Company had two joint ventures with P&O. For some reason, Eller didn't want DPW to take over that relationship. So it hired Joe Muldoon, a retired lobbyist and polo player, to get the deal overturned somehow. In the end, Muldoon turned out to be one of the

great overachievers in history. Not since Andrew Jackson fought the Battle of New Orleans has anyone won such an influential victory after the war was over. And never has such a public fuss been unleashed on behalf of such a tiny commercial interest.

Muldoon had never handled a CFIUS matter, and he probably didn't know that the approval was already final. He also didn't know—or perhaps didn't care—that terminal operators don't have much to do with port security. He just started telling anyone who would listen that national security was somehow at stake in the transaction. Finally, two weeks after the deal had been approved, someone heard him.

On Sunday, February 12, a story by the Associated Press claimed for the first time that the deal raised security issues, a twist raised by Senator Charles E. Schumer, who said that the transaction would “outsource . . . sensitive Homeland Security duties.”

I assumed he was repeating things he heard from Muldoon. They weren't true.

But that didn't matter.

By the end of Sunday, the blogs were buzzing. And the administration's rapid response team was silent. For one good reason. They had no idea what Senator Schumer was talking about. The transaction had set off no alarms as it wended its slow way through CFIUS. The lobbyist and politicians now complaining had said nothing while the deal was being reviewed.

And the review process had been over and done with for a month. For policymakers, that might as well have been eternity. Whatever whisper of worry they might have heard at the time of approval had long ago been crowded out by more pressing matters. Not until the working staff who had dealt with the case came to work on Monday were we able to gather the information we needed to respond.

By then it was too late. On Tuesday, February 14, the press had launched a story line that treated the transaction as the sale—lock, stock, and barrel—of six large American ports to an Arab company. That was the line taken that day by the Associated Press, which headlined the story “Arab firm may run 6 U.S. ports.”² Soon, the *Washington*

Times had the same slant: “Some of the country’s busiest ports—New York, New Jersey, Baltimore and three others—are about to become the property of the United Arab Emirates.”³ By Friday morning, a *Washington Post* writer channeling administration critics was frothing: “The management of major U.S. ports taken over by an Arab-owned company? What was the Bush administration thinking when it allowed such a thing?”⁴

For a couple of weeks, that was the nicest thing anyone said about us. No one listened when we tried to explain that port security is the job of the port authority and DHS, not the terminal operators.

It was a full-fledged Washington panic, of a kind seen only rarely, when a brand-new issue breaks suddenly and politicians have to wing it, with only their jangling switchboards for guidance.

The talk shows and blogs had a field day. So did the partisans. The issue let Democrats get to the president’s right on national security by demanding that Arabs not be allowed to run the security of American ports. Congressional Republicans, who couldn’t afford to seem soft on national security, rushed to condemn the deal as well. Congressmen of both parties launched crude attacks on Dubai and the United Arab Emirates to which it belongs. Congress held hearing after hearing to condemn the administration and to demand that the deal be overturned. In the end, the company buckled, promising to sell off its U.S. port properties.

A Washington panic is a funny thing. It seems to take Washington by the throat. No one can think or talk about anything else. Congress is suddenly ready to enact legislation in days, not weeks or months.

And then, like a tropical monsoon, the panic lifts. The clouds part. Politicians blink a bit shamefacedly in the sun. And everything goes back to normal.

That’s what eventually happened with the DPW case. Though you couldn’t have guessed from the hearings, our message slowly got through: DPW wasn’t buying American ports, we patiently repeated. It wasn’t going to be responsible for security. It had signed an unprecedented

mitigation agreement that addressed any reasonable security concerns. And Congress's noisy performance was undermining the U.S. reputation as a good place to invest—as well as CFIUS's reputation for raising only serious national security concerns.

Behind the bluster, Congress started to get nervous. It began looking for an exit. When DPW finally bowed to political reality and agreed to get rid of P&O's U.S. facilities, Congress was eager to claim the scalp and move on. Muldoon had earned his fee, at great cost to America's credibility in world financial markets.

The monsoon had passed. The sun was out again. But the DPW affair would hang over CFIUS for the rest of President Bush's second term. For a time, fear of another CFIUS eruption would allow DHS to turn the committee into a powerful bulwark against new computer and telecommunications insecurities. In the end, though, it would create a business backlash that showed the limits of security regulation even in a time of great and growing vulnerability.

For DHS, the fight over Dubai ports was a distraction from the real security risks posed by globalization of telecommunications and networks. The insecurity of U.S. networks wasn't just an organized crime problem. It was the result of deliberate policies adopted by countries that viewed us as an intelligence target. If they could get their companies to compromise U.S. networks, they'd do it in a heartbeat. So allowing foreign companies to take up critical positions in U.S. computer and telecommunications networks, either as suppliers or as service providers, raised serious national security issues. At the same time, globalization was relentless. The old days, when AT&T provided local and long distance service—and made all the equipment on the network—were long gone. And the collapse of the high-tech bubble had transformed the industry that emerged from AT&T's breakup. The Baby Bells were consolidating; long distance was disappearing as a separate business; wireless was displacing land-lines; and the equipment companies that had dominated North America for a century were in trouble. We couldn't just say no when foreign companies came

courting. In that context, mitigation agreements became a way to say yes to globalization without completely surrendering to foreign espionage. The agreements became a kind of company-specific network security regulation. We began to insist on a mitigation agreement in any transaction that posed even a modest threat. Each agreement created an ad hoc regime designed to curb foreign government infiltration of U.S. telecommunications and information technology.

The toughest agreements created a wall between the foreign owner and U.S. production facilities. This was common where CFIUS wanted to approve a deal in which the acquired company had sensitive government contracts. The wall was meant to keep the contracts free from foreign influence. The same thing was occasionally done for highly sensitive commercial contracts.

Another common security measure was to insist that the government (or an approved third party with technical skills) be guaranteed the right to inspect the buyer's hardware designs and processes, its software source code and testing results, and any other part of the production process that might reveal a deliberate compromise. To make sure that data was not shipped abroad and compromised there, some mitigation agreements required that data about Americans be kept in the country; sometimes the agreements required special security measures for the data.

The agreements also established a host of procedural security safeguards. These often included a government-approved security officer with broad powers and an obligation to report any suspicious incidents to the U.S. government. They also included regular audits by the government or a third party designated by the government. Personnel with access to sensitive data typically had to be screened; this sometimes included limits on outsourcing. Workers usually had to be trained in the security requirements and encouraged to report violations; and whistleblowers had to be protected from retaliation.

We were acutely aware that these measures weren't perfect. The substantive requirements were at best a mixed bag as far as security went. In theory, access to source code and hardware designs would

allow our experts to find any Trojan horse built into the product. But few government workers have the expertise to find these needles in a haystack of products. Unless we insisted that the companies pay for very expensive outside experts to check their work, or we received an intelligence tip about corporate misbehavior, we had only a modest chance of catching a really clever compromise.

The same was true of the procedural safeguards. Reporting obligations and whistleblower protections couldn't guarantee that we'd hear about an attempt at compromising U.S. products. They just increased the chances that someone would blow the whistle.

Still, imperfect as they were, mitigation agreements were well ahead of whatever was in second place. They were in fact our only good tool for policing foreign efforts to build insecurity into U.S. networks.

There was just one difficulty. The law didn't actually authorize mitigation agreements. No one knew how to enforce them, or even whether they could be enforced. If we were going to turn mitigation agreements into a kind of regulatory regime, we'd have to make sure they got the same respect as other regulatory measures.

Practically the first case I saw when I came on board was a small transaction that raised just this concern. The confidentiality of the process prevents me from providing details unless the companies have made them public, so I will not name the foreign buyer or the U.S. target. But both sold computer security products, so trust was critical. If you can't count on the loyalty of the company that provides your security, you have no security. *Quis custodiet ipsos custodes*—who will guard the guards themselves?—and all that.

As it happened, the foreign buyer of the security company had already entered into a mitigation agreement with DHS. An earlier transaction had been flagged for review, but the company had persuaded the government that negotiated safeguards would protect the national interest. The new case was tougher, but it became easier as we looked more closely. It turned out that no one had closely followed up as the company implemented the earlier agreement. The company

had sent the government letters putting forward self-serving interpretations of the agreement, and no one in government had responded. Now, as we took a close look, we didn't like what we saw. We were sure that the company had deliberately misread—and then violated—the mitigation agreement.

That was that. Why would we trust the company a second time if it hadn't lived up to the first set of promises? DHS took the lead in fighting the transaction. We ruled out another mitigation agreement. The transaction had to be rejected, we insisted. After a long period of disbelief that DHS truly intended to block the deal, the foreign buyer ultimately withdrew from the transaction.

That was the right result. The risk of foreign ownership can hardly be higher than in the area of security services. If we couldn't rely on the company's promises we couldn't find a middle ground.

I knew that the decision would enhance compliance with mitigation agreements. Before this, lawyers could tell their foreign clients that compliance with mitigation agreements was, if not optional, at least negotiable. After all, they might not even be enforceable, and for sure the government would have to sue to get compliance. If so, what was the harm in adopting an unreasonably narrow reading of the agreement? As long as its reading sounds plausible to a judge, the client would suffer no harm from defying the intent of the agreement.

But we didn't want to be forced to go to court over every misreading of the agreement, as though a security agency was just another party with a contract claim. Now we wouldn't have to. We had made it clear that companies would suffer very severe consequences indeed if they failed to live up to a reasonable reading of their mitigation responsibilities. We had taken a big step toward making CFIUS mitigation agreements a credible regulatory regime.

Still, I wasn't completely happy with DHS's performance. Not one member of CFIUS had taken responsibility for making sure the mitigation agreements that protected our security were actually being followed. How could we expect companies to take these mitigation

agreements seriously, I asked, if the government agencies that negotiated them didn't seem to care?

In one sense, DHS was the last agency that should have been responsible for enforcement of mitigation agreements. We were brand-new members of CFIUS, and the Policy office, which had been assigned to handle CFIUS, didn't exist until late 2005 and had not yet been staffed. Even so, we decided to take the lead in reviewing and auditing all of the mitigation agreements that DHS had signed. I hired Stephen Heifetz, a lean, sharp lawyer whose instincts and work habits had been honed in private practice. He could handle anything that the big-firm lawyers on the other side of the table threw at him.

Once he had his team assembled, I sent Heifetz out to audit the companies that had signed mitigation agreements with DHS. The team gave notice that they were coming, but not too much. When they arrived, they demanded records showing compliance and also insisted on reviewing all emails relating to the agreement. If the companies had been deliberately skirting their obligations it would have been hard to hide.

As we expected, most companies were complying, but we also saw clearly that they had become less than vigilant over the years. Heifetz said that email records told the same story in almost every company. Once the deal was done, months might go by without any special attention to the mitigation requirements. Then, suddenly, there would be a spike in high-level attention to compliance. The companies would launch internal reviews to make sure their performance was up to snuff. The spike almost always occurred a day or two after we had sent notice that our audit team was coming out for an inspection.

That was exactly what we hoped to achieve. It is human nature not to follow inconvenient rules when no one is watching. Every regulator knows that. If you want your rules followed, you have to remind companies that you're watching. That's what our audits did. Never again would the companies feel that DHS didn't care whether they complied with mitigation agreements. We were on our way to creating a successful cybersecurity enforcement regime.

This was not our only step to ensure that mitigation agreements were respected. We began to include financial penalties in the agreements. And to make sure that the buyer could never treat fines as simply a cost of doing business, we tied the size of the penalties to the value of the target company. The bigger the transaction, then, the higher the price would be for violating the agreement.

We soon had an opportunity to show we meant business when it came to assessing fines. One buyer of highly sensitive equipment had agreed to spin off a particular portion of the business within a few months of the closing. As the deadline grew nearer, though, the company began coming in regularly, explaining how hard it was working to find a buyer, and how much trouble it had encountered. It was clearly angling for an extension. We agreed, but we also declared that we'd begin imposing fines if the next deadline was missed. What's more, the fines would get bigger every month.

After agreeing to those terms, the company missed the next deadline, too. It asked us to forgo the fines. We refused. The penalties kicked in. As they began to mount, the company quickly found a way to spin off the business.

In a handful of cases, where the national security stakes were very high, we went even further. As the North American equipment market collapsed, the dominant supplier, Lucent, began to hemorrhage. The company put itself up for sale, and Alcatel won the bidding. For us, the stakes could not have been higher.

Alcatel manufactures telecommunications equipment and has been quite close to the French government for years. The French government had frequently been accused of carrying out espionage against U.S. targets. Lucent may have fallen on hard times, but it still manufactured and maintained the switches that carry most of North America's telephone calls. It was the home of the storied Bell Laboratories, whose Nobel-winning research had developed technologies from the transistor and the laser to the Unix operating system. Even the slightest risk that Lucent's capabilities might be turned against the United States was unacceptable.

I thought hard about saying no to the transaction, but the more we looked at the market, the more convinced we became that Lucent couldn't survive on its own. Vetoing the deal would put Lucent on a road to rapid decline. (That judgment still looks correct in hindsight; at the time, Nortel, the other North American telecom manufacturer, looked a bit healthier and chose to stay independent as the industry consolidated. That strategy turned out worse than Lucent's. In 2009, Nortel declared bankruptcy and was sold off in pieces.)

To salvage what we could from a bad set of options, DHS and other national security agencies decided to approve the deal and negotiate the toughest security measures ever imposed under CFIUS. We wanted above all to make sure that there would be no cheating on the deal. To make sure that the agreement would be scrupulously observed, the committee decided on the harshest penalty for breach that had ever been proposed.

If Alcatel breached the agreement in a way that threatened U.S. security, we insisted, the committee could reopen the acquisition. In other words, if there was a breach, the United States could require that Lucent be disgorged and restored to independence. This was called the "evergreen" provision because CFIUS's right to disapprove the transaction would remain in effect forever.

Alcatel and Lucent were nearly slack-jawed when we put this proposal on the table. How could that possibly work, they wanted to know. Five or ten years after the transaction had closed, Lucent would be deeply integrated into Alcatel; undoing the merger at that point could be a death sentence for both companies.

They weren't wrong. No one was sure how the companies could be pried apart at that stage. For that reason, some doubted that the United States would ever invoke the remedy. But the committee members believed that the risk was enormous—a compromise of Lucent's switches could disclose all of the government's wiretaps and make Americans subject to foreign wiretaps at home. If those were the stakes for U.S. national security, we needed to do everything possible to deter a violation of the network security measures.

A death sentence, we thought, should provide a measure of corporate deterrence.

In the end, Alcatel and Lucent accepted the agreement, including the evergreen clause. They decided that the risk created by the clause was material to their future prospects and disclosed it publicly to their investors (which is why I can discuss it publicly). In some ways, the Alcatel-Lucent deal was a high-water mark in the effort to make CFIUS a bulwark against subversion of U.S. information and telecommunications networks. It was public, it was demanding, and it was clearly going to be enforced. Indeed, other agencies, particularly Justice and the Treasury, began imitating DHS and bulking up their audit and enforcement capabilities at about the time we signed the Alcatel-Lucent agreement.

The tough new CFIUS regime benefited from the fallout from the Dubai port debacle. No policymaker wanted to be caught asleep at the switch if another transaction raised national security concerns. Agencies that had shown little interest in CFIUS before DPW now understood its importance, and they were reluctant to second-guess the security agencies. At least at first.

The same was true of investors, who had come to think of CFIUS as something of a paper tiger. CFIUS filings had hit an all-time low in 2003, but by 2006 and 2007 they had rebounded to levels not seen since Exon-Florio was enacted. (Part of that was DHS's doing; we began actively monitoring new transactions and requiring the parties to bring their deals—no matter how small—to CFIUS for review.)

Mitigation agreements also increased. DHS had signed seven such agreements in 2004 and 2005. In 2006 and 2007, after DPW, DHS signed an average of fifteen mitigation agreements a year. And many of the strongest enforcement measures for mitigation agreements were adopted in the same time frame.

For all the value we got from mitigation agreements, we weren't kidding ourselves that we'd solved the cybersecurity problem. CFIUS and its mitigation agreements were an unsatisfying way to address a

broader problem. CFIUS made it harder to compromise U.S. networks by buying a U.S. company. But foreign governments have other ways to compromise U.S. networks. They can provide subsidies so their own companies can underbid U.S. suppliers. If their price and quality are right, sooner or later the foreign companies will end up with a big share of the U.S. market—without ever making an investment that CFIUS can review. And if a company never makes a CFIUS filing, it will never have to sign a mitigation agreement, leaving some markets half-regulated.

Even more difficult to police is the supply chain. IT hardware and software are assembled from components made all over the world. A foreign government seeking to compromise U.S. computers doesn't need to buy Dell, or Intel, or Microsoft. It could buy a hard drive maker, a motherboard assembler, a modem supplier, even a keyboard manufacturer. Any of those components can be the source of computer security compromises. Again, without an investment in a U.S. company, CFIUS can do nothing about a "supply chain attack."

Even so, we had made a start, and a good one. Partial as they were, CFIUS mitigation agreements were still the best tool in our toolkit. They helped to close off the quickest and most obvious route that foreign governments might follow to compromise U.S. communications and data. Best of all, we seemed to have strong popular support for careful scrutiny of foreign acquisitions. If anything, the public had been convinced by the Congressional and media flap over DPW that CFIUS review was too lax.

In the end, though, DPW was poisoned fruit. The unjustified abuse that Congress had heaped on DPW eventually spurred a backlash. But when it came, it was aimed not at the worst Congressional offenders but at DHS.

Using CFIUS to reduce cybersecurity vulnerabilities was DHS's key strategy. As we turned mitigation agreements into a regulatory tool, we were drawing fire. And from some of the same forces that opposed us when we used new tools to deal with the risk of jet travel—business and the international community.

These forces slowly turned the DPW case into a millstone around the necks of the security community. At first, they concentrated on stopping the congressional effort to enact legislation that would make CFIUS tougher. Business groups quietly communicated their concern about the bill's effect on investment. After the initial burst of enthusiasm, work on the bill slowed. Nothing had been enacted by the mid-term elections of 2006, in which Democrats took control of both the House and Senate. Although they had been loud in condemning the DPW deal while out of power, by the time they took control, those calls had muted.

Congress was now hearing from other governments as well as business. Other governments have no reason to encourage the United States to protect its national security through CFIUS. In fact, some governments have a direct interest in precisely the opposite. But even for our friends, there's no reason to praise CFIUS. The safest—and most conservative—stance was disapproval, and the DPW case certainly offered plenty of fodder for that position.

Many governments claimed to see a protectionist motive in CFIUS. For some, the accusation of protectionism was clearly a projection of their own inclinations. France, famously, had decided in 2005 that a French yogurt company was a "jewel" of French industry and therefore could not be sold to Pepsi. The Germans had refused to let foreigners buy into their auto industry. But the best defense is a good offense, the Europeans had learned; so European and other trade negotiators began to criticize U.S. CFIUS practice, hinting that it would have to be negotiated away in the next round of trade talks.

In the United States, unease about CFIUS spread to businesses that depended on foreign companies—from Wall Street investment banks to K Street lawyers. They too began quietly campaigning against the new regulatory push. They didn't want to see the United States opened further to espionage or sabotage, of course. But couldn't we do that without cutting off their deal flow?

The Alcatel-Lucent "evergreen" clause added to the tumult. From a foreign investor's point of view, the one good thing about CFIUS was

its certainty; once a deal cleared, it was cleared for good. It was a safe harbor against future storms. By adding an evergreen clause to the mitigation agreement, though, we had torn down the breakwater, leaving Alcatel and perhaps others exposed to future national security storms.

For foreign investors and their lawyers, the evergreen clause offered a second issue to rally round. In our view, the furor over the provision was out of all proportion to how often it was likely to be used. The fines and other enforcement measures that DHS had introduced were almost always tough enough to keep companies on the straight and narrow. Evergreen clauses were worthwhile only when normal incentives might not be enough to ensure compliance (usually when we feared that a foreign government could force the foreign company to take actions without regard for the company's own financial interests).

Part of the problem was perception. We couldn't talk about individual cases, and we didn't tell the parties to the transaction what our intelligence said about the buyer. So from the outside, our decisions did not look consistent or predictable. Sometimes we'd oppose a deal fiercely because intelligence revealed dangers that weren't obvious to outside observers, or even the parties. To outsiders, the role of intelligence in CFIUS was deeply frustrating, because it deprived them of the opportunity to rebut the charges.

They weren't wrong to be concerned. Intelligence is never perfect, and it should always be challenged before it is relied upon. Some of the CFIUS agencies didn't have a broad understanding of intelligence, and they sometimes gave it too much credence. (From time to time, I would propose audits or inspections of foreign buyers as a way of checking what the intelligence agencies were saying, but it was not easy to get the buyers to agree. Perhaps they didn't understand that an inspection might help them by providing a check on the intelligence—or perhaps they feared that it would confirm what the intelligence was telling us.)

The backlash against CFIUS was also aided from within. CFIUS has a peculiar structure that is almost guaranteed to spur bitter conflict. Originally established as a committee of cabinet members

headed by the treasury secretary, the committee has gradually added members from the White House bureaucracy. So, in addition to cabinet departments like Defense, State, DHS, Justice, and Commerce, the table is cluttered with representatives from the U.S. trade negotiating office, the Office of Science and Technology Policy, the National Security Council, the National Economic Council, and so on and so on. I say cluttered because these offices could talk but did not vote on transactions.

The White House offices are, in theory, at the table to protect the president. The idea is that their advice will be conveyed confidentially to the president if and when the committee makes a formal recommendation. That's the theory. In fact, the White House agencies all have turf struggles with each other, and they're often at the table to fight their rivals, or in the hope of influencing the debate before it reaches the White House.

White House staff didn't vote in CFIUS cases. But that hardly mattered, because votes were rarely useful. CFIUS is not an agency, and the Treasury Department, though it chairs the committee, is little more than first among equals. The purpose of the committee is to make a recommendation to the president. If one cabinet secretary wants to say something to the president, and another secretary is adamant that something else must be said, the treasury secretary will not be able to resolve the dispute. Both messages will be delivered.

And, given the institutional interests of the departments, it was almost inevitable that disagreements would arise. The State Department, for example, is always concerned about the reaction of foreign nations to our CFIUS decisions, and foreign nations never welcome a tough CFIUS regime. So State invariably opposed for as long as it could any effort to put conditions on transactions. The Office of the U.S. Trade Representative (USTR) was, if anything, even more predictable in opposing the use of CFIUS for cybersecurity purposes.

DHS, in contrast, was among the most likely to propose mitigation agreements or outright vetoes of risky deals. With Justice and the Defense Department, DHS formed the heart of CFIUS's national

security wing. Less predictable, at least over time, were Treasury and Commerce. Treasury is deeply sensitive to the mood of foreign investors, and that tended to push it toward the State Department. But it also had a large security role in stopping terrorist finance, and it was constrained by the need to act as chair of the committee, muting its natural sympathies. The Commerce Department speaks for U.S. business interests; some of its leaders thought that was enough to determine their position in CFIUS cases. Under other leadership, though, Commerce would sometimes give weight to its own national security arm, the office that oversees export controls on high technology and recognizes the risk posed by potential compromises.

There was a deep divide between the “national security” agencies and the “trade” or “economic” agencies. And, because Treasury could never force a decision over an impassioned dissent, arguments at CFIUS, particularly at the lowest levels, had a kind of well-worn vitriol to them. Everyone knew that the dispute would go higher. The only reason to pull back was fear that your boss wouldn’t support you. At DHS, that was never a problem. We had short lines of communication and decisive leaders at the top. The secretary and deputy secretary could absorb new information and pass judgment on a course of action in minutes. Other agencies with less certainty of their boss’s views were less willing to hold firm, and that sometimes helped us advance our cybersecurity agenda.

In the long run, though, the DPW flap hurt us badly. As the panic wore off, policymakers all across the government began to realize that they had been foolish to make such an issue of the DPW investment. That had been DHS’s view all along. We thought the DPW case was a distraction from the greater dangers in telecom and information technology. For the business and international interests that opposed those measures, though, DPW was a godsend. Everyone knew that DPW had been a serious overreaction, and it was easy to lump everything together and argue that CFIUS was being abused.

As that idea took hold, CFIUS meetings grew more divided. Decision makers at the top of the Commerce Department or the U.S.

Trade Representative's office might not know much about cybersecurity, but they were happy to take a stand against CFIUS abuse. They began to back their lower-level officials more frequently on that basis. And so a logjam of unresolved conflict over CFIUS issues began to creep higher up the decision chain.

Deadlock became the norm. Trade agencies in particular would exercise a "bureaucrat's veto" by insisting that nothing could be done without their agreement and then asking for more paper, more process, and more debate before that agreement could be granted. They didn't say no, they just asked for more time.

Everyone in government is familiar with this tactic. The power to delay is often the power to prevent a policy decision. It was one more weapon in the arsenal of the institutional conservatives trying to prevent new policies from being adopted.

But delay had unexpected costs. The thirty-day deadline for decisions on most transactions was increasingly ignored. Too often, CFIUS would launch forty-five-day "investigations" simply to give the contending agencies more time to resolve their differences. Or it would strong-arm companies into "withdrawing" their applications and refilling them, starting the clock over again.

Much of this delay was caused by a growing determination on the part of the trade agencies to fight over the terms of mitigation agreements. But from the outside, all that the parties knew was that CFIUS was slowing their deal. For the trade agencies, this was a twofer. Their stalling tactics made it harder to get tough new mitigation agreements. And the delays brought the entire CFIUS process into disrepute, which increased the business backlash against strong CFIUS review.

Of course, the delay was hard on the companies involved in the transaction, but the trade agencies only occasionally seemed bothered by that. In fact, while DHS was viewed from the outside as the principal source of CFIUS scrutiny, and thus of delays, we were often the only voice arguing that the process should move more quickly to protect investors' need for certainty and promptness.

Officials who joined the administration after DPW also brought with them views shaped by the public debate but not informed by the intelligence that had driven our decisions in particular cases. Without access to that information, they tended to assume that all CFIUS decision making had been as irrational as the DPW case.

The National Security Council in particular suffered from this effect, and by 2007 it had abandoned any pretense of being an honest broker in CFIUS disputes. It became instead the principal combatant, working relentlessly to cut back the tough new security regime that we had introduced to CFIUS.

The critical showdown would come over who could negotiate and sign mitigation agreements. There was a long tradition of agency autonomy in this area. For years, mitigation agreements had been viewed as agreements with individual CFIUS members, not with CFIUS as a whole.

“This proposed mitigation agreement is between you and DHS,” we used to tell companies when we tabled a draft. “It is meant to address the concerns that DHS has about your transaction. If we negotiate a satisfactory agreement, DHS will not oppose the transaction. We’re not speaking for CFIUS, so there’s always a possibility that the committee will disapprove the deal notwithstanding this agreement. And if we don’t reach agreement, your deal may still be approved. You are simply taking the risk that DHS will oppose the deal and that we’ll be able to persuade CFIUS not to approve it.”

Because we were negotiating only for DHS (and sometimes for other agencies with similar concerns), it was easy for us to agree on tactics, priorities, and reach agreement on a deadline. That autonomy and flexibility is what allowed DHS to sign the quick mitigation agreement with DPW that was the administration’s best defense during the Washington panic over the case.

For the trade agencies, though, that was all history. As far as they were concerned, DHS’s authority to sign mitigation agreements had to be taken away. First, they argued that DHS and other agencies negotiating mitigation agreements should keep the rest of CFIUS

informed about the progress of the talks; then they argued that DHS should take their views into account in negotiating the agreements. Both of those positions sounded perfectly reasonable, but we accepted them with foreboding.

In theory, consultation with other agencies may provide useful new perspectives or avoid problems. In government practice, however, a consultation requirement is just a first step; it allows the consulted agency to second-guess and interfere, because it gives the agencies a chance to probe for weak spots. That is what happened in CFIUS. The trade agencies had little interest in helping the security agencies improve their mitigation agreements. Their principal interest was gaining enough information to argue that no mitigation agreement was necessary. Some of the more extreme agencies even violated the spirit and perhaps the letter of the CFIUS confidentiality requirements by “coaching” parties, suggesting arguments they should make when negotiating with DHS and then seconding those arguments in internal debates.

Eventually, the trade agencies began to insist that they weren’t being consulted in good faith if DHS reserved the right to sign an agreement while the trade agencies were still asking questions. Consultation, in other words, couldn’t end until the trade agencies agreed that all their questions had been answered. Of course, that formulation simply meant that the trade agencies could stall an agreement for as long as they could think up new questions. Or, more commonly, for as long as they could find new ways to ask the same old questions.

Impatient with this effort to undercut DHS’s authority in a back-door fashion, DHS simply continued to sign mitigation agreements. The parties usually were happy to do the deals, and the quicker the better. They had no interest in the ideological issues being raised by the trade agencies; they just wanted to move on. The trade agencies believed that the willingness of the parties to accept DHS’s terms was irrelevant. They thought the parties were simply knuckling under because they needed to get their deals done quickly. They thought it was bad policy to use the leverage of CFIUS approval to extract

security agreements that would not apply to everyone in the industry. And they found DHS's refusal to be cowed more and more frustrating. Even at the deputy secretary level, conflict grew intense as CFIUS pressed DHS to give up its traditional authority to execute mitigation agreements.

By early 2007, the trade agencies and Treasury decided to take their frustration to Congress. The new Congress was led by Democrats, and they had made CFIUS reform a priority. But they were also listening to the business groups and foreign countries that had begun complaining about DHS's tough scrutiny. At a hearing to which DHS was not invited, its mitigation agreements were roundly criticized. Witnesses repeatedly bemoaned the fact that the number of mitigation agreements required by DHS tripled in 2006 from the previous three-year average (up from 4.5 to 15).

One witness expressed concern "that some agencies are taking undue advantage of the leverage inherent in CFIUS. CFIUS should not be a fishing expedition for a single agency to address comprehensive industry objectives on a "catch-as-catch-can" basis merely because they have leverage over one industry participant. ... [I]f the Department of Homeland Security perceives a vulnerability in our telecommunications infrastructure, it should address that vulnerability across the sector, without regard to the ownership of firms."⁵ Others made similar complaints.

Congress continued to insist that it wanted to make CFIUS tougher, but its actions said something else. Throughout the hearings and debates, congressmen touted the new bill as strengthening CFIUS and security. But when the television lights were turned off, the drafters sat down with the Treasury Department, and the committee leadership added language designed to undercut the authority of any agency to enter into a mitigation agreement on its own.

The new bill took the long overdue step of acknowledging the need for mitigation agreements, and it called for a "lead agency" in each case to negotiate the mitigation agreement. At the last moment, though, the House financial services committee leadership slipped in

an amendment to the bill, requiring that any mitigation agreement be negotiated “on behalf of the committee.”⁶ The effect of this modest phrase was dramatic. It would allow Treasury and the trade agencies to insist that they had to supervise the negotiation of any mitigation agreement now that the talks were being conducted “on behalf of” the entire committee. That meant that no negotiation could occur without a consensus among all CFIUS members. And that in turn meant that the trade agencies could use the “bureaucrat’s veto” of endless delay to kill mitigation agreements even over the objection of the agency negotiating them.

The new CFIUS law⁷ also contained a provision requiring that mitigation agreements be based upon a “risk-based analysis” of the threat to national security of the proposed transaction. The same manager’s amendment described above also added language to this provision to specify that this analysis must be “conducted by the committee.” This amendment gave the trade agencies a hand in analyzing national security threats and determining the level of appropriate mitigation. Once again, the committee leadership had reduced the security provided by CFIUS.

The House didn’t exactly advertise the fact that it was weakening the hand of the security agencies. That wouldn’t have been consistent with the dominant narrative in the press, where Congress was still loudly proclaiming the need to strengthen CFIUS because the administration hadn’t given enough weight to security in the DPW case. Still, it seems likely that Congress knew exactly what it was doing. The business witnesses had asked that agency autonomy be abolished or constrained in the name of encouraging foreign investment, and as the Congressional Research Service noted, the amendment was adopted because the earlier bill, which lacked it, “could have delayed and discouraged foreign investment.”⁸

International and business groups, in short, seem to have persuaded the committees that the real problem with CFIUS was not that it was too weak but that it was too tough. Needless to say, that wasn’t a change of mind that Congress was eager to shout from the rooftops.

After the bill was enacted, the National Security Council wasted little time turning the “on behalf of” language into precisely what DHS had feared—a radical restriction in the authority of the security agencies. In fact, it built an entire edifice of obstruction on those few words. Under the executive order⁹, the lead agency must achieve consensus within CFIUS before it can even propose a mitigation agreement. To do this, the agency must prepare a written statement that (1) identifies the national security risk posed by the transaction, including potential threats, vulnerabilities, and consequences, and (2) sets forth the mitigation measures, which must be “reasonably necessary” to address the risk.

After jumping through these hoops just to propose a mitigation measure, the lead agency must also get committee approval before negotiations can begin. It must keep the committee fully informed of its activities and must notify the Secretary of the Treasury in advance of any proposed major action, allowing time for the committee to consult and direct the lead agency about how it should act.

By the time the order was fully written, the lead agency was less a leader than an indentured servant. It might sit in the driver’s seat, but every member of CFIUS would have a hand on the steering wheel and a foot on the brakes of the negotiations. The trade agencies were happy to use the brakes. No negotiations could occur, they would insist, until a final position had been agreed to by all agencies. This made the old tactic of delay and refusal to agree a potent weapon again. Security agencies were ordered not to even tell the parties to the transaction what their concerns were until they had the consent of the other agencies.

This quickly led to absurd results. In one case, when DHS expressed concerns about what might happen after the merger, the parties promised to take action after the merger that would completely resolve the worry. DHS suggested to the committee that the promise be put in writing so that the assurance was binding. Some members objected and the assurance was never formalized.

In another case, the deadlock in the committee went on so long that the parties wrote letters to all members of the committee begging

to be told what DHS wanted, arguing that it would much rather agree to reasonable mitigation conditions than wait for the committee to finish its internecine bureaucratic war. Nothing doing. The trade agencies were determined to make the United States safe for foreign investment no matter how many foreign investors they had to hurt in the process.

The most ironic note was sounded toward the end of the administration, when another foreign purchase of port facilities was submitted for approval. DHS proposed a modest mitigation agreement, similar to the DPW agreement that so many in Congress had condemned as inadequate during the panic. This time, though, under the law that Congress had enacted in reaction to DPW, even this modest agreement could not be imposed. The trade agencies refused to accept it, and Congress had made their consent a necessary condition to any mitigation agreement.

The counterattack on behalf of business and the international community had come a long way against heavy odds. The new law had been so trimmed and twisted that in the end the one part of the DPW affair that could not be repeated was the one part that contributed to security—the mitigation agreement.

The effect was felt quickly. In 2008, the number of mitigation agreements fell dramatically, and they became even more rare in 2009.

In the end, though, much of what DHS did to make CFIUS a force for network security endured. Even in the waning days of the administration, long after the new CFIUS law and executive order took effect, a new transaction raising severe security concerns came to CFIUS. Working with the other security agencies, DHS made the case against the deal. Faced with evidence of grave risk, the trade agencies folded; they did not oppose our recommendation that the transaction be rejected. Had the security agencies been willing to execute a mitigation agreement, they would have accepted that recommendation as well.

The lesson of that transaction was that the trade agencies would not fight the security agencies when the chips were down. Security is

the mission of DHS, Defense, and Justice. If those agencies say with confidence that a transaction will raise serious security concerns, it is hard for an agency like USTR to second-guess them. And at the highest levels, each agency tends to take a broader view than simply its own bureaucratic interest. This means that, for transactions that raise the greatest concern, the new law is not fatal to the reforms that DHS pioneered.

Still, the story shows how hard it is to regulate even the most dangerous cybersecurity threats. CFIUS dealt with the particularly overt and troubling threats, and in most cases it had found a way to allow investments to go forward, though with safeguards.

Even so, the nations and companies that opposed any regulation had successfully advocated for a law and executive order that undermined the security agencies, at least somewhat. That they accomplished their mission in the teeth of noisy public demands for tougher CFIUS security standards is a testament to their formidable clout.

Smallpox in the Garage

In January 1970, a German electrician fell ill after a trip to Pakistan. He was hospitalized with what appeared to be typhoid fever. He had been isolated for several days when the doctors realized that he didn't have typhoid fever.

It was smallpox.

Fear ruffled through the hospital, and the community beyond. Smallpox has probably killed more human beings than any other disease. And it kills them with particular cruelty. After starting out like a bad flu, after a few days the disease attacks the victim's skin. Tiny spots appear, spread, and then harden into pus-filled blisters. Gradually, with excruciating pain, the blisters pull the outer layer of skin away from the under-layers. Sometimes the skin pulls loose in sheets. Sometimes the blisters attack not just the skin but the eyes, the throat, and every other orifice, ripping loose skin inside the body as well. Desperate with thirst, the victims can't drink; swallowing is just too painful.

Throughout it all, the victim remains fully conscious. A third or more of the victims die. Those who survive are often permanently scarred, or blind or both.

The electrician lived. But many who came into contact with him were infected. Several died.

What was most frightening was how the virus spread. One victim spent only fifteen minutes in the hospital. All he did was ask directions, briefly opening a door that led to a corridor thirty feet from the patient's room. That was enough. He came down with smallpox.

Three other victims were even farther away—two floors above the electrician’s isolation ward. It was January, but tests revealed that opening the hospital windows just a crack allowed currents of air to drift between rooms on different floors. The virus had floated out the patient’s window and along the outside wall; it then slipped into three different rooms two stories above, infecting patients in each room.

Seven years later, in 1977, Ali Maow Maalin also fell ill with smallpox. This time, though, it turned out to be good news.

Maalin was a cook from Merca, Somalia—where smallpox was making its last stand. Vaccination was slowly tightening a noose around the disease. Because smallpox reproduces only in humans, widespread vaccination left fewer and fewer places for the virus to reproduce and spread.

The first vaccination for smallpox—or indeed for any disease—came in 1796. That was when Edward Jenner realized that milkmaids who caught cowpox seemed to be protected from smallpox, to which cowpox was related. Jenner’s vaccine based on cowpox marked the beginning of man’s counterattack on smallpox. By the 1970s, vaccinations had gradually reduced the disease’s natural range to the wilds of Somalia and Ethiopia.

The World Health Organization hoped to make Ali Maow Maalin the last victim of smallpox in history. It quickly vaccinated everyone who had been in contact with him, then held its breath. Would other cases flare up?

WHO waited.

A year.

Two years.

Three.

At last, after three years with no natural cases of smallpox, the World Health Assembly declared victory. It triumphantly called a special 1980 meeting.

“[T]he world and all its peoples have won freedom from smallpox,” the assembly declared. This was “an unprecedented achievement

in the history of public health.” Together, the nations of the assembly had “freed mankind of this ancient scourge.”¹

Copies of the virus were locked away in Atlanta and Moscow for research purposes, but the disease was gone from nature. Vaccinations stopped. Few Americans born after the 1960s have the dimpled scar on their arm that is the last trace of mankind’s worst nightmare.

It had taken a bit less than two centuries for vaccination to free the world from “this ancient scourge.”

Today, the likelihood that the world will remain free from this ancient scourge is close to zero.

Smallpox is back, or nearly so.

Within ten years, any competent biologist with a good lab and up-to-date DNA synthesis skills will be able to recreate the smallpox virus from scratch. Millions of people will have it in their power to waft this cruel death into the air, where it can feed on a world that has given up its immunity.

How can I be so sure? Easy. I’ve seen the same thing happen already, and so have you. The very same revolution that made possible the explosion of information technology—and set the table for network attacks—is now transforming biology, with consequences that are both exalting and frightening.

The same relentlessly exponential improvement in technology that gave us Moore’s Law and that democratized the computer is now democratizing the technology of life. It is empowering an army of biologists to tinker with biology in ways that will help us all live longer and more comfortable lives.

And then, unless we do something, it will kill us in great numbers.

“Synthetic biology” blends biology, chemistry, and engineering. The field really began to take off when it moved from laboriously replacing a single gene to building whole stretches of the genome from scratch.

DNA is organized like a spiral staircase, and each step on the stairs is called a base pair. Linking base pairs together into longer sequences

allows researchers to make more complex genes—and ultimately more complex organisms. So progress in synthetic DNA is measured by how many base pairs have been successfully strung together. In recent years, progress has been exponential.

In 2002, after a two-year effort, a team of researchers announced that they had assembled the entire polio virus. To do that, the team had to assemble 7,500 base pairs of DNA, precisely in order. The next year, scientists managed to knock years off the process, assembling a bacteriophage with 5,300 base pairs in just two weeks.

Two years later, in 2005, researchers' capabilities had tripled. A team managed to synthesize an influenza virus with 14,000 base pairs. Just a year later, they had surpassed that mark by a factor of ten, synthesizing the Epstein-Barr virus, with 170,000 base pairs.

Smallpox has 180,000.

By 2005, whether smallpox would be synthesized was simply a matter of choice, not of capability.

The following year, the outgoing secretary general of the United Nations, Kofi Annan, grew alarmed. He pointed to researchers' successes in building an entire virus from scratch and said, "In the right hands, and with the appropriate safety precautions, these are sound scientific endeavours that increase our knowledge of viruses. But if they fall into the wrong hands, they could be catastrophic."²

Too late. By 2009, the state of the art had left 180,000 base pairs in the dust. A team of researchers announced that it had assembled a bacterial genome with 583,000 base pairs. Creating smallpox from scratch was no longer even an interesting challenge.

Nor were these capabilities confined to a few specialty laboratories. Foundries sprang up to sell made-to-measure DNA, at ever-declining prices that put Moore's Law to shame. Synthesizing DNA cost \$10 per base pair when George W. Bush ran for president in 2000. By the time of his second inauguration, the price was \$2 per base pair. When he left office in 2009, the price was down to about 25 cents. For those who don't want to use a foundry, DNA synthesizers are available for sale on eBay.

Kofi Annan was wrong. This technology isn't going to fall into the wrong hands. Just like jet travel and powerful computers, it's going to fall into *everybody's* hands. The Mayo Clinic. Hezbollah. Pfizer. Al Qaeda. Apple. Ted Kaczynski, Timothy McVeigh, and the Fort Hood shooter.

They won't need their own labs to build bugs to order. Even today, it's possible to obtain long sequences of synthetic DNA simply by sending a message to the private "foundries" that assemble DNA to order.

Struggling to survive in a new market with thin margins, the foundries' sense of responsibility for what they make is, well, limited. The *Guardian* newspaper in Great Britain demonstrated this when one of its journalists successfully ordered a lightly modified piece of the smallpox genome over the web. The order was mailed to his home, no questions asked. When a dozen foundries were asked whether they screened DNA orders to see whether they were providing sequences that terrorists could turn into weapons, only five answered "yes."

As many as half the foundries questioned by journalists did not routinely screen their orders to make sure that they were not helping terrorists construct a dangerous virus. The order came in, and they filled it, often with no questions asked.

If current trends continue, anyone who can get his hands on a computer virus today will soon be able to get his hands on a custom-built biological virus.

And who can get his hands on a computer virus today? In an age of drop-down-menu malware attacks, the answer is simple.

Anyone who wants to.

Perhaps it isn't completely fair to assume that exponential growth in biotechnology will democratize biological terror in the same way that computer technology democratized computer crime. After all, unlike computer hackers, bio-hackers can't pretend that releasing pathogens is a good way to demonstrate their skills or to dramatize the need for better biosecurity. So perhaps biological malware will arrive more slowly than its computer counterpart. That's good.

So far, the terrorists who've tried to use biological weapons have turned out to be more hapless than terrifying. A cult that wanted to win an election in rural Oregon poisoned the local salad bar to suppress turnout. A Japanese group experimented with anthrax and ended up spreading a harmless, non-virulent vaccine strain around Tokyo. The anthrax-laced letters sent to prominent journalists and politicians in 2001 included a warning to take antibiotics and thus dramatically reduced casualties. Al Qaeda tried to acquire biological weapons before 9/11, but its efforts never really got off the ground.

Maybe large-scale bioterrorism is harder than it seems. Or maybe we're just in that golden era we also experienced in computer technology; maybe the bad news just hasn't caught up with the good news. Much the same thing happened with jet travel for that matter. Apart from some Brazilian military officers who commandeered a civilian flight in 1959 to further their coup attempt, there were no notable hijackings of a commercial flight before 1968, even though they had been possible since at least the 1950s. Early that year, though, an El Al plane was seized by Palestinian terrorists and a U.S. flight was hijacked and diverted to Cuba. Then the deluge began. By the end of 1968, there had been half a dozen hijackings to Cuba alone, and the stage was set for decades of ever more spectacular hijackings.

The lag between good news and bad owes something to the surprisingly conservative nature of terrorism. Terrorists don't like to fail; failure doesn't inspire fear. But once a new tactic has been pioneered, and it has become clear that governments don't know how to respond to it, everyone piles on. Suicide bombings were virtually unknown until the early 1980s, when they were used in the Lebanese and Sri Lankan conflicts. The tactic is now widely used by terror groups in many countries. We may be only one or two successful attacks away from a similar wave of bioterrorism.

When those attacks will occur, however, is anyone's guess. All we can say is that every year biological attacks become more probable, just as biotechnology becomes ever more democratized. And, of course, if disaster becomes more probable every year, then sooner or

later disaster will happen, though it may show up late. That's a lesson financial markets learned again in 2008 (as did New Orleans residents in 2005). Sooner or later, the inevitable does happen.

One cabinet-rank official summed it up a little differently after I gave him a briefing on the topic.

"Maybe," he said, "the human race isn't meant to survive."

I understood how bad the threat was. I had been briefed on it while investigating U.S. intelligence agencies' work on Iraq's WMD program. The agencies were eager to tell us how much they knew about other nations' nuclear weapons programs. We got briefing after briefing. Nukes were a major concern, and the agencies had scored many successes in penetrating other nations' programs.

On biological weapons, the intelligence community was noticeably less voluble. Everyone acknowledged that biological weapons were a terrible threat. Worse than nuclear weapons in some ways. They could kill as many people. And the aftermath would be worse. The day after a nuclear weapon goes off in an American city, a hundred nations will order their airlines to fly to the United States, carrying assistance until the crisis has passed. The day after a biological weapon is used in an American city, a hundred nations will order their airlines to stop flying to the United States until the crisis has passed.

But, with a few exceptions, intelligence operatives and analysts seemed almost to have lost hope of understanding other nations' biological weapons programs. The programs are easier to hide and require less in the way of investment than nuclear weapons. The equipment and training that supports them have many innocent commercial uses in the pharmaceutical and pesticide industries.

And the agencies' track records were not good. The Soviet Union—and Russia thereafter—had maintained a truly loathsome biological weapons program for decades after the United States gave up its program. It treated the disappearance of smallpox, and the worldwide end of smallpox vaccinations, as an invitation to devise more potent weapons using its stores of the pathogen. The Soviet program was

discovered only when defectors began to talk about their work on artificial new diseases that were proof against existing countermeasures, or that responded to treatment by changing into something even worse.

The same was true in Iraq. Saddam Hussein maintained a biological weapons program for years, hidden from both U.S. intelligence and UN inspectors. (If you're wondering why no such program was found after the U.S. invasion, the answer is that Saddam Hussein finally dismantled the program after his son-in-law defected and disclosed it to the West in 1995. Saddam admitted the existence of the program and announced that it had been shut down; intelligence agencies, shocked by what they had missed, credited Saddam's admission but doubted his claim that the easy-to-hide program had ended.)

Intelligence gaps on biological weapons raised our concern about anonymous attacks. Like computer malware, biological agents are hard to tie back to an individual or group. Ambiguity about attribution has already prevented the United States from taking effective retaliatory action against computer attackers. It's quite possible that we won't do any better against attackers armed with biological weapons. The best test of our capabilities came in the 2001 anthrax attacks. The FBI used great ingenuity and massive resources to question, search, and investigate all the likely suspects. It finally announced, to some skepticism, that it had identified the guilty man in 2008—seven years after the attack.

When I got to DHS, I asked my staff what we could do to cut the risk of biological terrorism. They described two new programs launched after the 2001 anthrax attacks. The first was to develop countermeasures—vaccines, treatments, etc.—for the most threatening pathogens. The second was to get a better picture of who actually had access to such pathogens inside the United States. These were large programs, funded by a Congress that feared another attack was imminent. But as the years went by without an attack, the programs had slowly been bent to fit the institutional inclinations of the agencies that got the money.

Take the countermeasures program. This is an absolutely essential step. Unlike nuclear weapons, biological weapons can be defeated even after the attack. That is, if we have a smallpox vaccine and can distribute it quickly, we can stop an infection in its tracks, greatly limiting the harm done by the disease. We could take the weapon out of terrorists' hands. A biological attack that is met by quick, effective countermeasures is like a bomb that has been defused before the blast.

But our countermeasures strategy has serious flaws. It requires a massive investment in medicines that often have no civilian use. We will never have a need for smallpox vaccine except to defend ourselves against attack. The doctors and researchers of the National Institutes of Health (NIH) were not used to battling human adversaries. They were scientists who wanted to do pure research, not something that felt like military work. Like any industry facing a market change, the traditional research community resented the funding that went to countermeasures research, and they didn't have much trouble turning that resentment into an ideological and personal campaign against the program. (The debate broke into the open when traditional NIH researchers launched a smear campaign against Tara O'Toole, the Obama administration's nominee to head DHS's science office. Her success at building a countermeasures research program led to her being labeled an alarmist and a female Dr. Strangelove by traditional researchers, delaying her confirmation for months.)

More troubling was the way business as usual in other parts of the Department of Health and Human Services threatened our ability to actually use the countermeasures that had been developed at such great cost. For example, getting approval for such countermeasures is staggeringly expensive. A host of regulatory hurdles has been set up for new drugs. The regulations assume that the drugs are being championed by private companies hoping to make billions in profits if they are approved. But the private sector will not spend billions to get regulatory approval for a product that may never be deployed.

Even if government pays that cost, most countermeasures, such as vaccines, have side effects that may be rare but can be quite serious.

Even faced with the threat of an occasionally deadly H1N1 influenza in 2009 and 2010, many Americans refused to be vaccinated. It would be nearly impossible to persuade them to be vaccinated against anthrax or smallpox on the chance that these pathogens would be unleashed by terrorists.

So the countermeasures will sit in warehouses, waiting for an event. Once smallpox or anthrax is released in a vulnerable population, the countermeasures will have to be deployed on a massive scale in a matter of days, even hours. At DHS we knew that this would be a logistical nightmare. After all, we'd lived through the errors and delays as government tried to improvise in the wake of Hurricane Katrina. An incident of biological terrorism would create the same problems, except the victims might be desperately sick, not just hungry and thirsty, and the rescuers would be delayed longer by fears for their own safety.

Imagine a biological attack in which terrorists release a large cloud of anthrax in an urban area without telling anyone. Even with air sampling equipment in place it might take a day or two to confirm the attack. If everyone who'd been exposed took antibiotics within three days, practically all of them could be saved. The weapon could be defused. But if it took five or six days to start antibiotics, we could lose half the population. That's an enormous difference, making every hour of logistical delay a matter of life and death.

So how were we planning to deliver antibiotics? The postal service. That's right. The aggressively unionized postal service workforce would be asked to show up and drive into anthrax-contaminated areas to distribute antibiotics. Of course, they would want armed protection, so law enforcement agents would somehow meet up with the postal workers and they'd both go around delivering antibiotics. To me, this sounded, well, unlikely. Getting the workers to show, hooking them up with their armed escorts, making sure they and their escorts had started antibiotics, verifying the routes, making sure they weren't swamped by people who couldn't stay home for their antibiotics, keeping others from trailing them to steal antibiotics from mailboxes, all of

this would have to be done for the very first time under unbelievable time pressures.

There was a way to cut through this mess. If everyone had their own medical kit of antibiotics at home, all they'd have to do is open it and start taking antibiotics as soon as the attack was discovered. We'd save days of delay and avoid the chaos of distribution. Even if only one-fourth of the exposed population had antibiotics, that would take a load off the distribution system. And in a pinch, people could share their antibiotics, so they wouldn't need government distribution until a week into the course of treatment. That would buy us time and ease the crisis no matter how many people had the home med kits. Not only that, it would leave people in charge of their fate. Instead of being helplessly dependent on government action, they could actively plan for and assist in the emergency.

That's also why the bureaucrats of Health and Human Services hated it. Government officials rarely doubt their own capacity to direct the lives of ordinary citizens. Doctors too seem to have vast confidence in their own judgment, at least as compared to patients. So it shouldn't be a surprise that government doctors have no faith whatsoever in the great unwashed mass of citizens. The Public Health Service has, basically, one piece of advice for the public in any health emergency: sit tight and wait for our instructions. We'll decide who should get vaccines or antibiotics, and in what order. If it's a close question, we'll send you to your family doctor, and he or she will tell you what to do. On no account should you do anything to help yourself. If you try to buy antibiotics, you'll be "hoarding" medicines that are needed more by others, like, uh, medical professionals.

When the first anthrax attacks occurred, that's exactly what government doctors said, and their guidance was posted on government and American Medical Association Web sites. Anyone trying to obtain Cipro or other antibiotics was seen as ignorant or selfish or both. In addition to the fear that medicines wouldn't be rationed in accord with government priorities, medical professionals were understandably concerned about the overuse of antibiotics, which has

encouraged the evolution of antibiotic resistance. So letting ordinary people have antibiotics in their homes was considered too risky. They might take it for a headache.

So the med kit idea met a wall of medical and bureaucratic resistance, even though both the secretary and deputy secretary of Health and Human Services eventually became supporters of the idea. Unable to defy their superiors, the bureaucrats who worked for them slow-rolled the idea. Eager to prove that you and I can't be trusted, and to wait out their bosses, they insisted on a large-scale test, putting emergency kits in the hands of citizens and telling them not to open the kits except in a government-announced emergency. I was delighted when they had to report back to the interagency that only one person had opened the kit improperly—an elderly woman who heard an official tornado emergency announcement and opened her package in the hope that it might offer some guidance.

Since the study hadn't turned out quite the way the bureaucrats expected, it was clear that what we needed was, well, more studies. The leaders of DHS and Health and Human Services pushed hard for a better set of plans to distribute med kits and use other methods to avoid the postal service option. In the month before the election, despite concerns that we'd look as though we were spreading fear, the two departments announced a number of steps that would make med kits possible. But time had run out; the efficacy of med kits was still being studied (in a Minnesota pilot project) when the Bush administration left office.

A year later, the bureaucrats won. An unimaginative bioterror strategy was released by the White House in December 2009.³ It contains an inevitable section, beloved of bureaucrats, setting out everyone's "Roles and Responsibilities." Such documents are beloved of bureaucrats because that's where all the turf wars are fought.

Now, you and your family probably didn't hire anyone to participate in those turf wars on your behalf.

Believe me, it shows.

Because when the document sets out your roles and your responsibilities (*i.e.*, the roles and responsibilities of “Individuals and Families”), here’s what it says:

There is a critical role for families and individuals in reducing the risks from biological threats. Individual contributions to community resilience can undermine motivations for biological threats by reducing their effectiveness. We will encourage individuals and families to undertake the following:

- ✦ Following general guidance for disaster preparedness, such as keeping supplies of food and other materials at home—as recommended by authorities—to support essential needs of the household for several days if necessary;
- ✦ Being prepared to follow public health guidance that may include limiting their mobility throughout the community for several days or weeks, or utilizing designated evacuation routes; and
- ✦ Informing appropriate authorities when they encounter or observe suspicious or unusual activities.⁴

This language was surely meant to resolve the bureaucratic battle conclusively against do-it-yourself preparedness. It says individuals are supposed to “follow guidance” about keeping food and other materials at home. But in case you didn’t understand the first time that you’re only supposed to do what the government tells you, the bit about keeping materials at home gets an added and quite redundant qualifier. While you’re following government guidance about keeping materials at home, remember that you’re only to keep materials “as recommended by authorities.”

And how will you get, say, antibiotics in an emergency? That shoe dropped a few weeks later. The Obama administration decided to make a big bet on the postal service’s nimbleness, sense of urgency, and dedication to duty. In a Christmas week executive order⁵, it announced plans to bet your life on the postal service having all those qualities and more.

Stop for a moment to imagine the scene. Postal workers will be asked to drive into contaminated neighborhoods even though they can't be sure their countermeasures will work against whatever strain has been spread there. The neighborhoods will be full of people desperate to get antibiotics, so for protection, the postal workers will first have to meet up with guys with guns whom they've never seen before. They also have to collect antibiotics from pickup points that they may or may not have seen before. They'll meet the guys with guns there, or someplace else that may have to be made up at the last minute. Then they'll start out on routes that almost certainly will be new to them. As they go, they will be expected to seamlessly and fairly make decisions about whether to deliver the antibiotics to homes where no one is present, to rural mailboxes that may or may not be easily rifled, to people on the street who claim to live down the way, to the guys with guns who are riding with them and have friends or family at risk, and to men in big cars who offer cash for anything that falls off the truck.

And this will put antibiotics in the hands of every single exposed person within forty-eight hours, from a no-notice standing start?

No way. It will be a nightmare. And that's not a knock on the postal service, which may, in fact, be as good a public agency as any for getting antibiotics into the hands of an exposed population.

That said, no one but an idiot would bet his life or his children's lives on flawless execution from a public agency doing something it's never done before.

So here's what I did—and what you should do, too. I asked my doctor for an emergency supply of antibiotics that would get me through the first week or so of a crisis. I promised not to take the antibiotics irresponsibly for colds or other viral infections. And I was ready to change doctors over the issue.

I got the prescription.

Some public health officials may try to make you feel guilty about "hoarding" antibiotics or contributing to antibiotic resistance. Poppycock. If you buy while supplies are plentiful, you're actually creating a bigger market for these products and contributing to the maintenance

of production capability. And if you don't take them in response to a tornado warning, you won't affect resistance.

In fact, you're being socially responsible. If we do suffer an anthrax attack and the postal service has trouble keeping up, a sure bet if ever there was one, you can defer your delivery in favor of someone who has no stash. You'll take a bit of strain off a system that is going to need all the relief it can get.

And for those who'd like to recapture their youth, in addition to the glow of virtue, you might even feel a bit of leftover sixties civil disobedience thrill. When I tried to give this home stockpile advice in a speech toward the tail end of the last administration, the lawyers at Health and Human Services told our lawyers that I'd be violating the law—because advocating an unapproved use of prescription medicine is a criminal offense under the federal food and drug laws. And, while taking antibiotics for an anthrax attack is an approved use, getting antibiotics in case of an anthrax attack is not. If the Health and Human Services lawyers were right, then this part of the book would be a felony. I think they're full of it, or I wouldn't be writing this. But if I'm wrong, well, power to the people.

The new policy is a throwback to an era of government-knows-best. There's a big role for government in countering terrorism, but this isn't it. This is like telling passengers that the best response to an air hijacking is to sit tight and wait for the authorities to arrive.

It's insufferably paternalistic. And it's bad advice.

So the bad news is that the administration isn't going to help you prepare a home med kit. No standard packaging and labels, no encouragement for doctors to prescribe the kits responsibly, no sober discussion of the risks. You're officially discouraged from worrying your sweet head about such things.

The good news is, no one will listen.

At least, not if I can help it. In fact, since no one in government has followed through on the claim that my advocacy of home med kits is illegal, you've got an easy response if government doctors try to discourage you from getting a home stash. Just tell them you're

adhering to the roles and responsibilities in the administration's bioterrorism strategy: You're keeping material at home "as recommended by authorities"—two of them, the authority of this book and of your own common sense as an independent citizen.

The other government program to thwart biological terrorism is based on the Willie Sutton principle. Sutton robbed banks "because that's where the money is." If you want to prevent the release of pathogens, probably the best place to start is where they are. And the people who ought to get the earliest scrutiny are those who have regular access to those pathogens. Because history tells us that bugs in the lab have a way of ending up in the wild.

In February 1978, Christmas break was a distant memory for the cadets of the U.S. Air Force Academy near Colorado Springs, Colorado. They were grinding their way through the bleakest stretch of the academic year. Suddenly, in less than three hours, five hundred of them had lined up outside the academy's clinic. They had the flu, and within days, three-fourths of the student body had fever, sore throats, headaches, and weakness.

Yet the faculty suffered no ill effects. They lectured to nearly empty rooms. Later, researchers pieced together the flu's origins. It was an H1N1 virus, very like one that had circulated in 1950. That explained why the cadets fell ill while the faculty did not. The older instructors had already been exposed. The younger ones had not. Still, the older faculty's resistance seemed surprisingly complete.

The reason for that soon became clear. The virus that hit the academy wasn't just similar to the 1950 version. It was identical. Now, nature doesn't usually repeat herself so precisely. But human researchers do. Many scientists think the 1977-78 influenza was released from a store of the 1950 strain—in error or otherwise. We still don't know.

Twenty-three years later, though, there wasn't much doubt that someone could release a pathogen from an existing store. According to the FBI, Bruce Ivins exploited his status as a biodefense worker at Fort Detrick in Maryland to obtain enough anthrax to kill seven people.

Fears of an inside job led Congress to adopt the “select agent” program in 2002. Its purpose was to keep the worst pathogens out of the wrong hands. It called on the Department of Health and Human Services to identify truly dangerous pathogens such as Ebola, plague, and anthrax. Researchers who wanted to work with these agents had to register their facilities, name an officer who was responsible for security, and prepare both a security and a safety plan for the agents. Those who worked with the agents had to undergo background checks; they were to be listed in a database and checked against criminal and immigration records. Foreigners who passed a background check could work with the agents if they did not come from a country that sponsors terrorism. All shipments and handling of these materials had to be tracked, and exports were subject to control.

DHS didn't exist when the select agent program was created. But we thought we had something to offer. The program was trying to solve a problem that looked a lot like the problem we faced at the border. Most lab workers, like most travelers, are entirely innocent; we want them to keep doing exactly what they're doing. So we needed a way to separate the great mass of ordinary researchers from a few risky ones. In the travel arena, the key was good data about travelers. If we knew who was coming to the United States, and we had a good idea who was risky, we could concentrate our attention on the tiny minority of risky travelers.

The same was true of researchers. In fact, that was the theory behind the select agent rules already enacted. Anyone with access to highly dangerous pathogens would be identified and investigated by the FBI. If the bureau had reason to think the researcher was a risk, access to the pathogens could be denied. But the FBI is at heart a criminal investigative enterprise. It doesn't make the kind of screening decisions DHS has to make every day at the border.

So DHS maintained electronic databases that offered up-to-date information about who was coming to the United States and who was a security concern. The select agent records, in contrast, were kept in paper files, or at best were frozen electronic pictures of documents

rather than easily searched electronic data. This meant that the FBI performed a one-time check on each researcher, using this paper record. Once that person was cleared, there was no good way to go back and look at his or her record without doing a paper search. As a result, the records simply sat in file cabinets for years. If a new fact showed up that made a researcher seem more risky—calls to a known terrorist, for example, or a decision to overstay his visa illegally—the federal government might never know that he also had access to an extraordinarily dangerous biological agent, at least not without getting out the paper files and checking names.

That didn't seem sufficient to us; we thought that researchers with access to the most deadly biological agents on the planet should get at least as much scrutiny as sleepy tourists arriving from Munich or Bangkok. We offered to put the files into a modern database or spreadsheet format so that they could be cross-checked automatically on a regular basis. We knew that even this would not be a foolproof system. A well-organized terrorist group could recruit people with clean records to work at pathogen research facilities. But it's almost always a mistake not to do something about terrorism risks just because you don't have a 100 percent guarantee of success. Terrorists are human, too. Sometimes they can be discouraged by measures that might not hold up to extended testing. And sometimes their efforts to evade and test your systems will backfire, drawing attention to the plot. The more information you have, the more likely you are to spot these efforts.

Since our approach to the problem of biotechnology involved learning more about researchers, we could expect privacy objections. But all we were proposing was to digitize records that had already been given to the government for purposes of background checks. You wouldn't think that privacy groups would object to government doing a better job with data it already had. At least that's what DHS thought. But in the end we didn't get a chance to find out how they'd react.

DHS was the new boy. The FBI and Health and Human Services had been given responsibility for the select agent program by

Congress before DHS was even created. They didn't get along particularly well, but they agreed on this much: They didn't need a third agency involved in the program, no matter what improvements the agency was willing to pay for. When we asked HHS which research labs held select agents, something we'd have to know to perform any review—or to plan a rescue if a flood, hurricane or earthquake struck the laboratory—HHS staff simply refused to provide the data. Even after the secretary of HHS twice promised our secretary that the data would be sent, his staff refused.

To justify their stonewalling, both the FBI and HHS played the privacy card. They told us they couldn't give DHS access to the background check data because, conveniently, they hadn't mentioned such information sharing when they wrote the privacy statement explaining how the data would be used. They'd have to publish a new privacy statement, then take comments on the change, then respond to the comments, they said, and maybe, maybe then, they could give us access.

We'd been down that road before. Even routine changes to a privacy statement take a year-and-a-half. And that's assuming the agency wants to make the change. If the agency didn't want to do it, the opportunities for delays and detours were endless. The FBI began the process, but I wasn't surprised that it hadn't been completed by the time we left office.

Maybe it never will be. One of the open secrets of the federal government is that privacy concerns can often be a useful way to advance bureaucratic interests without sounding parochial. ("We're not turf; we're civil libertarians.") No agency likes to share information with another. The other agency may use the information successfully but not share credit. Or it may use the information to second-guess the operations of the agency that gathered it. That's one reason the wall was so difficult to eradicate. Privacy claims simply reinforced a natural bureaucratic instinct to hold information close. In 2001, that mix of turf and privacy constraints had cost us dearly. For a while, it had receded as we counted the cost. But this was a different threat, and as we turned the reins over to a new administration, all the old instincts had revived.

And just like the fight over the wall in 2001, privacy groups had won this fight without even having to show up. The rest of us had lost.

That was frustrating; it was also just the beginning of our difficulties. The select agent program was based on an assumption that wouldn't be true much longer. Congress had assumed that we knew where the pathogens were. It hadn't prepared for a world where pathogens could be assembled from the blueprints of life.

History had already demonstrated that even the workers in government labs couldn't be fully trusted to keep pathogens under lock and key. What were we going to do when anyone with access to the DNA sequence of a pathogen could simply build it—or, even more simply, order it from a foundry?

We probably had a few years to find a way to head off this nightmare, but we needed a plan. I began to consult biotech experts, looking for someone who understood the technology, the risks, and perhaps some of the opportunities.

Craig Venter is a bald man with a beard and the tanned, bulky fitness of a sixty-year-old defying his years. He leans across the DHS conference room table as though he owns it. But the meeting isn't going quite as smoothly as Venter expected.

Venter is used to government meetings. He'd been a government researcher himself, long ago. But now he is a kind of biotech rock star, famous for sequencing the human genome in a bitter, elbow-throwing race between the National Institutes for Health and an upstart private company he created. Venter's company caught the NIH from behind, and the drama of the chase helped Venter raise a billion dollars for his company.

Venter learned then that sizzle sells, and he's a master at creating a narrative that catches journalists' imagination. In a second biotech undertaking, he sailed around the world, dipping into the ocean and parsing the DNA he found there. Now he's launched on his third—a private effort jump-started with government funds that has already

assembled nearly 600,000 base pairs to make the chromosome of a bacterium. He hopes to create an artificial organism that will make hydrogen or ethanol for industrial fuels.

If anyone represents the promise of biotech, it is Venter. He sees engineered organisms as the key to progress and riches on a vast scale. So he can't be comfortable with the theme of the meeting.

I am pressing him on risks, not promise. Venter knows more about biotech than almost anyone. If there's a way to avoid the dangers that come with democratizing genetic engineering, Venter should have it at his fingertips.

"What will stop terrorists from inventing new diseases?" I ask. Even if they're afraid of blowback that infects their supporters, plenty of pathogens affect different ethnic groups differently; and some viruses cause genetic mutations. Won't we see groups or individuals trying to engage in a kind of DIY eugenics—improving the species by killing off disfavored racial or ethnic groups or by introducing new genetic material to make future generations more peaceful and compliant?

They wouldn't even have to succeed to cause a disaster, it seems to me. A badly coded biological virus probably won't act like a badly coded computer virus. Bad computer code usually does more or less nothing. The computer's default state is inactivity. But in the biological world, the easiest way to build a new organism is to start with one that already exists and then change a few genes. That means using one that's been honed by billions of years of evolution to survive—to feed and breed at all costs. Even if the new gene turns out to be defective, the resulting organism could find a way to keep on feeding and breeding. We don't know what it will feed on or how quickly it will breed, but any surprises on this front are likely to be bad ones.

I'm thinking of what happened in 2001, when an Australian research project went frighteningly wrong. The researchers were trying to create a rodent contraceptive from the mousepox virus. They spliced a gene into the mousepox virus. They didn't want to hurt the mice, so they injected the engineered virus only into mice bred for resistance to mousepox. And, adding suspenders to

their belt, they vaccinated some of the mice for mousepox before administering the injection.

As a contraceptive, it turned out, the new virus was an overachiever. Dead mice don't have sex, and dead mice were what the virus produced. The new gene turned the formerly mild mousepox virus into a killer, overriding the genetic resistance of every unvaccinated mouse. And then it turned on the vaccinated mice, killing half of *them* for good measure. If just one researcher made just one mistake as bad as that with human subjects, I tell Venter, even nations that had stockpiled vaccines would be destroyed. How do we know, I say, that well-intentioned hobbyists, not to mention hapless terrorists, won't produce pathogens that are far more lethal and contagious than they intended?

Truth be told, this is turning into a bit of a rant, but I'm still not done. I'm not going to have another chance to get biotech advice from a rock star. Perhaps mistakes and terrorism aren't even the worst we have to fear, I offer. Computer viruses became ubiquitous only when hackers realized that they could make money from the infections. They had invented a new form of organized crime. Why couldn't the same thing happen in biotech? If we don't know who has released a pathogen, couldn't some crooked business, somewhere in the world, be tempted to design a disease, patent a cure, and then let the disease loose upon the world? Even if others suspected wrongdoing, the sick would still pay whatever it costs to get well, and with the proceeds, a company could buy a lot of protection from its government. What can we do to keep foreign businesses from trying such a tactic?

I pause. That's a lot to put on the table. But at least I've laid out all my concerns. I'm hoping Venter can see something I've missed, some reason why democratizing this technology won't ultimately empower the worst in human behavior as well as the best. Or at least some way to keep his beloved technology from putting humanity at risk.

I wait. Venter leans in, clears his throat. He smiles the winning smile that has charmed reporters and government funders for more than a decade.

"My, my, don't *you* have an imagination," he beams.

That's how it goes with many of the biotech leaders I consult. They know what the risks are. They just don't like to talk about them.

Rob Carlson is a principal at Biodesic and one of the industry's most astute observers. A physicist by training, he's spent years studying biotechnology as a business and a human undertaking.

Carlson has close-cropped hair and a genial, wonkish air. He's an eager teacher. But he grows distinctly uncomfortable when I turn the conversation to bioengineered pathogens.

Carlson wants to talk about where the industry is going. Biotech has already produced enormous improvements in productivity, he says. Drugs developed with recombinant DNA already have sales of \$65 billion a year, and biotech products already account for 2.5 percent of GDP growth. One company has modified yeast into a bug that can transform sugar into everything from malaria drugs to jet fuel and gasoline. Production will begin in 2010. And many companies expect to build bugs that can produce other chemicals out of petroleum. The chemical industry could be transformed by bioengineering, Carlson argues, but these changes cannot be achieved without making the tools for bioengineering cheaper and more efficient.

So, cheaper they will get. And bioengineers everywhere will benefit. Already, the foundries that assemble small bits of DNA into large stretches have been driven by competition into fully automating the process from code to gene sequence. Even so, the biggest bottleneck in industry is the time engineers spend waiting around for foundries to send back the sequences they've ordered. The engineers don't want to wait. Carlson thinks the chemical industry's need to experiment quickly with many different genes and organisms will continue to force the pace of automation until the process can be performed in a single machine that can be run by the engineers on premises. That machine will grow cheaper and smaller at an exponential rate because of the returns and the integration of semiconductor processes. The result will be desktop DNA synthesis, Carlson predicts, and perhaps very soon.

When that happens, he sees a golden age of bioengineering. Bugs will eat our waste—literally, feasting on municipal sewage—producing

raw materials that other bugs will turn into plastics and chemicals. Energy independence may come to any nation with modern sewers. The opportunities are astonishing.

I interrupt. Yes, I know. Biotech is irresistible. But that desktop DNA synthesizer—who's going to use it besides chemists? What about all the bad things that will come from putting this power into everyone's hands?

Carlson blinks. Well, sure, there could be bad things. Terrible things, maybe. But with technology like this in our hands, we can devise countermeasures faster and make them more effective than we ever dreamed possible. A revolution is coming. Why do you insist on looking at the downside?

He pauses and returns to the emerging economic opportunities. The industry is already global, and the business logic of bioengineering is already established. It's a fantastic new technology that will transform our lives for the better. Surely we'll be able to handle the risks in that transformed world.

After all, I think, who wants to be the voice of doom when everyone else is hoping to be the Steve Jobs and Steve Wozniak of biotech, playfully hacking genomes and starting a global empire in the garage?

Silicon Valley and the computer revolution is exactly what Rob Carlson and the rest of his generation hope to emulate. A growing "DIY bio" movement shows bio-hackers how to extract and modify DNA on their own, using household equipment. There's a *Biotech Hobbyist* magazine with a "series that will show you how to grow your own skin culture and suggest some very cool projects you can do with it."

There's even a biotech version of the Linux open source operating system. "Biobrick" prizes are awarded to teams that create standardized open-source DNA parts that perform predictable biological functions and can be combined in new ways.

Today, colleges hold lighthearted competitions for the best biological design. MIT's winning team in 2006 re-engineered *Escherichia coli*—an organism that lives in the human gut and helps to give our waste its distinctively foul smell. When the students were done, the redesigned *E. coli* smelled like wintergreen.

Biotech: it's cute, it's fun; and then you get rich.

I remember when the computer software geeks first came to Washington in the early 1990s. They were shocked to hear that the government wouldn't let them offer strong encryption to the world. The government feared that unbreakable encryption would allow criminals, terrorists, and pedophiles to hide evidence and communicate without fear of wiretaps. The technologists dismissed the fears. Encryption would be necessary to do business on the Internet, a development that was inevitable, they said, sounding a lot like Rob Carlson. Government would just have to get out of their way.

Carlson and other biotech industry representatives have none of the software industry's in-your-face contempt for government. After all, many of them are funded by NIH and hope to develop treatments that will pass muster with the Food and Drug Administration. Instead of defiance, they offer deflection, simply gliding past the risks and averting their gaze. It's the way most of us deal with the animal experiments that make new drugs possible: They're unfortunate, tragic even, but that's the price of progress; now, can we talk about something else, please?

Sixty-five years ago, with a bright flash and a mushroom cloud, the nuclear age was born in the New Mexico desert. Robert Oppenheimer was a prime mover in the first nuclear test, and he later told how the scientists reacted:

We knew the world would not be the same. A few people laughed. A few people cried. Most people were silent. I remembered the line from the Hindu scripture, the *Bhagavad Gita*. Vishnu is trying to persuade the prince that he should do his duty and to impress him takes on his multi-armed form, and says, "Now I am become death, the destroyer of worlds." I suppose we all thought that, one way or another.⁶

Nuclear technology came into the world burdened by a sense of original sin. Before it became a source of cheap, carbon-free energy, it would kill and wound two hundred thousand people in Hiroshima

and Nagasaki. For nuclear scientists even their most satisfying work was alloyed with tragedy.

It's a long way from that sober sense of guilt to the spirit that gave the world *E. coli* that smells like wintergreen. That's because, with nuclear technology, the deaths came first. With biotech, as with jet travel and computer networks, it's the delight, and the profits, that have come first.

It's odd. No one in the industry denies the risks, and some can be eloquent about the need to address the problem. But a curious disconnect remains between their intellectual acceptance of the danger and their response to it. At a visceral level, many of the biological and medical researchers who are leading the revolution simply cannot believe their technology may end up causing more harm than good. Some of them seem convinced that doctors, or at least medical researchers, just aren't the kind of people who would do such a thing. And so they fight restrictions on their work with the fervor of men and women who are determined to make the world a better place—no matter what the bureaucrats say.

DHS had no authority to force the foundries to screen their orders. Many of them were overseas, and none were subject to direct regulation. But we decided to press them anyway. We might not have regulatory authority, but we could make noncompliant foundries uncomfortable. We met with some of the DNA synthesis companies and told them they had a responsibility to prevent misuse of their products. They should know each customer and whether the customer was a legitimate business. And they should make sure the string of code they were building was not dangerous—the string of code that gives a pathogen its virulence, say, or the insertion of a toxin into the gene for an edible plant. If they got a suspicious order, they should report it to the government.

The purpose of this screening wasn't just to keep terrorists from building pathogens. We were also thinking about attribution after an attack. If we are attacked with an agent that might have been engineered, we will quickly find the resources to review

every synthetic DNA order in recent years—and to interview every purchaser whose orders resemble the pathogen. But if the foundry doesn't keep records, we can't review them later. Quickly identifying the attacker is one of the great challenges of biological terrorism; if we can do that, we will deter many future acts and we will reassure our citizens that their government is not helpless in the face of what could be a devastating attack.

Measured against the horrors and risks that come with exponential biotechnology, that may not seem like much of a response. But it was a start; it reflected a core strategy of expanding the information needed to identify risky people, either before or after an event. And if it seems like too little too late to you (as it does to me), there were plenty of officials who were prepared to fight even these modest steps.

Some of the American and European foundries were responsive. A few had already begun screening customers and keeping records. They were in business for the long haul, and they couldn't afford to acquire a reputation for irresponsibility. That was worth something, but if other foundries refused to screen orders, then we'd just be moving the risky customers to the irresponsible suppliers.

DHS's proposal to press the foundries to engage in screening met with a tepid reaction at the lower levels of HHS. The NIH, in particular, was so sure that basic research in biology was a boon to mankind that it refused even to keep track of who was accessing the research on dangerous pathogens that it published on the Internet. Researchers who blithely published work that could be used both for weapons development and energy production would be widely condemned as dangerously irresponsible; but unrestricted publication of biological research is still an article of faith, even though such research can also be used both for commercial and military purposes.

Only after members of the industry and two independent biosecurity boards had made similar recommendations did NIH agree in principle to do something about foundry screening. NIH proposed to tell its grantees that they should send orders only to foundries that engaged in screening.

For other countries, controlling biotechnology was simply not on the agenda. Biotech expertise had spread throughout the world. Nations that missed the information technology boom were rushing to stake a claim in the next hot field. Commercial DNA foundries can be found in California, New York, and Massachusetts, of course, but also in Pretoria, Moscow, Dalian, and Tehran. Where we saw a global risk requiring oversight, these capitals saw a chance to catch and pass the United States in the exploitation of biotechnology. They still chafed at the role that Intel and Microsoft played in information technology. Why couldn't the Microsoft of biotech be Chinese or Singaporean or Dutch, they asked? If the United States wanted to hobble its researchers with elaborate restrictions, well, fine. That was an opportunity not to cooperate with the United States but to steal a march on it.

If pressed for cooperation, international diplomats argue that the key is enforcing the Biological Weapons Convention. This is an example of just how wedded to the status quo international diplomacy can be. The Biological Weapons Convention is modeled on treaties to control nuclear weapons that can trace their roots back to the 1940s, when U.S. policymakers hoped to move from nuclear weapons to the peaceful production of nuclear power. The nuclear weapons convention adopted in the 1970s seeks to follow the same pattern; it offers a simple bargain to countries that lack nuclear weapons: Abandon military use of nuclear technology and the countries that have weapons will teach you how to use nuclear technology for peaceful purposes. Every five years, the nuclear haves and have-nots get together in Geneva. There, the have-nots press the haves to abandon nuclear weapons before they get down to the less high-minded task of demanding more aid and more technical assistance in using nuclear technology.

The Biological Weapons Convention more or less borrowed the same model when it was adopted in the 1970s, even though it was never a good fit. The nuclear convention makes at least some sense because there is a vast difference between building a nuclear power plant and building a nuclear weapon. Information about the peaceful

uses of nuclear technology is not easily used in a weapons program. So it's possible to transfer peaceful-use technology without dramatically increasing the risk of weapons proliferation.

That's not true for biological technology. There's no real difference between a bioengineering facility meant to cure disease and one meant to cause it. Facilities can be switched from one purpose to another with little more than a long weekend and a few gallons of bleach. Inspections to catch cheaters would have to be deeply intrusive, could easily become a cover for the theft of intellectual property, and would almost certainly fail to catch countries that were serious about maintaining an illicit program. The advent of synthetic DNA, with its radical empowerment of all researchers, makes the model even less relevant.

If ever there were a doubt about the dysfunctional conservatism of international forums, the persistence of the Biological Weapons Convention surely should put an end to it. The risks of biotech are novel and pressing. But the solution posed by internationalists is to draw on a model that was adopted for nuclear weapons in the 1970s and hasn't been a notable success in the forty years since. Finding a new response to a new problem seems to be simply beyond the capability of the international community.

In short, we were on our own. DHS kept pressing for action on foundry screening. A year after we left office, five of the biggest DNA foundries agreed on a common screening protocol that they would apply to every synthetic gene order; they also agreed to keep customer records for eight years.

This was progress, if it actually survived scrutiny by the European privacy bureaucracy. (European members of the group did not explain how they would square this new practice with the EU requirement that order data be destroyed when no longer needed for commercial purposes.) But at best, it covered only 80 percent of the foundries by market share.

Domestically, in 2009, HHS issued voluntary guidelines meant to encourage and set standards for screening of foundry orders. But

the incentives to follow the guidelines remained limited. Exponential growth in the market has made NIH's standards less important. Today, NIH grantees probably account for no more than 10 percent of the foundries' business. Foundries that find the standards constraining can simply limit their sales to customers who aren't using NIH money. And if the United States tries to make the rules mandatory, they can take their facility elsewhere; biotech firms are likely to be welcomed in other countries with open arms and less demanding laws.

In a globalized world, where regulations may be put on the block to get an edge in the international competition for new industry, is there any way to prevent a race to the bottom on synthetic DNA? Perhaps, but only over the opposition of privacy, business, and other governments. If the United States really wants to ensure that biotechnology researchers and developers meet biosafety and biosecurity standards, it can use the one piece of government leverage that still counts in that world.

For biotech firms, the road to riches is intellectual property. A patent entitling firms to a royalty on the exploitation of some new biotech technique or drug is the key to most startups' business plans. And U.S. patents are particularly important because, in the absence of government medical price controls, the U.S. market probably pays a disproportionate share of the development costs for new drugs.

If all companies seeking patents derived from biotech research were required to demonstrate compliance with reasonable safety and security measures, the requirements would likely be observed globally, since even companies located in deeply hostile nations, such as Cuba, have sought U.S. patents for their research. (Despite sanctions and a bitter war of words between the two countries, Cuba has been granted more than seventy-five U.S. patents in the last thirty-five years.)

Of course, the governments that would be bypassed by such a measure can be counted on to protest, as will the business interests

that want intellectual property protection without regard to their security record. And, since the most obvious biosecurity measures include detailed records of who is performing what kinds of research, we can expect other nations and the business community to cloak their interests in a cloud of privacy objections.

Requiring biotech companies to demonstrate that they have met biosecurity standards in order to get patent protection might well work, but it's guaranteed to trigger hostility from business, privacy, and international interests, and that's why it probably won't happen, at least not until the ever-steepening curve of biotechnology produces a disaster.

