Testimony

It Takes More than Offensive Capability to Have an Effective Cyberdeterrence Posture

Martin C. Libicki

CT-465
Testimony presented before the House Armed Services Committee on March 1, 2017.

Published by the RAND Corporation, Santa Monica, Calif.

It Takes More than Offensive Capability to Have an Effective Cyberdeterrence Posture

Testimony of Martin C. Libicki[1]
The RAND Corporation[2]

Before the Committee on Armed Services
United States House of Representatives

March 1, 2017

Good morning, Chairman Thornberry, Ranking Member Smith, and distinguished members of the committee. My name is Martin Libicki; I hold the Maryellen and Richard Keyser Chair of Cybersecurity Studies at the U.S. Naval Academy, and am also an adjunct management scientist at the nonpartisan, nonprofit RAND Corporation. The following represents my own viewpoint and not the viewpoint of the U.S. Naval Academy, the federal government, or the RAND Corporation.

I thank you for the opportunity to testify today about some issues associated with deterrence of cyberattacks.

Two years ago, Admiral Michael S. Rogers, commander of the U.S. Cyber Command, argued in Congressional testimony that he needed a greater ability to conduct offensive cyber operations, stating that this greater ability was needed to be able to deter cyberattacks against the United States.[3]

Clearly, greater capability would not hurt—but would it do much to help achieve deterrence, much less completely deter attacks?

A successful posture of deterrence—that is, the use of threats to compel others to restrain themselves—has many prerequisites. Four of them merit note. First, the United States has to be able to *attribute* cyberattacks in order to punish the correct party and convince others that the punishment is justified. Second, the United States needs to have and communicate its *thresholds*—that is, which actions will lead to reprisals. Third, U.S. promises to retaliate need

[1] The opinions and conclusions expressed in this testimony are the author's alone and should not be interpreted as representing those of the RAND Corporation or any of the sponsors of its research.

[2] The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest.

[3] Ellen Nakashima, "Cyber Chief: Efforts to Deter Attacks Against the U.S. Are Not Working," *Washington Post,* March 19, 2015.

credibility—so others believe that punishment will, in fact, follow crossing such thresholds. Fourth, the United States needs the *capability* to carry out reprisals.

There are also other considerations, but they are not prerequisites as such. One is that carrying out reprisals affects the *broader* relationship between the United States and the attacking country; there may be larger issues in the ongoing relationship that may modulate or exacerbate the reprisal, in turn affecting the credibility and even legitimacy of the threat. For instance, however annoying the Iranian distributed denial-of-service (DDOS) attacks on U.S. banks were in late 2012, efforts to halt Iran's nuclear program clearly had higher priority. Therefore, had reprisals been on the table, their impact on such efforts had to be taken into account. Another consideration is the extent to which the attacker feels justified in its original cyberattack (which may have been prompted by a perceived injury). This, in turn, will color the attacker's view of the legitimacy of U.S. reprisals— which, in turn, may influence the likelihood of counter-reprisals.

Returning to the prerequisites, the U.S. *capability* to retaliate in cyberspace is least in doubt among the four (even if United States need not respond in kind, Admiral Rogers' argument assumed that we needed to be able to do so). Any country credited with Stuxnet and the system penetration techniques described by Edward Snowden has demonstrated a very impressive capability (whether or not the credit is deserved is secondary).[5] As long as other countries believe we can do magic, what we can *actually* do matters less for deterrence purposes. That noted, however, countries vary in their susceptibility to cyberspace reprisals. North Korea is a good example, as its economic primitiveness and paranoia about the outside world mean that computers and connectivity are far less important to national well-being than in other countries. Note that susceptibility consideration had only a modest effect on the efficacy of the nuclear deterrent. Furthermore, while the U.S. attention to the laws of armed conflict (specifically *jus in bello*) is laudable, the effect of following them is to take certain targets off the list. Such prohibitions loom larger if people are worried that cyberattacks on some targets may yield unacceptable collateral damage. Lastly, for those who believe that reprisals delayed are reprisals denied, note that even a very capable United States is limited in its ability to respond from a cyberattack from a country that it did not consider a threat and whose systems it did not scope in advance. Otherwise, U.S. capability is more than sufficient for purposes of reprisals.

The other three prerequisites are what hobble the ability to develop a coherent deterrence policy.

Attribution, to be fair, has improved considerably over the past ten years. Roughly a decade ago, difficulties in attribution were recognized as an important barrier to establishing a deterrence posture, and considerable time and attention was invested in improving the intelligence and science behind attribution. By late 2012, the Secretary of Defense was able to claim that two-thirds of all incidents could be traced. Furthermore, several private cybersecurity companies—starting most publicly with Mandiant in early 2013—started making their own

[5] In the last year, Israel has publicly declared that it and the United States together authored Stuxnet. Belfer Center, "Deterring Terror: English Translation of the Official Strategy of the Israel Defense Forces," August 23, 2016, p. 48.

attribution claims.[7] This allowed the U.S. government to make cases against other countries without having to reveal its own sources and methods (even if some government officials believe private attribution claims force their hands when the evidence is less-than-overwhelming or decisions on reprisals need time to make correctly). Although the consonance between intelligence community knowledge and private cybersecurity claims is less than perfect, the two efforts remain complementary. It is quite plausible that China's perception that the U.S. ability to attribute acts of economic cyberespionage to the Chinese sufficed to inhibit further economic espionage after the Xi-Obama agreement to foreswear such activity.

Nevertheless, a few cautions are in order.

First, the ability to attribute and the ability to evade attribution are a measure-countermeasure game. Until the consequences of being caught are severe enough, hackers may simply not feel that they have to hide their origins (as opposed to their tracks) very well. Yet, if the point of having a deterrence policy is to inhibit cyberattacks, then presumably consequences have to be severe. If the prospects of reprisals are daunting enough, hackers can be expected to take pains to keep from getting *caught* carrying out cyberattacks. Hence countermeasures to attribution can be expected. Another way of putting it is that attribution will be good until it becomes useful, at which point it will cease being good.

Second, the U.S. government has made less progress in *explaining* why it believes its attribution is correct. After the Sony attack, the publicly released Federal Bureau of Investigation (FBI) statement on North Korean attribution devoted just 140 words to justifying its conclusion.[8] The public justification of Russian attribution for the Democratic National Committee (DNC) hack is even more problematic. The two public documents released on the matter—one by the Department of Homeland Security (DHS), the other by the Director of National Intelligence (DNI)—were generally deemed far from satisfactory.[11] Granted, it may not be obvious why the United States has to convince others that it is right about attribution; by this argument, as long as the attacker knows that it could get caught and punished for what it did—and knows it did—then the opinion of third parties is irrelevant. But is it? To skeptics, U.S. retaliation against a country that could be innocent may strike them not as punishment but aggression. Worse, if potential attackers come to believe that innocence is no guarantee against reprisals, what is the point of being innocent? The accused country could easily maintain its innocence, and having done so credibly (for lack of a good case against it), could justify its responding to retaliation as if it were responding to unprovoked aggression. Thus, what started as an attempt to make other countries conform to standards of responsible behavior becomes a tit-for-tat exchange, where no one can easily claim the high ground.

[7] Mandiant, "APT1: Exposing One of China's Cyber Espionage Units," March 2013.

[8] FBI, "Update on Sony Investigation," December 17, 2014.

[11] DHS National Cybersecurity and Communications Integration Center and FBI, "GRIZZLY STEPPE – Russian Malicious Cyber Activity," December 29, 2016; Office of the DNI, "Background to 'Assessing Russian Activities and Intentions in Recent US Elections': The Analytic Process and Cyber Incident Attribution," January 6, 2017.

Credibility also remains an issue when it comes to *cyber* deterrence. Put simply, the United States has yet to retaliate to any cyberattack with truly serious consequences that the rest of the world can see.

The U.S. retaliation against North Korea involved sanctions on a handful of individuals. The only quasiserious response was a DDOS attack on North Korea's thin Internet connection to the rest of the world—and the United States, if anything, distanced itself from taking credit for that act.[12] There are reports that the United States carried out reprisals against North Korea that did not make the news; although I have no way of evaluating that claim, suffice it to say that hidden reprisals lack effectiveness in persuading *other countries* of the folly of carrying out cyberattacks on the United States.

The United States also retaliated against Russia for the DNC hack by increasing some sanctions and expelling some Russian diplomats; there may have also been reprisals not visible to the public. Since the Russians probably believe that their contribution to defeating a presidential candidate they disliked exceeded the pain of having to replace a few diplomats, it is difficult to see how the threat of such future punishment would deter them. Does anyone think the Russians will hereafter refrain from injecting themselves into other countries' elections? And what does it say for the credibility of the U.S. Government when representatives of an incoming administration delegitimize the reprisals levied by an outgoing administration?

After two weak *public* responses, the credibility of U.S. reprisals cannot be ranked very high. Perhaps the failure to respond with anything harsher was wise, given the relatively limited harm associated with the Sony and DNC hacks and the possibility that a major confrontation would have raised much higher levels of risk. But it would now take a serious response to raise the credibility of a *possible* U.S. response off its current floor—and several serious responses to convert the possibility into a likelihood. These hypothetical responses to as-yet-potential cyberattacks would carry their own risks. Put another way—if the United States wanted to achieve credibility for a cyberspace deterrence policy, the costs of doing so would not be small at this point.

That leaves *thresholds*, which I want to focus on, partly because they seem to get the least attention. Here is the relevant question: Which cyberattacks merit cranking up the machinery of U.S. retaliation? The term "machinery" is deliberately used: The decision on whether and how to retaliate would certainly involve the President and the National Security Council, and it would have to be followed up by policy adjustments throughout the bureaucracies to reconcile retaliation with other actions involving the attacking country. Retaliation, after all, is an unfriendly act. By contrast, foreign individuals can be indicted in U.S. court as cybercriminals, based on decisions taken at the level of a U.S. district attorney and without much reference to the U.S. relationship with their countries of origin. Although the indictment of five members of China's People's Liberation Army and seven Iranian nationals doubtless required greater

[12] See Nicole Perlroth and David Sanger, "North Korea Loses Its Link to the Internet," *New York Times*, December 22, 2014. However, two weeks later, sanctions were described as a "first response," suggesting that the DDOS attack was not a U.S. response (BBC, "Sony Cyber-Attack: North Korea Faces New US Sanctions," January 3, 2015).

coordination, these moves were announced by someone no higher than an assistant attorney general.

The need for a threshold is obvious. Objectionable acts in cyberspace range from a network hiccup to a major catastrophe. Not all of them merit Presidential attention. By contrast, in the nuclear realm, even the detonation of the smallest nuclear weapon on U.S. soil would be catastrophic.

Finding a defensible threshold is, alas, a problem not easily solved. Let's consider some definitions that have been bruited about.

Perhaps something is actionable if it violates the U.S. Computer Fraud and Abuse Act (CFAA). Three problems arise. First, using a national law as a red line sets a precedent that can be easily abused by countries whose laws criminalize behavior that is acceptable, even normal, in the United States, such as posting material critical of the government on the Internet. In other words, if we use our domestic laws as a basis for international reprisals, what keeps others from using their domestic laws in the same way? Second, the CFAA has been violated literally millions of times—notably, every time a computer is infected as part of an effort to build a botnet, or every time some teenager wants to go exploring in someone else's machine. Third, such a law makes cyberespionage generally actionable, but the United States relies on cyberespionage techniques to protect itself from terrorists and hostile countries. Another good reason not to make all cyberespionage actionable is that cyberespionage penetrations can often go undetected for months or years, unlike disrupting operations or corrupting information, which are harder to hide. The less likely a violation is to be caught, the more problematic it is to punish violations that are found.

Another alternative threshold is to use a metric of size to determine whether a cyberattack is actionable. As one Assistant Secretary of Defense has argued, the United States cares primarily about the top 2 percent of all cyberattacks.[13] The problem with that formulation is that the criterion for membership in the set of cyberattacks has no obvious lower bound. Two percent of something unmeasurable is itself unmeasurable. Insofar as the effects of cyberattack can almost always be measured in terms of dollars, an economic threshold might make sense—until it comes time to measure impacts. If Sony's statement to the Securities and Exchange Commission is indicative, the North Korean attack cost only $35 million (in the financial quarter that took place plus the quarter afterwards). However, that dollar metric may not capture intangible costs, such as damage to the reputation of Sony's executives, the hassle of shifting communications from email to phones, and anxiety among employees. Furthermore, the administration justified its responses to the Sony and DNC attacks not with economic criteria, but by appealing to transcendent values—the attack on Sony contravened freedom of speech, and the attack on the DNC contravened U.S. political sovereignty. Meanwhile, there was no U.S. response to the Iranian attack on Las Vegas Sands Corporation, which wreaked damages approximately as large as those suffered by Sony.

Another criterion for judging a cyberattack actionable is if it hurts some part of the U.S. critical infrastructure. One would think such a threshold had sufficient clarity, since the key

[13] David Sanger, "Pentagon Announces New Strategy for Cyberwarfare," *New York Times,* April 23, 2015.

elements of that infrastructure had been publicly enumerated by DHS (admittedly in response to physical terrorism, which generates a somewhat different list than a focus on cyberspace would). But following the attacks on Sony and the DNC, some have tried to stretch the definition to include such attacks. There were desultory attempts to note that, technically, Sony Entertainment was part of the U.S. critical infrastructure, but they were not taken seriously.[14] The DNC hack, however, did persuade the government to declare the U.S. election system to be critical infrastructure, and properly so.

Perhaps a criterion is needed that offers a parallel with physical attack. Perhaps then, something is actionable if it violates the laws of armed conflict, specifically *jus ad bellum*. The law of armed conflict has the benefit of being established international law. But the various laws of armed conflict, having been established for physical combat, focus on destruction and injury. They do not cover economic loss from hostile activity, perhaps because one country can make many types of decisions that cost other countries money without using force at all. In the decades-long history of cyberwar, physical destruction has occurred twice: with the use of Stuxnet, and during a putative Russian cyberattack on a German blast furnace. (In many other cases, attackers altered information in machines, making the machines unusable until reformatted, but that is not physical destruction).[15] No one has yet been harmed as a direct consequence of a cyberattack. Instead, the effects of cyberattacks are usually felt in terms of lost productivity: e.g., time lost restarting systems or recovering data. It is unclear whether an attack that, say, bankrupts a trading house would be actionable by such criteria—and a willingness to declare cyberattacks actionable after the fact is not a basis for deterrence.

To complicate matters further, the reliance on precedents like the laws of armed conflict fosters the notion that cyberattack, like physical attack, is actionable, while cyberespionage, like non-cyberespionage, is acceptable. But accepting *all* cyberespionage as acceptable state behavior is *not* U.S. policy. When the United States successfully pressed China to stop its economically motivated cyberespionage, it established a norm that was adopted by the G20.[17] Given the G20's membership, this is now close to a universal norm. If the information taken from the Office of Personnel Management (OPM) had been sold into the black market—the possibility of which was implied by OPM's offer of credit-monitoring services to potential victims—then it is quite plausible that the United States would have strongly objected that the acceptability of cyberespionage did not imply the acceptability of every use of taken information. Fortunately, there is scant evidence that such information was transferred to criminals. Lastly, it helps to remember that the DNC hack was actually cyberespionage, the results of which would not have led to a U.S. response if the Russians had kept what they took to themselves, rather than using it to influence the outcome of a Presidential election.

[14] Kim Zetter, "Hacker Lexicon: What Counts as a Nation's Critical Infrastructure?" Wired, February 16, 2016

[15] Kim Zetter, "A Cyberattack Has Caused Confirmed Physical Damage for the Second Time Ever," *Wired*, January 8, 2015.

[17] For a copy of the communique and a discussion thereof, see Cody Poplin, "Cyber Sections of the Latest G20 Leaders' Communiqué," *Lawfare*, November 17, 2015.

These three examples may not be the only occasions where cyberespionage rises to the point where it is as obnoxious as cyberattack. It is very difficult to distinguish between cyberespionage against a system and the preparations made for a cyberattack on such systems. In some cases, the motivation for cyberespionage is so plausible that countries caught penetrating systems with valuable information can be assumed to have done so to gain information, not to disable or destroy systems. But it may be hard to give others the benefit of the doubt when they are caught carrying out cyberespionage against certain elements of a country's critical infrastructure—notably, the machine-control systems associated with transportation, energy production, distribution, and manufacturing in general. The information such systems contain is of modest value in and of itself, but the potential for its use in mischief is substantial. Here, too, certain types of cyberespionage may be plausibly deemed actionable if detected, characterized, and attributed.

In the face of these many issues, ensuring that countries do not convince themselves that there is a threshold below which that they can operate with impunity entails deliberately maintaining a threshold so low that the United States can afford to be indifferent to cyberattacks that fall beneath that level. This is hardly a panacea. First, it forces inordinate attention to above-threshold attacks, even if they are low level, because a failure to respond erodes the credibility of the U.S. promise to respond (although for some observers, the failure to respond will only erode their belief that the stated threshold is the real one). Second, if there is no difference between responses to low-level and high-level attacks, potential attackers may reason that if they are going to get caught and punished (again, no sure prospect), they might as well try to achieve a bigger attack. Third, too low a threshold, coupled with a fixed minimum cost associated with cranking up the machinery of retaliation, may strike others as disproportional, expensive, and even arbitrary.

A broader issue in all this is whether any country, even the world's most powerful, can arbitrarily establish redlines, as opposed to first creating some consensus on norms to achieve deterrence. To be fair, redlines are not the worst option; at least they have the advantage of needing to be declared beforehand. One of the problems with responding to the DNC hack—apart from its inherently political nature—was that few anticipated that the United States would need to declare against other countries hacking political organizations, extracting their emails, and posting them online. To react to injury solely after the fact assumes that a reasonable presumption could have been made by the attacker that something so injurious could not go unanswered. Such thinking is far from easy even in the physical domain, where precedents to almost every conceivable action abound. In the cyber domain, such precedents are absent, and the best one can resort to are inexact analogies between something that has merited objection in the past and some objectionable act in the present. Deterrence, after all, only works when the potential attacker knows *in advance* where the redlines are, at least approximately. A country's willingness to respond based on *post facto* redlines presupposes the willingness of others to give the aggrieved country a wide berth.

Redlines have had their place in U.S. history; the Monroe Doctrine, which stated U.S. intolerance for any establishment of new colonies in the Americas, could not possibly have been a norm. It was geographically delimited to one hemisphere, and the prevailing norm in those days actually allowed colonization in general. Russia's concern over activities in its near abroad, or China's concern over activities within its self-defined first island chain, are also geographically defined. But cyberspace, as oft observed, does not have the same geography and, to an important extent, has no geography at all. Therefore, redlines cannot be stated in geophysical terms very easily—and the justification that redlines are needed to defend the *physical* basis of a country's sovereignty does not apply.

Redlines and norms differ in several key respects. A country can establish redlines without having to abide by them; when a country establishes exclusion zones for others, it hardly signals its intention to exclude itself. But a norm implies mutual constraint. Every UN member, by dint of its membership, has pledged adherence to norms against carrying out an armed attack on others. Clearly, redlines are less constraining than norms—but that may be exactly why arbitrary redlines sit poorly with long-standing U.S. ideals.

At issue is how rules should govern the world. Until the mid-20th century, international relations could be said to be taken from Thucydides' Melian Dialogue: The strong do as they will, and the weak suffer what they must. Redlines bespeak a world in which strong countries—and the United States is the strongest—can compel others to live by their rules, even if they have no intention of living by such rules themselves. But U.S. leadership in the post-war era allowed a different notion to take root. International stability and world peace result when everyone follows the rules, just as domestic stability and safety follow when everyone obeys the law. To achieve legitimacy, that meant that the United States and its friends had to obey the same laws. And much of the history of the Cold War was an attempt—one that was largely successful—to define these laws and use the muscle of the United States and its allies to see that such laws were largely obeyed. The end of the Cold War made that task easier and spread the rule of law wider, but the effort remains nontrivial.

This theoretical difference has a practical consideration. Reconsider the OPM hack. Should the United States have responded? The attack transferred information of great value to China. It embarrassed the U.S. government. U.S. officials were angry with the Chinese, and there is evidence that Chinese officials were at least somewhat abashed at having been associated with the hack (they subsequently announced an arrest for having carried out the hack).[18] But the DNI and a former Central Intelligence Agency director admitted that what the Chinese did was something that the United States would have done if it could have (and it may well have done similar things).[19] The United States could easily declare that it would regard a repeat as having

[18] Ellen Nakashima, "Chinese Government Has Arrested Hackers It Says Breached OPM Database," *Washington Post*, December 2, 2015.

[19] "Don't blame the Chinese for the OPM hack," former NSA and CIA Director Michael Hayden said, arguing that he "would not have thought twice" about seizing similar information from China if he had the chance. (Matthew Ferraro, "On the OPM Hack, Don't Let China Off the Hook," The Diplomat, July 14, 2015,). Director of National Intelligence James Clapper echoed the sentiment, saying at a conference, "you have to kind of salute the Chinese for what they did. . . . If we had the opportunity to do that [to them], I don't think we'd hesitate for a minute."(Jim Sciutto, "Director of National Intelligence blames China for OPM hack," June 25, 2015; http://www.cnn.com/2015/06/25/politics/james-clapper-china-opm-hacking/.).

crossed a red line; it might even be able to enforce its dictum. But if the United States would not foreswear doing likewise, it could not argue that a repeat would have violated a norm. One of the reasons that the United States could persuade China to abjure economic cyberespionage is that it could make a reasonable case that this was behavior that the United States would not conduct—and, indeed, had not conducted (or at least no one has proved the contrary). By the same token, one of the difficulties of dealing with Russia's politically motivated cyberespionage/doxxing was the lack of a norm that made it easy to argue that such activity was out of bounds. Because countries, even the United States, seek to influence the elections of other countries all the time, mere unwarranted influence is a poor guide to norms-writing—but a norm condemning the use of cyberespionage coupled with doxxing (for political ends) would be more precise and consistent with U.S. behavior.

A norms-based deterrence posture has its issues. One is determining how large a consensus is required to establish a norm. One advantage of working from the UN Charter is that UN membership is universal. However, translating the words of the charter into the new fields of cyberspace is hardly obvious. The European Convention on Cybercrime, aka the Budapest Convention, counts almost every advanced country as a signatory—but Russia is not one. Treating, for example, Russian sanctuary for major cybercriminals as an actionable violation of universal norms is an iffy proposition. Conversely, waiting until North Korea signs up to norms before deeming them universal means waiting indefinitely. A best guess is that a norm can be deemed universal if it wins adherence from either Russia or China. The other issue is holding others to norms. A country that has declared a redline has put the onus on itself—and only itself—to respond to a violation. Responding to a norms violation, however, is a collective responsibility. This can be positive, because many countries join together in responding, but there is a negative side, as each country can shift enforcement responsibility to the others. In the past, it has fallen to the United States to enforce norms of international behavior, picking up other countries as active allies or passive supporters as politics dictated. But it is fair to note that despite the lip service that the United States pays to its mutual-defense alliances, it is more likely to react to a cyberattack on itself than to an ally. The best indicator of future response comes from comparing the U.S. response to the Sony attack to the U.S. nonresponse to a longer series of more damaging incursions into South Korean systems.

Conclusions

Using the threat of reprisals to dissuade cyberattacks introduces multiple issues that need far more careful attention than they have received to date. The notion that having the best offensive capability suffices for deterrence is simplistic, to say the least. Granted, weak countries cannot deter, and in there is a basis for Admiral Rogers's argument. But the United States is by no means weak, especially in cyberspace. If the U.S. deterrence policy has problems, they are not ones of weakness but wisdom, notably in determining where to draw the line between cyberattacks that are actionable at the national level and those that can either be ignored or responded to via judicial processes.

In the interim, we should understand that certain types of cyberattack—e.g., attacks that plunge the country into a blackout—clearly cannot go unanswered, while other ones are simply too trivial to bother with. It is the in-between that is the problem. As a general rule, the United States should develop its thresholds by working towards a regime of norms, creating a consensus on the difference between acceptable government actions and those that are unacceptable and actionable.

I appreciate the opportunity to discuss this important topic, and I look forward to your questions.