STRATEGIC STUDIES QUARTERLY

Expectations of Cyber Deterrence

Martin C. Libicki

Abstract

Cyber capabilities and national attitudes toward their use continue to evolve. As long as other countries understand that the United States is capable of impressive cyberspace operations, then the threat that it will use them in reprisal is inevitably part of the US deterrence package. The policy question is whether the United States should emphasize the possibility by impressing adversaries with what cyberattacks can do, reminding adversaries that the United States would be willing to do it, and investing in making cyberattacks more reliable and even more painful. Two factors, though, vitiate cyberattack as part of an overall deterrence menu—especially compared to similar kinetic threats. One is the difficulty of finding acts we wish to deter against for which a cyberattack is a just-right response: neither too weak nor too harsh. The other is the great uncertainty associated with cyberattack and thus the great difficulty of making the threat dissuasive.

✵✵✵✵✵

While the United States may seem peerless when measured in terms of its military, economic, financial, or cultural power, this does not prevent other countries from challenging US interests. If this is to change, the United States must be ready and willing to impose costs on those who work against those interests.[1] Such is the theory behind deterrence by punishment.[2] Clearly, the United States has the means to impose costs; the challenge is to develop a strategy that persuades other nations and actors that the United States can and will impose costs effectively and assuredly.

Martin C. Libicki is the Maryellen and Dick Keyser Distinguished Visiting Professor at the US Naval Academy's Center for Cybersecurity Studies and author of Cyberspace in Peace and War. His research has focused on the impact of information technology on domestic and national security. Previously, he worked for RAND (where he is still adjunct), the National Defense University, and the Navy staff. He holds a master's degree and a doctorate from the University of California–Berkeley.

In the last decade a new way of imposing costs has emerged through cyberspace, notably by cyberattacks on information systems. Cyber capabilities and national attitudes toward their use continue to evolve. What was deemed fantastical or science fictional 10 years ago now gets serious attention. Past assessments that "they couldn't do that" or "they wouldn't do that" are proving unreliable guides for assessing which capabilities may be used in the future. Many countries, not least of which the United States, can hurt other countries through cyberattacks as well as cyberespionage, cyber sabotage, or cyber subversion. Arguably, this additional capability should therefore add to the deterrence posture of its possessors—but in what way and by how much? This article argues that cyber capabilities can make a contribution to the broader deterrence framework. It seeks to establish a reasonable set of expectations that should inform where the *existence* of cyber capabilities, coupled with the threat they might be used as punishment, may help deter others. While additional capability to punish should improve the ability to deter, two factors vitiate cyberattack as part of an overall deterrence menu. The first is the existence of a deterrence scale associated with the response to a particular bad act, below which any particular reprisal threat may be too weak and above which any particular reprisal threat may be too costly or noncredible. Second, the consequences of suffering a retaliatory cyberattack are so uncertain that they are easy for opposing leaders to deprecate.

The requirements for reasonable and well-communicated thresholds as well as for credibility have been part of deterrence theory since its inception. Furthermore, the notion that certain types of punishment cannot be credibly threatened if they are disproportionate to the crime has been well understood in the literature on nuclear deterrence.[3] The difference here is that using cyberattacks for punishment raises issues whose considerations are not so salient in other domains.

To be clear, the focus here is on those cyberattacks that are punitive rather than carried out to support kinetic operations, largely because the contribution of a cyberattack capability to an overall kinetic *reprisal* capability is usually indirect. To illustrate as much, assume US airpower is a tool of reprisal against aggressive actions. The potential *aggressor* may conclude that it may suffer an air raid in response; it is therefore discouraged from aggression. If the aggressor could nullify the air threat because of its own air defenses, such leaders would be emboldened—that is, until their own cyber warriors caution that US cyberattack capabilities can nullify the aggressor's air defenses. Thus, because of cyberspace capabilities, a US air raid is more likely to succeed and its prospect is thus more likely to deter the aggressor. Here is a case where a cyberattack capability *within* the context of kinetic operations can add to deterrence. But, as the story suggests, it is a second- or third-order consideration and, for that reason, unlikely to factor strongly into the other side's decision making.

As for the deterrent value of an unaccompanied cyberattack threat against another country's military, it would seem to matter only if such a military is at war or is expected to be at war before the effects of any cyberattack can be eliminated; otherwise, a disruptive cyberattack threat against an idle military leaves little for its owner to fear. Also, in fairness, specific reprisal actions need not be explicitly threatened to contribute to deterrence. The president need not say "if you do this, we will take down your power grid." It suffices that another country understands that acting in a particular way will anger the United States and that, in its anger, the United States might use any of many instruments of reprisal. If a country is more inhibited because these instruments *now* include cyberattack, then cyberattack capabilities can be said to contribute to deterrence.

A Deterrence Scale

Deterrence by punishment must satisfy at least two criteria. On the one hand, the threatened punishment must be sufficiently painful that potential attackers believe they will be worse off after punishment has been delivered even after factoring in the benefit from the bad act. On the other hand, the threat of punishment must be credible; having the requisite capability means nothing if the other side thinks it will not be used. Therein lies the dilemma. Some threatened reprisals may be perceived as too disproportionately harsh to be credible. Others may be too small to merit notice on their own or too limited to add appreciably to the deterrent effect of other larger reprisals—for example, a cyberattack simultaneous with a kinetic attack. In other words, for every action there is a just-right deterrence range of reprisals: large enough to be effective but small enough to credibly threaten. Conversely, for every reprisal there is a corresponding range of actions that can be deterred by it.

Consider the risks to credibility of threatening a disproportionate response. If the potential aggressor is powerful, it could counter-retaliate, perhaps even in an escalatory manner. In evaluating the credibility of a US reprisal threat, the aggressor could ask itself whether the United States would still retaliate even in the face of a possible counter-retaliation. If confident that a counter-retaliation would make a difference, an aggressor is likely to present as daunting—that is, painful and credible—a counter-threat as possible. Furthermore, if it believed that its counter-threat had registered with the United States, the aggressor would deem a disproportionately large reprisal threat from the United States as simply *not* credible because the United States would hesitate to get into a large tit-for-tat to make its point. But the US might be willing to get into a smaller tit-for-tat to preserve its credibility, particularly with third parties it would like to deter. In other words, the US threat to use a large cyberattack to punish a small misdeed could be discounted; it would provide scant deterrence.

Deterrence calculations necessarily presume a high order of rationality and calculability. When the subject is cyberspace, it also requires a mindset capable of inferring effects and costs from threatened cyberattacks. Those being deterred may well impute some rationality and a set of reasonable objectives to the United States and thereby figure that the United States will abjure using a capability if using it is costly or risky.

That noted, for some leaders, such rationality (as well as a well-grounded confidence in its assessment of the United States) may be asking a lot. Leaders tend to mirror potential adversaries. Someone whose perspective sees aggression as a country's self-expression, for instance, may not necessarily believe that the United States uses a rationality that would be foreign to the aggressor, itself. One important reason why Iran settled with Iraq in 1988 to end their eight-year war was that its leaders saw the shoot down of the Iranian airliner as prelude to US lethal and pro-Iraq intervention (rather than the mistake it was).[4] In this particular case, the use of a weapon that the United States might deem overkill may be exactly what the other side would have used were tables turned. That noted, if, given the opportunity and the United States fails to use a retaliatory capability that the aggressor would have used were it available, the assessments of aggressors are likely to be adjusted toward realism. The more such foregone opportunities, the more realism.[5]

An additional barrier to translating a cyberspace capability into deterrence is whether or not others conclude that the US would actually use cyberattacks as reprisals. True, the United States is presumed to have employed Stuxnet and may have used similar capabilities to retard the

Democratic People's Republic of Korea (DPRK) missile program.[6] Yet in both cases, the United States did not take credit for doing so, something it might have done if such acts were consciously meant to be part of a deterrence package. Attribution is something a country might have wanted if it was hoping to leverage such capabilities for deterrence. The United States may also have used cyber operations to respond to the DPRK Sony hack or Russia's Democratic National Committee (DNC) hack, but there is no public indication of this actually happening (which does not prove the US did nothing). The United States has yet to carry out overt cyberattacks in response to any insult or injury, to cyberspace or otherwise. By contrast, one would expect that if cyberattacks *were* used for reprisals, the United States *would* own up to them. But the United States may be unique in its reticence. China, for its part, has carried out or at least condoned low-level cyberattacks as retaliatory responses. It fired its "Great Cannon" at Github when the latter provided a path for people inside China to access the *New York Times*.[7] The Chinese government has condoned its hackers defacing the website of a South Korean firm (Lotte) in response to that country's acceptance of a Thaad missile defense system.[8] The North Korean cyberattack on Sony was carried out presumably in retaliation for the imminent release of Sony's movie *The Interview*. Iran almost certainly retaliated with distributed denial-of-service (DDOS) attacks against US banks for what it believed to be the US sponsorship of Stuxnet—but it also carried out cyberattacks on energy companies of Saudi Arabia (Saudi Aramco) and even Qatar (RasGas) with no single specific Saudi provocation or any discernable Qatari provocation. Russia has also carried out doxing-motivated cyberespionage (DNC, World Anti-Doping Agency, Soros Foundation), DDOS attacks on multiple countries (Estonia, Georgia, Lithuania, Kyrgyzstan, Ukraine), and cyberattacks against electric power (Ukraine).

Another inhibition to the US use of a cyberattack is the precedent that doing so may set for others. In contrast to every other country, the United States has traditionally seen itself as a global leader whose actions are the gold standard by which another country will judge its own actions—or at least its ability to justify its own actions. Of course, if all plausible others have already crossed that threshold, and particularly if that fact were widely acknowledged, there may be no precedent to set, no inhibition to using cyberattack capabilities, and no basis for the United States to justify its diffidence in using cyberattack for retaliation. Some might argue that the widespread association of the United States with Stuxnet may mean that Rubicon has long been crossed even if Stuxnet were not itself a reprisal and even if the United States has not admitted authorship of it.

Nevertheless, many of the cyberattack capabilities putatively possessed by the United States may be those it said it will not use in peacetime and thus cannot be used for deterrence outside a setting in which warlike operations are also taking place. The United States has signed on to the UN's Group of Government Experts (GGE) convention that abjures cyberattacks on critical infrastructure[9]—and many of the targets of a reprisal cyberattack can be classified thusly. This promise could limit what the United States can threaten in any situation short of armed conflict. And in situations of armed conflict, the expectation of additional pain created by a cyberattack may be too modest to weigh heavily in the deliberations of the potential aggressor. Again, that noted, other countries may not necessarily rest easy by relying on US promises to limit US behavior. If they themselves give such norms no more than lip service, they may see the United States acting similarly, either violating the norm outright, deeming something a war-level cyberattack, or operating covertly so as not to take responsibility. If so, a cyberattack may be a credible reprisal threat even where kinetic response is not.

Can a cyberattack capability add punch to other retaliatory capabilities? This is unlikely if the "other" threat is nuclear, territorial occupation, or an air campaign (although, in the latter case, it depends on how thorough an air campaign). A cyberattack may pale beside what NATO did to Serbia in 1999 but might stand out beside a single cruise missile strike (e.g., such as that against Sudan in 1998). Similarly, since the damage from cyberattacks is largely denominated in dollars rather than lives, it may line up with a full-fledged blockade or even an embargo, but prospects of the latter, if serious, are likely to dominate the other side's calculations. If potential aggressors are already under embargo, they may feel cut off from the United States and thereby discount the threat from cyberspace operations that require connectivity to be implanted or activated—and it may not matter that their optimism is belied by the many ways cyberattacks can get into a system that lack obvious connectivity.

Not everything about the deterrence threshold casts doubt on the value of a cyberattack capability as a deterrent, in large part because retaliation is not homogeneous. Cyberattacks are high among retaliatory options that can be targeted at leadership without affecting the average person very much.[10] Leaders can be hurt without creating so much public pain as to be judged a disproportionate response. Leaders would pay the cost of their aggression, but they could well lack the evidence, hence the basis for a narrative, to justify their counter-response. Third parties may simply not believe that leaders have been hit unduly or unfairly. Leaders who fear as much may persuade themselves that targets of their aggression can credibly retaliate without the concerns other equally painful—but public and unambiguous—reprisals would raise. It is not necessary that their conclusions be warranted, only that they be plausible. Conversely, many of the more effective cyberspace operations, such as emptying a dictator's foreign bank account, may harm third parties such as the bank itself more than the dictator and thus make poor reprisal options.

Finally, adding even usable arrows to country's reprisal quiver is not always helpful. A potential aggressor facing a contingent risk of a US reprisal may nevertheless discount many of these reprisal options: some because they are too weak to matter and others because they present too many downsides for the United States. Assume, then, a new arrow is added. The other side may be more deterred because the United States now has another usable option—or even better, that the United States has one usable option to respond to a bad act before which it had no way to respond. But the opposite could be true. The other side could conclude that the existence of a cyberattack option *reduces* the odds of having to face a more forceful option. Indeed, the insistence with which the United States brandishes a cyberattack option may convince potential aggressors that threats by the United States to use costlier and/or riskier options are edging off the table. In that case, raising the credibility of one retaliation option would lower the likelihood of more painful ones. The substitution of something that *might* be painful for something that *would* be painful may *reduce* the overall deterrence posture.

Insofar as retaliation should bear some relationship to the transgression—the whole point of the deterrence scale—it may be hard to create an explainable equivalence between the initial event and the reprisal if the former takes place in the physical world and the latter takes place in the virtual world. When the potential of cyberattack has been likened to nuclear attack *and* where cyberattacks, as oft observed, have yet to kill anyone, making a credible case for equivalence *beforehand* is fraught.

And although it is explicitly not US policy that the response to an attack in one domain is a reprisal in another, it should be fairly clear that the only form of aggression for which the deterrence scale is not a major problem is a cyberattack itself. Unfortunately, the difficulties of any deterrence policy to govern the wilds of cyberspace are many.[11]

In practice, countries rarely lay out a specific set of options as part of their deterrence policy. At best, they may indicate what actions they would take issue with and perhaps point to capabilities they think others are unaware of or not paying enough attention to, as Pres. Barack Obama did with US cyberattack capabilities.[12] Doing so lets others draw their own inferences and calculate their own risks from aggression. What they conclude may be very different from what the United States wants them to conclude.

Does the Uncertainty of Effects Matter?

Uncertainty affects all form of warfare—but with kinetic weaponry others have a good idea what an individual weapon can do, how vulnerable its assets are to attack, what the United States has, and what defenses it has to counter US weapons. With physical attack, the cost of recovery is also predictable in that it is largely based on the cost of replacing destroyed items. However, cyberattacks have effects that are particularly unpredictable for several reasons.[13] First, the victim of a cyberattack will not know precisely what capabilities the perpetrator of a cyberattack possesses, although between Stuxnet and Snowden's revelations, the US arsenal appears to be impressive.[14] Second, in contrast to kinetic weapons where having a capability at all implies having a capability against anyone within range, having a cyberattack capability in the abstract does not mean that one has a cyberattack capability against a particular target or even a particular country, something that cannot be completely known even by its possessor before its cyber warrior cadres penetrate potential targets. Defenders themselves also may not know how vulnerable their systems are until tested. Third, the perpetrator of a cyberattack has only a partial insight into the victim's defenses. Fourth, neither may have a good understanding of how long it takes the victim to recover—which matters because, to a first-order approximation, the pain from a cyberattack is proportional to recovery time (i.e., the time to restore operations after a disruption attack, and the time to detect and eradicate induced errors from a corruption attack). To this uncertainty must be added the difficulty of understanding the direct effects a capability may have on things potential aggressors care about, notably regime legitimacy and survival.

Unpredictability, in turn, has two types of effects. First, the victim, not understanding how painful reprisals can be, can exaggerate or, alternatively, deprecate the direct effects of reprisals. As argued below, that last unknown may be influenced by the narrative that the targeted regime thinks it can create before and after reprisals hit. Secondly, the aggressor, unsure of how painful reprisals can be, can deem such reprisals too weak to make a difference or too strong to be used confidently and hence unlikely to be used.

Let us start with the doubts of the potential aggressor, which is to say, the potential victim of a cyberattack reprisal. When faced with a threat of unknown size, people can be pessimists and exaggerate it or optimists and deprecate it. As a deterrent threat, cyberattacks would have to impress only those leaders tempted to do something aggressive enough to call forth a reprisal. It would seem that anyone with a bias for aggression is presumptively forced to be optimistic about its own chances in a confrontation, even though, as with 1914-era Germany or 1941-era Japan, leaders can be simultaneously pessimistic about their prospects in the world if they do not act. That being so, given two reprisal threats of equal expected size, the one of more predictable effect would seem to deter more than the one of less predictable, albeit potentially greater, effect. In the latter case, the optimistic potential aggressor can tell itself that the retaliatory cyberattack will not work or will not cause much damage if it does. In the case of a kinetic attack such as an air raid there is less psychological basis for insouciance.

If the potential aggressor wishes to counter the fears of its policy, it can bluff with a narrative that it has nothing to fear from a particular capability. The extreme version of denial was Mao Tse-Tung's statement early in the Cold War era that the nuclear capabilities of the United States were those of a "paper tiger" because hundreds of millions would survive a nuclear war in China, a vast agricultural society.[15] Although this remark was derided as mad, its political purpose was to encourage the Chinese not to be cowed by the threat of nuclear reprisals. And, indeed, China had intervened against the United States in the Korean War without its territory being bombed, much less nuked.

With a cyberattack threat, the defensive narrative is both harder and easier to craft. If the threat of cyberattack reprisals goes unspoken and the targeted population is generally unaware that, for instance, access to their money could be denied, then creating a counter-narrative that says such a threat is empty may introduce a new fear (it is akin to "protest[ing] too much"). If the threat is explicit or if it is mentioned often enough by third parties, then there may be a point to introducing a new narrative. With cyberattacks, this new narrative could range from dismissal ("paper tiger") to blame-shifting ("only the feckless will suffer") and moralizing ("look how heinous the United States is by militarizing cyberspace"). Unless there have been enough incidents to make cyberattack threats tangible, the newness and the non-obviousness of the threat suggest that such a countering narrative is plausible.

Will others buy such narratives? The United States needs to deter many actors, and each is different. For some, the threat of cyberattacks may not register. For others, one or another counter-cyberattack narrative can work. Thus, it is not obvious that the United States would be better off crafting a narrative ("cyberattacks are really painful"), particularly if the effect of exaggeration is to scare the US public so that it palls from a confrontation lest their own systems falter. Two conditions help justify mounting such a narrative. If there is just one particular potential aggressor that needs to be deterred, the narrative can be focused in ways that scare others without being scary per se. It also helps to have a home front confident that it will not suffer from a counter-retaliatory cyberwar, perhaps because it believes in US superiority in cyberspace irrespective of whether it actually exists or would matter even if it did. Even then, the US developing a narrative that its cyberattacks can hurt others is not trivial. The best argument would come from carrying out cyberattacks and then pointing to the results—but the only cyberattacks that count are those against real foreign targets even though the whole purpose of deterrence is that if everyone behaves, no one gets to experience such attacks. And for cyber weapons, in large contrast to kinetic weapons, confidence in creating effects is limited to effects on a particular target system with its particular vulnerabilities, being administered in a particular way. Even then, the amount of damage that can be credibly threatened can and will change over time as the target adjusts to the threat.

As a rule, a target's adjustment to cyberattack threats has more effect on the efficacy of a cyberattack than when the threat is kinetic. The threat of a cyberwar deterrent in the long run is no greater than the costs—as measured in labor, time, resources, and decreased usability—of managing the risk to networks and systems down to tolerable levels. If the threat is sufficiently fearsome, a rational country would pay upfront for cybersecurity and resilience.[16] Forgoing such pains may be a clue that such cyberattacks do not frighten people enough for them to ward it off. This keeps the threat of cyberattacks from being a particularly persuasive deterrent. No such confidence can be reasonably expressed in the face of, say, a nuclear threat.

The impact of uncertainty on the United States reinforces the cyberattack argument. If the United States is uncertain about the effects of a retaliatory cyberattack it may not judge with requisite confidence that it falls within the deterrence scale: the reprisal may fail to impress, or the reprisal may be overkill. Indeed, if uncertainty is great enough, a cyberattack reprisal option may be rejected because the likelihood that it falls below the range *and* the likelihood that it falls above the range may *both* be disqualifying. To be sure, US doubts, on their own, do not matter because the thoughts that matter are those of the potential aggressor. But if potential aggressors convince themselves that this is how the United States thinks, they may conclude that no weapon whose impacts are so uncertain will actually be used by the United States regardless of what US leaders want others to believe.

Conclusions

As long as other countries understand that the United States is capable of impressive cyberspace operations then the threat that it will use them in reprisal is inevitably part of the US deterrence package. The policy question is whether the United States should emphasize the possibility: by impressing adversaries with what cyberattacks can do, reminding adversaries that the United States would be willing to do it, and investing in making cyberattacks more reliable and even more painful. Central to any answer is an assessment of how cyberwar might fit into an overall US conventional deterrence posture. Its contribution to deterrence is likely to be modest compared to the level of punishment cyberattacks promise or compared to similar kinetic threats. Several reasons exist for this.

First, the effective deterrence window for cyberattacks is narrow. Compared to the capabilities that the United States is likely to bring for a high-end scenario—ranging from nuclear war to a major conventional aggression—cyberwar, at this stage, does not add much. The United States has abjured low-end cyberattacks such as DDOS attacks or website defacements. It has forsworn attacks on critical infrastructure in peacetime; in wartime, as noted, the threats of cyberattack may add little to the threat of conventional force. The United States can credibly threaten retaliatory cyberattacks to punish an aggressor's cyberattacks, but credibility for other scenarios is problematic.

Second, the effects of cyberwar are highly uncertain. Thus, they are easy for optimistic aggressors to downplay, and most who would go up against the United States would have to be optimistic to do so at all. They are also relatively easy to deprecate if aggressors need to worry about bringing others (e.g., political and military elites) along before striking. Indeed, it is difficult to find a narrative in which the threat is both scary and legitimate. Lastly, a potential aggressor savvy to how the United States deals with uncertainty may conclude that no weapons of such uncertain effects will actually be used in a retaliatory package that must pass some sort of proportionality test. Thus, when it comes to cyberwar and cyberattacks, the US must carefully consider its options.

Because offensive cyberspace operations make unimpressive deterrents when used by the United States,[17] expectations about their efficacy should be tempered. They should probably be brandished only against aggression that, itself, comes via cyberspace. Retaliation in-kind limits the problem of the deterrence scale, and although the effects of retaliatory cyberspace operations are uncertain, the benefits of an aggressor's cyberspace operations suffer from similar uncertainly.

Notes

[1] Deterrence by punishment is a subset of coercion: the threat of force to persuade others to do other than what they might. This subject has generated a rich literature, including Thomas Schelling, *Arms and Influence* (New Haven, CT: Yale University Press, 1966); Robert Pape, *Bombing to Win*, (Ithaca, NY: Cornell University Press, 1996); and Dan Byman, Matthew Waxman, and Eric Larson, *Air Power as a Coercive Instrument* (Santa Monica, CA: RAND, 1999). Deterrence ("do not do this") is generally held to be easier than compellence ("do this or else"), but it is difficult to know whether deterrence works in any one case without knowing whether the prohibited act would have taken place in absence of a threat.

[2] This article uses deterrence to mean deterrence by punishment. Granted, making threats is not the only way to persuade people not to do something bad. They may also think twice if obstacles are put in their path or the reward for success is diminished. This latter path is called deterrence by denial. Although a valid consideration in persuasion, it is not the topic here.

[3] The history of the Cold War suggests that the vague threat of using nuclear weapons to change another country's nonnuclear behavior does not work very often, largely because it is not seen as credible; see Alexander George and Richard Smoke, *Deterrence in American Foreign Policy* (New York: Columbia University Press, 1974).

[4] See, for instance, Max Fisher, "The Forgotten Story of Iran Air Flight 655," *Washington Post*, 16 October 2013, https://www.washingtonpost.com/news/worldviews/wp/2013/10/16/the-forgotten-story-of-iran-air-flight-655/.

[5] And leaders are constantly evaluating the credibility of their counterparts: e.g., Kirk Semple and Elisabeth Malkin, "A Calmer Mexico Sees Trump Anew: As a 'Bluffer' at the Poker Table," *New York Times*, 27 April 2017, https://www.nytimes.com/2017/04/27/world/americas/mexico-trump-nafta-trade.html.

[6] David Sanger and William Broad, "Trump Inherits a Secret Cyberwar Against North Korean Missiles," *New York Times*, 4 March 2017, https://www.nytimes.com/2017/03/04/world/asia/north-korea-missile-program-sabotage.html.

[7] Dan Goodin, "Meet 'Great Cannon,' the Man-in-the-Middle Weapon China Used on GitHub," Ars Technica, 10 April 2015, http://arstechnica.com/security/2015/04/meet-great-cannon-the-man-in-the-middle-weapon-china-used-on-github/.

[8] Kang Seung-woo, "Cyber Warfare Rising in China's THAAD Retaliation," *Korea Times*, 8 March 2017, http://www.koreatimes.co.kr/www/tech/2017/03/133_225311.html.

[9] UN General Assembly (UNGA), *Group of Governmental Experts on Developments in the Field of Information and Telecommunications in the Context of International Security*, UNGA, 70th sess. (New York: United Nations, 22 July 2015), http://www.un.org/ga/search/view_doc.asp?symbol=A/70/174.

[10] See, for instance, Daniel W. Drezner, "Economic Sanctions in Theory and Practice: How Smart Are They?," in *Coercion: The Power to Hurt in International Politics*, ed. Kelly Greenhill and Peter Krause (Oxford, UK: Oxford University Press, 2018). He argues that although targeted sanctions against leaders are no more effective in compelling behavior than general ones, they are politically easier to maintain. Correspondingly, they may be more credible.

[11] For examples, see Martin Libicki, *Cyberdeterrence and Cyberwar* (Santa Monica, CA: RAND, 2009).

[12] As President Obama declared in 2015, "Among states, there has to be a framework [for behavior in cyberspace] that is analogous to what we've done with nuclear power because nobody stands to gain. And, frankly, although the Chinese and Russians are close, we're still the best at this. And if we wanted to go on offense, a whole bunch of countries would have some significant problems." Pres. Barack Obama, "Remarks by the President to the Business Roundtable" (address, Business Roundtable Headquarters, Washington, DC, 16 September 2015), https://www.whitehouse.gov/the-press-office/2015/09/16/remarks-president-business-roundtable.

[13] The difficulty of knowing how much one country can hurt another in cyberspace is key to arguments that cyberspace operations may have weak coercive power. See, for instance, Erica Borghard and Shawn Lonergan, "The Logic of Coercion in Cyberspace," *Security Studies* 26, no. 3 (2017): 452–81, https://doi.org/10.1080/09636412.2017.1306396; and Jon R. Lindsay and Erik Gartzke, "Coercion through Cyberspace: The Stability-Instability Paradox Revisited," in Greenhill and Krause, *Coercion: The Power to Hurt*. The latter argues that the

coercive potential of a cyberspace threat is generally limited because the target can disconnect threatened networks from the outside if the threat is unbearable. That noted, trends in digitization and connectivity mean that the cost borne by disconnection is rising every year.

[14] Henry Farrell, "The Political Science of Cybersecurity IV: How Edward Snowden Helps U.S. Deterrence," *Washington Post*, 12 March 2014, https://www.washingtonpost.com/news/monkey-cage/wp/2014/03/12/the-political-science-of-cybersecurity-iv-how-edward-snowden-helps-u-s-deterrence/.

[15] See, for instance, Jeffrey Lewis, *Paper Tigers: China's Nuclear Posture*, IISS Adelphi Series, no. 446 (New York: Routledge, 2014).

[16] This formulation assumes states to be unitary actors that can transfer resources as needed. A country's leaders may pay $2 to avoid the blame from a cyberattack on a private network when the network owners could not or would not spend $1 to protect the network and thereby void the threat—but there may be no reliable way to move the $2 into the necessary cybersecurity investment.

[17] Unfortunately, that does not mean that the United States itself may not fear cyberattacks as reprisals or counter-reprisals. Self-deterrence played a role in stymying the US reactions to Russian cyberspace operations during the 2016 election; see Michael Isikoff and David Corn, *Russian Roulette: The Inside Story of Putin's War on America and the Election of Donald Trump* (New York: Hachette, 2018); and Ben Rhodes, *The World as It Is* (New York: Random House, 2018).

Disclaimer

The views and opinions expressed or implied in SSQ are those of the authors and are not officially sanctioned by any agency or department of the US government. We encourage you to send comments to: strategicstudiesquarterly@us.af.mil