H-Diplo | ISSF

Author's Response to Article Review 83
issforum.org

H-Diplo/ISSF Editors: **Seth Offenbach** and **Diane Labrosse**
H-Diplo/ISSF Web and Production Editor: **George Fujii**

Author's Response to H-Diplo/ISSF Article Review by Brandon Valeriano of Rebecca Slayton, "What Is the Cyber Offense-Defense Balance? Conceptions, Causes, and Assessment," *International Security* (2017): 72-109, and Joseph S. Nye Jr., "Deterrence and Dissuasion in Cyberspace," *International Security* 41.3 (2017): 44-71.

Review URL: https://issforum.org/articlereviews/83-cyber
Response by **Rebecca Slayton**, Cornell University

Published by ISSF on 9 August 2017

tiny.cc/ISSF-AR83-Response
https://issforum.org/articlereviews/83-response
https://issforum.org/ISSF/PDF/ISSF-AR83-Response.pdf

I would like to begin by thanking Brandon Valeriano for reviewing my article "What is the Cyber Offense-Defense Balance?" I am also grateful to H-Diplo for publishing the review of my article and providing the opportunity to respond.

For those who did not read the original article, I will briefly summarize its three main points. First, it articulates distinctive concepts of cyber offense-defense balance that are currently in use, arguing that a useful conception should consider the benefits as well as the costs of cyber operations. Second, the article theorizes the sources of offensive or defensive advantage in terms of the relative skill with which adversaries manage complex information technology, and the relative complexity of their goals. Accordingly, the balance is a dyadic, not systemic, variable. Third, it provides an empirical analysis of the costs and benefits of Stuxnet, the U.S.-Israeli designed worm that attacked Iranian nuclear enrichment facilities in 2010. This analysis counters the conventional wisdom of offense-dominance by showing that the offense likely spent hundreds of millions of dollars on Stuxnet, while the defense likely spent only tens of millions. However, the value that both sides place on Iran's nuclear program appears to be on the order of $10 billion per year, making it unlikely that either side was fixated on costs.

Brandon Valeriano's review helpfully situates the article in its scholarly context and accurately summarizes some of the article's key points. I agree with his conclusion that further research is warranted. However, Valeriano also makes four unfounded criticisms of my article. Here I will briefly respond to each of these criticisms.

First, Valeriano states that my article "focuses too much on the platform of the offensive-defense balance," a concept that is "problematic" because of the "complicated question of just how to measure what an offensive weapon is versus a defensive weapon." However, my article moved beyond this critique by underscoring the impossibility of separating information technology from the skills needed to create, use, and repurpose it. Accordingly, I propose focusing not on technology, but on skills. While offensive and defensive skills may be very similar, the resources required for a successful offense or defense may be different because of the different levels of complexity involved in offensive and defensive information operations; as I explain, defense is usually (but not always) more complex. Additionally, whether acting offensively or defensively, more skilled actors can accomplish the same objectives with fewer resources. In short, since my article did not focus on "weapons" as traditionally construed, this criticism does not address the arguments I actually made.

The review advocates "moving beyond handed-down questions" such as offense-defense balance. However, the questions addressed in my article—what distinctive conceptions of cyber offense-defense balance are currently in use, what factors influence the costs and benefits of cyber-operations, and what empirical analysis can teach us about these costs and benefits—are not handed-down. They had not been previously addressed, despite the fact that divergent notions of offense-defense balance have influenced both scholarship and military policy.[1] Developing a clear understanding of the meaning, value, and limitations of the concepts we employ is crucial to both rigorous scholarship and sound policy.

Second, the review criticizes the selection of Stuxnet as a case study, arguing that it "is likely an outlier" and this choice is "not supported by any social science justification." Many experiences are worthy of study, not because they are typical, but because they are influential or potentially influential. Analysts study the Cuban missile crisis, Chernobyl, and other disasters or near-disasters in order to avoid repeating them.[2] Conversely, Silicon Valley's innovation system and other "successes" are studied by those who wish to emulate them. Stuxnet is worthy of study because it achieved a goal of many states: exerting kinetic effects through cyberspace. Such attacks are on the rise; Russia successfully caused blackouts in the Ukrainian electric grid in December 2015 and 2016.[3] Additionally, my article notes that Stuxnet "has received extensive analysis in the scholarly and trade press" (95) and has been used to support scholarly arguments "for either offense

[1] Three different concepts can be identified: the relative costs of offense and defense; the relative efficacy (payoff per unit of effort) for offense and defense; and first mover advantages. On costs, see Keir Lieber, "The Offense Defense Balance and Cyberwarfare," in Emily O. Goldman and John Arquilla, eds., *Cyber Analogies* (Monterey: Naval Postgraduate School, 2014); Ilai Saltzman, "Cyber Posturing and the Offense-Defense Balance," *Contemporary Security Policy* 34:1 (2013): 40-63, DOI: http://dx.doi.org/10.1080/13523260.2013.771031. On efficacy, see Andrew Krepinevich, "Cyberwarfare: A "Nuclear Option?," (Washington, D.C.: Center for Strategic and Budgetary Assessments, 2012); John Arquilla, "Cyberwar Is Already Upon Us," *Foreign Policy* 192 (March/April 2012), http://foreignpolicy.com/2012/02/27/cyberwar-is-already-upon-us/. On first mover advantages see Lani Kass, "A Warfighting Domain; Air Force Cyberspace Task Force Presentation, 26 September 2006," U.S. Air Force, http://www.au.af.mil/info-ops/usaf/cyberspace_taskforce_sep06.pdf.

[2] James March, Lee S. Sproull, and Michal Tamuz, "Learning from Samples of One or Fewer," *Organization Science* 2:1 (1991): 1-13. DOI: https://doi.org/10.1287/orsc.2.1.1.

[3] Andy Greenberg, "'Crash Override': The Malware That Took Down a Power Grid," *Wired*, 12 June 2017, https://www.wired.com/story/crash-override-malware/.

dominance or defense dominance (96)."[4] These are all good social scientific reasons for empirically assessing Stuxnet's costs and perceived benefits.

Third, Valeriano claims that my article does not address several issues that it does in fact address. He misrepresents the article when he writes:

> As Slayton notes, leaders were "unlikely to focus on costs" (75) of operations. This begs the question as to what, if not on costs, are they focused on.

This misquotes my argument. The complete quotation is:

> …the value that the United States, Israel, and Iran all attach to Iran's nuclear program appears to be much greater than the cost of either cyber offense or cyber defense, making it unlikely that leaders were focused on costs (75).

As is clear, my original statement does *not* beg the question of what focused leaders' attention. It suggests that they were focused on the value they attach to Iran's nuclear program. My article argues for including considerations of the benefits of cyber operations in any discussion of the balance, precisely because these benefits focus leaders' attention. It also evaluated the perceived benefits and costs of Stuxnet, providing empirical support for this claim.

The review goes on to state: "The questions of resolve, interest, and capability are left unexamined…" This is untrue. The article notes that "Stuxnet likely increased Iran's resolve to enrich uranium and spurred the development of both defensive and offensive cyber operations" (106). The article discusses revenues that both the United States and Iran have willingly lost in order to undermine or maintain Iran's nuclear program; this is clearly about national resolve. It focuses on the interest of Israel and the United States in preventing Iran from developing nuclear weapons. It also notes the U.S. interest in averting a conventional bomb strike on Natanz, as was proposed by Israel (95). Capability is a central element in the proposed model of offense-defense balance, and is discussed on five pages (87-91). The conclusion summarizes these points: "widespread claims about the offense dominance of cyberspace are fundamentally flawed; the offense-defense balance can

[4] Those who have used Stuxnet to demonstrate the difficulty of attack include: Jon Lindsay, "Stuxnet and the Limits of Cyber Warfare," *Security Studies* 22:3 (2013): 365-404, DOI: http://dx.doi.org/10.1080/09636412.2013.816122; Thomas Rid, *Cyberwar Will Not Take Place* (New York: Oxford University Press, 2013); Adam Liff, "Cyberwar: A New 'Absolute Weapon'? The Proliferation of Cyberwarfare Capabilities and Interstate War," *Journal of Strategic Studies* 35:3 (2012): 401-428, DOI: http://dx.doi.org/10.1080/01402390.2012.663252; David Betz, "Cyberpower in Strategic Affairs: Neither Unthinkable nor Blessed," *Journal of Strategic Studies* 35:5 (2012): 689-711, DOI: http://dx.doi.org/10.1080/01402390.2012.706970.Others use Stuxnet to demonstrate the ease of attack: Lucas Kello, "The Meaning of the Cyber Revolution: Perils to Theory and Statecraft," *International Security* 38:2 (Fall 2013): 7-40, DOI: https://doi.org/0.1162/ISEC_a_00138; Gary McGraw, "Cyber War Is Inevitable (Unless We Build Security in)," *Journal of Strategic Studies* 36:1 (2013): 109-119, DOI: http://dx.doi.org/10.1080/01402390.2012.742013; Dale Peterson, "Offensive Cyber Weapons: Construction, Development, and Employment," *Journal of Strategic Studies* 36:1 (2013): 4-41, DOI: http://dx.doi.org/10.1080/01402390.2012.732015.

be understood only in the context of specific adversaries with distinctive goals and levels of capability in managing complex information technology” (106).

Valeriano accurately notes that a “key point to be taken from the article is that the talent and skill of cyber security professionals in the field is the variable that requires more study,” but then claims that “Slayton does not really follow up on this trend relating to skill and talent…” This ignores the substantial attention that the article devotes to discussing both individual and organizational skills (pages 82-91). This includes discussion of the close relationship between offensive and defensive skills and the capability maturity model for organizational competence. While I agree that more research is warranted, I did devote more than one-quarter of a relatively long article to these subjects.

Fourth, Valeriano disagrees with my view that analysts are unlikely to establish usefully precise predictions of the probability of success for a cyber operation, or of the minimum resources required for success. My skepticism about making precise predictions stems from the paper’s core argument: “the offense-defense balance is shaped primarily by the relative skill with which adversaries manage complex information technology, and the relative complexity of their goals” (74). The complexity of information technology and operations is “arbitrary” because it emerges from human organizations and activities, and is unique to each system and operation (88). Thus, the basis for generalization about the costs of offense and defense is limited and uncertain.

The review does not acknowledge or counter the article’s core argument. Instead it states that good probabilistic data can be obtained because of the “evident public nature of cyber breaches.” It is true that some cyber breaches are public. It does not follow that all breaches are public. There is reason to believe that many are not, as there are many circumstances in which corporations and states do not want to publicize a successful breach. Additionally, even when breaches are public, much of the information needed to assess the costs and perceived benefits of the offense and defense is not available. Successful intruders rarely reveal all the false starts, blind alleys, and other futile avenues that they pursued before succeeding, yet all of these efforts are part of the cost of an intrusion. My own analysis of a single, well-publicized case—Stuxnet—was laborious and produced only retrospective order-of-magnitude estimates. Those who read the details of this analysis will appreciate the uncertainties intrinsic to any estimate. Indeed, I conducted the analysis in large part to demonstrate these uncertainties.

Valeriano also suggests that probabilities might be estimated from “wargames and cyber ranges at the classified and unclassified level.” However, those responsible for organizing such games generally do not gather the data that is needed to develop probabilistic predictions, nor do they agree that wargames are representative of real-world experience. On 11 March 2015, I queried an organizer of the National Security Agency’s Cyber Defense Exercise (CDX) about the number of intrusion attempts necessary for a successful breach. He explained that they did not collect data because categorizing the many activities that go into an intrusion as separate “attempts” and then recording those attempts would require attention from the people attempting the intrusion, distracting them from their work. He further added that exercises were not sufficiently realistic to be a good source of data; they were artificial situations designed to teach.[5]

[5] E-mail exchange dated 11 March 2015. The person who answered my e-mail requested that neither his name nor his position be used in any publication, and that he not be quoted directly.

In short, wargames and simulations are not a good means of gathering data on the effort required to succeed in offense or defense, or the probability of success at a given level of effort. In the interests of shortening a very long article, I deleted discussion of the limitations of wargames from earlier drafts. I am glad for the opportunity to discuss these limitations here.

Although we disagree on the prospects for making precise predictions, Valeriano and I agree that research should focus on what cyber operations accomplish strategically and politically, rather than simply focusing on whether systems were successfully breached or not. We also agree on the need for further research into the skills and organizational capabilities associated with cyber offense and cyber defense. Significant questions include: can offensive and defensive skills be distinguished, and if so how? What factors, if any, distinguish organizational capabilities for offense and defense, respectively? What roles might transnational communities of experts play in encouraging defensive activities and enabling international cooperation? These kinds of questions may eventually be addressed by quantitative methods, but there is also a significant prior need for qualitative historical and ethnographic research to characterize these problems. If my article helps to move this research forward, it will have been a worthwhile contribution.

Rebecca Slayton is an Associate Professor jointly appointed in the Department of Science & Technology Studies and the Judith Reppy Institute for Peace and Conflict Studies, both at Cornell University. She is author of *Arguments that Count: Physics, Computing, and Missile Defense, 1949-2012* (MIT Press, 2013), as well as many articles in science and technology studies. She is currently working on a second book project, *Shadowing Cybersecurity*, which examines the emergence of cybersecurity expertise through the interplay of innovation and repair.