What Is the Cyber Offense-Defense Balance?

Conceptions, Causes, and Assessment

Rebecca Slayton

"In cyberspace, the offense has the upper hand."[1] These words, written in 2010 by Deputy Secretary of Defense William Lynn, reflect conventional wisdom among military officers, policymakers, and scholars.[2] A minority of scholars disagree.[3] Nonetheless, the prevalent belief that cyberspace favors the offense has major consequences for international security. According to offense-defense theory, state perceptions that technology favors the offense increase fears of attack encourage arms races, and through interactions between fears and capabilities, increase the likelihood and consequences of war.[4] Overconfidence in the

Rebecca Slayton is Assistant Professor at Cornell University with a joint appointment in the Science and Technology Studies Department and the Judith Reppy Institute for Peace and Conflict Studies.

The author thanks Matthew Evangelista, Trey Herr, Drew Herrick, Herb Lin, Jon Lindsay, Jason Reinhardt, Fred Schneider, participants in Cornell University's Science Studies Research Group, and the anonymous reviewers for critiques and suggestions that greatly improved this article.

1. William J. Lynn III, "Defending a New Domain: The Pentagon's Cyberstrategy," *Foreign Affairs*, Vol. 89, No. 5 (September/October 2010), p. 98.
2. Lucas Kello, "The Meaning of the Cyber Revolution: Perils to Theory and Statecraft," *International Security*, Vol. 38, No. 2 (Fall 2013), pp. 7–40; John Arquilla, "Cyberwar Is Already upon Us," *Foreign Policy*, February 27, 2012, http://foreignpolicy.com/2012/02/27/cyberwar-is-already-upon-us/; Kenneth Lieberthal and Peter W. Singer, "Cybersecurity and U.S.-China Relations" (Washington, D.C.: Brookings Institution, 2012); Martin C. Libicki, *Conquest in Cyberspace: National Security and Information Warfare* (New York: Cambridge University Press, 2007); Martin C. Libicki, "Cyberdeterrence and Cyberwar" (Santa Monica, Calif.: RAND Corporation, 2009); Joseph S. Nye Jr., "Cyber Power" (Cambridge, Mass.: Belfer Center for Science and International Affairs, John F. Kennedy School of Government, Harvard University, 2010); Adam P. Liff, "Cyberwar: A New 'Absolute Weapon'? The Proliferation of Cyberwarfare Capabilities and Interstate War," *Journal of Strategic Studies*, Vol. 35, No. 3 (June 2012), pp. 401–428; Chris C. Demchak, *Wars of Disruption and Resilience: Cybered Conflict, Power, and National Security* (Athens: University of Georgia Press, 2011); Nazli Choucri, *Cyberpolitics in International Relations* (Cambridge, Mass.: MIT Press, 2012); Keir Lieber, "The Offense-Defense Balance and Cyber Warfare," in Emily O. Goldman and John Arquilla, eds., *Cyber Analogies* (Monterey, Calif.: Naval Postgraduate School, 2015), pp. 96–107; and Timothy J. Junio, "How Probable Is Cyber War? Bringing IR Theory Back in to the Cyber Conflict Debate," *Journal of Strategic Studies*, Vol. 36, No. 1 (2013), pp. 125–133.
3. Thomas Rid, *Cyber War Will Not Take Place* (Oxford: Oxford University Press, 2013); Jon R. Lindsay, "Stuxnet and the Limits of Cyber Warfare," *Security Studies*, Vol. 22, No. 3 (2013), pp. 365–404; and Jon R. Lindsay, "The Impact of China on Cybersecurity: Fiction and Friction," *International Security*, Vol. 39, No. 3 (Winter 2014/15), pp. 7–47.
4. Robert Jervis published the foundational work in contemporary offense-defense theory. See Jervis, "Cooperation under the Security Dilemma," *World Politics*, Vol. 30, No. 2 (January 1978), pp. 167–214. Additional important works include George H. Quester, *Offense and Defense in the International System* (New York: John Wiley, 1977); Robert Gilpin, *War and Change in World Politics* (Cambridge: Cambridge University Press, 1981); Stephen Van Evera, "Offense, Defense, and the Causes of War," *International Security*, Vol. 22, No. 4 (Spring 1998), pp. 5–43; Stephen Van Evera, "The Cult of the Offensive and the Origins of the First World War," *International Security*, Vol. 9,

International Security, Vol. 41, No. 3 (Winter 2016/17), pp. 72–109, doi:10.1162/ISEC_a_00267

advantages of offense can create a "cult of the offensive" with potentially tragic results.[5]

Many of these dynamics appear to be at work with cyber conflict. Military leaders perceive cyberspace as favoring the offense and are seeking more discretion to conduct offensive cyber operations.[6] Cyberattack features prominently in the U.S. intelligence community's list of global threats.[7] In 2012 Defense Secretary Leon Panetta warned of a potential "cyber-Pearl Harbor."[8] Fears of being hacked, optimism about hacking others, or both have spurred massive investments in military cyber operations around the world, suggesting a cyber arms race.[9] Because cyber operations can blur lines between espionage (or "cyber exploitation") and use of force (or "cyberattack"), they create a "cybersecurity dilemma," wherein network intrusions undertaken for defensive purposes are easily misunderstood as preparation for an attack, creating the risk of escalation and use of force.[10]

No. 1 (Summer 1984), pp. 58–107; and Charles L. Glaser and Chaim Kaufmann, "What Is the Offense-Defense Balance and Can We Measure It?" *International Security*, Vol. 22, No. 4 (Spring 1998), pp. 44–82. For reviews of offense-defense theory, see Jack S. Levy, "The Offensive/Defensive Balance of Military Technology: A Theoretical and Historical Analysis," *International Studies Quarterly*, Vol. 28, No. 2 (June 1984), pp. 219–238; and Sean M. Lynn-Jones, "Offense-Defense Theory and Its Critics," *Security Studies*, Vol. 4, No. 4 (Summer 1995), pp. 660–691.

5. Van Evera, "The Cult of the Offensive and the Origins of the First World War."

6. Lynn, "Defending a New Domain"; and Ellen Nakashima, "Cyber Chief: Efforts to Deter Attacks against the U.S. Are Not Working," *Washington Post*, March 19, 2015 https://www.washingtonpost.com/world/national-security/head-of-cyber-command-us-may-need-to-boost-offensive-cyber-powers/2015/03/19/1ad79a34-ce4e-11e4-a2a7-9517a3a70506_story.html. See also James Cartwright statement, *National Defense Authorization Act for Fiscal Year 2008 and Reauthorization of Existing Programs*, Public Law 110-186, 110th Cong., 1st sess. (2007), p. 65.

7. James Clapper, "Statement for the Record: Worldwide Threat Assessment of the US Intelligence Community," Senate Select Committee on Intelligence, 114th Cong., 2nd sess., February 9, 2016, https://www.dni.gov/files/documents/SSCI_Unclassified_2016_ATA_SFR%20_FINAL.pdf.

8. Elisabeth Bumiller and Thom Shanker, "Panetta Warns of Dire Threat of Cyberattack on U.S.," *New York Times*, October 11, 2012, http://www.nytimes.com/2012/10/12/world/panetta-warns-of-dire-threat-of-cyberattack.html?_r=0.

9. The operating budget for the U.S. Cyber Command has more than quadrupled in the past five years, at a time when most other Defense Department budgets have been slashed. See Joe Gould, "Constructing a Cyber Superpower," *Defense News*, June 29, 2015, http://www.defensenews.com/story/defense/policy-budget/cyber/2015/06/27/us-cyber-command-budget-expand-fort-meade-offensive/28829321/. For reactions from China and Russia, see Bill Gertz, "China Sharply Boosts Cyber Warfare Funding," *Free Beacon*, April 1, 2015, http://freebeacon.com/national-security/china-sharply-boosts-cyber-warfare-funding/; and Eugene Gerden, "$500 Million for New Russian Cyber Army," *SCmagazineuk.com*, November 6, 2014, http://www.scmagazineuk .com/500-million-for-new-russian-cyber-army/article/381720/.

10. Ben Buchanan, *The Cybersecurity Dilemma: Hacking, Trust, and Fear between Nations* (London: C. Hurst, 2016). Buchanan develops Jervis's conception of the security dilemma. See Jervis, "Cooperation under the Security Dilemma." For more on escalation, see William A. Owens, Kenneth W. Dam, and Herbert S. Lim, *Technology, Policy, Law, and Ethics Regarding U.S. Acquisition and Use of Cyberattack Capabilities* (Washington, D.C.: National Academies Press, 2009); Vincent Manzo, "Deterrence and Escalation in Cross-Domain Operations: Where Do Space and Cyberspace Fit?" Strategic Forum No. 272 (Washington, D.C.: National Defense University, December 2011); Herb

These risks raise urgent questions. What is meant by cyber offensive advantage? Can the offense-defense balance be measured, and if so how? What are the causes of offensive and defensive advantage, and can they be distinguished? Under what conditions do cyber operations favor the offense, and what are the implications of this for policy?

This article provides a framework for reasoning about the offense-defense balance of cyber operations and shows that sweeping claims about offensive advantage in cyberspace are deeply misguided. I make four arguments. First, a useful conception of offense-defense balance must include consideration of the value that decisionmakers place on cyber operations, and not merely the costs. I identify three distinctive notions of cyber offense-defense balance, showing that all can be framed in terms of relative utility (i.e., the benefits and costs of offense relative to the benefits and costs of defense). This approach makes the stakes of cyber conflict more explicit and exposes some false assumptions in the dominant conception of offensive advantage.

Second, I analyze the sources of cyber offensive or defensive advantage, arguing that they are determined not by technology alone, but by the organizational processes that govern interactions between technology and skilled actors—processes such as software updating, vulnerability scanning, and access management. Although in theory the complexity of information technology offers the offense advantages, in practice information technology, skills, and organizations are not easily separable variables. As a result, 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. One can assess the offense-defense balance of cyber operations between two adversaries, but not of cyberspace; the balance is a dyadic, not a systemic, variable.[11] Additionally, the advantages that complex information technology offers the offense do not extend to the physical world, making

Lin, "Escalation Dynamics and Conflict Termination in Cyberspace," *Strategic Studies Quarterly*, Vol. 6, No. 3 (Fall 2012), pp. 46–70; and Avery Goldstein, "First Things First: The Pressing Danger of Crisis Instability in U.S.-China Relations," *International Security*, Vol. 37, No. 4 (Spring 2013), pp. 49–89.

11. I define "cyber operations" to include efforts to compromise the confidentiality, integrity, or availability of digital electronic computers or communications, or to prevent such compromises. I define "cyberspace" as the infrastructure that enables digital electronic computing and communications. This includes elements that are "air gapped" from the internet, as these computers are nonetheless connected by the "sneaker net" (i.e., people who move data on flash drives or other media). This infrastructure is physical, even though its purpose is informational. For alternate definitions, see Brandon Valeriano and Ryan C. Maness, *Cyber War versus Cyber Realities: Cyber Conflict in the International System* (Oxford: Oxford University Press, 2015), pp. 22–24; Richard A. Clarke and Robert Knake, *Cyber War: The Next Threat to National Security and What to Do about It* (New York: HarperCollins, 2010), p. 70; Choucri, *Cyberpolitics in International Relations*, p. 8; and David J. Betz and Tim Stevens, *Cyberspace and the State: Toward a Strategy for Cyber-Power* (New York: Routledge, 2011), pp. 35–38.

cyber offenses much more expensive for achieving physical effects than for conducting espionage and deception.[12]

Third, by analyzing the distinctive sources of offensive or defensive advantage, I contribute to efforts to empirically distinguish and assess the offensive and defensive capabilities of particular actors. Some scholars have assessed such capabilities for different nation-states, assigning scores for cyber-defense capabilities, cyber-offense capabilities, and cyber dependence, and treating the sum of the scores as an indication of their overall capability in cyber conflict.[13] Such work typically assumes that defensive capabilities correlate with state control over civilian networks—an assumption that has been challenged—and scores for offensive capabilities are largely based on the authors' judgments of available evidence.[14] The scores are thus descriptive, but not explanatory. The explanation for offense-defense balance advanced here, based on the relative skill of organizations and level of complexity they must manage, suggests that the success of the offense to date has largely been the result of poor management and the relatively limited goals of offense, rather than a technologically determined offensive advantage.

Fourth, I conduct a cost-benefit analysis of Stuxnet—the U.S.-Israeli cyberattack on Iran's uranium enrichment facilities. This analysis suggests that the defense was likely less costly than the offense in the Stuxnet attack, contrary to dominant assumptions about cyber offense dominance. Perhaps most significantly, 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.

In the remainder of this article, I first outline distinctive conceptions of cyber

12. This point supports others' observations that cyberspace seems to have the greatest impact on espionage. See Rid, *Cyber War Will Not Take Place*; Erik Gartzke and Jon R. Lindsay, "Weaving Tangled Webs: Offense, Defense, and Deception in Cyberspace," *Security Studies*, Vol. 24, No. 2 (2015), pp. 316–348; Erik Gartzke, "The Myth of Cyberwar: Bringing War in Cyberspace Back Down to Earth," *International Security*, Vol. 38, No. 2 (Fall 2013), pp. 41–73; Lindsay, "Stuxnet and the Limits of Cyber Warfare"; Lindsay, "The Impact of China on Cybersecurity"; Maness and Valeriano, *Cyber War versus Cyber Realities*; Derek S. Reveron, ed., *Cyberspace and National Security: Threats, Opportunities, and Power in a Virtual World* (Washington, D.C.: Georgetown University Press, 2012); and Jon R. Lindsay, Tai Ming Cheung, and Derek S. Reveron, eds., *China and Cybersecurity: Espionage, Strategy, and Politics in the Digital Domain* (New York: Oxford University Press, 2015).

13. Clarke and Knake, *Cyber War*; and Maness and Valeriano, *Cyber War versus Cyber Realities*.

14. The assumption that countries such as China have better defenses because they attempt to censor their networks is not necessarily well founded. Scholars have shown that the "great firewall of China" can be easily circumvented and even manipulated into a denial-of-service attack. See Richard Clayton, Steven J. Murdoch, and Robert N.M. Watson, "Ignoring the Great Firewall of China," paper presented at the Sixth Workshop on Privacy Enhancing Technologies, Robinson College, Cambridge, United Kingdom, June 28–30, 2006; and David J. Betz and Tim Stevens, "Analogical Reasoning and Cyber Security," *Security Dialogue*, Vol. 44, No. 2 (April 2013), pp. 147–164. Additionally, China struggles with cybercrime just as the United States does. See Lindsay, Cheung, and Reveron, *China and Cybersecurity*.

offensive advantage and reframe them in terms of relative utility. The second section articulates five different issues that shape the relative utility of offensive and defensive cyber operations. The third section presents an empirical cost-benefit analysis of Stuxnet. In conclusion, I summarize and consider policy implications of these findings.

Conceptualizing Cyber Offense-Defense Balance

The idea that some technologies make either offense or defense easier is at least as old as the 1930s League of Nations discussions on reducing or limiting arms, but the formalization of offense-defense theory came much later.[15] Some analysts argue that offense-defense balance should be defined narrowly in terms of the effects of military technology on the costs of offense and defense for the entire international system, while acknowledging that other factors (such as military skill or resources) affect battle outcomes between any specific dyad of states.[16] Others argue that the balance should be defined more broadly to include military technology and doctrine, geographic advantages, social and political order, and diplomatic behavior.[17] Most concepts share the idea that individual technologies or systems of technologies possess intrinsic features that tend to make either offense or defense easier.[18]

Although offense-defense theory has been critiqued on several grounds, it retains its intuitive appeal and has been extended to cyber conflict.[19] Most scholars argue that cyberspace favors the offense. Joseph Nye writes, "Because the internet was designed for ease of use rather than security, the offense currently has the advantage over the defense."[20] Kenneth Lieberthal and Peter Singer similarly argue that offense has the advantage in cyberspace because

15. Foundational works include Jervis, "Cooperation under the Security Dilemma"; and Quester, *Offense and Defense in the International System*.
16. See, for example, Lynn-Jones, "Offense-Defense Theory and Its Critics," p. 665.
17. Van Evera, "Offense, Defense, and the Causes of War"; and Glaser and Kaufmann, "What Is the Offense-Defense Balance and Can We Measure It?"
18. For a good review of concepts, see Levy, "The Offensive/Defensive Balance of Military Technology." Some formulations of offense-defense theory deemphasize military technology and focus more on the operations and strategies of attackers or defenders. See Stephen Biddle, *Military Power: Explaining Victory and Defeat in Modern Battle* (Princeton, N.J.: Princeton University Press, 2004); Stephen Biddle, "Rebuilding the Foundations of Offense-Defense Theory," *Journal of Politics*, Vol. 63, No. 2 (August 2001), pp. 741–774; and Jonathan Shimshoni, "Technology, Military Advantage, and World War I: A Case for Military Entrepreneurship," *International Security*, Vol. 15, No. 3 (Winter 1990/91), pp. 187–215.
19. Many of the original works and critiques can be found in Michael E. Brown et al., eds., *Offense, Defense, and War* (Cambridge, Mass.: MIT Press, 2004). See also Levy, "The Offensive/Defensive Balance of Military Technology"; and Lynn-Jones, "Offense-Defense Theory and Its Critics."
20. Nye, "Cyber Power," p. 5.

"the Internet was designed to share information easily, not prevent its flow," and because "there are many vulnerabilities that can be exploited."[21] Others echo these views, largely focusing on conflict within cyberspace.

A few scholars have instead presented evidence suggesting that cyberspace may favor the defense.[22] Some accept that offense may dominate within cyberspace, but underscore the distinction between cyber conflict and territorial or strategic conflict.[23] For example, Erik Gartzke notes that if command and control is more crucial for offense than defense, as is commonly believed, then offense dominance in cyberspace would actually translate into territorial defense dominance, and vice versa.[24] A sizable literature on cyber threat inflation implicitly supports the thesis that cyber offenses may have been exaggerated for political or other purposes, but does not engage explicitly with offense-defense theory.[25]

Debate over the offense-defense balance is tied to a larger dispute about the impact of cyberspace on international relations. Some analysts argue that cyberspace favors the offense because it offers anonymity to the attacker, preventing meaningful deterrence.[26] In this view, offense dominance gives asym-

21. Lieberthal and Singer, "Cybersecurity and U.S.-China Relations," p. 14.
22. Lindsay, "Stuxnet and the Limits of Cyber Warfare"; Lindsay, "The Impact of China on Cybersecurity"; and Rid, *Cyber War Will Not Take Place*. Rid argues that cyber conflict favors defense for three reasons: sophisticated attacks (such as those that target industrial control systems) require significant resources; cyberweapons are usually easy to defend against once discovered and therefore cannot be reused reliably; and the defensive market is bullish. Rid writes, "[C]ompetition between various computer security companies has heated up, red-teaming is steadily improving, active defense is emerging, and . . . consumers are becoming more security aware." Ibid., p. 169.
23. See, for example, Lieber, "The Offense-Defense Balance and Cyber Warfare."
24. Gartzke, "The Myth of Cyberwar," p. 67. Gartzke and Lindsay also argue that the ease of attack has been mistaken for ease of deception, and that the ease of deception in cyberspace has unclear implications for offense-defense balance. See Gartzke and Lindsay, "Weaving Tangled Webs."
25. Ralf Bendrath, "The Cyberwar Debate: Perception and Politics in U.S. Critical Infrastructure Protection," in "The Internet and the Changing Face of International Relations and Security," ed. Andreas Wenger, special issue, *Information & Security: An International Journal*, Vol. 7 (2001), pp. 80–103; Ralf Bendrath, "The American Cyber-Angst and the Real World—Any Link?" in Robert Latham, ed., *Bombs and Bandwidth: The Emerging Relationship between Information Technology and Security* (New York: Free Press, 2003), pp. 49–73; Myriam Dunn Cavelty, *Cyber-Security and Threat Politics: UA Efforts to Secure the Information Age* (New York: Routledge, 2007); Sean Lawson, "Beyond Cyber-Doom: Assessing the Limits of Hypothetical Scenarios in the Framing of Cyber-Threats," *Journal of Information Technology & Politics*, Vol. 10, No. 1 (2013), pp. 86–103; and Lene Hansen and Helen Nissenbaum, "Digital Disaster, Cyber Security, and the Copenhagen School," *International Studies Quarterly*, Vol. 53, No. 4 (2009), pp. 1155–1175. For a comparison of constructivist approaches to cybersecurity with those of other international relations theories, see Johan Eriksson and Giampiero Giacomello, "The Information Revolution, Security, and International Relations: (IR)Relevant Theory?" *International Political Science Review*, Vol. 27, No. 3 (July 2006), pp. 221–244.
26. For work emphasizing the difficulty of deterrence, see Libicki, "Cyberdeterrence and Cyberwar"; Stephen Blank, "Can Information Warfare Be Deterred?" in David S. Alberts and Daniel S. Papp, eds., *Information Age Anthology*, Vol. 3: *The Information Age Military* (Washington, D.C.:

metric advantages to weaker actors.[27] Others argue that deterrence is still possible, and that even though offense may dominate within cyberspace, this will not significantly alter the current global distribution of power.[28]

THREE CONCEPTIONS OF CYBER OFFENSE-DEFENSE BALANCE

The debate remains muddy in part because the offense-defense balance is rarely defined, let alone empirically assessed. Below I outline three distinctive ways in which scholars and military thinkers have implicitly defined the cyber offense-defense balance.

RELATIVE COSTS OF OFFENSE AND DEFENSE. Within traditional offense-defense theory, the balance is most often defined as the ratio of the cost of attacking territory to the cost of defending it; this balance may vary with the size of the offense and defense.[29] Many scholars extend this approach to cyberspace by focusing on the relative costs of offense and defense within cyberspace.[30] One

Department of Defense C4ISR Cooperative Research Program, 2001), pp. 125–158; and Arquilla, "Cyberwar Is Already upon Us."

27. Gregory J. Rattray, *Strategic Warfare in Cyberspace* (Cambridge, Mass.: MIT Press, 2001); Clarke and Knake, *Cyber War*; Nye, "Cyber Power"; Demchak, *Wars of Disruption and Resilience*; David C. Gompert and Phillip C. Saunders, *Paradox of Power: Sino-American Strategic Restraint in an Age of Vulnerability* (Washington, D.C.: National Defense University, 2011); Choucri, *Cyberpolitics in International Relations*; Andrew Krepinevich, "Cyberwarfare: A 'Nuclear Option?'" (Washington, D.C.: Center for Strategic and Budgetary Assessments, 2012); Kello, "The Meaning of the Cyber Revolution"; and Ilai Saltzman, "Cyber Posturing and the Offense-Defense Balance," *Contemporary Security Policy*, Vol. 34, No. 1 (2013), pp. 40–63.

28. Libicki, *Conquest in Cyberspace*; Richard L. Kugler, "Deterrence of Cyber Attacks," in Franklin D. Kramer, Stuart H. Starr, and Larry K. Wentz, eds., *Cyberpower and National Security* (Washington, D.C.: Potomac, 2009), pp. 309–342; Libicki, "Cyberdeterrence and Cyberwar"; Will Goodman, "Cyber Deterrence: Tougher in Theory Than in Practice?" *Strategic Studies Quarterly*, Vol. 4, No. 3 (Fall 2010), pp. 102–135; Liff, "Cyberwar"; Gartzke, "The Myth of Cyberwar"; Lindsay, "Stuxnet and the Limits of Cyber Warfare"; Lieber, "The Offense-Defense Balance and Cyber Warfare"; Jon R. Lindsay, "Tipping the Scales: The Attribution Problem and the Feasibility of Deterrence," *Journal of Cybersecurity*, Vol. 1, No. 1 (2015), pp. 53–67; Maness and Valeriano, *Cyber War versus Cyber Realities*; and Thomas Rid and Ben Buchanan, "Attributing Cyber Attacks," *Journal of Strategic Studies*, Vol. 38, Nos. 1–2 (2015), pp. 4–37. Others argue for a new framework for deterrence in cyberspace. See Richard B. Andres, "The Emerging Structure of Strategic Cyber Offense, Cyber Defense, and Cyber Deterrence," in Reveron, *Cyberspace and National Security*, pp. 89–104; Jeffrey R. Cooper, "A New Framework for Cyber Deterrence," in Reveron, *Cyberspace and National Security*, pp. 105–120; and Uri Tor, "'Cumulative Deterrence' as a New Paradigm for Cyber Deterrence," *Journal of Strategic Studies* (online edition), 2015, http://www.tandfonline.com/doi/pdf/10.1080/01402390.2015.1115975?needAccess=true.

29. For a list of alternate measures, see Glaser and Kaufmann, "What Is the Offense-Defense Balance and Can We Measure It?" Karen Ruth Adams includes deterrence in the balance. See Adams, "Attack and Conquer? International Anarchy and the Offense-Defense-Deterrence Balance," *International Security*, Vol. 28, No. 3 (Winter 2003/04), pp. 45–83. Biddle argues that minimum costs needed to defend or attack are not good measures because they are not observable. See Biddle, "Rebuilding the Foundations of Offense-Defense Theory."

30. Patrick J. Malone, "Offense-Defense Balance in Cyberspace: A Proposed Model," Naval Postgraudate School, 2012; Saltzman, "Cyber Posturing and the Offense-Defense Balance"; and Lieber, "The Offense-Defense Balance and Cyber Warfare."

report states, "The cyber competition appears to be an offense-dominant competition. If both the attacker and defender are given equal resources, the attacker will prevail."[31] The most detailed analysis to date, a master's thesis completed by Patrick Malone at the Naval Postgraduate School, models the balance in terms of the total costs of software, hardware, and wages for personnel on both sides.[32]

EFFICACY OF CYBER OFFENSE. A second conception of cyber offensive advantage presumes that offensive operations are low cost and have a high payoff for the offense, whereas defensive operations are expensive and ineffective. This notion of offensive advantage differs from the first in focusing more on the offensive payoff. One report states, "It is unlikely that cyber criminals would persist in their activities if the competition favored the defense. In that case, crime would 'not pay.'"[33] John Arquilla notes that the 2007 cyberattacks on Estonia cost the attackers very little but had a high payoff in terms of disruption to the Estonian government.[34] The notion of efficacy also appears in the traditional offense-defense theory literature. Gilpin writes, "The defense is said to be superior if the resources required to capture territory are greater than the value of the territory itself; the offense is superior if the cost of conquest is less than the value of the territory."[35] A focus on efficacy, however, has more continuity with the literature on expected utility than with the traditional offense-defense theory literature, as discussed further below.[36]

FIRST-MOVER ADVANTAGES. The 2006 Air Force Cyberspace Task Force states that "[o]ffensive capabilities in cyberspace offer both the US and our adversaries an opportunity to gain and maintain the initiative," suggesting a focus on first-mover advantages.[37] This focus also has roots in traditional offense-defense theory.[38] Some argue that cyberattackers do not enjoy a first-mover advantage for "purely" cyber conflicts because a first-strike cyberattack could not render an enemy's cyber capabilities completely ineffective.[39] Even if this argument is accepted, however, first-mover advantages need not be confined to cyberspace; the ability to disrupt the command and control capabilities of a

31. Krepinevich, "Cyberwarfare," p. 40.
32. Malone, "Offense-Defense Balance in Cyberspace."
33. Krepinevich, "Cyberwarfare," p. 40 n. 141.
34. Arquilla, "Cyberwar Is Already upon Us."
35. Gilpin, *War and Change in World Politics*, p. 63.
36. Gilpin uses multiple definitions (both cost based and efficacy based), a point that Levy critiques. See Levy, "The Offensive/Defensive Balance of Military Technology," p. 227.
37. Lani Kass, "A Warfighting Domain," Air Force Cyberspace Task Force Presentation, U.S. Air Force Headquarters, Washington, D.C., September 26, 2006, http://www.au.af.mil/info-ops/usaf/cyberspace_taskforce_sep06.pdf.
38. Jervis, "Cooperation under the Security Dilemma," p. 188.
39. Libicki, *Conquest in Cyberspace*; and Buchanan, *The Cybersecurity Dilemma*.

military force could be seen as enabling a decisive territorial victory. Scholars and policymakers have suggested, for example, that Russian cyberattacks on Georgia helped pave the way for Russian tanks to roll over the border in the 2008 Russo-Georgian war.[40] Others argue that the effects of cyberattacks were minimal in this case.[41] The question of whether such cyber operations provide a first-mover advantage and how it should be measured is currently irresolvable, because there is no consensus on how to calculate first-mover advantages for either cyber operations or cross-domain operations.

A common framework for measurement: utility

Although multiple conceptions of offense-defense balance have been implied, the balance has rarely been explicitly defined or operationalized to enable empirical measurement. This article begins that task by noting that all of the preceding conceptions of offensive advantage can be expressed in terms of utility, a notion with roots in economics that has been adopted in game-theoretic and bargaining theories of war.[42] The utility of cyber offense is the value of the offensive goal (e.g., taking territory, stealing secrets, or gaining control of a computer) less the minimum costs of achieving it; the utility of the defense is the value of the defensive goal (e.g., holding territory, maintaining secrecy, keeping control of a computer) less the minimum costs of defense. When the utility is greater for the offense than the defense, the situation may be said to favor the offense, and vice versa.

Both sides aim to maximize their expected utility from outcomes that may be achieved with varying degrees of probability.[43] This article, however, does

40. Lieberthal and Singer, "Cybersecurity and U.S.-China Relations"; Arquilla, "Cyberwar Is Already upon Us"; Krepinevich, "Cyberwarfare"; Kello, "The Meaning of the Cyber Revolution"; Libicki, "Cyberdeterrence and Cyberwar"; Clarke and Knake, *Cyber War*; and Jon R. Lindsay and Lucas Kello, "Correspondence: A Cyber Disagreement," *International Security*, Vol. 39, No. 2 (Fall 2014), pp. 181–192.

41. Lindsay, "Tipping the Scales," p. 62; and Martin C. Libicki, *Cyberspace in Peace and War* (Annapolis: Naval Institute Press, 2016), p. 12.

42. Bruce Bueno de Mesquita, *The War Trap* (New Haven, Conn.: Yale University Press, 1981); Bruce Bueno de Mesquita, "The War Trap Revisited: A Revised Expected Utility Model," *American Political Science Review*, Vol. 79, No. 1 (March 1985), pp. 156–177; Bruce Bueno de Mesquita, "The Contribution of Expected Utility Theory to the Study of International Conflict," *Journal of Interdisciplinary History*, Vol. 18, No. 4 (Spring 1988), pp. 629–652; James D. Fearon, "Rationalist Explanations for War," *International Organization*, Vol. 49, No. 3 (Summer 1995), pp. 379–414; and Dan Reiter, "Exploring the Bargaining Model of War," *Perspectives on Politics*, Vol. 1, No. 1 (March 2003), pp. 27–43.

43. Expected utility theory has been used in the economics of information security, but it generally aims to identify optimal strategies under conditions of uncertainty, rather than to empirically determine the offense-defense balance. For examples of the information security literature on this topic, see Sankardas Roy et al., "A Survey of Game Theory as Applied to Network Security," paper presented at the Forty-third Hawaii International Conference on System Sciences, Koloa, Kauai, Hawaii, January 5–8, 2010; Ross Anderson and Tyler Moore, "Economics and Internet Security: A Survey of Recent Analytical, Empirical, and Behavioral Research (Technical Report Tr-03-11)"

not attempt to analyze expected utility for two reasons. First, the offense-defense balance has predominantly been defined as a comparison between two different outcomes that are achieved with certainty: one in which the offense expends the minimum necessary to succeed (e.g., to offset the defensive investment), and another in which the defense spends the minimum necessary to succeed (e.g., to offset the offensive investment).[44] Probabilistic considerations complicate comparison with traditional offense-defense theory. Second, data for predicting the probability of successful cyber operations do not exist, and there is reason to doubt that precise probabilistic data will be produced, as discussed further below. Nonetheless, both retrospective empirical analysis of the utility of cyber operations and theoretical analysis of what drives the costs and value of cyber operations can inform expectations. This article provides such analyses.

The dominant conception of the offense-defense balance in the academic literature—relative cost of offense and defense—is identical to relative utility under conditions in which the offense and defense value their goals equally. This approach, however, fails to account for the variable ways in which offense and defense may value their respective goals, a critique Richard Betts makes of territorially based offense-defense theory.[45] Furthermore, cyber operations have even more highly variable goals than territorial conflict, including sabotage, espionage, and military operational advantages.

Considering the variable valuation of the goals of cyber operations reveals logical flaws in some of the dominant arguments for offensive advantage. For example, contrary to what some have suggested, the fact that cybercrime pays does not necessarily mean that it is more expensive to defend than to attack.[46] It could simply mean that criminals value their gains more than companies value the loss, a situation in which the defense is unlikely to allocate the resources needed to stop crime. In fact, game-theoretic analyses have shown that cyber defenders acting under conditions of uncertainty may rationally choose not to fully defend themselves in order to maximize their gains and minimize their losses.[47] More importantly, cybercrime typically refers to individuals

(Cambridge, Mass.: Harvard Computer Science Group, 2011); and Tyler Moore, Allan Friedman, and Ariel D. Procaccia, "Would a 'Cyber Warrior' Protect Us? Exploring Trade-Offs between Attack and Defense of Information Systems" (Concord, Mass.: New Security Paradigms Workshop, 2010).

44. This approach has been critiqued by Biddle, who argues that minimum costs needed to defend or attack are not observable; he suggests that probability of victory is one of multiple better metrics. Biddle, "Rebuilding the Foundations of Offense-Defense Theory," p. 749 n. 12.

45. Richard K. Betts, "Must War Find a Way? A Review Essay," *International Security*, Vol. 24, No. 2 (Fall 1999), pp. 166–198.

46. Krepinevich, "Cyberwarfare," p. 40 n. 141.

47. Rainer Böhme and Tyler Moore, "The Iterated Weakest Link: A Model of Adaptive Security Investment," paper presented at the Workshop on the Economics of Information Security, London,

seeking personal gain, rather than to nation-states seeking strategic advantage; the costs and benefits of success are likely to be very different in these different contexts, but discussions of cyber offensive advantage often conflate them.

The second conception of offensive advantage noted above, efficacy, explicitly includes the payoff of a successful offense, and thus is more appropriate to an analysis of cyber conflict. The third conception outlined above, first-mover advantage, is explicitly focused on a particular kind of utility. First-mover advantages are typically considered for cyber operations used in conjunction with other domains; for example, a cyberattack might disrupt an air defense system, allowing bombers to enter enemy airspace with no attrition, as Israel reportedly did in its 2007 bombing of suspected nuclear materials sites in Syria.[48] Although such a cyber operation has a physical payoff, it is analytically important to distinguish the utility of the cyber operation and the utility of the bombing operation because they function by different mechanisms that have different cost structures, as discussed further below.[49]

In short, a relative utility conception of the offense-defense balance is better than alternatives because it explicitly includes valuations of the goals of cyber operations, and thus better reflects the considerations of decisionmakers. In what follows, I analyze factors that will shape the relative utility of cyber operations (i.e., their valuation and costliness).

What Drives the Value and Costs of Cyber Operations?

I argue that the cost of cyber operations will depend not on the features of technology alone, but instead on the skills and competence of the actors and organizations that continually create, use, and modify information technology. This argument stands in stark contrast to most existing theories of the offense-defense balance, which focus on technological affordances such as mobility or firepower.[50] Some analysts similarly focus on the technological affordances of

June 24–25, 2009. See also Lawrence A. Gordon and Martin P. Loeb, "Budgeting Process for Information Security Expenditures," *Communications of the ACM*, Vol. 49, No. 1 (January 2006), pp. 121–125.

48. Sharon Weinberger, "How Israel Spoofed Syria's Air Defense System," *Wired*, October 4, 2007, https://www.wired.com/2007/10/how-israel-spoo/.

49. For example, a calculation of the utility of a cyber operation that disables Syria's air defense system would include the costs of the cyber operation and the payoff in terms of a more effective bombing operation. The utility of the bombing, however, could be considered separately, in terms of the cost of bombing and the value of success (e.g., the value Israel placed on destroying the facilities). Although the utility of the bombing operation would be influenced by the utility of the cyber operation, the two can be analytically distinguished, provided that they have a common framework for valuing goals.

50. For a review of such arguments, see Levy, "The Offensive/Defensive Balance of Military Technology," pp. 225–227; and Van Evera, "Offense, Defense, and the Causes of War," pp. 16–18.

cyberspace, suggesting that speed, "ease of use," or "versatility" of information technology favors offense.[51] Such properties are not intrinsic to information technology, however; they arise out of the interactions between technology and skilled actors. As the discussion below shows, a cyberweapon may operate very quickly, but building and deploying it is a slow, labor-intensive process.[52]

The approach advocated here shares some commonality with territorial theories of the offense-defense balance that have focused on skills and operations.[53] However, it differs from virtually all offense-defense theories in a key respect: I argue that technology, organizations, and skill are not easily separable variables, particularly not in cyberspace.[54] As work in science and technology studies has shown, information technology is constantly structuring and being structured by organizations and skilled actors.[55] Thus, cyberspace cannot be treated as a "black box," something that can be understood primarily in terms of its external effects.[56] Understanding both the costs and the value of

51. For an example citing "ease of use," see Nye, "Cyber Power." Ilai Saltzman proposes to "cybernate" offense-defense theory by replacing territorial mobility (which is presumed to favor offense) with versatility in cyber operations, and by replacing "kinetic firepower" (which is presumed to favor defense) with "byte power." It is unclear, however, how "versatility" or "byte power" would be operationalized. See Saltzman, "Cyber Posturing and the Offense-Defense Balance."

52. For work highlighting the time-intensive nature of designing and deploying advanced cyberweapons, see Liff, "Cyberwar"; Rid, "Cyber War Will Not Take Place"; and Lindsay, "Stuxnet and the Limits of Cyber Warfare."

53. See Shimshoni, "Technology, Military Advantage, and World War I"; and Biddle, "Rebuilding the Foundations of Offense-Defense Theory." For similar arguments that are focused not on offense-defense theory, but on military effectiveness and technological innovation and use, see Brian A. Jackson, "Technology Acquisition by Terrorist Groups: Threat Assessment Informed by Lessons from Private Sector Technology Adoption," *Studies in Conflict & Terrorism*, Vol. 24 (2001), pp. 183–213; and Stephen Biddle and Robert Zirkle, "Technology, Civil-Military Relations, and Warfare in the Developing World," *Journal of Strategic Studies*, Vol. 19, No. 2 (June 1996), pp. 171–212.

54. The distinction between skill and technology is implicit in most accounts, although Glaser and Kaufmann make this distinction explicit. See Glaser and Kaufmann, "What Is the Offense-Defense Balance and Can We Measure It?" pp. 48–49. Biddle, though he underscores the importance of how technologies are used, nonetheless treats technology as a systemic variable that is separate from force employment (a dyadic variable). See Biddle, "Rebuilding the Foundations of Offense-Defense Theory," p. 749.

55. This approach is informed by work in organizational sociology and science and technology studies that examines the mutual shaping of information technology and organizations. See, for example, Wanda J. Orlikowski, "The Duality of Technology: Rethinking the Concept of Technology in Organizations," *Organization Science*, Vol. 3, No. 3 (August 1992), pp. 398–427; Wanda J. Orlikowski, "Using Technology and Constituting Structures: A Practice Lens for Studying Technology in Organizations," *Organization Science*, Vol. 11, No. 4 (July/August 2000), pp. 404–428; and Paul N. Edwards, "From Impact to Social Process," in Shiela Jasanoff et al., eds., *Handbook of Science and Technology Studies* (Thousand Oaks, Calif.: Sage, 1995), pp. 257–285.

56. For more on black-boxing science and technology, see Trevor J. Pinch and Weibe E. Bijker, "The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other," *Social Studies of Science*, Vol. 14, No. 3 (August 1984), pp. 399–441; Bruno Latour, *Pandora's Hope: Essays on the Reality of Science Studies* (Cambridge,

cyber operations requires a close examination of the processes of creating and using information technology.

The following sections discuss five dynamics that shape the relative utility of offensive and defensive cyber operations. First, the tight coupling of individual skills and information technology make the economics of producing cyberweapons fundamentally different from those of producing physical weapons. Second, the costs of cyber operations increase with the complexity of the requisite information technology and decrease with organizational competence in managing complex technology. Third, the relatively greater cost of defensive cyber operations stems from the greater complexity of the defensive goals; when the offensive goals grow more complex, offensive costs also rise. Fourth, the advantages that complexity offers the offense diminish at the "edges" of cyberspace, where digital control systems meet physical equipment. Fifth, the value of cyber operations must be considered in relation to the complex political objectives of adversaries.

THE INTERTWINING OF SKILL AND INFORMATION TECHNOLOGY

The intertwining of skill and technology is not unique to cyberweapons. The word "technology" has historically been used interchangeably with phrases such as "the useful arts" and words such as "manufacturing," and "invention" (i.e., the making of technology rather than the final artifacts).[57] Furthermore, users of technology can be innovators who repurpose technology in ways their inventors never imagined.[58] Combat histories are full of innovations made in the exigencies of the moment, suggesting that technological effects are not easily separable from the skills of users.[59] Nonetheless, it may be that on

Mass.: Harvard University Press, 1999); and Edwin T. Layton, "Conditions of Technological Development," in Ina Spiegel Rosing and Derek de Solla Price, eds., *Science, Technology, and Society: A Cross-Disciplinary Perspective* (London: Sage, 1977), p. 198.

57. Ruth Oldenziel, *Making Technology Masculine: Men, Women, and Modern Machines in America, 1870–1945* (Amsterdam: Amsterdam University Press, 1999); Eric Schatzberg, "*Technik* Comes to America: Changing Meanings of Technology before 1930," *Technology and Culture*, Vol. 47, No. 3 (2006), pp. 486–512; and Leo Marx, "Technology: The Emergence of a Hazardous Concept," *Technology and Culture*, Vol. 51, No. 3 (July 2010), pp. 561–577.

58. Ruth Cowan, "The Consumption Junction: A Proposal for Research Strategies in the Sociology of Technology," in Wiebe Bijker, Thomas P. Hughes, and Trevor Pinch, eds., *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology* (Cambridge, Mass.: MIT Press, 1987), pp. 253–272; Nelly Oudshoorn and Trevor Pinch, *How Users Matter: The Co-Construction of Users and Technology* (Cambridge, Mass.: MIT Press, 2003); Janet Abbate, *Inventing the Internet* (Cambridge, Mass.: MIT Press, 1999); Ronald R. Kline, *Consumers in the Country: Technology and Social Change in Rural America* (Baltimore, Md.: John Hopkins University Press, 2000); and Nelly Oudshoorn and Trevor Pinch, "User-Technology Relationships: Some Recent Developments," in Edward J. Hackett et al., eds., *The Handbook of Science and Technology Studies* (Cambridge, Mass.: MIT Press, 2008).

59. John H. Hay, *Vietnam Studies: Tactical and Materiel Innovations* (Washington, D.C.: Department

the physical battlefield, technology and skills can be treated as approximately separate variables.

By contrast, cyberweapons are inseparable from skills.[60] Alan Paller, a founder of the SANS Institute (an information security research and training organization), notes that in cyberspace "the skills are the weapon."[61] Skill levels vary widely. An early study found that the speed of programmers varied by a factor of nearly 30, whereas the speed with which their programs ran on the computer varied by more than a factor of ten.[62] There are also order of magnitude differences in programmers' abilities to find bugs in code—a key skill for both offensive and defensive cyber operations.[63] After the size of the software development project, personnel factors are seen as the second-largest cost driver for software.[64]

The key point here is that skill is not neatly separable from the technologies of cyber conflict. Any user of an office computer knows that after a few years, the technology has changed. Unless updating features are disabled, software is continually modified by skilled programmers who develop and distribute patches. Computer users frustrated with disruptive automatic updates may disable them, however. It is thus organizational processes—including software companies' bug hunting, patch development, and distribution processes as well as user organizations' adoption or rejection of such updates—that keep computer technology tightly coupled to skilled actors. I return to this point below.

---

of the Army, 1974); Timothy T. Lupfer, *The Dynamics of Doctrine: The Change in German Tactical Doctrine during the First World War* (Fort Leavenworth, Kans.: Combat Studies Institute, U.S. Army Command and General Staff College, 1981); Center of Military History, *Improvisations during the Russian Campaign* (Washington, D.C.: United States Army, 1986); Michael D. Doubler, *Closing with the Enemy: How GIs Fought the War in Europe, 1944–1945* (Lawrence: University Press of Kansas, 1995); James Jay Carafano, *GI Ingenuity: Improvisation, Technology, and Winning World War II* (Mechanicsburg, Pa.: Stackpole, 2006); and Jon R. Lindsay, "'War upon the Map': User Innovation in American Military Software," *Technology and Culture*, Vol. 51, No. 3 (July 2010), pp. 619–651.

60. For works emphasizing the close relationship between skilled users and information technology, see Eric von Hippel, *The Sources of Innovation* (New York: Oxford University Press, 1988); Abbate, *Inventing the Internet*; Eric von Hippel, *Democratizing Innovation* (Cambridge, Mass.: MIT Press, 2005); and Barbara Van Schewick, *Internet Architecture and Innovation* (Cambridge, Mass.: MIT Press, 2010).

61. Quoted in Anna Mulrine, "Cyber Security: The New Arms Race for a New Front Line," *Christian Science Monitor*, September 15, 2013, http://www.csmonitor.com/USA/Military/2013/0915/Cyber-security-The-new-arms-race-for-a-new-front-line.

62. H. Sackman, W.J. Erikson, and E.E. Grant, "Exploratory Experimental Studies Comparing Online and Offline Programming Performance," *Communications of the ACM*, Vol. 11, No. 1 (January 1968), pp. 3–11.

63. Bill Curtis, "Substantiating Programmer Variability," *Proceedings of the IEEE*, Vol. 69, No. 7 (July 1981), p. 846.

64. Barry Boehm et al., *Software Cost Estimation Using Cocomo II* (Upper Saddle River, N.J.: Prentice Hall, 2000).

Skills are particularly important because cyberweapons, unlike physical weapons, are readily defeated once they are revealed as weapons. Once defenders recognize that they have been compromised, they can usually fix the vulnerability that the attacker exploited, and help others do the same, rendering the attack obsolete. Malicious code is a "use and lose" weapon.[65] Poorly defended targets may be easy to compromise with recycled malware, but high-value targets are likely to be protected against known threats. Without skilled hackers to continually develop new surprises, cyberweapons are readily made obsolete.

Defensive cyber technologies are similarly inextricable from skilled practices. For example, intrusion detection systems may help analysts identify and mitigate attacks, but these technologies continue to rely on the judgment of skilled defenders. People, rather than technology, remain the weakest link in computer security. For example, tough password rules cannot prevent (and may even encourage) computer users from writing down their passwords and leaving them in visible places.

The close coupling of information technology and skill makes the economics of cyberweapons fundamentally different from physical weapons. After an initial period of research and development, most of the cost of physical weapons comes from physical production. By contrast, most of the cost of cyberweapons comes from the skilled workers who research, develop, and deploy the technology. The cost of reproducing cyberweapons—that is, the cost of copying bits from one location to another—is nearly zero. Software cost-estimation frameworks do not even include it.[66] Although John Arquilla and others suggest that such low-cost replication contributes to offense dominance, such arguments neglect the fact that replicating defensive code is also nearly free.[67] In either case, the costly part of cyber operations, offensive or defensive, is the skilled person behind the keyboard who must develop and deploy the technology.

Offensive and defensive skills can be difficult to distinguish. Ed Skoudis, a trainer at SANS and founder of Counter Hack, explains: "Offensive techniques can be used to achieve defensive ends; defensive means can be used to achieve offensive ends; and, sometimes, the inherent technical skills of offense and defense are actually identical."[68] Skoudis adds that he is "not say-

65. Gartzke, "The Myth of Cyberwar," p. 60.
66. For such frameworks, see Boehm et al., *Software Cost Estimation Using Cocomo II*.
67. Arquilla, "Cyberwar Is Already upon Us."
68. Ed Skoudis, "SANS Penetration Testing: When Offense and Defense Become One" (Bethesda, Md.: SANS Institute, April 8, 2013), http://pen-testing.sans.org/blog/2013/04/08/when-offense-and-defense-become-one. For example, endpoint security suites, which interconnect operating system calls to give a security administrator control over a machine, are technically identical to rootkits—but the latter term is typically associated with an attacker taking malicious control of the

ing that offense and defense always merge, or that they are always the same thing," but that "the techniques behind each have more in common than is typically observed."[69]

ORGANIZATION AND MANAGING ARBITRARY COMPLEXITY

Cyber offense and defense depend not only on the skills of individuals, but also on skilled managers and the organization of workers. A Defense Science Board task force on resilient military systems notes that although offensive teams are typically able to disrupt defensive teams in exercises, most "successful attacks reaching DoD [Department of Defense] networks today result from a personnel failure or out-of-date software in firewalls and detection systems. Most of these attacks are understood and preventable through known signature management (patching), yet DoD defensive systems don't keep up, and attacks continue to penetrate DoD's networks."[70] In other words, the success of offense is largely the result of a poorly managed defense.[71] Additionally, contrary to the image of the lone hacker, the success of offense also depends on organizational capability, a point discussed further in the next section.

The difficulty of organizing for cyber operations stems from a fundamental characteristic of information technology: complexity. Bruce Schneier, a cryptographer and cybersecurity consultant, states that "complexity is the worst enemy of security."[72] Similarly, a 2003 report notes: "Prevention of insecure operating modes in complex systems is difficult to do well and impossible to do cheaply: The defender has to counter all possible attacks; the attacker only has to find one unblocked means of attack."[73]

---

system. The difference is only intent. Similarly, some defenders in capture-the-flag exercises have used traditionally offensive tools such as Metasploit to better control their own systems, detecting intruders and removing their access to the defenders' machines.

69. Ibid.

70. Defense Science Board, "Resilient Military Systems and the Advanced Cyber Threat" (Washington, D.C.: U.S. Department of Defense, 2013), p. 65.

71. It is easy to find dozens of similar statements in the trade literature. For example, one computer security textbook declares that "security is a managerial problem, not a technical one." Michael E. Whitman and Herbert J. Mattord, *Principles of Information Security* (Boston: Cengage Learning, 2011), p. 282. See also Daniel Messler, "Security Is Not a Technology Problem: Why Companies Need to Be Looking at Organizational Issues Instead of Products," August 3, 2007, https://danielmiessler.com/blog/security-is-not-a-technology-problem-why-companies-need-to-be-looking-at-organizational-issues-instead-of-products/; and Julia H. Allen, "Security Is Not Just a Technical Issue" (Washington, D.C.: U.S. Department of Homeland Security, May 13, 2013), https://buildsecurityin.us-cert.gov/articles/best-practices/governance-and-management/security-is-not-just-a-technical-issue.

72. Bruce Schneider, testimony, *Overview of the Cyber Problem—a Nation Dependent and Dealing with Risk: Hearing before the House Select Committee on Homeland Security*, 108th Cong., 1st. sess., June 25, 2003.

73. Daniel Geer et al., "Cyber*in*security: The Cost of Monopoly—How the Dominance of Microsoft's Products Poses a Risk to Security," September 27, 2003, http://cryptome.org/cyberinsecurity.htm.

Complexity gives the offense advantages for purely probabilistic reasons.[74] Imagine a race in which offense and defense go hunting for randomly distributed vulnerabilities—the offense to exploit those vulnerabilities, and the defense to patch them. If there is only one vulnerability, the offense has no intrinsic advantage. But with a vast number of vulnerabilities, it is unlikely that the defense will be able to find and patch every vulnerability before the offense finds and exploits it. The number of vulnerabilities grows with the size and complexity of the computer system, and so do the technological advantages of offense—at least in principle.

In practice, information technology is only useful when it is embedded in social organizations, and this embeddedness can either increase or decrease the advantages of the offense. The organizations that develop software can check for common errors before making hardware-software systems available for use.[75] The defender has complete access to the computer system, whereas the attacker has a more limited set of attack vectors. Organizations that establish good cybersecurity processes, such as continually scanning for vulnerabilities and updating software, help skilled defenders manage complexity.

Ultimately, the complexity of software is created and managed by social organizations. Whereas the complexity of physics is assumed to stem from universal laws, the prominent software engineer Frederick Brooks notes that the complexity of software is "arbitrary . . . forced without rhyme or reason by the many human institutions and systems to which [interfaces] must conform."[76] Arbitrary complexity, changeability, and invisibility are "essential" aspects of software that make it difficult to develop correctly.[77]

Although no silver bullet will ever make trustworthy software easy to develop, good organizational processes can reduce the costs of creating and maintaining software. Software engineers use a "capability maturity model" to characterize an organization's software development and maintenance processes. Mature organizations have been shown to decrease costs while increas-

74. Complexity can also pose different kinds of challenges for the attacker. See Drew Herrick and Trey Herr, "Combating Complexity: Offensive Cyber Capabilities and Integrated War Fighting," George Washingotn University and Harvard Univeristy, February 2016, https://ssrn.com/abstract=2845709

75. For an overview, see Gary McGraw, "Software Security," *IEEE Security & Privacy*, Vol. 2, No. 2 (2004), https://www.cigital.com/blog/software-security/.

76. Frederick Brooks, "No Silver Bullet: Essence and Accidents of Software Engineering," *IEEE Computer*, Vol. 20, No. 4 (April 1987), pp. 11–12.

77. On the twentieth anniversary of Brooks's seminal article, a panel discussion generally agreed that his original argument remained valid. See Frederick Brooks et al., "'No Silver Bullet' Reloaded: Retrospective on 'Essence and Accidents of Software Engineering,'" paper presented at the Twenty-second ACM SIGPLAN Conference on Object-Oriented Programming Languages and Applications, Montreal, Canada, October 21–25, 2007.

ing the quality of software.[78] The Cybersecurity Capability Maturity Model has been developed by the U.S. Department of Energy to help organizations improve their cyber defenses.[79]

In short, just as arbitrary complexity is rooted in human organizations and institutions, the relative advantages of offense or defense depend on the ways in which complex computer systems and skilled actors are integrated and organized. This is one reason that information security training institutes do not recommend security products, but best practices—such as keeping a complete inventory of all devices on a network, monitoring user accounts, and engaging in continuous vulnerability scanning and remediation.[80] The fact remains that most attackers today succeed not because of fundamental technological asymmetries, but because of poor management. One computer scientist with experience in both offensive and defensive cyber operations argues against the notion that defense is more intrinsically expensive than offense: "The supposed asymmetry of cost is actually just lack of defensive coordination."[81] Organizations with mature processes may not be able to keep attackers out, but they can increase the attackers' costs significantly.

More mature organizations should also have a higher probability of successfully defending against a particular threat than less mature organizations, presuming that defensive expenditures are equal. Analysts have yet to develop metrics that can predict such probabilities, however.[82] Similarly, there are no good statistics that would enable analysts to predict the probability of successfully compromising a target at a specified offensive cost. In general, analysts lack good empirical data about the costs of cyber offense and defense, as well

78. Dennis R. Goldenson and Diane L. Gibson, "Demonstrating the Impact and Benefits of CMMI®: An Update and Preliminary Results" (Pittsburgh, Pa.: Carnegie Mellon University Software Engineering Institute, 2003).

79. U.S. Department of Energy, "Cybersecurity Capability Maturity Model (C2M2)" (Washington, D.C.: U.S. Department of Energy, 2014). Similar models have been developed for specific sectors such as the energy sector. See also Center for Infrastructure Assurance and Safety, "The Community Cyber Security Maturity Model" (San Antonio: University of Texas at San Antonio, 2016), http://cias.utsa.edu/the-ccsmm.html.

80. "CIS Critical Security Controls—Version 6.0" (Bethesda, Md.: SANS Institute, n.d.), https://www.sans.org/critical-security-controls.

81. Matthew Monte describes most asymmetries between offense and defense in terms of motivation, knowledge, and access—not technology. See Monte, *Network Attacks and Exploitation: A Framework* (Indianapolis: John Wiley and Sons, 2015), p. 56.

82. Although organizations using the capability maturity model report reduced costs associated with developing and maintaining software, the net benefits of cybersecurity process maturity are uncertain. Proponents of cybersecurity maturity models acknowledge that mature processes incur new costs, and suggest that organizations choose what processes to implement based on risk assessment. Such models are designed to be independent of resources, however, so that organizations of any size can achieve the highest "maturity indicator level." See U.S. Department of Energy, "Cybersecurity Capability Maturity Model (C2M2)," p. 9.

as the probability of success.[83] Nonetheless, understanding that organizational skills and process maturity shape the costs of cyber operations provides a basis for predicting the conditions that are likely to increase or decrease the relative costs of cyber offense and defense.

Distinguishing offensive and defensive organizational competencies

Are the organizational competencies required for offense and defense distinguishable? It is tempting to imagine that offenses require little or no organization; popular legends depict hackers as creative, individualistic geniuses. Lynn suggests that "a couple dozen talented programmers wearing flip-flops and drinking Red Bull can do a lot of damage."[84]

It would be a mistake, however, to imagine serious offense as a spontaneous, undisciplined activity. The offense must also navigate the complexity of software, while trying not to make missteps that lead to detection.[85] Serious attackers keep research notebooks.[86] One computer scientist with experience in offensive cyber operations argues that the costs of offense are similar to the costs of building an infrastructure: "Breaking into a particular network may be cheap after the tools and infrastructure are in place," but "building and maintaining the infrastructure for a program of sustained operations requires targeting, research, hardware engineering, software development, and training. This is not cheap."[87] A well-organized offense can attack at relatively low cost, but the infrastructure must be built first. The cost of developing and maintaining that infrastructure grows with its complexity.

The importance of organization for cyber defense is implicit in discussions of the need for better public-private partnerships and information sharing across many agencies of government, but much less has been written about organizing for offense.[88] Chris Demchak provides perhaps the most detailed

83. For discussion of problems with cybercrime statistics, see Dinei Florencio and Cormac Herley, "Sex, Lies, and Cyber-Crime Surveys" (Richmond, Wash.: Microsoft Research, 2011).
84. William Lynn, "Remarks on Cyber at the RSA Conference," San Francisco, California, February 15, 2011 (Washington, D.C.: U.S. Department of Defense, 2011), http://archive.defense.gov/speeches/speech.aspx?speechid=1535.
85. Herrick and Herr, "Combating Complexity."
86. Clifford Stoll, who stalked a hacker whom he discovered in Lawrence Berkeley National Laboratory's networks, slowly realized that the hacker must be keeping a research notebook, just as Stoll was. The notorious hacker Kevin Mitnick was also known to keep a research notebook. See Clifford Stoll, *The Cuckoo's Egg: Tracking a Spy through the Maze of Computer Espionage* (New York: Doubleday, 1989); and Katherine Hafner and John Markoff, *Cyberpunk: Outlaws and Hackers on the Computer Frontier* (New York: Simon and Schuster, 1991).
87. Monte, *Network Attacks and Exploitation*, p. 56.
88. Philip E. Auerswald et al., *Seeds of Disaster, Roots of Response: How Private Action Can Reduce Public Vulnerability* (Cambridge: Cambridge University Press, 2006); Myriam Dunn-Cavelty and Manuel Suter, "Public-Private Partnerships Are No Silver Bullet: An Expanded Governance Model

discussion of how to organize for cyber operations, proposing the "atrium model," which has been treated as appropriate for both offense and defense.[89] Here I argue that the primary difference between organizational requirements for offense and defense is the relative complexity of their respective tools and tasks and corresponding coordination costs. The size and associated complexity of the asset to be defended is usually much higher than the size and associated complexity of malicious code; as a result, the coordination costs associated with defense organizations are likely to be higher. The defense must manage a large complex infrastructure that includes many users; the offensive team can often be kept relatively small.

Nonetheless, the offense must learn enough about its highly complex target to ensure success. Ultimately, the smaller size and lower coordination costs of the offense result from its narrower goal. While the defense must provide high levels of complex functionality to many users who may have relatively little security expertise, the offense often only needs to achieve limited functionality in its malware to be disruptive. By contrast, when the offense must control the target precisely to achieve success, its task and coordination costs become much more significant; mistakes are more likely because the attacker must understand the complex target more completely. The size of the effort thus scales not only with the size of the code used to hijack the system, but also with the need to understand the complex system being attacked.

The Edges of Cyberspace: Costs of Achieving Physical Effects

Those concerned about cyberwar focus on the potential for a cyberattacker to take control of the industrial control systems that run electric power grids, nuclear power plants, water systems, chemical facilities, and other physical facilities. Although digital industrial control systems (ICS) have been exploited by cyberweapons, achieving a desired physical effect is much more expensive than simply exploiting software vulnerabilities.

The potential advantages that complexity offers an attacker in software—a plethora of vulnerabilities that can be readily discovered and exploited—are

for Critical Infrastructure Protection," *International Journal of Critical Infrastructure Protection*, Vol. 2, No. 4 (December 2009), pp. 179–187; Lynn, "Defending a New Domain"; and Madeleine Carr, "Public-Private Partnerships in National Cyber-Security Strategies," *International Affairs*, Vol. 92, No. 1 (2016), pp. 43–62.

89. Demchak, *Wars of Disruption and Resilience*. The book primarily discusses the context of the defense, but several comments make it clear that offense is also being considered. The model was originally proposed as an antidote to failures to respond to crises, including military offensive failures. See Chris C. Demchak, "Lessons from the Military: Surprise, Resilience, and the Atrium Model," in Louise C. Comfort, Arjen Boin, and Demchak, eds., *Designing Resilience: Preparing for Extreme Events* (Pittsburgh, Pa.: Pittsburgh University Press, 2010), pp. 62–83.

much less prevalent in the physical world. Although unanticipated failure modes are common to all complex systems, they are less common in physical systems than they are in software.[90] Additionally, using cyberweapons to achieve physical effects typically requires knowledge not only of information systems, but also of the physical processes and systems that are targeted, increasing the range of required skills for attack. This information and the associated skills are not readily available in computers. Indeed, much of the knowledge needed for an attack on a control system is very likely tacit knowledge that gets passed from one person to the next, and becomes embedded in organizations, but is never codified.[91]

Dale Peterson, an expert in ICS security, has described the challenges of such attacks in terms of creation, deployment, and persistence. He explains that the creation of a simple ICS weapon is straightforward, but "complex attacks require process and automation talent in addition to Information Technology (IT) security capabilities."[92] The deployment of an ICS cyberweapon is "difficult" because critical infrastructure ICS systems are typically "located behind multiple layers of firewalls and other security-perimeter devices."[93] Although deployment is "difficult but possible, the main challenge is to maintain contact with it once it is deployed."[94] The attacker will likely not be able to deploy the cyberweapon immediately before wanting to use it; triggering it at a specific time that achieves the desired strategic effects requires persistence of communications. These are relatively easy for defenders to detect and stop, because ICS systems are generally not designed to communicate with the outside world.[95]

90. This does not mean that there are no vulnerabilities in the computerized industrial control systems that operate physical machinery; instead it means that the failure modes of physical machinery are easier to model using continuous mathematics and therefore are easier to identify and correct. By contrast, software is a discrete state system that cannot be modeled using continuous mathematics; thus, its behavior is more difficult to predict in advance of real-world operation. See David Lorge Parnas, "Software Aspects of Strategic Defense Systems," *Communications of the ACM*, Vol. 28, No. 12 (December 1985), pp. 1326–1335.

91. The literature on tacit knowledge is extensive. For examples of work relevant to international security, see Kathleen M. Vogel, "Expert Knowledge in Intelligence Assessments: Bird Flu and Bioterrorism," *International Security*, Vol. 38, No. 3 (Winter 2013/14), pp. 39–71; Donald MacKenzie and Graham Spinardi, "Tacit Knowledge, Weapons Design, and the Uninvention of Nuclear Weapons," *American Journal of Sociology*, Vol. 101, No. 1 (July 1995), pp. 44–99; Sonia Ben Ouagrham-Gormley, *Barriers to Bioweapons: The Challenges of Expertise and Organization for Weapons Development* (Ithaca, N.Y.: Cornell University Press, 2014); and Alexander H. Montgomery, "Ringing in Proliferation: How to Dismantle an Atomic Bomb Network," *International Security*, Vol. 30, No. 2 (Fall 2005), pp. 153–187.

92. Dale Peterson, "Offensive Cyber Weapons: Construction, Development, and Employment," *Journal of Strategic Studies*, Vol. 36, No. 1 (2013), p. 122.

93. Ibid., p. 123.

94. Ibid.

95. Ibid.

Although none of these difficulties make it impossible to successfully attack an industrial control system, they increase the costs significantly. At the 2015 DEFCON conference, researchers presented an analysis of the challenges of hacking a chemical plant, noting that a competitor might use such an attack to inflict economic damage. They summarized, "The target plant is not designed in a hacker friendly way. . . . An attacker targeting a remote process is not immediately gifted with complete knowledge of the process and the means to manipulate it. . . . The cost of attack can quickly exceed damage worth."[96] Another DEFCON presentation similarly emphasized the difficulty of creating physical effects.[97]

In short, doing physical damage through cyber operations is much more difficult and costly than simply gathering, distorting, or denying access to information. Some evidence of this difficulty emerges from an analysis of the cyberattacks on Ukraine's electricity distribution stations in December 2015. The attackers made several mistakes in their efforts to control a physical system, including tripping the power bus of a substation before completing operations and thereby cutting themselves off, as well as forgetting to enable an option that would have deprived station operators from locally controlling their equipment (which allowed an operator to fight back).[98] Although the attackers succeeded in being disruptive, the relative efficacy of the offense and defense remain uncertain.

VALUING OFFENSIVE AND DEFENSIVE CYBER OPERATIONS

There are a multitude of potential goals and associated payoffs associated with offensive cyber operations. They can sabotage an enemy's military-industrial capabilities, as the Stuxnet attack on Iran's nuclear facilities did. Offensive cyber operations can also be used to subvert military command and control, as the United States has reportedly done in counterinsurgency efforts.[99] Offensive cyber operations are also important for espionage. The offensive payoff depends not only on the nature of the payoff (e.g., information, an incapacitated

96. Marina Krotofil and Jason Larsen, "Rocking the Pocketbook: Hacking Chemical Plants for Competition and Extortion," DEFCON 23 (conference), Las Vegas, Nevada, August 6–9, 2015, https://www.defcon.org/html/defcon-23/dc-23-speakers.html#Krotofil.
97. Jason Larsen, "Physical Damage 101: Bread and Butter Attacks," Black Hat (conference), Las Vegas, August 1–6, 2015, https://www.blackhat.com/docs/us-15/materials/us-15-Larsen-Remote-Physical-Damage-101-Bread-And-Butter-Attacks.pdf.
98. Michael Assante and Tim Conway, "Ukraine Workshop," North American Electric Reliability Corporation (NERC) Grid Security Conference, October 18, 2016, Quebec City, Canada.
99. Fred Kaplan, *Dark Territory: The Secret History of Cyberwar* (New York: Simon & Schuster, 2016), pp. 156–161.

military, or reduced productivity), but also on the offense's subjective assessment of the value of its own goals. This limits precise prediction.

Nonetheless, the potential payoff of cyber offense and defense can be expected to increase with the value of activities that depend on computer and network technology. For example, if payoff is measured in terms of damage to the target, the payoff of a successful North Korean cyberattack on the U.S. electrical power grid is much greater than the payoff of a successful U.S. cyberattack on North Korea's power grid. This is not only because North Korea's grid supports much less economic activity, but because its grid is less dependent on digital computers. Greater dependence on information and communications technologies typically comes with more opportunities for attack, creating asymmetric vulnerabilities.[100] Additionally, greater dependence on cyberspace increases the potential consequences and thus increases the potential payoff for an adversary.

How do adversaries value the various objectives of cyber operations? Any answer to this question must consider the deeply political nature of cyber offense and defense, which is driven not only by concerns for security or territorial control, but also by the pursuit of prestige and distinctive political values. For example, the value that some states place on stability may directly oppose the value that other actors place on justice.[101] Although an economic value can in principle be assigned to such goals, the method of calculation is by no means obvious, and will include uncertainties. Nonetheless, rough estimates can be obtained in retrospect and can inform expectations for the utility of cyber operations, as shown below.

An Empirical Cost-Benefit Analysis of Stuxnet

A cost-benefit analysis can be a useful way to discipline understanding of cyber conflict and the conditions under which cyberspace may encourage attacks. The following sections provide such an analysis for a celebrated case—Stuxnet—showing that the offense very likely cost more than the defense, but that considerations of benefits probably dwarfed considerations of cost for both offense and defense.

100. Rattray, *Strategic Warfare in Cyberspace*; Clarke and Knake, *Cyber War*; Nye, "Cyber Power"; Gompert and Saunders, *Paradox of Power*; Krepinevich, "Cyberwarfare"; Saltzman, "Cyber Posturing and the Offense-Defense Balance"; Kello, "The Meaning of the Cyber Revolution"; and Joel Brenner and Jon R. Lindsay, "Correspondence: Debating the Chinese Cyber Threat," *International Security*, Vol. 40, No. 1 (Summer 2015), pp. 191–195.

101. This is also a problem for territorial conceptions of the offense-defense balance. See, for example, Betts, "Must War Find a Way?"

A BRIEF HISTORY OF STUXNET

Stuxnet has received extensive analysis in the scholarly and trade press, and here a brief summary will suffice.[102] President George W. Bush learned about the possibility of using cyber operations to undermine the Iranian nuclear program in 2006, not long after negotiations stalled and Iran renewed enrichment at the Natanz nuclear facility.[103] The Stuxnet worm was developed by computer experts from Israel and the U.S. Defense Department; work had begun by 2005, and early versions were found in the wild by November 2007.[104] In 2008, when Israel requested U.S. assistance for a military strike on Natanz, Bush could deflect the request by emphasizing the potential for cyberweapons to slow enrichment.[105]

The worm did not aim to eliminate the enrichment capability immediately, but rather to slowly destroy centrifuges; it was a way of buying time. Early versions would have undermined enrichment by increasing the pressure in the centrifuge cascade, thereby stressing the centrifuge rotors.[106] In 2009 Stuxnet was modified for a more aggressive approach to destroying centrifuges.[107] Rather than targeting valves, it would drive the rotation frequency of the centrifuges to very low and high speeds. The 2009 attack plan accepted higher levels of risk that the attack would be detected.[108]

The 2009 attack seems to have succeeded in destroying some centrifuges. As figure 1 shows, the number of operational cascades and associated centrifuges declined during the summer of 2009, from a high of 4,920 centrifuges in May 2009 to a low of 3,772 centrifuges in January 2010—the point at which Iran began to replace large numbers of broken centrifuges. The number of centrifuges

102. For previous accounts, see Kim Zetter, "How Digital Detectives Deciphered Stuxnet, the Most Menacing Malware in History," *Wired*, June 11, 2011, http://www.wired.com/2011/07/how-digital-detectives-deciphered-stuxnet/; Kim Zetter, *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon* (New York: Crown, 2014); David E. Sanger, *Confront and Conceal: Obama's Secret Wars and the Surprising Use of American Power* (New York: Random House, 2012); Ivanka Barzashka, "Are Cyber-Weapons Effective? Assessing Stuxnet's Impact on the Iranian Enrichment Programme," *RUSI Journal*, Vol. 158, No. 2 (2013), pp. 48–56; and Ralph Langner, "To Kill a Centrifuge: A Technical Analysis of What Stuxnet's Creators Tried to Achieve" (Munich: Langner Group, 2013).

103. Sanger, *Confront and Conceal*.

104. Geoff McDonald et al., "Stuxnet 0.5: The Missing Link" (Mountain View, Calif.: Symantec, 2013).

105. David E. Sanger, "U.S. Rejected Aid for Israeli Raid on Iranian Nuclear Site," *New York Times*, January 10, 2009, http://www.nytimes.com/2009/01/11/washington/11iran.html?pagewanted=all&_r=4&; and Ewen MacAskill, "Stuxnet Cyberworm Heads Off US Strike on Iran," *Guardian*, January 16, 2011, http://www.theguardian.com/world/2011/jan/16/stuxnet-cyberworm-us-strike-iran.

106. McDonald et al., "Stuxnet 0.5."

107. Ibid.

108. Langner, "To Kill a Centrifuge."

Figure 1. Installed and Operational Cascades and Operational Centrifuges at Natanz

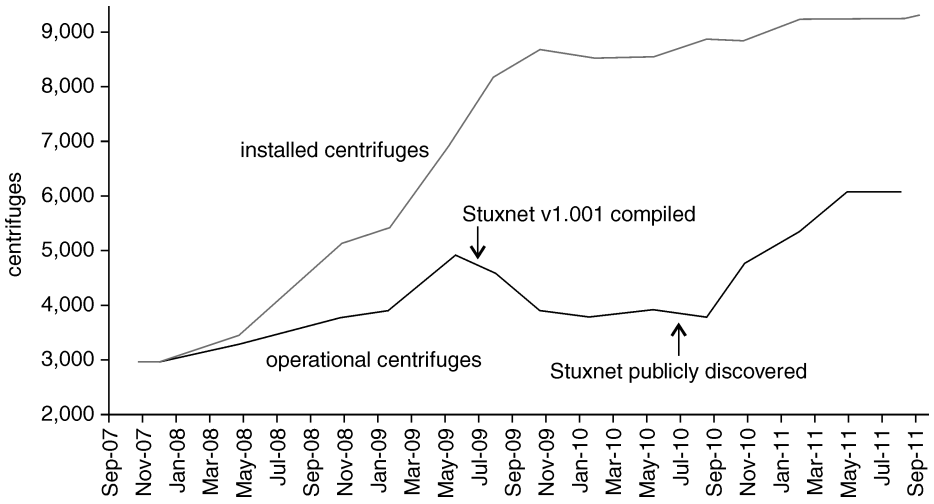

SOURCE: Data were collected from individual reports available on the International Atomic Energy Agency website: https://www.iaea.org/newscenter/focus/iran/iaea-and-iran-iaea-reports.

began to rise again not long after the public discovery of Stuxnet in the summer of 2010, and surpassed the previous peak in November 2010.

Although Stuxnet was aimed at Iran, it spread around the world. After a Belarusian antivirus firm identified and published findings about the worm in June 2010, computer professionals and companies in at least three different countries became involved in deciphering and helping mitigate Stuxnet. Researchers at the U.S.-based Symantec noted that even after they began to realize that the attack was most likely a U.S.-Israeli collaboration, they felt obligated to publish their findings online.[109] The public analysis of Stuxnet likely benefited Iran and illustrates the difficulty of drawing boundaries around the offensive and defensive players, a point discussed further below.

Scholars have used Stuxnet to support their case for either offense dominance or defense dominance.[110] I am aware of only one previous effort to

109. Zetter, "How Digital Detectives Deciphered Stuxnet, the Most Menacing Malware in History."

110. Those who have pointed to Stuxnet as demonstrating the difficulty of attack include Lindsay, "Stuxnet and the Limits of Cyber Warfare"; Rid, *Cyber War Will Not Take Place*; Liff, "Cyberwar"; David Betz, "Cyberpower in Strategic Affairs: Neither Unthinkable nor Blessed," *Journal of Strategic Studies*, Vol. 35, No. 5 (October 2012), pp. 689–711; and Rid and Buchanan, "Attributing Cyber Attacks." Others point to Stuxnet as demonstrating the ease of attack. See Kello, "The Meaning of the Cyber Revolution"; Gary McGraw, "Cyber War Is Inevitable (Unless We Build Security In),"

define and empirically assess the offense-defense balance in cyberspace—the aforementioned thesis by Malone.[111] Although this thesis does an admirable job of cataloging various costs, it does not assess the value of offense or defense, and as the following analysis shows, some of the assumptions used to derive costs are problematic. Additionally, the underlying theoretical model does not recognize the high variability of skills (it assumes that the skills of the offense and defense are equal), and neglects the interactions between technology and people, treating them as separable costs.

The analysis provided in this article is based solely on public information. Table 1 summarizes the results with a comparison to Malone's results. These figures are somewhat uncertain, for reasons discussed below, but they are unlikely to be inaccurate by a factor of 10 or more. Two results are immediately apparent. First, despite significant uncertainties, it is likely that the Stuxnet attack cost the United States and Israel more than it cost Iran. In fact, public reports suggest that the United States allocated $300 million to "joint covert projects aimed at Iran," of which Stuxnet was "a priority."[112] If this is true, the offense almost certainly spent more than the defense. Second, the value that each country places on Iran's nuclear program dwarfs the likely costs of these attacks, suggesting that leaders were unlikely to have quibbled over costs.

VALUE OF ATTACKING AND DEFENDING IRAN'S NUCLEAR PROGRAM

Although the value that the United States and Israel place on undermining Iran's nuclear program is uncertain, one proxy measure is the economic value of sanctions designed to stop the program. In 2014 the National Iranian American Council estimated that sanctions cost the United States $135–$175 billion from 1995 to 2012, or $7.5–$9.7 billion per year, not including lost employment and similar costs.[113] Similarly, Iran's losses due to nuclear-related sanctions are likely tens of billions of dollars per year.[114]

Journal of Strategic Studies, Vol. 36, No. 1 (2013), pp. 109–119; and Peterson, "Offensive Cyber Weapons."

111. Malone, "Offense-Defense Balance in Cyberspace."

112. MacAskill, "Stuxnet Cyberworm Heads Off US Strike on Iran."

113. Jonathan Leslie, Reza Marashi, and Trita Parsi, "Losing Billions—The Cost of Iran Sanctions to the U.S. Economy" (Washington, D.C.: National Iranian American Council, 2014). These estimates are very rough, because the types of sanctions imposed and their effects varied considerably over this period.

114. Estimates of losses to Iran as a result of sanctions vary, but were in the range of tens of billions of dollars per year during the period that Stuxnet was developed. Losses are typically estimated primarily in terms of lost foreign investment and lost oil revenue. In 2006, tensions over Iran's nuclear program resulted in tens of billions of dollars of capital moving out of Iran. Neil King Jr. and Marc Champion, "Nations' Rich Trade with Iran Is Hurdle for Sanctions Plan," *Wall Street Journal*, September 20, 2006. In 2008, a letter from Iranian economists stated that sanctions had cost Iran "many billions of dollars" by forcing many of its imports to run through middlemen. See Orde F.

Table 1. Costs and Benefits of Stuxnet to the Offense and Defense

United States and Israel (offense)	Iran (defense)
Value of Offense and Defense	
Value of undermining the program, as measured by sanctions: on the order of $10 billion per year, or $30 billion from 2007 to 2010 Not considered by Malone	Value of protecting the program, as measured by losses resulting from sanctions: on the order of $10 billion per year, or $30 billion from 2007 to 2010 Not considered by Malone
Costs of Offense and Defense	
Intelligence $260 million from 2007 to 2010 Not included in Malone's estimate	Ongoing security expenditures $7 million from 2007 to 2010 Malone's estimate: $34 million
Labor required to develop the Stuxnet worm $7 million from 2007 to 2010 Malone's estimate: $3.5 million	Manufacturing new centrifuges $1.8 million from 2007 to 2010 Not included in Malone's estimate
Cost of zero-day exploits $2 million Not included in Malone's estimate	
Cost of stolen certificates $1 million Not included in Malone's estimate	
Infrastructure costs from 2007 to 2010 $40 million Malone's estimate: $327,000	Lost productivity at Natanz resulting from attacks $5 million Not included in Malone's estimate
Total costs (United States/Israel): $300 million Malone's estimate: $4.7 million	Total costs: $14 million Malone's estimate: $34 million

These are imperfect proxies for the value of Stuxnet and defense from it; the effects of sanctions are different from those of a covert cyberattack, and Iran's tolerance of sanctions may also reflect reputational concerns—that is, countries may be willing to suffer a loss simply to avoid the appearance of capitulation. Nonetheless, these estimates suggest that the United States is willing to forgo

Kittrie, "New Sanctions for a New Century: Treasury's Innovative Use of Financial Sanctions," *University of Pennsylvania Journal of International Law*, Vol. 30 (2008), p. 799. Asked about the costs of existing sanctions to Iran in 2010, the Treasury Department noted that Iran's oil ministry aimed to attract at least $25 billion of investment per year in its oil and gas sectors, but had attracted only a few billion dollars in foreign direct investment, implying losses of $20 billion per year. See *Iran Policy in the Aftermath of U.S. Sanctions: Hearing before the Senate Committee on Foreign Relations*, 111th Cong., 2nd sess., June 22, 2010, p. 55. Sanctions became stricter in 2010, leading to estimates that they cost Iran $60 billion in lost foreign investment that year, and a decline of $40 billion in oil revenue from 2011 to 2012. See Ali Vaez and Karim Sadjadpour, "Iran's Nuclear Odyssey: Costs and Risks" (Washington, D.C.: Carnegie Endowment for International Peace, 2013).

tens of billions of dollars per year to slow or stop Iran's nuclear program, and Iran (with a much smaller gross domestic product) is willing to take a similar loss to preserve its nuclear program. Israel views an Iranian nuclear weapon as an existential threat and is likely to place an exceedingly high value on eliminating that threat. The intense debate over the Joint Comprehensive Plan of Action reached in 2015 indicates the high stakes to the entire international community.

By contrast, the following analysis suggests that the cost of the Stuxnet worm to both Iran and the United States was likely closer to tens or hundreds of millions of dollars. Because such costs are two or three orders of magnitude smaller than the value of the offense and defense, they likely had very little impact on decisionmaking. Consideration of those costs, however, can provide important insights into the challenges of offense and defense.

costs of stuxnet to the offense

The costs of Stuxnet to the offense include technical expertise, zero-day vulnerabilities, stolen certificates, infrastructure costs, and intelligence.

technical expertise. Stuxnet was developed and deployed by a team of highly skilled computer, nuclear, and ICS experts for at least six years, though focused development likely took closer to three years.[115] Estimates of the labor required to develop Stuxnet have varied from five to thirty people, although it is likely that some of them worked on Stuxnet only part-time.[116] Based on published salaries for computer security professionals and nuclear engineers, employing an elite team of five to thirty such highly skilled individuals for three years would cost between $1.9 million and $11.4 million; the average of $6.7 million can be used as an approximate figure.[117]

115. Sanger reports that the project was started in early 2008. See Sanger, "U.S. Rejected Aid for Israeli Raid on Iranian Nuclear Site." However, Stuxnet was first discovered in the wild in 2007, and analysis suggests that work had started by 2005. McDonald et al., "Stuxnet 0.5."

116. Josh Halliday, "Stuxnet Worm Is the 'Work of a National Government Agency,'" *Guardian*, September 24, 2010, http://www.theguardian.com/technology/2010/sep/24/stuxnet-worm-national-agency; Michael Joseph Gross, "A Declaration of Cyber-War," *Vanity Fair*, April 2011, http://www.vanityfair.com/news/2011/04/stuxnet-201104; and David Kushner, "The Real Story of Stuxnet," *IEEESpectrum.com*, February 26, 2013, http://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet

117. In 2008, reported salaries for computer engineers ranged from $76,760 to $115,085. For nuclear engineers reported salaries ranged from $90,866 to $138,532. See Maureen Byko, "Engineering Salaries: The AAES Tracks the Trends," *JOM*, Vol. 60, No. 12 (2008), pp. 14–15. I assume that elite experts operating in secret would be paid at the high end of this scale; averaging over computer and nuclear engineering, this implies a salary of $126,808 per person. By contrast, Malone uses average (rather than top) U.S. government pay rates for information security, inflates salaries by an additional 25 percent to account for secrecy, and estimates that a ten-person team would cost approximately $3.5 million for a four-year effort (2006–10). See Malone, "Offense-Defense Balance in Cyberspace."

FOUR ZERO-DAY VULNERABILITIES. Stuxnet contained four zero-day vulnerabilities (i.e., vulnerabilities for which no patches existed when they were first exploited).[118] It is possible that the costs of these vulnerabilities were included in the work of the Stuxnet team. The U.S. Defense Department is the largest purchaser of zero-day vulnerabilities, however, with the National Security Agency budgeting $25 million to purchase vulnerabilities in 2013.[119] The cost of vulnerabilities varies widely, with particularly valuable vulnerabilities costing up to $1 million.[120] Thus, the cost of four zero-day vulnerabilities was likely between $0 (if they were developed by the Stuxnet team) and $4 million (if they were purchased separately); the average of these can be taken as a rough estimate.

TWO STOLEN CERTIFICATES. Stuxnet used two stolen certificates, that is, digital documents used to (falsely) assure the operating system that the worm was authorized to update software.[121] Prices for digital certificates are less openly discussed than prices for vulnerabilities, because it is difficult to conceive of a legitimate market for stolen certificates. In the past several years, Trojan malware has been developed that is capable of stealing certificates from personal computers, but certificates were harder to obtain during the period in which Stuxnet was being developed. Keys to authenticate software (rather than an individual computer) are typically held by software developers on private networks, and thus should be more difficult to steal.[122] This suggests that the price of a stolen key for a software developer such as JMicron and RealTek may have been comparable to that of an expensive zero-day vulnerability. Accordingly, the cost of two certificates is estimated at anywhere from $0

118. Two of the vulnerabilities had appeared in earlier exploits, but remained unpatched, whereas the other two vulnerabilities had not appeared in the wild before Stuxnet. See Zetter, "How Digital Detectives Deciphered Stuxnet, the Most Menacing Malware in History."

119. Brian Fung, "The NSA Hacks Other Countries by Buying Millions of Dollars' Worth of Computer Vulnerabilities," *Washington Post*, August 13, 2013, http://www.washingtonpost.com/blogs/the-switch/wp/2013/08/31/the-nsa-hacks-other-countries-by-buying-millions-of-dollars-worth-of-computer-vulnerabilities/.

120. One vulnerability broker was paying $60,000 to $120,000 for Windows vulnerabilities in 2012. See Andy Greenberg, "Shopping for Zero-Days: A Price List for Hackers' Secret Software Exploits," *Forbes*, March 23, 2012, http://www.forbes.com/sites/andygreenberg/2012/03/23/shopping-for-zero-days-an-price-list-for-hackers-secret-software-exploits/. Endgame offers a subscription to twenty-five zero-day vulnerabilities per year for $2.5 million, suggesting that $100,000 is a typical cost per zero day. See Stefan Frei, "The Known Unknowns: Empirical Analysis of Publicly Unknown Security Vulnerabilities" (Austin, Tx.: NSS Labs, 2013).

121. Pierre-Marc Bureau, "Win32/Stuxnet Signed Binaries," *welivesecurity.com*, July 19, 2010, http://www.welivesecurity.com/2010/07/19/win32stuxnet-signed-binaries/.

122. Hiroshi Shinotsuka, "How Attackers Steal Private Keys from Digital Certificates," Symantec official blog, February 22, 2013, http://www.symantec.com/connect/blogs/how-attackers-steal-private-keys-digital-certificates.

(if they were stolen by the Stuxnet team) to $2 million, if they were purchased. The average of these is used as a rough estimate.

INFRASTRUCTURE COSTS. Creating software requires infrastructure such as computers and machines on which to do development, testing, and debugging. In the case of Stuxnet, some of these activities required access to a uranium enrichment testbed (the Dimona facility in Israel), and equipment in nuclear energy laboratories across the United States.[123] These infrastructure costs are likely similar to those of the National Nuclear Security Administration (NNSA), whose facilities were involved in the Stuxnet project. In 2010 the NNSA spent an average of $744,000 per permanent employee on infrastructure.[124] With five to thirty people involved over three years, this cost would range from $11 to $67 million; I use the average of these as a rough estimate.[125]

COST OF INTELLIGENCE. The cost of intelligence for Stuxnet was likely very high. The first version of Stuxnet relied on precise knowledge about the Iranian enrichment cascade and protection mechanisms. Ralph Langner, the industrial control system security expert who deconstructed the Stuxnet code, noted that the "detailed pin-point manipulations of these sub-controllers indicate a deep physical and functional knowledge of the target environment; whoever provided the required intelligence may as well know the favorite pizza toppings of the local head of engineering."[126]

Intelligence may be the most difficult expense to estimate because intelli-

123. J. Broad William, John Markoff, and David E. Sanger, "Israeli Test on Worm Called Crucial in Iran Nuclear Delay," *New York Times*, January 16, 2011, http://www.nytimes.com/2011/01/16/world/middleeast/16stuxnet.html?pagewanted=all&_r=0.

124. In 2010 the infrastructure budget was $2.1 billion, and the number of permanent NNSA employees was 2,823. See NNSA, "Department of Energy FY 2012 Congressional Budget Request: National Nuclear Security Administration" (Washington, D.C.: U.S. Department of Energy, February 2011), p. 53, https://nnsa.energy.gov/sites/default/files/nnsa/inlinefiles/FY%202012%20NNSA%20Congressional%20Budget%20Submission_0.pdf; and NNSA, "Federal Equal Opportunity Recruitment Program Accomplishment Report Fiscal Year 2011" (Washington, D.C.: NNSA, November 21, 2012), http://nnsa.energy.gov/sites/default/files/nnsa/10-12-inlinefiles/2012-10-05%20FY%202011%20FEORP%20&%20Relevant%20Hispanic%20Employ%20Plan.pdf.

125. Malone instead estimates a cost of only $327,000 over a four-year development effort. He considers only the cyber-infrastructure budget of NNSA, plus the cost of one centrifuge sold by URENCO. This is likely a low estimate, because Stuxnet involved testing many kinds of equipment that are not typically included in cyber infrastructure, and testing different versions of Stuxnet likely required the destruction of more than one centrifuge. Malone also divides the cyber-infrastructure cost by employees and contractors rather than just permanent employees. See Malone, "Offense-Defense Balance in Cyberspace." Malone states that NNSA employs 33,000 permanent employees and contractors, citing the NNSA jobs webpage: http://nnsa.energy.gov/federalemployment/ourjobs. That number no longer appears on that webpage, however, or in any documentation about the NNSA workforce. Reports for 2010 discuss only the total number of permanent employees.

126. Langner, "To Kill a Centrifuge," p. 10.

gence budgets tend to be secret. Documents leaked by Edward Snowden show that $6.9 billion of the 2013 U.S. intelligence budget was requested to counter the proliferation of nuclear, chemical, and biological weapons.[127] The program description emphasizes nuclear weapons more than either chemical or biological weapons, so a low estimate is that one-third of the counterproliferation budget ($2.3 billion) goes to nuclear weapons counterproliferation. Israeli intelligence on the Iranian nuclear program was particularly important for Stuxnet, but its budget is much smaller than that of the United States, and thus is not considered further.[128]

If the U.S. nuclear counterproliferation intelligence budget is further divided three ways to reflect the three primary states of concern (Pakistan, Iran, and North Korea), and divided again by the nine nuclear facilities in Iran that are safeguarded by the International Atomic Energy Agency (IAEA), the estimate is $85 million annually spent on counterproliferation at Natanz, or $260 million over three years.[129] Some analysts might object that this cost is too high, because the marginal cost of intelligence needed for Stuxnet was lower than the total cost of intelligence infrastructure. But because intelligence analysis is intrinsically inefficient—the entire "haystack" must be searched to find the "needle"—the marginal costs are not necessarily meaningful.

costs of defense against stuxnet

There are at least four components to Iran's defense: measures to prevent a cyberattack from affecting Natanz; the immediate productivity loss at Natanz resulting from the centrifuge failures and replacement; the cost of replacing centrifuges; and the costs of deciphering and mitigating Stuxnet. This final cost is negligible given the availability of free public analyses of Stuxnet and mitigation tools. The other three costs are detailed below.

ongoing security costs at natanz. One way to estimate Iran's expenses for securing Natanz from cyberattack is by analogy to the cybersecurity expenses of U.S. industrial control systems. A 2013 survey of U.S. firms in the electric power sector found that average spending on cybersecurity was $1.45 million, with an average firm size of 2,878 employees and spending scaling roughly with size.[130] With approximately 2,000 employees, Natanz

127. Director of National Intelligence, "FY 2013 Congressional Budget Justification" (Washington, D.C.: Office of the Director of National Intelligence, February 12, 2012), http://www.scribd.com/doc/164056434/FY-2013-Congressional-Budget-Justification.
128. Sanger, *Confront and Conceal*.
129. For a list of Iran's safeguarded nuclear facilities, see https://www.iaea.org/newscenter/focus/iran
130. Zpryme, "2012 Utility Cyber Security Survey" (Austin, Tex.: Zpryme, 2013).

would spend about $1 million annually to achieve approximately the same level of security.[131]

An alternate estimate comes from Malone and is based on the average of three different estimates for Iran's entire nuclear budget: the entire budget of the U.S. NNSA, adjusted downward proportionally to the smaller Iranian gross domestic product; 9 percent of Iran's defense budget (a study by Global Zero estimates that nuclear weapons states on average spend 9 percent of their defense budget on nuclear weapons); and an estimate based on a statement in a speech by Iranian President Mahmoud Ahmadinejad. This average is then multiplied by the fraction of the U.S. NNSA budget that goes to both physical and cybersecurity, producing an estimate of $34 million. This estimate is too high, because it includes security for the entire weapons program, of which Natanz is only a part. I divide Malone's figure by the nine nuclear facilities that are the subject of IAEA safeguards, considering the result ($4 million) to be an upper bound for annual costs. Table 1 shows the average of the lower and upper bounds, multiplied over three years, as a rough estimate.

COST OF LOST PRODUCTIVITY AT NATANZ. Iran continued to steadily increase its enrichment capacity in spite of the unreliable centrifuges.[132] This was possible because Iran continued to install new cascades, even as the damaged centrifuges were removed (see figure 1). Nonetheless, the facility lost productivity due to the centrifuges being removed and replaced. I estimate the cost of this reduced productivity as the fraction of centrifuges that were removed and subsequently replaced, multiplied by both the days that they were missing and the cost of wages paid to approximately 2,000 employees during this time.[133] An analysis of IAEA data suggests that Natanz lost approximately ninety days of productivity through February 2011, when it surpassed its previous peak number of operational centrifuges; based on public information about salaries and manufacturing workforce composition, the total cost of lost productivity would be roughly $5 million.[134]

131. A report by an Iranian physicist estimates that there were about 2,000 workers at Natanz in 2012. See Khosrow Semnani, "The Ayatollah's Nuclear Gamble" (Salt Lake City: University of Utah Hinckley Institute of Politics, 2012). Gholamreza Aqazadeh, head of the Atomic Energy Organization of Iran, stated that 3,000 people were working in two shifts at Natanz. See "Developments at Natanz," *Nuclear Engineering International*, January 29, 2007, http://www.neimagazine.com/news/newsdevelopments-at-natanz. The former estimate is the more reliable of the two, as Aqazadeh was attempting to make the case that Iran's nuclear program was too institutionalized to be stopped.

132. David Albright, Paul Brannan, and Christina Walrond, "Did Stuxnet Take Out 1,000 Centrifuges at the Natanz Enrichment Plant? Preliminary Assessment" (Washington, D.C.: Institute for Science and International Security, 2010); and Barzashka, "Are Cyber-Weapons Effective?"

133. Other expenses would include electricity to power Natanz, but this is relatively inexpensive, and UF_6 should not have been lost because of the malfunctioning centrifuges.

134. A 2006 study suggests that approximately 25 percent of U.S. manufacturing jobs were high

CENTRIFUGE REPLACEMENT. Based on IAEA data, it appears that approximately 984 centrifuges were replaced in late 2009 or early 2010, most likely as a result of the Stuxnet attacks.[135] Iran's IR-1 centrifuge is estimated to have a capability similar to centrifuges manufactured in the United States in 1959; the latter cost about $1,800 in 2012.[136] This suggests a total replacement cost of about $1.8 million.[137]

LESSONS AND LIMITATIONS

Contrary to the prevailing assumption of offense dominance, Stuxnet ultimately had a relatively small payoff and high costs for the offense. Despite initial claims that Stuxnet set back the Iranian nuclear program by years, it appears to have had very little impact on Iran's enrichment capability; the analysis above suggests a delay of only about three months.[138] Additionally, Stuxnet likely cost the offense more than the defense. This is largely because of the cost-intensive nature of gathering intelligence on a physical target. Hackers can efficiently gain massive amounts of digital data once they have penetrated a computer network, but crucial kinds of knowledge about physical control systems are not stored digitally. Gathering such information can be very costly.

skilled in 2002 and 75 percent low skilled. See Richard Deitz and James Orr, "A Leaner, More Skilled U.S. Manufacturing Workforce," *Current Issues in Economics and Finance*, Vol. 12, No. 2 (February/March 2006), p. 3, https://www.newyorkfed.org/medialibrary/media/research/current_issues/ci12-2.pdf. For highly skilled workers, I extrapolate from the median U.S. salary of $90,866 for nuclear engineers with 13–16 years of experience. See Byko, "Engineering Salaries." Because U.S. median household incomes are roughly 3.6 times of those in Iran, a nuclear engineer with comparable experience would make $25,113 in Iran; I use this figure for the 25 percent of workers who are highly skilled. Since most Iranian households have multiple sources of income and women play an important role in the economy, I divide Iran's median household income by two and use the result of $6,023 to establish an average salary for low-skilled or mid-skilled workers at Natanz. Median household incomes can be found in a Gallup poll of 131 countries, including Iran, conducted from 2006 to 2012. Glenn Phelps and Steve Crabtree, "Worldwide, Median Household Income about $10,000," Gallup, December 16, 2013, http://www.gallup.com/poll/166211/worldwide-median-household-income-000.aspx.

135. Albright, Brannan, and Walrond, "Did Stuxnet Take out 1,000 Centrifuges at the Natanz Enrichment Plant?"

136. R. Scott Kemp, "Centrifuges: A New Era for Nuclear Proliferation" (Arlington, Va.: Nonproliferation Policy Education Center, 2012); and R. Scott Kemp, "The Nonproliferation Emperor Has No Clothes: The Gas Centrifuge, Supply-Side Controls, and the Future of Nuclear Proliferation," *International Security*, Vol. 38, No. 4 (Spring 2014), pp. 39–78.

137. Geoffrey Forden estimates the cost of Iran's centrifuges to be $20,000, the same as that of the much more sophisticated centrifuges that France is leasing to URENCO. This figure seems unreasonably high, so I do not use it here. See Geoffrey Forden, "What Does Natanz Cost?" *Arms Control Wonk*, June 12, 2009, http://forden.armscontrolwonk.com/archive/2363/what-does-natanz-cost.

138. For an example of claims of great efficacy, see William, Markoff, and Sanger, "Israeli Test on Worm Called Crucial in Iran Nuclear Delay." For work suggesting a minimal or even positive impact, see Barzashka, "Are Cyber-Weapons Effective?" A close study of IAEA reports from 2007 to 2011 confirms that enrichment capacity did not significantly decline. Reports can be found online at https://www.iaea.org/newscenter/focus/iran/iaea-and-iran-iaea-reports.

Even if the cost of intelligence is neglected, however, the offense likely spent more than the defense ($49 million versus $14 million).

Perhaps more importantly, when the potential value of offense and defense are considered, the costs of offense and defense come into clearer perspective. In the case of Stuxnet, the value that the United States, Israel, and Iran assign to undermining or defending the Iranian nuclear program appears to be at least two orders of magnitude greater than the cost of a cyberattack. The limited impact of Stuxnet suggests that the United States and Israel did not fully realize this value; nonetheless, the high potential value associated with undermining Iran's nuclear program is one reason that decisionmakers probably did not carefully assess the costs of offense when considering whether or not to develop and deploy Stuxnet. It also confirms the importance of including benefits in any theorization of offense-defense balance.

This analysis has largely ignored the impact of private-sector actors on the costs of Stuxnet. Several multinational companies helped with analysis that likely benefited Iran; the total costs of this help were relatively small.[139] Although U.S. companies were affected by Stuxnet, remediation tools were distributed freely by Siemens, and companies have not complained about high costs associated with remediation; hence the cost of removing Stuxnet is not considered in this analysis. Nonetheless, it is important to note that because attacks on one part of cyberspace have the potential to harm many other parts, nonstate actors that were not originally targeted have a strong incentive to diagnose and mitigate any cyberattacks. Security corporations can restrict some services to paying customers, but they often publicize their analyses freely on the internet, where it can help everyone, including the target of the original attack. One reason for providing such public resources is that a novel attack provides security companies and professionals an opportunity to develop their reputations and expertise, while also presenting a fascinating challenge. This merits further research, as the private sector's role in defense and its implications for international relations have received relatively little attention.[140]

This analysis has also not considered the long-term impacts of Stuxnet,

139. This is based on the previously cited accounts of the key actors involved with analyzing Stuxnet (approximately three Symantec researchers working for four months; three at Langner Group working for one month; and three in Belarus for one month). See Zetter, *Countdown to Zero Day*; "The Man Who Found Stuxnet—Sergey Ulasen in the Spotlight," *Nota Bene* blog, November 2, 2011, https://eugene.kaspersky.com/2011/11/02/the-man-who-found-stuxnet-sergey-ulasen-in-the-spotlight/. Based on these accounts, about eighteen months were spent on Stuxnet. The cost of this analysis is tiny compared to the cost for Iran of lost productivity.

140. As a recent historical analysis of cyber conflict notes, "[T]he principal difference between cyber and traditional conflicts" is "the primacy of the private sector" in cyber defense; yet this aspect is also "the one most often overlooked." See Jason Healey and Karl Grindal, *A Fierce Domain: Conflict in Cyberspace, 1986 to 2012* (Vienna, Va.: Cyber Conflict Studies Association, 2013), p. 22.

which may ultimately be more significant than the destruction of centrifuges. For example, it is possible that increased suspicion and paranoia will take a significant long-term toll on the Iranian nuclear project. The U.S. nuclear program incurred similar damage in the wake of allegations of espionage at Los Alamos National Laboratory.[141] Former Los Alamos Director Siegfried Hecker has argued that tightened security measures reduced the laboratory's productivity enormously. In fact, Hecker decided to stop doing classified research because it was too much hassle.[142] A similar reaction in Iran's nuclear program could impose significant costs.

Other long-term impacts of Stuxnet may or may not benefit the United States. Stuxnet likely increased Iran's resolve to enrich uranium and spurred the development of both defensive and offensive cyber operations. Since the Stuxnet revelations, Iran has been linked to multiple attacks against U.S. companies and government targets, and the National Security Agency states that Iran has learned from the U.S.-Israeli attacks.[143] On the other hand, it is also possible that Stuxnet will enhance the U.S. cyber deterrent.

Whatever the long-term impacts, it is clear that one of the most sophisticated cyberattacks ever conducted bore no resemblance to the "cyber-Pearl Harbor" that has drawn so much attention.[144] Stuxnet's immediate impact was small, and the most significant impacts of the attack will likely be felt over the longer term.

Conclusion

This article has shown that widespread claims about the offense dominance of cyberspace are fundamentally flawed; the offense-defense balance can be understood only in the context of specific adversaries with distinctive goals and levels of capability in managing complex information technology. In many cases, particularly those in which the goal is to achieve complex kinetic effects, cyber operations may well be less costly for the defense than the offense.

This is beginning to change. See, for example, Shane Harris, *@War: The Rise of the Military-Internet Complex* (New York: Houghton Mifflin Harcourt, 2014).

141. Joe Masco, "Lie Detectors: On Secrets and Hypersecurity in Los Alamos," *Public Culture*, Vol. 14, No. 3 (2002), pp. 441–467, http://ceas.iscte.pt/ethnografeast/Joseph%20Masco/%20Lie%20Detectors.pdf.

142. Author interview with Sigfried Hecker, Stanford University, February 17, 2015.

143. Glenn Greenwald, "NSA Claims Iran Learned from Western Cyberattacks," *Intercept*, February 10, 2015, https://theintercept.com/2015/02/10/nsa-iran-developing-sophisticated-cyber-attacks-learning-attacks/; and Evan Perez and Shimon Prokupecz, "U.S. Charges Iranians for Cyberattacks on Banks, Dam," *CNN.com*, March 23, 2016, http://www.cnn.com/2016/03/23/politics/iran-hackers-cyber-new-york-dam/.

144. See Bumiller and Shanker, "Panetta Warns of Dire Threat of Cyberattack on U.S."

I have presented this argument in four parts. First, I have argued that conceptions of offensive advantage need to include valuations of the goals as well as the costs of cyber operations. The goals of cyber operations are much more varied than are battles for territory and may include propaganda, espionage, counter-espionage, and sabotage, in addition to assistance with territorial military operations. This article has focused on the relative utility of offense and defense—the value of the goals of offense less the costs of offense, and the value of the goals of defense less the costs of defense. A more complete analysis would consider expected utility—that is, it would include the probability of success for offense or defense as a function of offensive and defensive expenditures. The statistics for empirically predicting the probability of success for offense or defense under various conditions do not exist, however, and there is reason to doubt that such calculations will reach a useful level of accuracy or precision. Nonetheless, decisionmakers' beliefs about the probability of success will shape behavior, and thus both theoretical and empirical analysis of the factors that make cyber offense and defense costly and valuable is crucial to informing policy.

Second, I have theorized what makes cyber operations costly and valuable, arguing that the resulting offense-defense balance is a characteristic not of cyberspace, but rather of the relationship between two adversaries; the balance is not systemic, but dyadic. Although software does have an essential characteristic—arbitrary complexity—which provides the offense with many vulnerabilities to exploit, technology alone does not determine the balance. Nor is technology one of several independent factors to be summed up in a net assessment. Instead, it is the processes that govern the interactions between skilled users and technology that determine an organization's readiness for offense or defense. Research on the capability maturity model shows that the cost of managing complex software decreases as the maturity of an organization's processes increases. Nonetheless, cost grows with complexity.

The skills and organizational capabilities needed for offense and defense are very similar, but offensive capabilities often require less coordination and therefore are less costly than defensive operations. Nonetheless, an offensive operation that aims to precisely control a complex system is much more difficult than one that merely aims to disrupt the system, and may be costlier than defense. Additionally, the advantages that complex software offers attackers diminish rapidly at the "edges" of cyberspace, where computers are used to control physical systems, because knowledge of the physical systems is needed to exercise careful control. Such knowledge is often tacit and therefore unavailable through cyber espionage. Thus, although information technology offers unprecedented efficiency for espionage, it is not the most

cost-effective means of destruction. Cyberweapons are primarily advantageous for their covertness, and they become expensive when physical systems are the target.

Third, this theorization contributes to efforts to empirically characterize offensive and defensive capabilities. Evidence suggests that the current ease with which offensive operations succeed is more a result of a poorly organized defense and the relatively simple goals of offense than it is a result of a fundamental technological advantage.

Fourth, by providing an empirical cost-benefit analysis of Stuxnet, this article counters the dominant assumption that cyberspace favors the offense, and it offers qualified support for the thesis that attacking physical targets may be costlier than defending them. The high cost of precisely controlling physical machinery, however, is not cause for complacency about the threat of cyberattacks on critical infrastructure or other physical assets, because the cost of an attack is not all that matters. If an actor assigns sufficiently high value to an objective that can be achieved only through cyberattack, the costs become irrelevant. The costs of Stuxnet are uncertain, but they are likely to be two orders of magnitude lower than the perceived value of Iran's nuclear program.

The shift in analytic focus proposed here—from technology to the organizational processes that integrate technology and skilled workers—opens up avenues for further research. For example, it has implications for theories of general deterrence and compellence in cyber operations. Several scholars have suggested that cyberweapons cannot be used for compellence or deterrence because they are "use and lose" weapons. To be a credible threat, the malicious code must be demonstrated, but this demonstration provides the basis for defense and thereby undermines the threat. If, however, the weapons are not computer programs, but rather skilled people working in well-oiled organizations, general deterrence and compellence become possible. Demonstration of past capabilities can make future threats seem much more plausible. The question of what kinds of demonstrations are persuasive for deterrence or compellence is a matter of future research.

The most important policy implication of these findings is that leaders should not presume that cyberspace favors the offense, but instead recognize that any such advantage will depend on the complexity of the offensive goals and the capabilities of the defender. The difficulty of distinguishing offensive and defensive cyber capabilities may exacerbate the security dilemma; even if nations are motivated primarily by defensive interests, investments in cyber capabilities will appear threatening to others.[145] The ease of switching

145. Jervis, "Cooperation under the Security Dilemma."

from defensive to offensive operations, however, should reduce the need to focus on offensive operations and create a space for investing in defense.[146]

Finally, shifting the focus from black-boxed technologies to organizational processes and skills suggests that a focus on stockpiling technological vulnerabilities, rather than developing the skills necessary to continually identify and mitigate vulnerabilities, is misguided. Defensive practices—such as improved software development and maintenance processes, better personnel training, and continual vulnerability scanning—will not produce invulnerable organizations, but they can increase the costs to attackers and decrease the costs of defenders. Similarly, innovation in software development processes and technologies can make attack much more difficult. Offensive advantages are not inevitable in cyberspace, and they cannot be eliminated by a technological fix. Instead, gaining defensive advantage will require persistent investments in technological management, innovation, and skill.

146. Maness and Valeriano also suggest that the "use it and lose it" nature of cyberweapons calls the applicability of the security dilemma into question. See Maness and Valeriano, *Cyber War versus Cyber Realities*, p. 36. Their analysis, however, also frequently invokes the security dilemma as a problem.