Hacking Democracy

Thomas Rid, Ben Buchanan

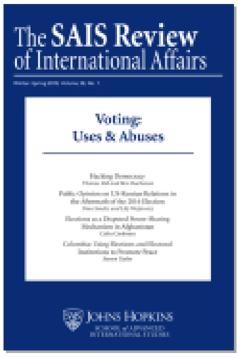

SAIS Review of International Affairs, Volume 38, Number 1,
Winter-Spring 2018, pp. 3-16 (Article)

Published by Johns Hopkins University Press
DOI: *https://doi.org/10.1353/sais.2018.0001*

Hacking Democracy

Thomas Rid and Ben Buchanan

This article outlines democracy's four main weak points against cyber operations. The first weak point is the cybersecurity debate itself. Adversaries are able to exploit fissures among the information security community to increase the effectiveness of their operations. Institutions, including political organizations, campaigns, civil society groups, etc. are the main target of cyber operations against democracy. Targeting institutions allows adversaries to undermine the population's trust in its political system. The third weak point is technical infrastructure, which includes voting information databases, voting machines and vote-counting machines. Operations targeting technical infrastructure can change the result of an election or undermine public confidence in the results. Democracies' intelligence agencies are the fourth major target. Cyber operations against their intelligence agencies weaken democracies' ability to defend themselves from further operations.

Introduction

Cybersecurity—the practice and the debate—is more than a quarter century old. Early on, military concepts dominated, with the US Air Force and the RAND Corporation among the earliest adopters.[1] The context for this pioneering work was command-and-control warfare. By the mid-1990s the declared goal was winning in network centric warfare, taking advantage of a new "revolution in military affairs," and achieving "information dominance." The utopian goal of turn-of-the-century military visionaries was striking: to win a war before it even started. Meanwhile, for twenty-five years, the corresponding dystopian vision of an "electronic Pearl Harbor" formed a counterpoint in the early cyberwar debate. The vision of winning swiftly by high-tech cyberattack dialectically nourished the fear of perishing in one.[2] Perhaps no idea was more critical at the extreme ends of the spectrum of computer network attack—and defeat—than the commonly accepted view that the internet, like airpower, affords advantage to the offense over the defense.[3] Whoever acts first, wins.

Then came the year 2016. Cybersecurity turned into the central issue of the US general election, gaining further in profile during the transition period and early 2017. Information operations helped mar the presidential ambitions of the losing Democratic candidate and undermined the legitimacy of the winning Republican president. Both of these assessments are highly-charged and animate a debate that is more political than technical—this deeply politicized

Thomas Rid is a professor of strategic studies at Johns Hopkins University SAIS. Ben Buchanan is a post-doctoral fellow at Harvard University Kennedy School of Government.

SAIS Review vol. 38 no. 1 (Winter–Spring 2018) © 2018 Johns Hopkins University

state of the art illustrates that cybersecurity has been elevated to a public profile and significance never seen before in its quarter-century history. Yet, despite it all, almost no serious commentators were ready to see the much-feared electronic Pearl Harbor in Russia's election interference, let alone a "cyberwar"; these military monikers were flawed. 2016 showed that the quarter-century-old debate was littered with broken ideas.

Yet, despite it all, almost no serious commentators were ready to see the much-feared electronic Pearl Harbor in Russia's election interference, let alone a "cyberwar"; these military monikers were flawed.

Cybersecurity was not and still is not ready for prime time. The field disappoints in practice, policy, and scholarship. Cybersecurity is under-delivering on the defense, because, even decades later, soft targets are still soft, and fruits are still hanging low aplenty. Cybersecurity is under-delivering in the public debate because facts are too often poorly shared, major incidents not revealed, with too many public commentators still struggling to distinguish firm forensics from flimsy. Finally, cybersecurity is under-delivering in theory, because 25 years after the first seminal works, core concepts are still wobbly and contested too easily. These flaws are not just impairing hacking victims, news stories, or research articles—this triple deficiency is subverting democracy itself. 2016 has brought the collective information security vulnerabilities of the American democratic process into sharp relief. The stakes could not be higher: liberal democracy has become a juicy target of ever more sophisticated Computer Network Attacks (CNA)—and CNA have become a preferred semi-covert action tool of the early 21st century.

In this article we take stock. We do so by asking three questions. The first is a conceptual one: what kind of internet-enabled offensive operations have proven to be most dangerous for liberal democracies? The second question is a historical one: how have these types of political computer network operations evolved, and what proved to be their most significant mechanisms of hacking democracy as they evolved? The third question is practical: what can be done to fix the problem as we move forward?

The well-established adage that the internet is "offense-dominant" is accurate, but for a widely misunderstood reason. The common, yet flawed, technical argument is that the architecture of the internet means that the offense dominates in computer network operations. This argument has become gospel in the cybersecurity debate. The offense has indeed often dominated in computer network operations thus far, but this dominance has been a result of flaws in the broader cybersecurity field, not just in specific technologies. The offense dominated in the 2016 election interference, in the 2012 Saudi Aramco incident, and in the 2007 Estonia case because of the field's triple deficiency, not because of long-suspected technical characteristics inherent in the architecture of the internet. To focus on only the technical, or only on the political, is to take too narrow a view.

The re-emergence of active measures fueled by hacking and leaking exacerbated this already existing trend. Active measures are intelligence operations designed to shape decisions and opinions of adversaries. Sometimes active measures exploit technical weaknesses, but more fundamentally these operations exploit fissures in political systems, societies, and communities, including in the cybersecurity research community. The oft-repeated claim that "cyberspace is offense dominant" both admits failure in the past and concedes defeat in the future, thus conceding weakness and limiting ambitions in theory and in operations. Instead, we argue that the cult of the online offensive is contingent, that it can be reversed through better cybersecurity practice, policy, and theory. Upon closer examination it becomes clear that open societies have already begun to harden their defenses against computer network operations and have mechanisms to continue this buildup into the future.

This paper is structured as follows. In four sections, we will briefly outline democracy's main potential weak points against cyber operations. The first weak point is the cybersecurity debate itself. The more direct political targets then fall into three sets: institutions, infrastructure, and intelligence. Within these sets, this paper will first present the most visible type of offensive activity and work toward the most clandestine and hard-to-examine form of operation. As the analysis moves from overt to covert, the challenges for researchers and investigators fast increase in scholarship, government, and industry.

Sometimes active measures exploit technical weaknesses, but more fundamentally these operations exploit fissures in political systems, societies, and communities, including in the cybersecurity research community.

A Fractured Field

The field of cybersecurity has been shaped by major incidents over the decades. These advances and attention have often fed into the policy discussion. Perhaps the first major case, thirty years ago, was the Morris worm, a highly infectious piece of malicious software. Later, in late 1996, Department of Defense investigators discovered *Moonlight Maze*, the first major state-on-state espionage campaign; "MM," as it came to be known, arguably never stopped.[4] Slammer, the uncontrolled computer worm that found its way into a nuclear power plant, made its mark fourteen years ago.[5] The first *Sofacy*/APT28 samples date back to 2004, but were identified publicly only years later. *Titan Rain*, the first large Chinese espionage campaign, was revealed in 2005, but had already been under way for several years by then.[6] The infamous "Denial of Service" attacks on Estonia, which were christened *Web War One* by *Wired* magazine,[7] occurred one decade ago roughly concurrently with the development of Stuxnet,[8] a potent Israeli-United States sabotage attack against Iran's nuclear enrichment facilities. Five years ago, in the summer of 2012, Saudi Aramco experienced what may still be one of the most devastating publicly known computer network intrusions, which incapacitated 30,000 workstations in the world's largest oil company.[9]

In February 2013, Mandiant published a landmark attribution report, "APT1." The document contributed to changing the notion that digital intruders could easily cover their tracks, and therefore attracted a great deal of attention.[10] Since 2013, both the pace of major incidents increased, as did their often-public investigations. In May 2014, the FBI filed its first indictment of adversarial hackers, of five Chinese People's Liberation Army intelligence operators.[11] Later in 2014, North Korea's high-profile hack of Sony Pictures Entertainment revealed what was then the politically most famous cybersecurity incident to date.[12] This increase in the volume of operations and their visibility is the result of a number of factors. In particular, it is likely that the publicity that followed Stuxnet as well as the National Security Agency (NSA) leaks led a number of countries to make investments in offensive capabilities—the leaks, counterintuitively, also showed that the West's advanced technical intelligence agencies have superb attribution capabilities and apparently better operational security than their Russian or Chinese counterparts.[13]

Meanwhile, the policy debate struggled to keep pace with incidents, especially after 2013.[14] Conferences like Defcon, Blackhat, RSA, or Infosec Europe have mushroomed from niche meetings into gargantuan conventions. A host of smaller annual conferences serves an ever more finely branched community of researchers, such as Kaspersky's Security Analyst Summit, BSides, Code Blue, ShmooCon, Troopers, S4, the new OffensiveCon, or the Chaos Computer Club conferences—the last with a more notable anti-establishment tinge. Graduate programs in cybersecurity are proliferating. Major news outlets have designated cybersecurity reporters, and the press coverage of major incidents has improved significantly.

Yet at closer examination, the community is disintegrating rather than consolidating. Entire subfields and sub-debates have emerged, for example the debates on vulnerabilities and exploits, on civil society targeting, on industrial control system security, on digital forensics and incident response—"DFIR" in jargon—or on threat intelligence and "threat hunting," as well as on privacy, encryption, and bulk collection and surveillance. The overlaps between these subfields are shrinking rather than growing, foreshadowing less fruitful dialogue even within the field. One example is the highly contentious issue of applying the Wassenaar export control (a multilateral agreement regulating the export of weapons) to hacking tools. This cacophony makes it harder for elected officials and policy makers to grasp important trends. Developments such as the emergence of politically motivated wiping attacks (deleting or destroying sensitive or key data), including a large number of undisclosed wiping incidents, have been lingering at the shadowy fringes of the debate for too long.

Developments such as the emergence of politically motivated wiping attacks (deleting or destroying sensitive or key data), including a large number of undisclosed wiping incidents, have been lingering at the shadowy fringes of the debate for too long.

The fissures in the debate are not just inconvenient for newcomers, journalists, and policy makers. The cracks in the field also make it easier for adversaries to exploit the cybersecurity debate itself. Just as journalists added value to the DNC leaks by sifting dumps and pulling out the gems, so did cybersecurity experts add value to the Shadow Brokers' recent dumps of NSA hacking tools by sifting files, testing them, and pulling out shiny objects—thus enabling the desired operational effect against the US government in the first place. The Shadow Brokers' operation was likely designed with this dynamic in mind. It cleverly distracted the cybersecurity community from larger political questions while weaponizing the community's collective sleuthing abilities.

Scholarship on cybersecurity faces unique and particularly difficult challenges. In recent years, every large international relations journal has published articles on cybersecurity. Yet some of the debate's most influential texts are conceptual in nature, for example the landmark 1993 piece "Cyberwar is Coming!"[15] the 1995 article "What is Information Warfare?"[16] or the 2013 retort "Cyber War Will Not Take Place".[17] Too many scholars find this old and tired debate about cyberwar to be irresistible—an especially glaring shortcoming considering the proliferation of incidents and the rich empirical material available today.

University-based academics, no matter their discipline, face an additional unique challenge in this field: most cutting-edge work requires access to insights and data from operations and incidents. Some of the most valuable sources for research, as a consequence, are either in the intelligence community (IC) or, more likely, in the vibrant private sector market for incident response, digital forensics, and threat intelligence. Yet most PhD students, post-docs, junior faculty, and even senior scholars have little or no access to these communities. The dearth of scholarship that effectively utilizes leaked NSA and GCHQ documents is an illustration of this problem. Insightful research that takes advantage of some of the more technical leaks is all too rare. The mirror image of this problem is reflected in the private sector, where a wide range of companies share reports under the so-called "Traffic Light Protocol" (TLP), a trust-based industry classification system. The number of TLP-coded reports is today reaching well into the five digits. Yet too few scholars have seen, let alone taken advantage of, these rich data sources.

Not even the name of the field is settled. The information security community still rolls its collective eyes at the vintage cringe-worthy "cyber"—which has quickly become a noun, especially in policy and military circles—although most use the moniker at the same time. "Imagine a stranger renamed your profession 'poop-mining'—and you had to start using the term yourself. This is how I feel about 'cyber,'" wrote Matthew Green, a prominent professor of cryptography at Johns Hopkins University, in May 2017.[18] His note received than 2,300 approvals on Twitter. To a non-specialist, "cyber" can feel at once ephemeral and intimidating.

Institutions

The main targets of computer network operations against democracy are institutions, not computer networks. Institutions are based on trust; they work because people trust them to work. That trust is one of the top targets of 21st century digital active measures. It is also one of the easiest targets. Undermining institutions that serve key democratic functions is a blow against democracy itself.

Two broader categories of institutions are being targeted. Political institutions comprise the first and most prominent type of target. These organizations, such as political committees or campaigns, offer a soft and large target surface. In many respects, election cycles are the most vulnerable moment for liberal democracies, akin to the period during which a lobster sheds its protective shell, and for a while remains covered in a soft, vulnerable shell, marking it out as particularly easy and attractive prey. Espionage during this vulnerable period is not new. Probably the most notable early examples of computer hacking assisting in this espionage effort come from the 2008 election cycle, in which Chinese operatives reportedly targeted both the Obama and McCain campaigns.[19] Even in 2016, the Democratic campaign was not targeted by only Russian actors. But while a democracy should attempt to thwart this kind of intrusion, the theft of documents does not pose the same kind of threat as more active operations.

In many respects, election cycles are the most vulnerable moment for liberal democracies, akin to the period during which a lobster sheds its protective shell, and for a while remains covered in a soft, vulnerable shell, marking it out as particularly easy and attractive prey.

Four points deserve attention, in light of the 2016 activities. First, the Russian operations targeted voters directly. Cold War active measures were mostly directed against politicians and journalists and reached the wider public only indirectly. By the early 2010s, social media introduced a new platform that allowed remote operators to reach voters directly, via ads, trolls, and personal messages. The Robert Mueller indictment brought forward against the Internet Research Agency on 18 February 2018 brought a number of case studies to wider public attention. Second, and related, the 2016 operations took place at a substantial scale. They played out in traditional mass media and on social media, reaching millions of people. The operators published tens of thousands of individual documents in more than 80 individual leaks. Most (although not all) active measures throughout the Cold War tended to be far smaller and narrower. Russian intelligence agencies successfully used the scalability of innovative hack-leak-amplify activities to their advantage.

Third, the 2016 influence operations were poorly disguised, perhaps even semi-overt by design. Not only was it obvious that a hacking-aided influence campaign was going on, but it was reasonably apparent which foreign power was conducting it. As early as mid-June 2016, a range of outside experts squarely placed the blame of the budding leaking operation on Russia, following the initial hacking attribution of the cybersecurity firm CrowdStrike.[20] Politically there may have been partisan incentives to call the evidence into question. Technically and historically, however, the evidence of a Russian hand was inescapable early on.[21] Fourth, the operation almost certainly worked better and longer than anticipated—independent of whether the active measures actually affected the vote or not. Candidate Donald Trump cited foreign hacking favorably on numerous occasions during the campaign.[22] He repeatedly cited hacked documents, professed his "love" for WikiLeaks,[23] and encouraged Russian computer network intrusions to recover Hillary Clinton's emails of interest. Taunted and provoked, Trump himself worsened damage to the legitimacy of his own presidency, both before and after the election. Moreover, after Trump was sworn in as President of the United States, he continued to call the evidence of Russian election interference into question, thus inadvertently straining the credibility of America's intelligence and law enforcement community, thus providing incentives and cover for follow-on active measures.

The other category of democratic institutions at exceptional risk of nation-state hacking is civil society groups, particularly those agitating for human rights or democratic policies. Activist movements often find themselves in at least partially threatening situations. Governments, often including their own domestic government, perceive grassroots activists to be threats to stability and order. Most non-democratic regimes target, with some semi-regularity, citizen activists. One of the clearest example of this campaign comes from the United Arab Emirates. Ahmad Mansoor, a noted activist there and a member of Human Rights Watch, found himself in the crosshairs of a conglomerate run by a leading member of the ruling family. Investigation by cybersecurity researchers revealed that—in addition to old-fashioned tools of repression like arrest, theft, and beating—Mansoor was the victim of more new-fangled hacking.

The spyware that targeted Mansoor was in place for more than a year before researchers found it. It was intrusive, gaining deep insight into his digital life and his organizing activities. The intimidating effects were real, even after the hacking operation was eventually thwarted. "It was as bad as someone encroaching in your living room, a total invasion of privacy, and you begin to learn that maybe you shouldn't trust anyone anymore," Mansoor said.[24]

Invoices from the company Hacking Team, which carries out or enables a great deal of government-sponsored hacking operations, show an extensive relationship with the United Arab Emirates.

The story, thus far, has no happy ending. The hacking effort against Mansoor garnered some media attention, and great notice within the cybersecurity community. The technical mechanics of it, including the use of significant hacking tools, indicate that Mansoor was a priority target. But the publicity did little to help him. On a later trip to the United Arab Emirates, he was arrested again and remains detained. Mansoor was not alone. Invoices from the company Hacking Team, which carries out or enables a great deal of government-sponsored hacking operations, show an extensive relationship with the United Arab Emirates. The Emirates, as of 2015, appear to be the company's second biggest customer. They have paid more than $630,000 for hacking operations targeting more than a thousand people. It is yet another reminder of how significant and scalable the hacking threat can be.

It is not just the United Arab Emirates that has used hacking tools against civil society. China, Mexico, Ethiopia, Morocco, and others are doing so as well.[25] Their targets vary, and some are more directly connected to movements for representative governments than others, but the overall trend is clear: hacking is a tool not just for non-democratic foreign governments to target mature liberal democracies, but for non-democratic governments to nip domestic threats to their regime in the bud.

Infrastructure

Another form of political targeting is directed not at the trust in institutions, but at the underlying technical infrastructure. This form of hacking goes one level deeper and interferes either with the integrity of key data or machines. While such operations can certainly also have an effect on trust—an attack on infrastructure can also be an attack on institutions—at its core it is something more direct than the last category of operations.

Some theorized attacks on infrastructure have been demonstrated in practice by researchers. Foremost amongst these are the myriad of ways in which voting infrastructure could be targeted.[26] The range of possible attack vectors is significant and can be illustrated by tracing the path a typical American voter takes in order to participate in an election. A citizen's first formal step of political engagement upon reaching adulthood is to register to vote. In doing so, they input their personal—and often private, sensitive, or identifying—data into a computer system. The maintenance and security of these computer systems varies enormously, but it is clear both that they are significant targets and that hackers have had success in breaching them and copying out the data within. These registration databases are ripe targets for compromise. There are numerous examples of hackers penetrating the registration systems and databases of American states and other entities. The total number of records affected often numbers well into the millions, sometimes into the hundreds of millions.[27]

There are numerous examples of hackers penetrating the registration systems and databases of American states and other entities.

On Election Day, in-person voters check in at their assigned polling places. Election workers typically verify the voter's registration in a poll book, which contains a list of all the voters eligible to vote at that place. Some states, such as Ohio, have digitized these books.[28] Digital records are more efficient, but can get hacked. One possible vector of attack is to manipulate the poll books in a given area, perhaps by removing five percent of the voters of a given party—a so-called data integrity attack. If voters of a particular political affiliation, gender, race, or age were disproportionately turned away, this could quickly fuel claims that the election was subject to illicit manipulation.

Next, the voter actually marks and casts their ballot. Unfortunately, the cybersecurity of voting machines is both important and flawed. Voting machine security is uneven. A series of audits turned up a wide range of problems in many of the major voting machines used in the United States. Due to a lack of funding, a sizable portion of voting machines is more than a decade old and depends upon outdated security models. Notorious examples abound, such as Wi-Fi connected systems that used a default password of "abcde,"[29] or the system that researchers were able to reconfigure in order to play popular video games like Pac Man. In 2017, the hacker conference DEF CON vividly highlighted the state of affairs with their "Voting Machine Hacking Village," in which attendees worked together to quickly and successfully demonstrate a substantial number of security flaws.

Before a Communist Party election, Josef Stalin famously said, "I consider it completely unimportant who in the party will vote, or how; but what is extraordinarily important is this—who will count the votes, and how."[30] This quote suggests a fourth step in the voter's journey, and another opportunity for hackers: tabulation. Many of the vulnerabilities with machines can permit the manipulation of vote-counting functions.[31] In more centralized systems, the computers at the core of the voting infrastructure are ripe targets. In Ukraine in 2014, hackers launched a wiping campaign against these vital machines just three days before the election.[32]

After vote tabulation, citizens find out the result of the election. Here, too, there is the possibility for trouble and worries of illegitimacy. No example is more prominent than the Bush-Gore controversy in 2000, in which major networks first called Florida for Al Gore, then retracted the claim amidst a vote count that came down to the wire. The uncertainty lasted into the next month before the United States Supreme Court awarded the state to George W. Bush. The case might provide inspiration for digital saboteurs.

Intelligence

Intelligence agencies are fundamental to protecting democracy. The NSA in particular takes on a direct role in engaging with and thwarting their foreign counterparts, including attempts to target democratic institutions and infrastructure. For this and other reasons, the US intelligence community and its allies find themselves in the line of fire. The most striking post-Cold War case of a counter-IC active measure is known as Shadow Brokers. First visible in a series of August 2016 messages, the operators behind various Shadow Brokers social media and developer accounts began posting evidence that they had obtained classified NSA tools (referring to the NSA in infosec-jargon as "Equation Group"). The messages escalated throughout the next twelve months, revealing more and more about NSA activities. The group revealed some of the agency's most potent penetration tools. This included one, ETERNALBLUE, which United States government operators had used for years; one source called the exploit “fishing with dynamite.”[33]

The Shadow Brokers’ move to publish the pilfered tools had a triple effect. First, it thrust secret NSA operations back into the media spotlight. The NSA faced questions about how these powerful tools got out of its secure confines, and whether it was right to retain such vulnerabilities in the first place, as others could have exploited them as well. Major companies, such as Microsoft, sharply criticized the NSA in this regard.[34] Even a hawkish former director of both the NSA and the CIA, Michael Hayden, said, “If American espionage cannot protect the special tools that it possesses, it doesn’t matter that they are good people working for good purposes under good oversight. If they cannot protect the tools, I just can’t mount the argument to defend that they should have them.”[35] Secondly, the leaks made it either technically impossible or politically much more difficult for the NSA to use the published tools.[36] The disclosure, highly likely an adversarial intelligence operation, sabotaged NSA’s collection capabilities. Third, the leaks enabled others to use the NSA’s insights for their own purposes. Others, less worried about getting caught, could use the publicized tools as blueprints for developing their own similar capabilities. The most prominent example of this is an attack, known as WannaCry, which occurred in May 2017. Likely conducted by North Korea, this attack rapidly spread to hundreds of thousands of computers around the world. Most prominently, this included many computers operated by Britain’s National Health Service. The attack code served as a broken form of ransomware, encrypting the machine’s files until a payment was made; in a sign of incompetence, an inadvertent start to the operation, or a desire to make noise and not money, the payment and decryption mechanism appeared broken.

The disclosure, highly likely an adversarial intelligence operation, sabotaged NSA’s collection capabilities.

A little more than a month later, Russian military intelligence launched a comparable attack that caused “billions” of dollars of damage to the world economy.[37] The operation, named NotPetya, limited its targets to users of a Ukrainian tax reporting software; but the attack spread quickly within the networks of a significant number of international companies. Though NotPetya spread using a variety of mechanisms, including credential theft, it in part took advantage of the same software vulnerabilities the NSA had exploited, again generating embarrassing and damaging publicity for the American intelligence agency. In effect, NotPetya served as a double sabotage tool: it crippled Ukraine and embarrassed the US.[38] One minister of the Russian Federation directly challenged US intelligence at the Munich Security Conference 2018 by bringing up NotPetya’s ETERNALBLUE re-use as an example of US government malpractice.

The likely goal of the Shadow Brokers operation was weakening Five Eyes (Australia, Canada, New Zealand, the United Kingdom, and the United States) intelligence. The NSA is the most potent signals intelligence agency of any liberal democracy, and its work has enabled many advance threat warnings to NATO allies, against terrorism as well as against ongoing espionage campaigns—weakening the West's wider intelligence community is therefore weakening the capacity of democracies to self-defend.

The technical intelligence community in the West especially faces another new problem: the loss of their near-complete monopoly on technical intelligence and classified information. A range of private sector companies are now competing with signals intelligence organizations not just for talent, but in collection and analysis as well. But in this newfound competition also lies a great opportunity for the defense of democracy from new threats.

In the early 2010s, more and more companies started to "publish" restricted reports under the so-called "Traffic Light Protocol." The traffic light protocol is an industry-wide, trust-based convention that governs the sharing of confidential information, both orally and in writing. The protocol is usually abbreviated as TLP. The protocol works as follows: TLP:WHITE means the document is public; TLP:GREEN means recipients can share the file with peers; TLP:AMBER means recipients can share within their own organization; and TLP:RED means recipients cannot share with any parties outside of the exchange or meeting in which the information was originally disclosed. The US and UK governments, and many others besides, have endorsed the system and sometimes use it themselves. TLP reports usually have a classification header, like a government document would have.

The private sector is taking the lead in intelligence because firms often have greater access to information and data through their products and services. This rise of private sector intelligence comes with a range of risks. One is imperfect quality assurance and vetting for reports and analysis. Another one is the rising number of proprietary, for-profit reports that never get published as TLP:WHITE at all, and thus remain out of view. But TLP reports have one major advantage over classified government documents: they can be shared more easily, and data, as well as insights, sometimes permeate out into the semipublic and public debate. Indeed, sometimes private sector reports may be used as a vehicle to publish previously classified information on high-profile cases.

Despite these challenges, the overall trend is a positive one: after decades of a predominantly military frame of mind, epitomized in the tired notion of "cyberwarfare" and "offense domination," the field of cybersecurity is beginning to right itself back to where it always belonged, which is at the intersection of public and private intelligence. Ever growing numbers of operators in the wider information security community have cut their teeth in intelligence. These individuals tend to be more comfortable talking about collection rather than about coercion, about interception rather than intervention. The notion that political aspects pollute a pure technical analysis has, thankfully, less and less currency. More and more terms of art from the study and the history of

The paradigm is shifting from a military mindset to an intelligence-led philosophy of information security—indeed the quiet rise of the term infosec over cyber highlights this trend.

intelligence find their way into the cybersecurity lexicon, for example active measures, active and passive collection, operational security, all-source in investigations, attribution, estimative language, assessment, analysis of competing hypotheses, aperture of analysis. The paradigm is shifting from a military mindset to an intelligence-led philosophy of information security—indeed the quiet rise of the term infosec over cyber highlights this trend. It is this area of public-private partnership, thus far mostly out of public view, where great opportunities to defend democracy have emerged.

Conclusion

The cybersecurity debate, and the wider infosec community, face a moment of reckoning. 2017 has exposed that one of the greatest vulnerabilities of liberal democracy in the 21st century is the information security of the wider political system, from personal email authentication to government routers, from boutique malware attacks to old-school infiltration, from voting machine compromise to shrewd forgeries, from hacking the grid to brazen active measures against the NSA as well as everyday voters.

The offense has indeed dominated for decades, not as a result of technical defects, but as a result of man-made defaults. It is still too easy not to use two-factor authentication; it is still too easy to hide bots at scale on some social media platforms; it is still too easy to breach networks and stay undetected for months; it is still too easy to lure journalists with a scoop on a flimsy intelligence story. Geeks and wonks face mirroring temptations: overestimating what they understand, underestimating what they don't understand, dismissing the silly mannerism and alien jargons of the other side, and secretly sneering at the inevitable mistakes of those who try to cross the divide. This is not a trite, smartass observation from two smug academics: 2016 has brought to the fore that adversaries are again getting better at driving wedges into the divisions that divide us. "Us" as liberal, open democracies. But also, "us" as a community of professionals dealing with information security. Spotting useful idiots is easy after all—just ask, "Am I still willing to adjust my view in response to new insights or new evidence and admit that I was wrong?"

> *Geeks and wonks face mirroring temptations: overestimating what they understand, underestimating what they don't understand, dismissing the silly mannerism and alien jargons of the other side, and secretly sneering at the inevitable mistakes of those who try to cross the divide.*

Notes

[1] These early adopters derived the term for "cybernetics," see Thomas Rid, *Rise of the Machines* (New York: WW Norton, 2016).

[2] These terms and the transition into this stark tension is best articulated by Richard Clarke and Robert Knake, *Cyberwar* (New York: HarperCollins, 2010).
[3] For further discussion of offense dominance and its central role in international relations scholarship, see Robert Jervis, "Cooperation under the Security Dilemma," *World Politics* 30, no. 2, 1978. For an analysis of sources arguing the offense dominance of the internet, see Ben Buchanan, *The Cybersecurity Dilemma* (New York: Oxford University Press, 2017), Ch. 5.
[4] For full discussion, see Rid, *Rise of the Machines*. See also Costin; Moore Raiu, Daniel; Guerrero-Saade, Juan Andrés; Thomas Rid, "Penquin's Moonlit Maze: The Dawn of Nation-State Digital Espionage," Kaspersky Lab, April 3, 2017.
[5] Kevin Poulsen, "Slammer Worm Crashed Ohio Nuke Plant Network," *Security Focus*, August 19, 2003.
[6] For discussion of this case, see Joel Brenner, *America the Vulnerable* (New York: Penguin, 2011).
[7] Joshua Davis, "Hackers Take Down the Most Wired Country in Europe," *Wired*, August 21, 2007.
[8] For the best history, see Kim Zetter, *Countdown to Zero Day* (New York: Crown, 2014).
[9] Christopher Bronk and Eneken Tikk-Ringas, "The Cyber Attack on Saudi Aramco," *Survival* 55, no. 2, 2013.
[10] "APT1," Mandiant, February 18, 2013.
[11] Ellen Nakashima, "Indictment of PLA Hackers Is Part of Broad U.S. Strategy to Curb Chinese Cyberspying," *the Washington Post*, May 22, 2014.
[12] Michael Schmidt, Nicole Perlroth, and Matthew Goldstein, "F.B.I. Says Little Doubt North Korea Hit Sony," *the New York Times*, January 7, 2015.
[13] Thomas Rid and Ben Buchanan, "Attributing Cyber Attacks," *Journal of Strategic Studies* 39, no. 1, 2015.
[14] Adam Segal calls this fateful moment "year zero" for cybersecurity. For a full discussion, see Adam Segal, *The Hacked World Order* (New York: Public Affairs, 2016).
[15] John Arquilla and David Ronfeldt, "Cyberwar Is Coming!" RAND Corporation, 1993.
[16] Martin C. Libicki, *What Is Information Warfare?* (Washington, DC: National Defense University, 1995).
[17] Thomas Rid, *Cyber War Will Not Take Place* (Oxford/New York: Oxford University Press, 2013).
[18] See @matthew_d_green, May 19, 2017, https://web.archive.org/web/20170724110256 /https://twitter.com/matthew_d_green/status/865666772241862656
[19] Ryan Naraine, "Newsweek: Obama, McCain Campaigns Hacked by 'Foreign Entity'," *Newsweek*, November 5, 2008.
[20] Dmitri Alperovitch, "Bears in the Midst: Intrusion into the Democratic National Committee," *CrowdStrike*, 2016.
[21] See Thomas Rid, "All Signs Point to Russia Being Behind DNC Hack," *Motherboard/Vice*, July 25, 2016.
[22] Ashley Parker and David Sanger, "Donald Trump Calls on Russia to Find Hillary Clinton's Missing Emails," *the New York Times*, July 27, 2016.
[23] Patrick Healy, David Sanger, and Maggie Haberman, "Donald Trump Finds Improbable Ally in Wikileaks," *the New York Times*, October 12, 2016.
[24] Nicole Perlroth, "Governments Turn to Commercial Spyware to Intimidate Dissidents," *the New York Times*, May 29, 2016.
[25] Nicole Perlroth, "Governments Turn to Commercial Spyware to Intimidate Dissidents," *the New York Times*, May 29, 2016.
[26] For a full discussion, see Ben Buchanan and Michael Sulmeyer, "Hacking Chads," Belfer Center for Science and International Affairs, October 2016.
[27] Sometimes these systems have shockingly little security. For the biggest example, see Nick Corasaniti and Rachel Shorey, "Millions of Voter Records Posted, and Some Fear Hacker Field Day," *the New York Times*, December 30, 2015.
[28] Karen Farkas, "Electronic Poll Books Will Be at Voting Locations across the State by November 2016," *Cleveland Plain Dealer*, August 28, 2015.
[29] "Security Assessment of Winvote Voting Equipment for Department of Elections," Commonwealth Security and Risk Management: Virginia Information Technology Agency, 2015.

[30] Boris Bazhanov, *Memoirs of the Former Secretary of Stalin* (Moscow: III Tysiacheletie, 2002).
[31] For a series of audits, see "Source Code Review of the Sequoia Voting System," Berkeley University of California: California Secretary of State, 2007; "Source Code Review of the Diebold Voting System," Berkeley University of California: California Secretary of State, 2007; "Source Code Review of the Hart Intercivic Voting System," Berkeley University of California: California Secretary of State, 2007.
[32] Mark Clayton, "Ukraine Election Narrowly Avoided 'Wanton Destruction' from Hackers," *Christian Science Monitor,* June 17, 2014.
[33] Ellen Nakashima and Craig Timberg, "NSA Officials Worried About the Day Its Potent Hacking Tool Would Get Loose. Then It Did.," *the Washington Post*, May 16, 2017.
[34] "Ransomware Attack 'Like Having a Tomahawk Missile Stolen', Says Microsoft Boss," *the Guardian*, May 14, 2017.
[35] Ken Dilanian, "Can the CIA and NSA Be Trusted with Cyber Hacking Tools?" *NBC News*, June 30, 2017.
[36] For two examples, see "ProjectSauron: Top Level Cyber-Espionage Platform Covertly Extracts Encrypted Government Comms," Kaspersky Lab, August 8, 2016; "Equation: The Death Star of Malware Galaxy," Kaspersky Lab, 16 February 2015. See also Ben Buchanan, "The Legend of Sophistication in Cyber Operations," Belfer Center for Science and International Affairs, January 2017.
[37] "Statement from the Press Secretary," White House, February 15, 2018, https://web.archive.org/web/20180315182022/https://www.whitehouse.gov/briefings-statements/statement-press-secretary-25/
[38] Nicole Perlroth, Mark Scott, and Sheera Frenkel, "Cyberattack Hits Ukraine Then Spreads Internationally," *the New York Times*, June 27, 2017.