A Tale of Two Cyber Conflicts

–

Civil Society and Commercial Threat Reporting[1]

Lennart Maschmeyer

Ronald Deibert

Jon Lindsay

University of Toronto

Paper to be presented at the 2019 ISA Annual Convention in Toronto

Abstract

Public and academic knowledge of cyber conflict relies heavily on data from commercial threat reporting, but this data provides a distorted view of cyber threat activity. Commercial actors are selective about the threats they focus on and what they report in public documents. This selection bias hampers not only scholarship on cybersecurity, but also has concerning consequences for democracy. Threat reporting can be conceived as a public good that is by and large provided by actors in the private sector. As collective action theory leads us to expect, accurate threat reporting is underprovided while the interests of commercial information security firms in high end threats to high-profile victims are overrepresented. Prevalent threats to civil society organizations, which lack the resources to pay for high-end cyber defense, tend to be ignored. Commercial threat reporting thus presents a truncated sample of cyber conflict. We present and analyze an original dataset of available public reporting by the private sector as well as independent research centers. We also present three in-depth case studies tracing reporting patterns on a cyber operation targeting civil society. Our findings strongly confirm the hypothesized selection bias in reporting.

[1] Acknowledgments: the authors would like to thank Max Smeets, Masashi Crete-Nishita and Irene Poetranto, the participants of the 2018 "Global Digital Futures" workshop at Columbia University's School of International and Public Affairs and the team at ETH Zurich's Center for Security Studies for their feedback on previous drafts of this paper. We are also grateful for the generous funding from the Carnegie Corporation of New York and the School of International and Public Affairs at Columbia University, the Ford Foundation, the John D. and Catherine T. MacArthur Foundation, the Sigrid Rausing Trust, the Oak Foundation and the Open Societies Foundation that helped make this project possible.

On October 1, 2018, the Citizen Lab at the University of Toronto published a research report indicating that the phone of Omar Abdulaziz, a prominent dissident of the Kingdom of Saudi Arabia, had been infected with sophisticated spyware (Marczak et al. 2018). Citizen Lab researchers established with a high degree of confidence that his phone was infected by an operator associated with the Saudi Arabian government, and identified the spyware as the 'Pegasus' suite manufactured by the Israel-based vendor NSO Group. Abdulaziz, a university student and Canadian resident, has a popular YouTube channel where he posts regime-critical videos. One day later, another high-profile dissident, the Washington Post journalist Jamal Khashoggi, was lured into the Saudi consulate in Istanbul, Turkey where he was murdered and dismembered. Soon thereafter it was revealed that Abdulaziz and Khashoggi were close confidants working together on a social media opposition campaign against the Crown Prince of Saudi Arabia, principally communicating over many weeks in what they mistakenly thought was private through a WhatsApp messaging application that the Citizen Lab discovered was being remotely monitored by Saudi intelligence. Although the reasons underlying the specific decision to murder Khashoggi are unknown, many have drawn connections to the surveillance uncovered in this operation, which revealed highly incriminating plans Khashoggi and Abdulaziz were discussing (Rogin 2018; Shezaf 2018).

This case illustrates several intersecting factors that are at the heart of our analysis presented in this paper. First, the tools of a sophisticated Israeli-based commercial spyware company, self-described as specializing in "cyber warfare", manufactured by a company valued at USD $1 billion, and selling its technology exclusively to government law enforcement, military and intelligence agencies, ended up being deployed by its Saudi client not to target an adversary government, or steal propriety business data, but instead to remotely monitor the iPhone of a critic of the regime. Second, the case underscores the relative asymmetry of threats faced by civil society, and their incapacity to deal with them. While governments and private companies can allocate vast sums of money to protect themselves from sophisticated electronic espionage operations, it is highly unlikely a university student can do so. Even a relatively well-resourced NGO typically lacks the capacity and resources to deal with threats on the order of a sophisticated Israeli-based surveillance company. Lastly, this case illustrates what might be described as a current reality of global cyber conflict – but one largely overlooked by both policy analysts and scholars alike. This is cyber conflict in action, but a very different kind of conflict than the one discussed in most scholarly and policy writing on cybersecurity. The latter focus almost exclusively on interstate conflict, threats to critical infrastructure and intellectual property theft for economic espionage; civil society targets like Omar Abdulaziz and Jamal Khashoggi receive only passing attention. What explains this mismatch?

We argue that incentives behind commercial threat reporting cause a selection bias in what is being reported, and what is not. The result is a distorted picture of cyber conflict, which hampers not only scholarship, but has concerning consequences for democracy.

Threat reporting is a public good provided by self-interested actors, which creates a collective action problem. Market failure leads to the underprovision of this good as the incentives that drive reporting create selection criteria that prioritize cases at the high end of conflict spectrum. There are three main criteria: unique tactics, techniques and procedures (TTP), high-profile victims, and high-profile threat actors. Selection across these criteria leads to underreporting of threats to civil society, which tend to score low across all three variables. Commercial threat reporting thus presents a truncated sample of cyber conflict.

To test our theory, we use a mixed methods approach to trace congruence between predicted and actual reporting patterns. We contribute an original dataset of all available public reporting by the private sector, 622 threat reports in total, as well as independent research centres, comprising 71 total reports. In addition, we obtained data from AccessNow's helpline, reflecting threat reporting from the perspective of civil society itself. Using content analysis, we code threat reports based on the selection criteria identified above. We conduct a quantitative analysis tracing reporting trends across these criteria and compare results between commercial and independent reporting. Findings confirm our expectations: only a low proportion of commercial threat reports discuss civil society, and those that do focus on high-profile victims and threat actors. Both the geographical distribution of reporting and attribution patterns are congruent with the hypothesized selection bias.

Finally, we select three cases of cyber operations targeting civil society from the dataset and conduct an in-depth qualitative analysis comparing independent and commercial reporting. We employ structured focused comparison and congruence testing to verify if threat intel reporting patterns on specific threat actors known to target civil society are consistent with our expectations. Cases are selected to cover the maximum range of values on the independent values (our hypothesized selection criteria). Hence, we select two most-likely cases with extreme values (e.g. highly unique TTP) and one least-likely case with intermediate values across all three variables. The findings of this plausibility probe strongly confirm our theory as even the least-likely case exhibits clearly selective reporting. In addition, we identify evidence for omission of civil society targeting by known actors in commercial reporting. The article concludes with a discussion of the dire implications of this selection bias, not only for scholarship, but more importantly civil society itself—and, by extension, democracy.

1. THE PROBLEM: WHAT DO WE KNOW?

Conventional wisdom considers the main 'cyber threats' to be disruption or destruction of critical infrastructure and intellectual property theft for economic espionage. For example, Demchak and Dombrowski warned that "future innovations might be able to destroy or disrupt other critical infrastructures upon which modern societies depend" (2011, 32). Meanwhile, former NSA director Keith Alexander asserted cyber espionage constitutes the "greatest transfer of wealth in history" (Rogin 2012). Betz and Stevens sum up prevailing threat perception as follows: "critical national infrastructures are said to be vulnerable to cyberattack; the vitality of the economy is said to be threatened with enervation through cyber-espionage and cyber-crime…and the property and well-being of citizens are claimed to be endangered by cyberwar" (Betz and Stevens 2011, 11). Similarly, Joseph Nye warns that "there is a wide range of cyber threats, including war, espionage, sabotage, and disruption" (Nye 2017, 47), and identifies possible means of deterring nation-state attacks on critical infrastructure and economic espionage (Nye 2017). In short, prevailing threat perception has focused primarily on cyber conflict as high-level interstate competition in a new domain, drawing on analogies with traditional warfare, strategy and intelligence.[2]

[2] See, among others: Arquilla, John, and David Ronfeldt. "Cyberwar Is Coming!" Comparative Strategy 12, no. 2 (April 1, 1993): 141–65; Gartzke, Erik. "The Myth of Cyberwar: Bringing War in Cyberspace Back down to Earth." International Security 38, no. 2 (2013): 41–73.Belk, Robert, and Matthew Noyes. On the Use of Offensive Cyber

Then came the discovery of the 2016 US Presidential Election interference campaign, now widely attributed to Russia, which demonstrated the effectiveness of cyber operations at the low end of conflict. None of the activity corresponded to prevailing threat models and prioritization, yet it was obviously significant. The discovery started with the detection of an intrusion into the systems of the Democratic National Convention and leaking of sensitive documents and emails via WikiLeaks. Subsequent investigations have revealed a large-scale, orchestrated influence campaign designed to sway voter opinions and foster divisions (Isaac and Wakabayashi 2017; Oremus and Newell 2017; ODNI 2017). The actual effects of the campaign on election outcomes continue to be hotly debated, but its significance as an instrument of cyber power is clear and reflected in threat perceptions of the campaign as an 'act of war' (Schleifer and Walsh 2017; Jamieson 2018).

Yet this campaign neither targeted critical infrastructure, nor did it cause any of the other major disruptions cyber warfare theorists have been warning about. Rather, its key targets were private individuals and civil society[3] organizations. Finally, it did not involve much sophisticated 'hacking. Instead, the campaign succeeded with a simple phishing email (Franceschi-Bicchierai 2016), online advertising services offered by Google and Facebook (Shane and Goel 2017), and fake accounts on popular social media sites (Swaine 2018). These revelations have sparked a wave of attention to threats to civil society for influence and interference (Landau 2017; Svetoka and Reynolds 2016; Nissen 2016), suggesting something new is afloat. As Jayamaha and Matisek put it, "no one expected that 'subversive instruments' would be used in such a way as to create intra-societal tensions through exploitation of civil society organizations" (Jayamaha and Matisek 2018).

Yet contrary to this collective surprise, civil society has been the target of cyber operations for a long time—and with many of the same TTP now employed in the election influence operation. Ronald Deibert warned as early as 2003 that "pressures from the security and commercial sectors to regulate and control the Internet are beginning to alter its basic material architecture in ways that may undermine not only the activities of global civic networks, but also the long-term prospects for an open global communications environment" (Deibert 2003). And while most scholars and policy-makers have

Capabilities. Belfer Center, 2012; Buchanan, Ben. The Cybersecurity Dilemma Hacking, Trust and Fear Between Nations. Oxford: Oxford University Press, 2017; Jensen, Eric Talbot. "CYBER DETERRENCE." Emory International Law Review, 2010; Lin, Herbert. "Offensive Cyber Operations and the Use of Force." Journal of National Security Law and Policy 4, no. 1 (2010): 63–86; Winterfeld, Steve. The Basics of Cyber Warfare: Understanding the Fundamentals of Cyber Warfare in Theory and Practice. Boston: Syngress/Elsevier, 2013.; Yannakogeorgos, Panayotis A., and Adam Lowther, eds. Conflict and Cooperation in Cyberspace: The Challenge to National Security. Boca Raton, FL: Taylor & Francis, 2014; Arquilla, John. "Cyberwar Is Already upon Us." Foreign Policy, no. 192 (2012): 84; Wirtz, James J. "The Cyber Pearl Harbor." In Cyber Analogies. Naval Postgraduate School, 2014; Rattray, Gregory J. Strategic Warfare in Cyberspace. Cambridge, Mass.: MIT Press, 2001; Nye, Joseph S. "Nuclear Lessons for Cyber Security?" Strategic Studies Quarterly: SSQ, Vol. 5, no. 4 (Winter 2011): 18–38; Kramer, Franklin D., Stuart H. Starr, and Larry K. Wentz, eds. Cyberpower and National Security. 1st ed. Washington, D.C: National Defense University Press : Potomac Books, 2009.

[3] We define these as independent organizations and individuals who engage in nonviolent political activity, including activists, dissidents, journalists, academics and nongovernment organizations. This is a heterogenous set of actors, but with one distinguishing characteristic: political engagement without being part of government, and without the use of coercion. Accordingly, political theorists define civil society as an "intermediate association realm between state and family" (White 1994, 379), and Mary Kaldor highlights the emphasis on rule-based behavior rather than coercion (Kaldor 2013, 1).

neglected this aspect of cyber conflict, independent research centers and non-profit organizations have documented the proliferation of targeted digital threats[4] to civil society for over a decade.[5] In fact, the Citizen Lab's report on *Gh0stNet* (Deibert and Rohozinski 2009), an international espionage campaign attributed to China targeting civil society and foreign governments across 103 countries, arguably started the practice of modern threat reporting.

These points suggest targeted threats to civil society have been proliferating for at least a decade, hence the use of the same methods in the election interference campaign should not have come as a surprise. Evidently, while scholarly and policy debates have focused mostly on speculative threats at the high end of conflict—preparing for a 'cyber pearl harbor' (ABC News 2011; Wirtz 2014)—in practice cyber conflict has thrived at the low end. This situation raises a crucial question: how do we know what cyber conflict is about: who is targeting whom and why?

2. THE CAUSE: MARKET FAILURE

We argue that the main cause for the mismatch between perception and practice of cyber conflict lies in academics and policy-makers' growing reliance on commercial threat reporting as a data source. Threat reporting is driven by business incentives, which shape a set of selection criteria that determine what is, and what is not, reported. The primary objective of reporting is to maximize attention, which results in selection bias towards operations at the high end of conflict. Consequently, commercial threat reporting presents a truncated sample of cyber conflict. Because they are situated predominantly at the low end, threats to civil society are underrepresented in this sample.

The information security industry is still relatively young, public threat reporting started only in 2010. The Gh0stNet report provided the template for in-depth reporting on targeted threats and the private sector picked up the practice, but changed the substance. Threat reporting is part of the threat intelligence and network defense business model in information security, and this subsector will be the focus of our analysis (from now simplified as 'threat intel').[6] The threat intel sector has been rapidly growing into a multibillion-dollar business, producing a regular stream of publicly available reporting on targeted digital threats. These reports not only form the largest body of available data on such threats, but they are also often the only publicly available source of data on cyber operations. For example, the 2016 intrusion into the DNC networks was revealed in reporting by private vendor CrowdStrike, and its reports remained the core source of public knowledge about this cyber operation for months after (CrowdStrike 2016; Sanger and Shane 2016). Hence, researchers in both the academic and policy worlds

[4] These are defined as "persistent attempts to compromise and infiltrate the networked devices and infrastructure of specific individuals, groups, organizations, and communities." (Crete-Nishihata et al. 2014, 5)

[5] These include Citizen Lab at the University of Toronto, Electronic Frontier Foundation, AccessNow, Human Rights Watch and Amnesty International. For an overview of their research, see: https://citizenlab.ca/tag/targeted-threats/, https://www.amnesty.ca/search/node/cybersecurity, https://www.accessnow.org/content-type/publications/, https://www.hrw.org/publications?keyword=cyber&date%5Bvalue%5D%5Byear%5D=, https://www.eff.org/search/site/civil%20society

[6] The information security sector provides three main types of services (vendor examples in brackets): targeted offensive capabilities and operations (e.g. Gamma Group, Hacking Team, NSO Group), mass data analytics and surveillance (e.g. Palantir), and threat intelligence and defense (FireEye, CrowdStrike, Kaspersky, etc.).

rely on commercial threat reports as an important, and often the main data source for threat assessments, theory development and policy-making on cyber conflict.

At face value, the emergence of public threat reporting is a good thing because it provides data on classified operations that would otherwise be inaccessible to the academic and policy communities. For reasons that lie beyond the scope of this paper, states have been unwilling or unable to provide cybersecurity services analogous to law enforcement or national defense. This vacuum has been filled by private actors, assuming a 'quasi-governmental' role by nature of the services they offer (Eichensehr 2017). Threat intel vendors offer three main services: public threat reporting, private threat reporting and customized protection and mitigation services.

Public threat reporting can be conceived as a public good. A pure public good is defined by two properties: it is non-exclusive (Samuelson 1954), and non-rivalrous (Ostrom and Ostrom 1977). In other words, no one can be excluded from the benefits of the good, while consumption of the good by one actor doesn't reduce the availability to others. Public threat reporting fulfils both these criteria: it is freely available online and reading a report does not reduce its readability for others.

Threat reporting is a public good that benefits all participants in the cybersecurity sector. Knowledge about the methods used by digital adversaries is an essential requirement to increase resilience of systems, detect and mitigate intrusions. Nascent community-driven initiatives among cybersecurity researchers to consolidate this knowledge from threat reporting attests to these potential benefits.[7] These initiatives reflect the value of this knowledge as a resource for both researchers and practitioners of cybersecurity, and they constitute efforts at self-governance of the shared resource of knowledge of targeted threats. Such shared resources are always subject to collective action problems, as highlighted in Hess and Ostrom's definition of "knowledge as a shared resource, a complex ecosystem that is a commons—a resource shared by a group of people that is subject to social dilemmas." (Hess and Ostrom 2007, 3). Public good provision by self-interested actors is a key collective action problem because it results in insufficient provision of the good (Olson 1971, 34–36). We argue that market failure results in insufficient reporting on civil society due to the underlying incentives driving threat intel reporting.

Threat intel vendors are not the only actors providing the good of threat information, as pointed out further above a group of independent research centers and nonprofits spearheaded by Citizen Lab has produced a body of rigorous reporting on targeted threats to civil society. Their resources are limited, however, and a far cry from those of leading threat intel vendors. Accordingly, as our dataset shows, the volume of commercial reporting dwarfs that of independent reporting at a ratio of about 9 to 1. In addition, some governments have commenced publishing reports on cyber operations. The authors are currently aware of such efforts only by the United States (US CERT n.d.), the United Kingdom (NCSC n.d.) and Germany (BfV n.d.). These nascent initiatives offer only a handful of reports, however, and, moreover, much of the weekly reporting offered by the United Kingdom's National Cyber Security Centre and the German quarterly 'Cyber Briefs' draw heavily from commercial reporting. In short, threat intel vendors are by far the largest producers of threat reports, and thus overwhelmingly shape the shared knowledge on targeted threats.

[7] There are two main consolidation projects: (1) APT Groups and Operations, a shared Google sheet consolidating naming schemes and operations of known threat actors; and (2) APTNotes, a repository consolidating threat reports by different companies.

Hence, it is important to examine why threat intel firms push out public reports. Self-interested actors typically fail to provide sufficient quantities of a public good because benefits are shared among the group while costs are borne by the individual actor(s) providing it alone (Olson 1971). Hence, public goods are often provided by governments, and a classic example is national defense (Murphy and Topel 2013). Yet threat reporting is a public good provided by private actors, which, as Olson has shown, only occurs if the individual benefits of provision outweigh the costs (Olson 1971, 34). Therefore, it is important to consider the individual and public benefits of reporting in more detail.

	Excludable	*Non-excludable*
Rival	**Private good** (food, clothes, laptops) - Protection Services	**Common good** (meadows, fish stocks)
Non-rival	**Club good** (cable television, Wi-Fi) - Private reporting	**Public good** (lighthouse, public parks, air) - Public Reporting

Threat intel firms are profit-seeking enterprises, and public reporting is one of three key products that comprise their business model. The other two products they offer are private reporting and customized protection and mitigation services, both of which are only available to paying customers (Guerrero-Saade 2015). Paul Rosenzweig has already pointed out the public goods characteristics of threat reporting versus customized protection services as an private good, which is both exclusive and rival since typically involves dedicated teams (Rosenzweig 2011). Rosenzweig's model still misses private reporting, however, which have now become a key part of the revenue stream. Private reports are offered to subscribers at a premium and contain more comprehensive information than public reports (Guerrero-Saade 2015, 2). These reports are available only to paying customers (exclusive), but consumption does not reduce the availability to others (non-rivalrous), hence they constitute club goods. How do these products relate?

Public reporting is foremost a marketing instrument, its purpose is increasing revenue from paid products offered by the vendor. According to Juan Andrés Guerrero-Saade, a researcher at the prominent threat intel provider Kaspersky, "the intended purpose is a PR-coup to both attract new customers for closed-release intelligence reports as well as garner brand recognition and industry respect for formidable findings." (Guerrero-Saade 2015, 4). The public good or reporting is thus tied to the private and club goods offered by the vendor, because the benefits of the latter outweigh the costs of producing the former—explaining provision of the public good (Demsetz 1970). Reporting is thus a means to increase profits, the individual benefit for the vendor. The diffuse benefits of contributing to public knowledge of threats are a byproduct of the pursuit of these individual benefits. There undoubtedly is a shared and genuine interests in furthering knowledge among security researchers, but as profit-seeking enterprises the vendors publishing reports depend on the revenue of the products this public good is tied to. Without these, the business could not exist. Threat reports inform potential customers about the problems they face, and the vendor offers (paid) solutions. These incentives lead to selection in bias in reporting.

As marketing instruments, threat reports need to maximize attention. Based on public statements and writing from industry insiders, informal conversations and one formal interview with a threat intelligence researcher at a prominent firm[8], we identify three key selection criteria.

First, a cyber operation exhibits some *unique characteristics*, typically in its techniques, tactics and procedures—the 'formidable findings' mentioned by Guerrero-Saade above. According to a threat intelligence researcher, to make it into a public report "it needs to be something unique, something that hasn't been reported before, for example a zero day, or some kind of unique tactic used" (Threat Intelligence Researcher 2018).

Second, it has a *high-profile victim*. Since threat reports are intended to sell protection products, the more significant the threat reported, and the more high-profile the targeted actor, the better. A 'high profile' is not a scientific measure, but a key indicator of public attention. Accordingly, the Cambridge English Dictionary defines it as "attracting a lot of attention and interest from the public and newspapers, television etc." (Cambridge English Dictionary n.d.). Hence, general public and in particular media attention to a given actor is a key indicator of a high-profile actor as perceived by the target audience.

However, from the perspective of threat intel firms, the highest profile actors are those with the greatest revenue potential. If threat reporting is intended to sell private reports and protection services, a rational profit-seeking actor can be expected to prioritize reporting on threats to the most lucrative targets.

Third, a *high-profile threat actor* is behind the campaign. We identify three key measures of high-profile actors: 1) attribution to strategic competitors of nation-state(s) in which the target audience resides, 2) previous coverage in general news outlets, 3) attribution to previous campaigns perceived as a national or international threat. The more of these apply, and the higher the score, the higher profile the threat actor.

What are the implications of reporting based on these selection criteria? Threats to civil society are underreported because they tend to score low on at least two of these variables.

First, operations targeting civil society tend to be technically unsophisticated, relying on generic tools and known vulnerabilities. The Citizen Lab's Communities @ Risk study, an in-depth investigation on the risks faced by civil society from targeted digital threats, showed that civil society actors typically lack resources and expertise necessary for strong cybersecurity, hence cyber operations targeting civil society can often succeed with generic and unsophisticated TTP (Crete-Nishihata et al. 2014, 3).

Second, their lack of resources render civil society actors unattractive potential clients since they promise little revenue. In contrast, they are typically cash-strapped (Crete-Nishihata et al. 2014, 2; CLTC 2018) and thus unable to afford the premium products offered by threat intelligence vendors. Some civil society actors receive a high level of public and media attention as a result of their activism, however,

[8] Structured interviews were a core part of our original research design, but despite multiple attempts at recruiting interview subjects across different levels of management at multiple firms, only one researcher was willing to be interviewed anonymously. As in any competitive industry, these firms carefully guard the details of their internal processes and organizational structures—which is understandable, but adds to the opaqueness of threat reporting as a data source.

which may promise enough potential attention to threat reporting on this actor to offset above disincentives.

Third, and in contrast, campaigns targeting civil society are regularly pursued by the same actors as those targeting governments, militaries and large corporations. The Citizen Lab's Communities @ Risk study (2014) conducted an in-depth analysis of the threat landscape faced by ten different civil society organizations over the course of four years, during which it identified ten distinct targeted threat campaigns. Importantly, half of these campaigns had "connections to threat actors, previously reported to have targeted government and private industries." (Crete-Nishihata et al. 2014, 2). Accordingly, based on these criteria, those campaigns pursued by high-profile threat actors have the highest probability of being reported on by private threat intelligence firms. However, if the trends identified in the Communities @ Risk study are representative, only a minority of cyber operations targeting civil society can be expected to fulfil these conditions—not all of the actors also targeting governments and private firms will fulfil the criteria of a high-profile actor as outlined above.

Based on these criteria, threat intel vendors can thus be expected to underreport threats to civil society compared to other threat types promising more attention and thus potential revenue. Consequently, the public good of threat information is insufficiently provisioned, constituting a classic market failure, where "the failure of a more or less idealized system of price-market institutions to sustain 'desirable' activities or to stop 'undesirable' activities" (Bator 1958, 351). This market failure has two key consequences.

First, it creates a truncated sample of cyber conflict that distorts perceptions of the priorities and methods of capable threat actors. Jack Levy has shown that the most important forms of misperception in explaining patterns of conflict are misperceptions of the capabilities and intentions of adversaries (Levy 1983). These misperceptions undermine both academic theory development on cyber conflict, and formulation of effective policy.

Second, it undermines the defenses of civil society organizations targeted by these actors. Civil society is notoriously short on resources to begin with, and many organizations lack even basic cybersecurity infrastructure—rendering them easy targets. Meanwhile, the consequences they face are potentially the most severe. The Citizen Lab study also tracked the impact of the identified operations on the targeted organizations, including "arrest and detention from groups in the study (and through other Citizen Lab projects) that appeared to be linked to electronic surveillance." (Crete-Nishihata et al. 2014, 117). Since they are already at a disadvantage, civil society organizations are most in need of accurate information on the threat landscape they face. Underrepresentation in threat reporting thus further undermines their already precarious security situation through underprovision of knowledge on threats. Finally, it likely results in a lack of prioritization of the problem among policy-makers who draw on threat reporting in formulating threat assessments and allocating resources.

3. HYPOTHESES AND RESEARCH DESIGN

This section lays out the research design. First, it outlines the mixed methods approach and specifies the division of labor between quantitative and qualitative analysis formulates a set of three hypotheses. Third, it introduces an original dataset of threat reporting and specifies the selection criteria.

We hypothesize that business incentives result in systematic selection bias in threat reporting. Because the selection and decision-making processes of threat reporting happen behind close doors, however, there is insufficient data to prove this causal relationship. Our original research design included structured interviews with researchers and management at a sample of threat intelligence firms to trace this decision-making process. We designed detailed questionnaires and robust anonymization procedures approved by our University's Research Ethics Board. However, multiple attempts at recruiting interview subjects remained fruitless and ultimately only one researcher was willing to be interviewed. This lack of responses itself constitutes a potential data point on the low prioritization of the issue. Regardless of the reasons, however, there is insufficient data to trace causal mechanisms in the decision-making process. Therefore, we examine whether the outcome of this process conforms to what our theory predicts.

For this analysis, we employ a mixed method research design of nested analysis. As laid out by Lieberman, this approach "combines the statistical analysis of a large sample of cases with the in-depth investigation of one or more of the cases contained within the large sample." (Lieberman 2005, 434–35). The aim is to make inferences about bias in threat reporting, and our research strategy involves two key steps. First, we formulate set of hypotheses. Second, we test them against evidence drawn from a quantitative analysis of all available threat reporting. Third, we proceed with a qualitative analysis of three case studies to verify whether our predictions are congruent with reporting patterns on these cases.

PART 1: QUANTITATIVE ANALYSIS

The quantitative part tests our theory against findings from an analysis of all available public threat reporting. This approach leverages the strength of large-N analysis in testing multiple hypotheses and testing robustness of our theory by identifying broad trends in a large body of data (Lieberman 2005, 436). Our main research question is straightforward and empirical: what threats are being reported by commercial vendors? Because we are comparing this sample to reporting by independent research centers, our auxiliary research question is how do reporting patterns differ among commercial and independent reporting? For this purpose, we contribute an original dataset built from content analysis of all available public threat reporting. Content analysis is a useful tool because it allows quantitative analysis of unstructured data to identify trends and potential biases (Mukherjee 2018, 29–30). Our dataset comprises 622 threat reports, 551 reports by threat intel firms and 71 reports by independent research centers. These reports were collected from two community-run websites as well as individual vendor websites. Selection criteria were straightforward: to be included, reports had to discuss (1) a targeted digital threat, and (2) be available publicly. Following established practice, we specified a set of categories and coded all reports across the same set of categories. We then use descriptive statistics to analyze three

of these variables: (1) whether civil society is mentioned[9], (2) the geographical location of the targeted organization, and (3) attribution of a campaign. This analysis aims to verify the following three hypotheses.

The preceding section has specified three sources of bias, based on which we formulate the following three empirical hypotheses:

H1: Threats to civil society are underreported.

This hypothesis needs little explanation: if business incentives shape reporting the way we hypothesize, only a low proportion of reporting will discuss threats to civil society. It would be confirmed if only a small proportion of reports discuss threats to civil society. Conversely, if a majority of reports discuss civil society threats, it is clearly false.

H2: Reporting is geographically skewed towards the Global North.

This hypothesis captures selection bias towards the most lucrative sectors and regions. We use geographical distribution as an indicator because it is the most uniformly reported metric. Many threat reports do not specify business sectors, and moreover, each firm tends to use different categories. Both commercial and independent reporting do specify the geographical location of victims, however, allowing a direct comparison. Since the Global North is more affluent, reporting can be expected to prioritize threats in this most lucrative region. Hence, if a majority of commercial reporting discusses threats targeting the Global North, it supports our hypothesis. If reporting is evenly distributed between the global North and South, or if a majority of reporting focuses on the global South, the hypothesis is invalidated. We then compare geographical distribution between commercial and independent reporting to identify diverging patterns. To ensure a fair comparison, we include only the subsample of commercial reporting that does discuss threats to civil society. If commercial reporting exhibits the expected geographical skew, and independent reporting doesn't, it further corroborates our hypothesis. If both commercial and independent reports exhibit identical or similar geographical bias, the hypothesis is invalidated.

In addition, we use a second source of data for comparison: AccessNow -- an international Internet rights advocacy group -- has generously provided us with anonymized data from their helpline. This helpline provides a resource for civil society organizations to get help if they are targeted by digital threats, and the helpline data thus offers a window into cyber conflict from the perspective of civil society itself. This data counts the incidents of malware, system compromise and phishing as self-reported by civil society organizations, categorized by country. It is globally accessible and in contrast to threat reporting—both commercial and independent—it is not shaped by intervening selection criteria. Therefore, this data provides the most neutral representation of targeted threats to civil society available, but it comes with one key caveat: the lack of relevant IT expertise among civil society that renders them vulnerable is also likely to result in mis-reporting of incidences. However, the classification of incidents is done by experienced staff at AccessNow based on a uniform categorization scheme. Combined, these measures

[9] Based on our definition (see further above), these include political activists, dissidents, journalists, non-government organizations as well as academics. Since the distinguishing characteristic of civil society is non-violent political engagement, we only include journalists and academics if this condition applies to them. For example, an academic calling out government corruption is included, while a physicist working on quantum theory is not.

ensure a reasonable degree of validity across these results. If geographical distribution between commercial reporting and self-reporting differs along our hypothesized North-South divide, this finding provides further strong support for our hypothesis. If, however, the geographical distribution of self-reporting and commercial reporting turn out to be congruent, with independent reporting showing a different distribution, the hypothesis is invalidated—and such findings indicate systematic bias in independent reporting instead.

H3: Reporting is skewed towards operations attributed to the target audience's main adversaries.

Our theory predicts that firms are more likely to report on high-profile threat actors, and in our quantitative analysis we measure this bias by tracking attribution patterns.[10] We chose this indicator to measure the profile of a threat actor because threat reporting does not provide data on the values for the other hypothesized determinants of profile (media attention, previous campaigns). However, our qualitative analysis tracks these additional indicators of actor profile as well.

If a majority of reporting discusses threats by the main strategic competitors of the target, the hypothesis is verified. If not, and reporting is evenly spread across different threat actors, or otherwise distributed, the hypothesis is invalidated. Again, we compare attribution patterns across commercial and independent reporting to spot divergences, which would further corroborate positive findings.

The quantitative analysis constitutes a hoop-test of our theory (Van Evera 1997, 31), meaning negative findings eliminate the theory yet positive findings do not invalidate rival explanations. There are several plausible alternate explanations. First, reported threat patterns may accurately reflect reality, and correlations may be purely coincidental. Second, divergences between commercial and independent reporting may reflect independent reporting's own bias(es), or other differences among these two slices of data. Third, the lack of reporting on specific regions is the result of lack of knowledge rather than a lack of prioritization. Threat intelligence firms often rely on telemetry data from their own security products to identify threats, and they may simply not have any visibility into some global regions or sectors. This constraint has been pointed out by several insiders, yet a lack of telemetry data itself likely reflects underlying business interests: telemetry data is drawn from existing customers, and if there are no customers in a region, that region has not been developed. In short, this analysis allows us to verify whether reporting patterns correspond to those one would expect if reporting was indeed biased across the hypothesized selection criteria.

Content analysis does not allow confirming causal relationships among variables, however (Mukherjee 2018, 36). Therefore, positive results of this analysis do not confirm whether there is systematic bias against threats civil society in commercial threat reporting.

[10] We focus on this measure because data on news reporting and previous campaigns by the same actor is not available from our dataset, and thus not feasible to include in the quantitative analysis. We do, however, include these additional two measures in our case studies.

PART II: QUALITATIVE ANALYSIS

Hence, in line with Lieberman's model, we shift the level of analysis to qualitative analysis of three cases from within the same dataset (Lieberman 2005, 440). These case studies serve as a plausibility probe for our argument. We use structured focused comparison and congruence testing to verify if threat intel reporting patterns on specific threat actors known to target civil society are consistent with our expectations.

The key problem is incomplete data on the decision-making process. Hence, we cannot rely on the strength of within-case analysis to uncover and trace causal processes (=process-tracing), because this requires complete, or at least adequate data (Waldner 2015; Bennett and Checkel 2015; George and Bennett 2005, chap. 1,7,9). Therefore, our qualitative strategy parallels that of the quantitative analysis in comparing observed patterns in testing the congruence between observed reporting trends and expected trends based on our predictions.

We are employing congruence testing and controlled comparison as laid out by George and Bennett to "match the predictions and expectations of the theory with the outcomes of the cases to see if they are consistent" (George and Bennett 2005, 227). To do so, we test the validity of the predictions that follow from our theory of selection bias in three case studies of campaigns targeting civil society taken from our sample of independent reporting.

To test the strength of our theory, we focus on most- and least-likely cases. To recall, most-likely cases are those where "independent variables posited by a theory are at values that strongly posit an outcome or posit an extreme outcome" while in a "least-likely case, the independent variables in a theory are at values that only weakly predict an outcome or predict a low-magnitude outcome." (George and Bennett 2005, 147). If our predictions fail in the most-likely cases, strong doubt is cast on our theory, while successful prediction in a least-likely case strongly supports the theory (Ibid.)

Our theory identifies three sources of selection bias that determine publication of reporting and prioritization of findings within published reporting: (1) unique TTP, (2) a high-profile target and (3) a high-profile threat actor. The dependent variable is whether a campaign is reported, and if so, how it is prioritized (both within individual reports and across the entire body of reporting). Our main assumption is that selection bias leads to systematic underreporting of threats to civil society, hence the null hypothesis states that no such bias exists. Our case studies thus test whether reporting patterns are congruent with expectations based on the selection criteria we have identified.

We select three cases of a campaign targeting civil society from Citizen Lab reporting to track and compare reporting patterns by threat intel firms across the hypothesized three variables of selection bias. We select two most-likely cases at extreme values on the independent variables, and one least-likely case with intermediate values. The use of only Citizen Lab reports minimizes variation between these cases since the lab uses a consistent methodological approach as well as classification scheme for sophistication.

4. QUANTITATIVE ANALYSIS

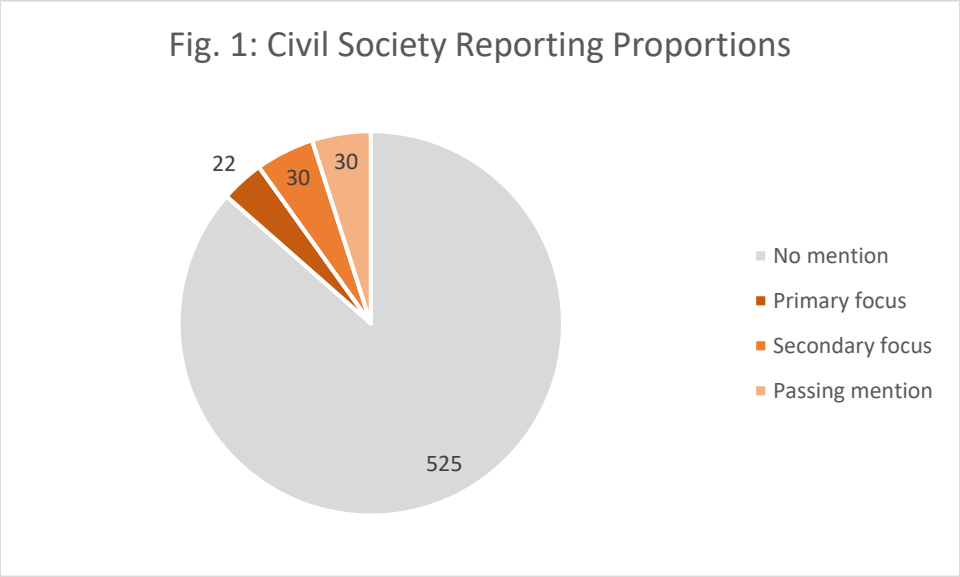

H1 predicts a low proportion of reporting prioritizing threats to civil society, and is verified by the findings. As shown in Figure 1, only a small minority 82 out of the 622 commercial reports analyzed do discuss a targeted threat civil society, a proportion of 14%. A deeper look at prioritization of the issue within this subset of commercial reporting revealed that only 22 out of these reports, or 3.4% of total reporting, place their primary focus on civil society. Meanwhile, 30 reports (5%) place a secondary focus on civil society targeting, with limited analysis, and 30 reports (5%) mention civil society in passing (including mentions of past targeting by the same threat actor). Findings thus provide strong support for H1.

H2 predicts a geographical bias towards the Global North, which is supported by our findings on the absolute distribution of reporting, and further corroborated by the relative distribution as compared to independent reporting.

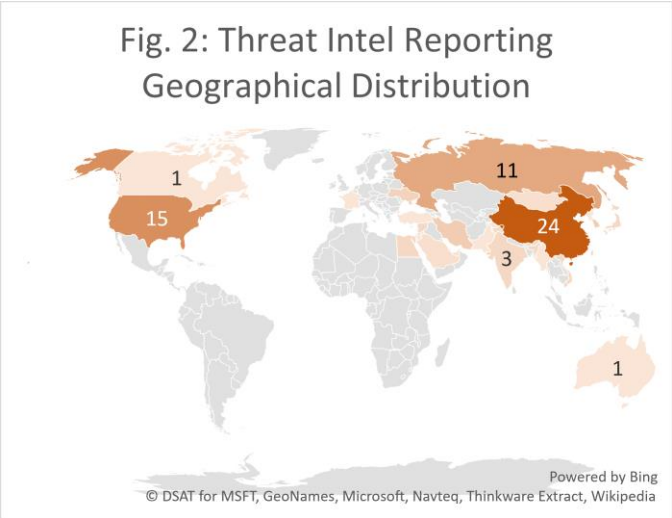

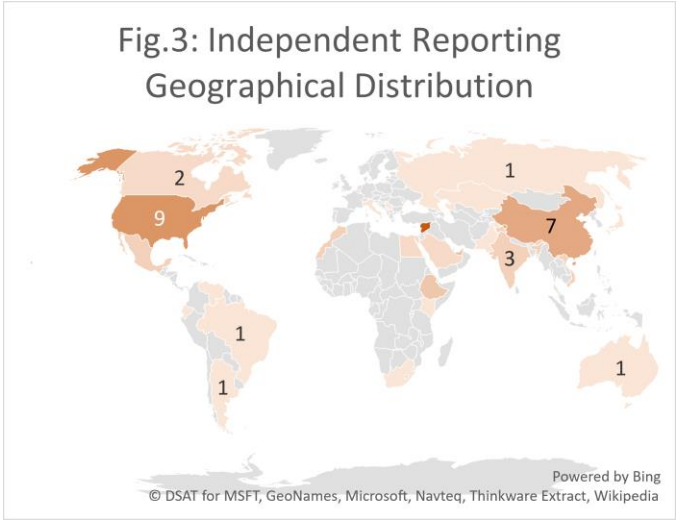

Figure 2 projects the geographical location of all threat intel reporting on targeted digital threats to civil society. Apart from three reported operations targeting civil society in Egypt, the Global South remains a blank spot on this map. In contrast, independent reporting reveals a host of activity in the Global South. There are six instances of civil society being targeted in four different countries in South and Central America (Mexico, Ecuador, Brazil, Paraguay), as well as twelve reported instances of targeted threats to civil society in five different countries in Africa (Morocco, Egypt, Ethiopia, Kenia, South Africa). The absolute distribution of threat intel reporting thus supports H2, and the relative distribution of reporting compared to independent reporting strongly supports the hypothesized geographical bias in commercial reporting. This difference becomes even more pronounced when comparing threat intel reporting to a second slice of data: self-reporting of incidents by civil society organizations themselves.

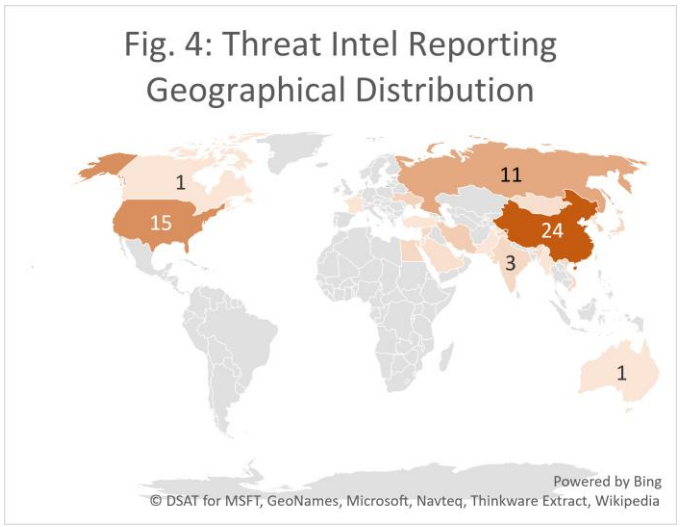

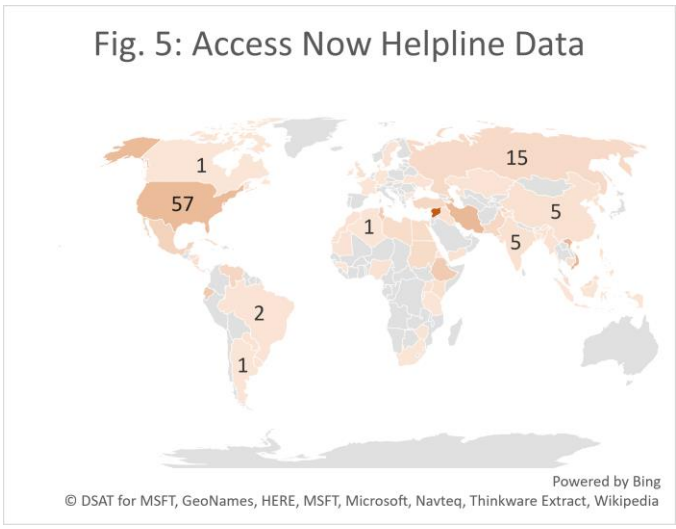

AccessNow shared a sample of their helpline data with us, capturing the number of reported incidents of account compromises, system compromises, malware and phishing as reported by member of civil society themselves. Although this sample only covers 17 months (January 2016 – May 2018), it shows a much wider and more even distribution. In particular, it shows a plethora of activity in the Global South in general, and in Africa in particular, that is missing from threat reporting. Granted, these are very different sources of data and thus not easily comparable—nonetheless, the divergence reinforces the trend identified in the comparison above, thus providing further support to our hypothesis.

H3 predicts reporting will be skewed towards operations attributed to perceived adversaries of the target audience. Our findings are strikingly clear.

Fig. 6: Threat Intel Attribution

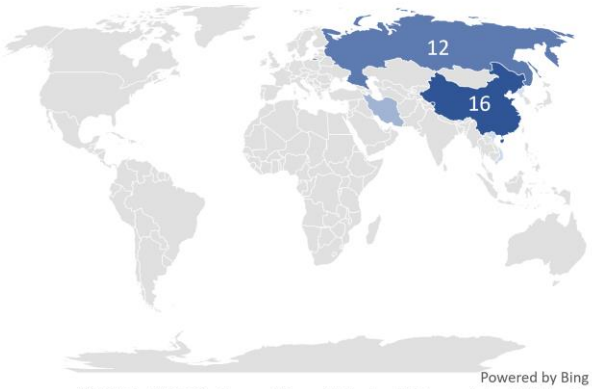

Fig. 7: Independent Attribution

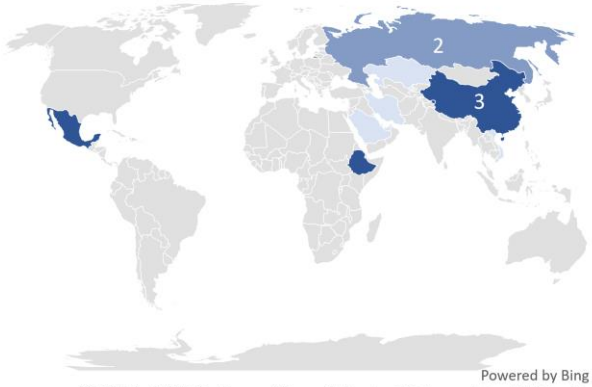

As depicted in figure 6, the vast majority of threats to civil society reported by threat intel firms are attributed to the main adversaries of the West, and particularly the United States: China, Russia and Iran. Apart from these, one campaign is attributed to North Korea and two to Vietnam. This pattern is strikingly congruent with what we would expect to see if reporting was indeed biased towards campaigns attributed to those countries perceived as most threatening by a North American audience, as determined in a recent survey (YouGov 2017). Russia, China and Iran are counted among the world's leading 'cyber powers', hence it is conceivable that they are simply the main perpetrators of threats to civil society. Independent reporting (figure 7) reveals six campaigns by these 'big three', but it also documents the use of targeted digital threats by a range of other governments that are missing from commercial reporting: Kazakhstan (1), Ethiopia (3), Kuwait (1), Saudi Arabia (1), United Arab Emirates (1) and Bahrain (1), and Mexico (3). Not only does independent reporting show a more evenly distributed attribution pattern, but the total number of operations by 'other' states is actually greater than those attributed to the 'big three'.[11] Combined these findings strongly support the hypothesized bias towards adversarial actors.

[11] It is interesting to note that the attribution rate of independent reporting is only about half that of commercial reporting: 25% vs. 51%. This finding likely reflects more stringent rules for attribution among academic and non-profit researchers.

The findings from our quantitative analysis thus provide strong support for hypothesis 1, 2 and 3. Our theory of bias in commercial reporting passes the hoop test.

5. CASE STUDIES

This section develops a plausibility probe of our theory by testing its assumptions against evidence from three case studies of cyber operations targeting civil society. We predict that publication of threat intel reports is subject to three selection criteria: unique TTP, high-profile victim, and high-profile perpetrator.

CASE 1: TAINTED LEAKS

Our first case is the Tainted Leaks operation, a campaign that targeted a journalist and involved the theft of personal emails and subsequent leaking of this data (Hulcoop et al. 2017). This operation fulfils two of the selection criteria for publication of threat reporting, and is linked to a larger campaign that fulfils all three criteria.

First, this operation involved a new and unique method. As the Citizen Lab investigation discovered, the data was obtained through 'credential phishing', a widely used method that uses a deceptive website made to look like their email login page to obtain their user name and password. While the method of obtaining data was thus fairly standard, when this data was subsequently leaked it had been carefully manipulated to include disinformation designed to discredit critics of the Russian regime. There were previous suspicions of this method of 'tainting' leaks of stolen data, but the case reported by Citizen Lab was the first instance involving concrete evidence.

Second, it involves high-profile targets. The operation targets an individual journalist—prominent Kremlin critic David Satter who has received some media attention, but a low-profile actor from the perspective of threat intel firms' commercial incentives. The subsequent analysis revealed, however, that this operation was part of a larger scale phishing operation that included high-profile targets such as a "former Russian Prime Minister, members of cabinets from Europe and Eurasia, ambassadors, high ranking military officers, CEOs of energy companies" (Hulcoop et al. 2017).

Third, the operation was pursued by a high-profile actor. The Citizen Lab report did not conclusively attribute this campaign, but noted that circumstantial evidence pointed to infamous APT 28 tied to Russian military intelligence service GRU. The day after publication of the Citizen Lab report, the magazine Forbes reported it had, "obtained evidence that indicated the group was responsible" (Fox-Brewster 2017). This degree of media attention and the attribution of several previous high-profile campaigns to APT 28—including the intrusion of the DNC in 2016—attests to its high-profile as a threat actor. Since this campaign includes a unique method by a high-profile actor, and is part of a larger operation targeting high-profile actors, it is a most-likely case where our theory strongly predicts reporting by threat intel firms.

How does this match up to reality? APT28 is attributed to Russia, one of the Wests main strategic competitors, and has received extensive media coverage—where it is consistently ranked among the most

dangerous threat actors due to its skill and apparent recklessness (Hackett 2017; Burgess 2017; NCSC 2018; Bartlett 2018). There are few, if any, higher-profile threat actors and therefore it is of no surprise that a significant proportion of threat reports discuss operations by this actor: 53 out of the 606 reports in our dataset (8.7%) discuss campaigns by this actor. Considering there are over thirty different known threat actors, this proportion is congruent with our predictions of reporting proportions being skewed towards high-profile threat actors.

Findings concerning targeting proportions are similar. Unfortunately, only a minority of threat reports include targeting proportions (known as 'verticals' in the industry) so it is difficult to make a general assessment. However, the Citizen Lab investigation of the wider spear phishing campaign to which the Tainted Leaks operation was part of, did include such proportions[12]: behind governments, comprising 24% of targets, civil society was the second-largest target group at 21%--even before the private sector (Hulcoop et al. 2017). These proportions suggest civil society is a priority target for APT28, and evidence from one of the few threat intel reports that mention such verticals confirms this assumption. SecureWorks' investigation of a previous spear phishing using similar methods to those reported by Citizen Lab found that while military targets were the largest group at 41%, the second largest target group were civil society actors at 36% (including NGOs, activists, and journalists) (SecureWorks 2016b). Meanwhile, a 2015 TrendMicro report found civil society to be the main target of domestic operations by APT28 (TrendMicro 2015b). Finally, the first dedicated threat report on APT28, published by FireEye in 2014, explicitly highlighted the actor's targeting of journalists as a means to "monitor public opinion, identify dissidents, spread disinformation or facilitate further targeting." (FireEye 2014). This evidence suggests civil society is a priority target of APT28, hence threat reporting free of selection bias should exhibit a corresponding prioritization of this targeting category.

Out of these 54 reports on APT28, however, only 15 mention targeting of civil society (28%) overall, and only two out of these prioritize a threat to civil society. Hence, the great majority of reporting (39 reports, 72% of all reporting on APT28) entirely omits threats to civil society. If one were to build an analysis of APT28 activity entirely based on commercial reporting, targeting of civil society would seem a low priority and more occasional event of this actor.

Meanwhile, congruent with our predictions, reporting prioritizes campaigns with high-profile targets. Most reports characterize APT28 as a highly sophisticated espionage actor that targets government and private sector entities to steal sensitive data. For example, the group is described by CrowdStrike as a "Russian-based threat actor, which has been active since mid-2000s, and has been responsible for targeted intrusion campaigns against the Aerospace, Defense, Energy, Government and Media sectors." (CrowdStrike 2016). Symantec summarizes activity by APT28 as follows: "The organizations targeted by APT28 during 2017 and 2018 include: a well-known international organization, military targets in Europe, Governments in Europe, a government of a South American country, an embassy belonging to an Eastern European country" (Symantec 2018). Notably, this periodization includes the Tainted Leaks campaign, yet both the *Tainted Leaks* operation and the linked spear phishing campaign

[12] Citizen Lab researchers identified these targets by investigating links created by the operators using the shortening service Tinyurl.cc. Because these links included a predictable feature, the researchers were able to trace the sequence of links created by the shortening service, which revealed their origins in the same spear-phishing campaign. This investigation revealed over 200 additional targets besides David Satter and showed a strong, though not conclusive, linkage to APT28 (Hulcoop et al. 2017).

widely targeting civil society is entirely bracketed from Symantec's summary of APT28 activity over the same timeframe. In short, the prioritization of government and military targets in the majority of reporting is congruent with our predicted selection bias in favor of high-profile targets. Moreover, the apparent omission of civil society targeting from the majority of reporting further supports this assumed prioritization. Recall that the pattern of targeting of dissidents was mentioned in, and hence known since, the very first report on APT28, while the available data on vertical targeting patterns suggest civil society is among the group's main targets. It would thus be highly surprising if this pattern is entirely absent from a majority of its reported activity.

Moreover, prioritization of high-profile targets is evident not only in the reports that do not mention civil society, but also in those that do. For example, a 2017 survey of APT28 activity by FireEye only mentions one civil society target in the period from 2014-2017: the widely covered (in Western media) dissident band PussyRiot (FireEye 2017, 4). This data point is taken from a TrendMicro report published in 2015, which also leads with the targeting of Pussy Riot, but actually reveals the extent of domestic spying by APT28 and its broad targeting of civil society (TrendMicro 2015b). Yet even this report is scant on the details of this operation, prioritizing high-profile victims in its analysis: "to illustrate one of the credential phishing attacks Pawn Storm sends to its targets, we will focus on a particular attack on high-profile Yahoo users" (TrendMicro 2015b). These findings are congruent with our expectations, supporting the hypothesized selection bias in favor of high-profile victims.

Finally, only a small proportion of reporting mentions the credential phishing operations preceding, and likely linked to, the Tainted Leaks campaign. This finding is congruent with the expected selection bias in favor of highly unique TTP. Most reports on APT28 focus on sophisticated methods and especially malware, while the relatively unsophisticated method of simply fooling users to reveal their access credentials is not as widely covered. The general technique of credential phishing used in the Tainted Leaks campaign is mentioned in nine reports, seven of which do discuss civil society. Two out of the latter discuss the specific technique of credential phishing used in the Tainted Leaks operation, which was also employed in the targeting of Hillary Clinton's campaign (SecureWorks 2016b, 2016a). It is worth noting that the first of these reports (on Hillary Clinton) was published a day after CrowdStrike revealed the intrusion into the DNC, which provided a high-profile target and window of media attention. The timing of the report thus provides additional support for the hypothesized prioritization of high-profile targets.

None of the above reports, however, mention the technique of 'tainting leaks' and its use to spread disinformation. Only one report briefly passes on the 'alleged' manipulation of data obtained through credential theft that is later published as a 'leak', noting how "we have credible information that CyberBerkut has published information which was stolen during Pawn Storm's [APT28] credential phishing campaigns. Prior to leaking the information, parts of the documents and emails were allegedly altered. " (TrendMicro 2017, 6). No further information is provided, however.

To sum up, reporting patterns on campaigns by APT28 are largely congruent with the predictions of our theory. Based on commercial reporting, APT28 is a highly sophisticated actor that focuses on espionage and primarily threatens governments entities and large private sector entities. There are references to the targeting of civil society, but only among a minority of the reports. While the targeting of journalists and the objective of spreading disinformation was indicated from the beginning of reporting on APT28, hence we would expect operations such as Tainted Leaks to be reported prominently. Yet the

body of threat intel reporting mostly brackets targeting of civil society, and does not discuss the use of targeted 'leaking' apart from a passing mention. This pattern is congruent with the hypothesized selection bias in favor of higher-profile targets.

CASE 2: SPYING ON A BUDGET

Our second case study lies at the opposite end of the spectrum concerning selection criteria. The *Spying on a Budget* report by Citizen Lab identified a spear phishing campaign using cheap and unsophisticated means to target Tibetan activists by a previously unknown threat actor using relatively sloppy methods (Crete-Nishihata et al. 2018). Because it scores low across all three variables identified as selection criteria, this operation also constitutes a most-likely case as our theory strongly predicts it will be largely or entirely missing from commercial threat reporting.

The Citizen Lab report tracked a phishing operation active from around January 2016 until July 2017 using a range of tactics to obtain the email credentials of members from the Tibetan activist community and potentially other social movements in China (Crete-Nishihata et al. 2018). These individual activists and civil society groups make for low-profile targets promising little business opportunities and media attention. However, while the decoy documents employed indicate the Tibetan community as the main target, additional documents used indicate the same campaign also targeted government agencies in South and Southeast Asia, namely the Sri Lankan Ministry of Defense and the Thai Ministry of Justice. Hence, there are some higher-profile victims that would increase the likelihood of threat reporting. Methods used are mostly generic and unsophisticated, hence the Citizen Lab estimate that the entire operation could have been done within a budget of only US$1000, and thus do not fulfill the uniqueness criterion. Finally, the threat actor involved is unknown, exhibits "only basic technical skills" and is sloppy, leading Citizen Lab researchers to conclude it is likely a 'low-level contractor'. In short, it is the opposite of a high-profile threat actor, and thus fails to fulfil the third selection criterion for publication in threat reporting. Hence, our theory would predict only scattered, passing mentions, or no mentions at all, of this campaign and/or actor in commercial threat reporting.

Results are congruent with our expectations. There are no preceding reports on this threat actor and/or the phishing campaign involved. Moreover, there are also no follow-up reports by firms in response to the publication of the Citizen Lab report. These findings are in line with our expectations. Half a year later, however, RecordedFuture published a thoroughly researched report on a campaign targeting the Tibetan community and attributes it to the same threat actor (Recorded Future 2018). This level of attention to the same low-level actor would challenge our theory, were it not for the fact that the overall "increased level of sophistication for the attacker" (Recorded Future 2018). Evidently, the threat actor had passed the necessary threshold in uniqueness and sophistication to be included in a threat report. The overall lack of commercial reporting on this threat actor, as well as the inclusion only following an increase in sophistication, are congruent with our theory's predictions.

CASE 3: FAMILIAR FEELING

Our final case also targets the Tibetan community, but it is pursued by a higher-profile actor, targets some additional higher-profile victims, and employs somewhat unique TTP. It is a least-likely case where our theory does not provide strong predictions either way since it scores intermediate values across the three selection criteria.

This campaign involves a known threat actor, called either Tropic Trooper or KeyBoy (Alexander et al. 2018), which was behind a previous campaign targeting the Tibetan community reported on by Citizen Lab in 2016 (Hulcoop et al. 2016). In contrast to APT28, this actor has not been conclusively attributed to a specific government, but is suspected to be associated with China, nor has it received attention in general news media. However, previous activity by this actor has been covered in dedicated information security media (Networks Asia Staff 2015; MacGregor 2015; Muncaster 2017). The campaign reported by Citizen Lab—dubbed *Resurfaced*—employed a new version of a previously known set of malware that exploited known vulnerabilities. Finally, while this campaign focuses on low-profile victims (an unnamed Tibetan NGO is mentioned), the Citizen Lab links it to previous campaigns by the same threat actor targeting government and large private sector actors in East and Southeast Asia. Hence, it involves both low and high(er)-profile targets.

In short, this campaign is right in the middle of the spectrum of selection criteria, with a medium level of sophistication, a medium-profile threat actor and a low-profile target but previous campaigns against higher profile targets. Because this case falls squarely 'in the middle', our theory only weakly predicts reporting outcomes: it is possible that some commercial reporting covers this campaign since it may cross the necessary thresholds for publication, but the opposite outcome is similarly likely. However, since this campaign includes both lower and higher-profile victims, any evidence for a prioritization of high-profile targets—and in particular omissions of lower profile targets—provides strong support for our theory.

Threat reporting patterns on this case provide surprisingly clear support for our predictions. Overall, five commercial reports discuss campaigns by Tropic Trooper / Keyboy. None of them mention civil society as a target, but exclusively focus on high-profile government and corporate targets. Rapid7 first reported on Keyboy in 2013, without conclusively identifying targets, but hypothesizing targeting of "either someone in the telecommunications industry or a representative of the local government" (Rapid7 2013). TrendMicro reported on a campaign targeting "major government sectors and corporations in both Taiwan and the Philippines" (TrendMicro 2015a). PwC reported on a campaign in 2017, quoting previous Citizen Lab reporting on the actor but does not mention civil society. Instead, while the PwC report notes the lack of "clear visibility" into targeting, it nonetheless highlights that it "does appear that this latest campaign targets at least some Western organisations, likely for corporate espionage purposes" (PwC 2017)—providing support for the hypothesized prioritization of victims in the Global North. Based on commercial reporting, one would thus conclude that this is an actor focusing exclusively on international espionage.

There are three plausible explanations for this exclusive focus on high-profile government and corporate actors: (1) the campaigns reported by Citizen Lab are the only ones targeting civil society, (2) commercial researchers were unaware of the targeting of civil society, or (3), commercial researchers were aware but did not include it in reporting due to the prioritization of high-profile victims. In the former two cases, we would not expect any evidence pointing towards civil society targeting commercial threat reporting. However, there are two key pieces of such evidence—and congruent with the third explanation.

First, the reference to Citizen Lab reporting on targeting of civil society in the PwC report shows the researchers were familiar with targeting of civil society. Second, the latest report on TropicTrooper / KeyBoy characterizes it as an actor "focusing on … government, healthcare, transportation, and high-tech industries" and reports on the evolution of its tradecraft (TrendMicro 2018). However, its 'indicators of compromise' section also includes the domain "tibetnews[.]today", which points directly towards the targeting of Tibetan community by the same actor (and was later shown to be relevant in the Resurfaced campaign by Citizen Lab). Yet the TrendMicro report does not address this overlap at all. To be sure, none of these findings provide conclusive 'smoking gun' evidence confirming hypothesized selection bias in favor of high-profile targets in commercial reporting. Both the overall reporting pattern and these anecdotal pieces of evidence pointing to the omission of civil society targeting are closely congruent with our predictions and thus strongly support our theory.

To summarize, reporting patterns in the most-likely cases do not invalidate the theory, while the least-likely case strongly supports it. Our hypotheses pass the plausibility probe.

6. DISCUSSION

Both the quantitative and qualitative analysis confirm our expectations. Reporting patterns overall, as well as on the specific cases we have selected, are congruent with predictions based on our theory of threefold selection bias. While the limitations of available data prevent a causal analysis, the unambiguousness of our findings—in particular the least-likely case—strongly indicate that reporting prioritizes sophisticated and unique campaigns by high-profile threat actors against high-profile targets. Conversely, the cybersecurity marketplace fails to provide sufficient reporting at the low-end of cyber conflict. This lack of attention is not a moral failure—each firm behaves exactly as it needs to in order to survive—but rather a classic collective action problem. The resulting market failure has three key implications.

First, it creates a distorted picture of cyber conflict as researchers base their analyses on a skewed sample of cases. There is growing evidence that cyber conflict thrives especially at the low end of the conflict spectrum (Lindsay 2017; Valeriano and Maness 2015), and in this conflict civil society is right at the frontlines (Deibert 2015). Yet, our findings suggest this portion of conflict is systematically sidelined in threat reporting.

Second, this distorted picture poses a risk for democracy by systematically underrepresenting the threats to civil society organizations that are vital for the functioning of democracy. This underrepresentation leads to a lack of awareness that can be expected to lead to (1) lower than necessary prioritization and resource-allocation in policy, as well as (2) insufficient preparedness by both policy-makers and civil society itself when it comes to detecting and mitigating these threats. The widespread surprise at the methods used in the 2016 election meddling campaigns attest to this lack of awareness.

Third, the solution cannot come from the market alone, yet states are simultaneously the key threats to civil society. Several threat intelligence firms are offering pro-bono services to civil society, which is a move to be welcomed, but these individual measures cannot override the market logic that dictates the prioritization of the sector as a whole. Since civil society organizations require independence

from governments, a government-driven solution to this problem—the classic answer to market failures—is not a viable option. After all, government-sponsored support to civil society organizations abroad challenging authoritarian regimes constitutes a form of interference that can trigger the same reaction as the election meddling in 2016. The best available solution to close the information gap is awareness of the limitations of commercial research, as well as increased independent research of targeted threats across the entire spectrum of cyber conflict.

This analysis points to the need for foundations and funders that are often the principal supporters of civil society to take notice of these targeted digital threats and take measures to mitigate them through their grant-making. There are signs of change, such as the Ford Foundation's digital security initiative (Brennan et al. 2017), but we are still far from a broad recognition and prioritization of this issue.

Finally, our findings highlight the need for a follow-up analysis with statistical methods to gauge significance of the selection criteria identified here. In particular, the impact of the level of sophistication as well as the profile of the threat actor on reporting volume needs to be analyzed more systematically. This task faces two key challenges: first, establishing a general measure of sophistication, including both technical and social aspects, and second, it requires consolidating the naming schemes to track reporting. Currently, each firm employs their own naming schemes, and there are no commonly accepted criteria for sophistication.

References

ABC News. 2011. “CIA Director Warns of ‘Cyber-Pearl Harbor.’” ABC News. February 11, 2011. https://abcnews.go.com/News/cia-director-leon-panetta-warns-cyber-pearl-harbor/story?id=12888905.

Alexander, Geoffrey, Matt Brooks, Masashi Crete-Nishihata, Etienne Maynier, John Scott-Railton, and Ronald J Deibert. 2018. “Familiar Feeling: A Malware Campaign Targeting the Tibetan Diaspora Resurfaces.” https://citizenlab.ca/2018/08/familiar-feeling-a-malware-campaign-targeting-the-tibetan-diaspora-resurfaces/.

Bartlett, Evan. 2018. “Fancy Bears: Who Are the Shady Hacking Group Exposing Doping, Cover-Ups and Corruption in Sport? | The Independent,” 2018. https://www.independent.co.uk/sport/football/news-and-comment/fancy-bears-who-are-hacking-group-doping-sport-football-russia-georgia-reedie-bach-a7906376.html.

Bator, Francis M. 1958. “The Anatomy of Market Failure.” *The Quarterly Journal of Economics* 72 (3): 351–79. https://doi.org/10.2307/1882231.

Bennett, Andrew, and Jeffrey T. Checkel. 2015. “Process Tracing: From Philosophical Roots to Best Practices.” In *Process Tracing: From Metaphor to Analytic Tool*, edited by Andrew Bennett and Jeffrey T. Checkel. Strategies for Social Inquiry. Cambridge ; New York: Cambridge University Press.

Betz, David., and Tim. Stevens. 2011. *Cyberspace and the State : Toward a Strategy for Cyber-Power*. Abingdon: Routledge.

BfV. n.d. “Spionage- und Proliferationsabwehr.” Bundesamt für Verfassungsschutz. Accessed February 12, 2019. https://www.verfassungsschutz.de/de/oeffentlichkeitsarbeit/publikationen/pb-spionage-und-proliferationsabwehr.

Brennan, Michael, Elizabeth Eagen, Bryan Nunez, John Scott-Railton, and Eric Sears. 2017. “Digital Security and Grantcraft Guide.” https://www.fordfoundation.org/media/3334/digital-security-grantcraft-guide-v10-final-22317.pdf.

Burgess, Matt. 2017. “Exposed: How One of Russia’s Most Sophisticated Hacking Groups Operates.” *Wired UK*, January 11, 2017. https://www.wired.co.uk/article/how-russian-hackers-work.

Cambridge English Dictionary. n.d. “HIGH-PROFILE | Meaning in the Cambridge English Dictionary.” Accessed February 7, 2019. https://dictionary.cambridge.org/dictionary/english/high-profile.

CLTC. 2018. “Center for Long-Term Cybersecurity Project on Protecting Politically Vulnerable Organizations Threat Landscape and Organizational Ecosystem.” UC Berkeley.

Crete-Nishihata, Masashi, Jakub Dalek, Ronald Deibert, Seth Hardy, Katharine Kleemola, Irene Poetranto, John Scott-Railton, et al. 2014. “Communities at Risk - Extended Analysis.” Citizen Lab. https://targetedthreats.net/media/1-ExecutiveSummary.pdf.

Crete-Nishihata, Masashi, Jakub Dalek, Etienne Maynier, and John Scott-Railton. 2018. “Spying on a Budget: Inside a Phishing Operation with Targets in the Tibetan Community.” The Citizen Lab. January 30, 2018. https://citizenlab.ca/2018/01/spying-on-a-budget-inside-a-phishing-operation-with-targets-in-the-tibetan-community/.

CrowdStrike. 2016. “Bears in the Midst: Intrusion into the Democratic National Committee ».” June 15, 2016. https://www.crowdstrike.com/blog/bears-midst-intrusion-democratic-national-committee/.

Deibert, Ronald. 2003. “Black Code: Censorship, Surveillance, and the Militarisation of Cyberspace.” *Millennium - Journal of International Studies* 32 (3): 501–30. https://doi.org/10.1177/03058298030320030801.

———. 2015. “Cyberspace Under Siege.” *Journal of Democracy* 26 (3): 64–78. https://doi.org/10.1353/jod.2015.0051.

Deibert, Ronald, and Rafal Rohozinski. 2009. “Tracking GhostNet: Investigating a Cyber Espionage Network.” *The Citizen Lab* (blog). March 28, 2009. https://citizenlab.org/2009/03/tracking-ghostnet-investigating-a-cyber-espionage-network/.

Demchak, Chris, and Peter Dombrowski. 2011. “Rise of a Cybered Westphalian Age.” *Strategic Studies Quarterly* 5 (1): 32–61.

Demsetz, Harold. 1970. “The Private Production of Public Goods.” *The Journal of Law & Economics* 13 (2): 293–306.

Eichensehr, Kristen E. 2017. “Public-Private Cybersecurity.” *Texas Law Review; Austin* 95 (3): 467–538.

FireEye. 2014. “APT28:A WINDOW INTO RUSSIA’S CYBER ESPIONAGE OPERATIONS?” FireEye. https://www2.fireeye.com/rs/fireye/images/rpt-apt28.pdf.

———. 2017. “APT28: At the Center of the Storm.” FireEye. https://www2.fireeye.com/rs/848-DID-242/images/APT28-Center-of-Storm-2017.pdf.

Fox-Brewster, Thomas. 2017. “Russian ‘Fancy Bear’ Hackers Tainted Their Huge Leaks With Fake Data.” *Forbes*, 2017. https://www.forbes.com/sites/thomasbrewster/2017/05/26/russian-dnc-hackers-planted-leaks-with-fake-data/.

Franceschi-Bicchierai, Lorenzo. 2016. “How Hackers Broke Into John Podesta and Colin Powell’s Gmail Accounts.” *Motherboard*, October 20, 2016. https://motherboard.vice.com/en_us/article/how-hackers-broke-into-john-podesta-and-colin-powells-gmail-accounts.

George, Alexander L., and Andrew. Bennett. 2005. *Case Studies and Theory Development in the Social Sciences*. Cambridge, Mass.: MIT Press.

Guerrero-Saade, Julian Andres. 2015. “THE ETHICS AND PERILS OF APT RESEARCH: AN UNEXPECTED TRANSITION INTO INTELLIGENCE BROKERAGE.” https://media.kaspersky.com/pdf/Guerrero-Saade-VB2015.pdf.

Hackett, Robert. 2017. “Cybersecurity: 5 of the World’s Most Dangerous Hacker Groups | Fortune.” 2017. http://fortune.com/2017/06/22/cybersecurity-5-hacker-groups/.

Hess, Charlotte, and Elinor Ostrom, eds. 2007. *Understanding Knowledge as a Commons: From Theory to Practice*. Cambridge, Mass: MIT Press.

Hulcoop, Adam, Matt Brooks, Etienne Maynier, John Scott-Railton, and Masashi Crete-Nishihata. 2016. "It's Parliamentary: KeyBoy and the Targeting of the Tibetan Community." https://citizenlab.ca/2016/11/parliament-keyboy/.

Hulcoop, Adam, John Scott-Railton, Peter Tanchak, Matt Brooks, and Ron Deibert. 2017. "Tainted Leaks: Disinformation and Phishing With a Russian Nexus." https://citizenlab.ca/2017/05/tainted-leaks-disinformation-phish/.

Isaac, Mike, and Daisuke Wakabayashi. 2017. "Russian Influence Reached 126 Million Through Facebook Alone." *The New York Times*, October 30, 2017, sec. Technology. https://www.nytimes.com/2017/10/30/technology/facebook-google-russia.html.

Jamieson, Kathleen Hall. 2018. *Cyberwar: How Russian Hackers and Trolls Helped Elect a President What We Don't, Can't, and Do Know*. New York, NY: Oxford University Press.

Jayamaha, Buddhika, and Frankie Matisek. 2018. "HYBRID WAR: ATTACKING THE 'CIVIL' IN CIVIL SOCIETY." *US Army War College War Room* (blog). April 13, 2018. https://warroom.armywarcollege.edu/articles/hybrid-war-attacking-the-civil-in-civil-society/.

Kaldor, Mary. 2013. *Global Civil Society: An Answer to War*. John Wiley & Sons.

Landau, Susan. 2017. "Russia's Hybrid Warriors Got the White House. Now They're Coming for America's Town Halls." *Foreign Policy* (blog). 2017. https://foreignpolicy.com/2017/09/26/russias-hybrid-warriors-are-coming-for-american-civil-society-hacking-trump-clinton/.

Levy, Jack S. 1983. "Misperception and the Causes of War: Theoretical Linkages and Analytical Problems." *World Politics* 36 (1): 76–99. https://doi.org/10.2307/2010176.

Lieberman, Evan S. 2005. "Nested Analysis as a Mixed-Method Strategy for Comparative Research." *American Political Science Review* 99 (03): 435–52. https://doi.org/10.1017/S0003055405051762.

Lindsay, Jon R. 2017. "Restrained by Design: The Political Economy of Cybersecurity." *Digital Policy, Regulation and Governance* 19 (6): 493–514. https://doi.org/10.1108/DPRG-05-2017-0023.

MacGregor, Alice. 2015. "Tropic Trooper Exploits Old Vulnerabilities to Unearth State and Corporate Secrets in Taiwan and Philippines." *The Stack* (blog). May 20, 2015. https://thestack.com/security/2015/05/20/tropic-trooper-exploits-old-vulnerabilities-to-unearth-state-and-corporate-secrets-in-taiwan-and-philippines-2/.

Marczak, Bill, John Scott-Railton, Adam Senft, Bahr Abdul Razzak, and Ron Deibert. 2018. "The Kingdom Came to Canada: How Saudi-Linked Digital Espionage Reached Canadian Soil." https://citizenlab.ca/2018/10/the-kingdom-came-to-canada-how-saudi-linked-digital-espionage-reached-canadian-soil/.

Mukherjee, S P. 2018. *Statistical Methods in Social Science Research*. New York, NY: Springer Berlin Heidelberg.

Muncaster, Phil. 2017. "Chinese KeyBoy Group Unlocks More Victim Networks." *Infosecurity Magazine*, November 6, 2017. https://www.infosecurity-magazine.com:443/news/chinese-keyboy-group-unlocks/.

Murphy, Kevin M., and Robert H. Topel. 2013. “Some Basic Economics of National Security.” *The American Economic Review* 103 (3): 508–11.

NCSC. 2018. “Indicators of Compromise for Malware Used by APT28 - NCSC Site.” 2018. https://www.ncsc.gov.uk/alerts/indicators-compromise-malware-used-apt28.

———. n.d. “Threats and Vulnerability Reports.” Accessed January 30, 2019. https://www.ncsc.gov.uk/index/report.

Networks Asia Staff. 2015. “Philippines, Taiwan Are Latest Targets of ‘Operation Tropic Trooper’ Malware.” *Networks Asia*, 2015. https://www.networksasia.net/article/philippines-taiwan-are-latest-targets-operation-tropic-trooper-malware.1432083840.

Nissen, Thomas Elkjer. 2016. “Cyber Warfare by Social Network Media.” In *Conflict in Cyber Space: Theoretical, Strategic and Legal Pespectives*, edited by Karsten Friis and Jens RIngmose, 1st ed. Routledge. https://doi.org/10.4324/9781315669878.

Nye, Joseph S. 2017. “Deterrence and Dissuasion in Cyberspace.” *International Security* 41 (3): 44–71. https://doi.org/10.1162/ISEC_a_00266.

ODNI. 2017. “ODNI Statement on Declassified Intelligence...” IC ON THE RECORD. January 6, 2017. http://t.umblr.com/redirect?z=https%3A%2F%2Fdni.gov%2Ffiles%2Fdocuments%2FICA_2017_01.pdf&t=ZGFkZWE0MjI4NDczNmI1NzU2ZDIzNGE1YzY4ODFlNjU2NjZjZjk1MCx4eWV3cW43NA%3D%3D&b=t%3ACeDO6NTe6pPkB8DydjGePw&p=https%3A%2F%2Ficontherecord.tumblr.com%2Fpost%2F155494946443%2Fodni-statement-on-declassified-intelligence&m=1.

Olson, Mancur. 1971. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Revised edition. Harvard Economic Studies, v. 124. Cambridge, Massachusetts ; London, England: Harvard University Press.

Oremus, Will, and Jim Newell. 2017. “Russia Used Fake News to Influence the Election, Says U.S. Intelligence Chief.” *Slate*, January 5, 2017. http://www.slate.com/blogs/future_tense/2017/01/05/russia_used_fake_news_to_influence_the_election_james_clapper_says.html.

Ostrom, Vincent, and Elinor Ostrom. 1977. *Public Goods and Public Choices*. Indiana University, Workshop in Political Theory and Policy Analysis.

PwC. 2017. “The KeyBoys Are Back in Town.” https://www.pwc.co.uk/issues/cyber-security-data-privacy/research/the-keyboys-are-back-in-town.html.

Rapid7. 2013. “KeyBoy, Targeted Attacks against Vietnam and India.” https://blog.rapid7.com/2013/06/07/keyboy-targeted-attacks-against-vietnam-and-india/.

Recorded Future. 2018. “RedAlpha: New Campaigns Discovered Targeting the Tibetan Community.” https://www.recordedfuture.com/redalpha-cyber-campaigns/.

Rogin, Josh. 2012. “NSA Chief: Cybercrime Constitutes the ‘Greatest Transfer of Wealth in History.’” *Foreign Policy*, July 9, 2012. https://foreignpolicy.com/2012/07/09/nsa-chief-cybercrime-constitutes-the-greatest-transfer-of-wealth-in-history/.

———. 2018. “Opinion | Washington Must Wake up to the Abuse of Software That Kills.” *Washington Post*, 2018. https://www.washingtonpost.com/opinions/2018/12/12/washington-must-wake-up-abuse-software-that-kills/.

Rosenzweig, Paul. 2011. “Cybersecurity and Public Goods.” http://media.hoover.org/sites/default/files/documents/EmergingThreats_Rosenzweig.pdf.

Samuelson, Paul A. 1954. “The Pure Theory of Public Expenditure.” *The Review of Economics and Statistics* 36 (4): 387–89. https://doi.org/10.2307/1925895.

Sanger, Eric Lipton, David E., and Scott Shane. 2016. “The Perfect Weapon: How Russian Cyberpower Invaded the U.S.” *The New York Times*, December 13, 2016. https://www.nytimes.com/2016/12/13/us/politics/russia-hack-election-dnc.html.

Schleifer, Theodore, and Deirdre Walsh. 2017. “McCain: Russian Cyberintrusions an ‘Act of War.’” CNN. January 6, 2017. http://www.cnn.com/2016/12/30/politics/mccain-cyber-hearing/index.html.

SecureWorks. 2016a. “Hillary Clinton Email Targeted by Threat Group-4127.” https://www.secureworks.com/research/threat-group-4127-targets-hillary-clinton-presidential-campaign.

———. 2016b. “Threat Group-4127 Targets Google Accounts.”

Shane, Scott, and Vindu Goel. 2017. “Fake Russian Facebook Accounts Bought $100,000 in Political Ads.” *The New York Times*, September 6, 2017, sec. Technology. https://www.nytimes.com/2017/09/06/technology/facebook-russian-political-ads.html.

Shezaf, Hagar. 2018. “Snowden: Israeli Firm’s Spyware Was Used to Track Khashoggi - Israel News - Haaretz.Com,” 2018. https://www.haaretz.com/israel-news/.premium-israeli-spyware-was-used-to-track-saudi-journalist-khashoggi-edward-snowden-says-1.6633745.

Svetoka, Sanda, and Anna Reynolds. 2016. *Social Media as a Tool of Hybrid Warfare*. Riga: NATO Strategic Communications Centre of Excellence. http://www.stratcomcoe.org/social-media-tool-hybrid-warfare.

Swaine, Jon. 2018. “Twitter Admits Far More Russian Bots Posted on Election than It Had Disclosed.” *The Guardian*, January 20, 2018, sec. Technology. https://www.theguardian.com/technology/2018/jan/19/twitter-admits-far-more-russian-bots-posted-on-election-than-it-had-disclosed.

Symantec. 2018. “APT28: New Espionage Operations Target Military and Government Organizations.” https://www.symantec.com/blogs/election-security/apt28-espionage-military-government.

Threat Intelligence Researcher. 2018Phone.

TrendMicro. 2015a. “Operation Tropic Trooper: Old Vulnerabilities Still Pack a Punch - TrendLabs Security Intelligence Blog.” http://blog.trendmicro.com/trendlabs-security-intelligence/operation-tropic-trooper-old-vulnerabilities-still-pack-a-punch/.

———. 2015b. “Pawn Storm’s Domestic Spying Campaign Revealed; Ukraine and US Top Global Targets - TrendLabs Security Intelligence Blog.” https://blog.trendmicro.com/trendlabs-security-intelligence/pawn-storms-domestic-spying-campaign-revealed-ukraine-and-us-top-global-targets/.

———. 2017. “Two Years of Pawn Storm.” https://documents.trendmicro.com/assets/wp/wp-two-years-of-pawn-storm.pdf.

———. 2018. “Tropic Trooper’s New Strategy - TrendLabs Security Intelligence Blog.” https://blog.trendmicro.com/trendlabs-security-intelligence/tropic-trooper-new-strategy/.

US CERT. n.d. “Incident Reporting System.” Accessed January 30, 2019. https://www.us-cert.gov/ncas/analysis-reports.

Valeriano, Brandon, and Ryan Maness. 2015. *Cyber War versus Cyber Realities: Cyber Conflict in the International System*. Oxford, New York: Oxford University Press.

Van Evera, Stephen. 1997. *Guide to Methods for Students of Political Science*. Ithaca, NY: Cornell University Press.

Waldner, David. 2015. “What Makes Process Tracing Good?” In *Process Tracing: From Metaphor to Analytic Tool*, edited by Andrew Bennett and Jeffrey T. Checkel. Strategies for Social Inquiry. Cambridge ; New York: Cambridge University Press.

Wirtz, James J. 2014. “The Cyber Pearl Harbor.” In *Cyber Analogies*. Naval Postgraduate School.

YouGov. 2017. “America’s Friends and Enemies | YouGov.” 2017. https://today.yougov.com/topics/politics/articles-reports/2017/02/02/americas-friends-and-enemies.