If it bleeps it leads? Media coverage on cyber conflict and misperception

Journal Article

Author(s):
Makridis, Christos; Maschmeyer, Lennart; Smeets, Max Willem Eline

Publication date:
2024

Permanent link:
https://doi.org/10.3929/ethz-b-000658766

Originally published in:
Journal of Peace Research, https://doi.org/10.1177/00223433231220264

Check for updates

Journal of Peace Research

If it bleeps it leads? Media coverage on cyber conflict and misperception

Journal of Peace Research
1–15

DOI: 10.1177/00223433231220264
journals.sagepub.com/home/jpr
Sage

Christos Makridis

Human-Centered Artificial Intelligence, Stanford University, USA

Lennart Maschmeyer

Center for Security Studies, ETH Zürich, Switzerland

Max Smeets

Center for Security Studies, ETH Zürich, Switzerland

Abstract

What determines media coverage on cyber conflict (CC)? Media bias fostering misperception is a well-established problem in conflict reporting. Because of the secrecy and complexity surrounding cyber operations (COs), where most data moreover come from marketing publications by private sector firms, this problem is likely to be especially pronounced in reporting on cyber threats. Because media reporting shapes public perception, such bias can shape conflict dynamics and outcomes with potentially destabilizing consequences. Yet little research has examined media bias systematically. This study connects existing literature on media reporting bias with the CC literature to formulate four theoretical explanations for variation in reporting on COs based on four corresponding characteristics of a CO. We introduce a new dataset of COs reporting by the private sector, which we call the Cyber Conflict Media Coverage Dataset, and media reporting on each of these operations. Consequently, we conduct a statistical analysis to identify which of these characteristics correlate with reporting quantity. This analysis shows that the use of novel techniques, specifically zero-day exploits, is a highly significant predictor of coverage quantity. Operations targeting the military or financial sector generate less coverage. We also find that cyber effect operations tend to receive more coverage compared to espionage, but this result is not statistically significant. Nonetheless, the predictive models explain limited variation in news coverage. These findings indicate that COs are treated differently in the media than other forms of conflict, and help explain persistent threat perception among the public despite the absence of catastrophic cyberattacks.

Keywords

bias, cyberattacks, cyber conflict, espionage, media, reporting, sophistication, zero-day

Introduction

Cyber threats remain a mysterious menace. Their secrecy and technical complexity hamper clear assessments and produce lingering uncertainty concerning the nature of the threat, its origins and its extent (Clarke and Knake, 2010; Kello, 2013; Lindsay, 2013). This situation fosters misperception and miscalculation – undermining stability and raising the risk of inadvertent or accidental escalation in crises (Buchanan, 2020; Buchanan and Cunningham, 2020; Jervis, 2017). An emerging consensus in cybersecurity scholarship holds cyber operations (COs) to be a low-intensity alternative to warfare (Buchanan, 2020; Harknett and Smeets, 2022; Maschmeyer, 2021). Accordingly, other articles in this special issue examine the use of COs for espionage as well as for influence operations (Akoto, 2024; Vicic and Gartzke, 2024).

Corresponding author:
lmaschmeyer@ethz.ch

And yet, a growing body of research using survey experiments indicates that exposure to cyberattacks evokes emotional responses and psychological distress among the public (Gomez and Villar, 2018; Gross et al., 2017; Shandler et al., 2022). These findings may explain why, although catastrophic attacks have not happened, most United States citizens still perceive cyber-attacks to be the greatest national security threat, even above nuclear weapons (Visé, 2023). Consequently, even if cyber-attacks themselves remain low in intensity, public threat perception may still increase the risk of escalation by pressuring governments to retaliate (Jardine and Shandler, 2024; Shandler et al., 2022, 2023). Moreover, how cyber-attacks are perceived can shape public opinion concerning security policy, such as attitudes towards surveillance (Arsenault et al., forthcoming). Importantly, media reporting is a key determinant of public perception of cyber threats (Gomez, 2019; Snider et al., 2021). How the media reports cyber threats, therefore, may contribute towards instability and escalation risks in cyber conflict (CC) as well as produce spillover effects on wider security policy.

In the ideal case, media reporting reduces uncertainty by providing the public (and decision-makers) with information on operations whose sponsors strive to keep secret. Previous research has shown that the prevalence of online media, in combination with well-oiled media ecosystems, can even reduce the likelihood of states opting in favour of covert operations against other states where these conditions are present due to perceived risk of exposure (Joseph and Poznansky, 2018). Yet, in practice not all cyber threats are created equal – nor does the media report about events consistently.

Consider the following example: in February 2014, cyber threat intelligence firm Kaspersky Lab published a report on 'The Mask,' a hacking group that was said to be likely backed by an unknown national government targeting a wide range of targets including government agencies, diplomatic offices and energy companies (Donohue, 2014). Kaspersky described the activity of The Mask as the 'The world's most sophisticated APT [Advanced Persistent Threat] Campaign' (Donohue, 2014). The threat intelligence company reported that the hacking group has been operating since at least 2007. In the subsequent months, several experts wrote about the hacking group – such as Bruce Schneier making the case that they were part of the Spanish intelligence services.[1] Some journalists also reported on the tradecraft of this group, and its targeting patterns (Leyden, 2014; Menn and Finkle, 2014). Yet, most media outlets paid scant attention. Reporting on the group did not reach the front page of major newspapers. In fact, most did not even write a short story for their online audience.

Four years later, in 2018, the same company, Kaspersky Lab, wrote a report on Olympic Destroyer, the hacking attack on the 2018 Olympics that 'temporarily paralyzed IT [information technology] systems, shutdown display monitors, crippled Wi-Fi and shuttered the Olympics website preventing visitors from printing tickets' (Spring, 2018). Other threat intelligence companies, such as Talos Checkpoint, also tracked this operation.[2] When the private sector released information about Olympic Destroyer it received wide coverage in the media. Over 2000 media stories were published.[3]

The two examples illustrate the significant variation in media coverage of COs. This article examines what explains this variation. Why do some COs receive more media attention than others? Answering this question is of great value not only for policy making, but also for understanding public perception of cyber threats – with important implications for stability.

Yet, there has been surprisingly little analysis addressing this question. Our current understanding of why some operations receive more news coverage than others is based solely on anecdotal evidence. Existing research highlights bias within news reporting, such as threat inflation towards 'cyber doom' scenarios and a hype of sophistication (Buchanan, 2017; Dunn-Cavelty, 2008, 2013; Lawson, 2013). Yet, no scholarship has examined which factors determine news coverage of COs. Consequently, the importance of quantitative data cannot be overstated in the study of biases in news coverage of COs. While case examples and anecdotes can provide valuable insights, they cannot offer the same level of generalizability and rigour that a quantitative analysis can.

In line with the goals of the special issue, 'cyber-conflict: moving from speculation to investigation' (Shandler and Canetti, 2024), we conduct the first quantitative analysis of the determinants of media attention in CC reporting. This study connects existing literature on media reporting bias with the CC literature to formulate four theoretical explanations of variation in reporting. These explanations focus on effect intensity, target, sophistication, and origin of a cyberattack, respectively.

To assess alternate explanations about variation in media reporting, we first introduce a new dataset of COs reporting, which we call the Cyber Conflict Media Coverage Dataset. This dataset is the most complete dataset of COs based on reporting from commercial threat intelligence providers, from which journalists get much of their information. Subsequently, we trace media reporting on each of these operations and conduct a statistical analysis to identify which of the characteristics in the previous step correlate with reporting quantity.

We obtain four main results. First, when a CO has used a zero-day exploit to gain or escalate access it is associated with 359–398 more stories, or 165–169% more, relative to operations that do not use zero-day attacks. Second, we find that disruptive and destructive COs generate significantly more news stories than their espionage counterparts. For example, when we use negative binomial models, we find that they generate 408–514 more news stories. Similarly, when we compress the distribution of stories by taking its logarithm, we find that they generate 84–100% more stories. We find statistically insignificant effects associated with the target country, for example, United States, the Group of Seven (an intergovernmental political and economic forum consisting of Canada, France, Germany, Italy, Japan, the United Kingdom and the United States) or Group of 20 (G20, an intergovernmental forum comprising Argentina, Australia, Brazil, Canada, China, France, Germany, India, Indonesia, Italy, Japan, Republic of Korea, Mexico, Russia, Saudi Arabia, South Africa, Türkiye, United Kingdom and United States, as well as the European Union and the African Union). Third, we find that certain sectors, namely healthcare and energy, receive much more coverage than operations targeting the military, government, finance, or media.

Fourth, we find interesting intertemporal patterns in coverage. Specifically, news stories after a six months period are generally positively predictive of news stories after 12 months, and news stories after 12 months are highly predictive of news stories after 18 months. While there could be some mean reversion in the short run, we do find long run persistence of stories. This is intuitive: given the same editor and momentum that builds up after a story, continued coverage will continue. Nonetheless, we recognize that some of our statistical estimates have large standard errors, so we caution that more data are needed to make fully definitive conclusions.

Prior research has shown that reporting by commercial threat intelligence firms is biased by their business interests, prioritizing high-profile threats and neglecting threats to weaker actors (Egloff, 2020; Maschmeyer et al., 2021; Work, 2020). Media reporting thus builds on a data source already subject to significant bias – and introduces its own bias by selecting and prioritizing material most likely to catch the attention of its target audiences. We measure its effects and show that the result is a 'double bias,' where only a fraction of a fraction of activity gets reported on, distorting academic and policy debates.

Theoretical propositions

We examine bias in media coverage of COs. Previous research has documented bias in private sector reporting on COs due to underlying business incentives, which privilege operations that score high on one or more of three key characteristics: (a) using unique tactics, techniques and procedures; (b) targeting of a high-profile victim; and (c), being sponsored by a high-profile threat actor (Maschmeyer et al., 2021). We assess whether the media reporting of such threats adds another source of bias, creating 'double bias.'

Media bias in general is a well-established and multifaceted problem, documented by a large body of research (Alterman, 2003; Baron, 2006; Groseclose and Milyo, 2005; Innis, 1951; Niven, 1999). We focus on print media, mainly for reasons of data collection – more on this further below. Within this media type, we examine a specific type of bias, namely selection bias, which is defined as 'the selection of but a few of the many possible events to observe and report' (McCarthy et al., 1996). Its main cause is straightforward: just like private threat intelligence firms are driven by incentives that shape their reporting, so are media organizations. Previous research identifies two key incentives: maximizing attention; and the sale of advertising space (Ellman and Germano, 2009). A large body of research has assessed the determinants of media selection bias in reporting international security events, such as civil conflict, drone attacks and terrorism (Berlemann and Thomas, 2019; Kearns et al., 2019; Moeller, 2006; Shoemaker and Cohen, 2006). We connect insights from these studies with the literature on CC to set up a novel theoretical framework on biases in CC media reporting. We identify four core characteristics that determine coverage, as laid out below.

Effects intensity

The more violent and grisly a news story, the more attention it tends to generate. This dynamic is captured by the well-worn trope, 'if it bleeds, it leads,' attributed to 19th century newspaper publisher William Randolph Hearst. Empirical research confirms this adage, with Miller and Albert (2015) finding clear statistical evidence that the quantity of news coverage of a given conflict increases with the quantity of fatalities. Apart from fatalities, research by Snyder and Kelley (1977) has shown that conflict intensity in general is a significant predictor of the quantity of newspaper coverage. Mueller (1997) has shown the same to be the case concerning the intensity of protests.

Accordingly, we expect the intensity of effects a CO produces to be similarly predictive of the quantity of news coverage it receives. Especially considering that critical researchers have identified the media as a key driver of heightened threat perception and fear of cyber doom (Dunn-Cavelty, 2013; Lawson and Middleton, 2019), we would expect media reporting to concentrate on the most intense and thus threatening effects. There are two types COs: those that passively monitor activity or steal information; and those that produce active effects against targeted systems, such as disruption, denial, degradation, or destruction. In technical circles, the former are commonly known as computer network exploitation (CNE), whereas the latter are referred to as computer network attacks (CNA) (Zetter, 2016).[4] For simplicity, and because there exists no common scale of effects intensity in COs, we use this binary distinction to measure intensity. Because we expect effects intensity to predict coverage quantity, we thus formulate the following first hypothesis:

> *Hypothesis 1*: CNA receives more news coverage than CNE.

Since more intense effects tend to receive more frequent news coverage, and since CNAs generate more intense effects than CNE, we expect CNAs to receive relatively more news coverage.

Target type

Apart from effects, the target of a CO is expected to be equally important. Shoemaker and Cohen (2006) have established a basic model of newsworthiness with a general measure of target significance as a key determinant of coverage quantity. In reporting on terrorism, the quantity of coverage tends to increase with the political or symbolic significance of the target (Nacos, 2016).

Similarly, Kearns et al. (2019) found target type to be a key predictor of news coverage, where terrorism against government targets received more coverage than non-government targets. Finally, an auxiliary source of bias is the target's distance to the media outlet. Berlemann and Thomas (2019) found systematic evidence of such distance bias in reporting of natural disaster. We expect the same types of biases to apply to reporting on cyber threats. This expectation is warranted since public perception and framing of cyber threats shows key parallels to both terrorism and natural disasters, particularly concerning recurring ‘cyber doom’ scenarios (Dunn-Cavelty, 2008, 2013; Lawson, 2013). Based on these expectations and previous research, we formulate the following hypothesis:

> *Hypothesis 2*: When an operation targets entities in the Global North it receives more news coverage.

Because most threat intelligence firms and their prospective customers are in the Global North, we expect COs in that area to be more likely to get reported than those in the Global South. Since the English-language news media we examine are also predominantly located in the Global North, and since news media have exhibited a distance bias, we expect them to be more likely to cover private sector reporting on COs in that area than in the Global South.

Sophistication

The third characteristic of COs we expect to correlate with media reporting quantity is perceived sophistication. The term is ubiquitous in private sector and media reports, yet rarely defined. Aitel introduced a five-level framework to score the sophistication of a CO (Aitel, 2016), which allows a nuanced classification. However, likely due to this complexity, in practice the term remains ill-defined to such a degree that Buchanan claims the idea of sophisticated hackers has assumed the status of a ‘legend’– ubiquitous, awe-inspiring and rarely questioned (Buchanan, 2017). The media is expected to pick up on easy to use and understand indicators of sophistication. The indicator that stands out are zero-days (Healey, 2016; Joyce, 2016; Smeets, 2022). Zero-days are vulnerabilities in software or hardware that are unknown to the vendor and the user(s) of the former (Zetter, 2014).

We expect COs with unique characteristics along the lines just outlined to receive more reporting for two reasons. First, there is a known media selection bias towards quirky and astonishing stories in disaster reporting (Moeller, 2006: 184). We expect the same to be the case with COs, hence the more sophisticated and thus astonishing a CO is, the likelier it is to be picked up. Second, and conversely, however, Moeller also argues that simplicity of natural disasters is a key predictor of reporting quantity (Moeller, 2006: 184–186). Accordingly, we expect that easily observable indicators of sophistication are more likely to predict reporting volume than complex investigations across the sophistication framework proposed by Aitel (2016) and Buchanan

(2017) (see further above). We thus formulate the following hypothesis:

> *Hypothesis 3*: COs with easily observable indicators of sophistication (i.e. use of zero-day exploits) receive more news coverage than those with less easily observable indicators of sophistication.

Because the legend of sophistication generates attention, and because past research indicates reporting tends to favour simpler incidents over more complex ones, we expect COs with easily observable (and explainable) indicators of sophistication, namely the presence of zero-day exploits, to receive more frequent coverage than those without such easily observable indicators.

Threat origin

The fourth and final aspect is threat origin. Existing communications research shows a media bias towards those not part of the audience in-group, and a corresponding overrepresentation of non-white people in news about terrorism (Kearns et al., 2019: 989). Applying this finding to the concept of a 'threat actor,' used in threat intelligence reporting to refer to a hacking group, we would accordingly expect those actors not part of the audience's 'in-group' to be overrepresented in media reporting. This expectation fits with previous research on threat intelligence reporting suggesting that threat group identity predicts reporting volume. Specifically, operations attributed to state-sponsored threat actors linked to perceived enemies of the Western alliance, namely Russia, China, Iran and North Korea, are more likely to be reported (Maschmeyer et al., 2021: 7). Accordingly, we expect COs originating within one of these countries to receive a higher amount of media reporting than those originating in other countries.[5]

> *Hypothesis 4*: Operations pursued by key adversaries of the Western powers (i.e. Russia, China, Iran and North Korea) are more likely to receive attention than operations pursued by other actors.

Because news reporting aims to maximize attention, it is more likely to report on COs by perceived enemies, which generate fear, than by other states.

Data and measurement

COs

We construct a new dataset of COs for the period 2000–2021, which we call the Cyber Operations Dataset.[6] The dataset is based on commercial cyber threat intelligence reporting on 'Advanced Persistent Threats' (APT) activity. Greg Rattray introduced the term in 2007 to characterize emerging adversaries in cyberspace that required a coordinated defence from the defence industrial base (Bejtlich, 2020). Today, however, the term more broadly refers to those actors that are advanced and/or persistent in their efforts to achieve certain objectives.

Not just the media, but also policymakers, scholars, and military professionals heavily rely on information from commercial threat intelligence companies to understand the activities of APTs (Maschmeyer et al., 2021; Work, 2020). Private sector reporting is driven by a mixed set of incentives. For one, gaining advanced knowledge about exploitation activity can help clients and vendors to fix vulnerabilities. There are also more indirect benefits, as the provision intelligence may help sell other products as well. The commercial cyber intelligence market was valued at $1.5 billion in 2018, excluding auxiliary activities (Work, 2020: 8).

The outcome of interest to this study is the coverage of a CO conducted by an APT. COs concern a set of activities that seek unauthorized access to computers, computer systems or networks to achieve a certain objective. The Cyber Kill Chain distinguishes seven phases of COs: reconnaissance; weaponization; delivery; exploitation; installation; command and control; and actions on objectives (Lockheed Martin, 2015). An APT can run multiple operations, each with different goals.[7]

Several datasets already exist that combine commercial cyber threat intelligence on APTs, the most comprehensive of which is a shared spreadsheet titled 'APT Groups and Operations,' developed by Florian Roth and maintained by multiple other researchers (Roth et al., 2015). It contains information on APT names, associated operations, and occasionally, tools used. However, it does not provide information on other relevant variables, such as the type of operation (CNA or CNE). Furthermore, this dataset is coded at the actor-level, rather than the operation-level. While some actors had operations listed in the dataset, this was not always the case. Consequently, to identify operations, we examined all sources listed for every actor in the spreadsheet, as the operations mentioned in the dataset were unlikely to represent the complete set of an actor's activities.

In addition, we did a Google search for all actors to identify threat intelligence reports that might have been missing, and if they covered an operation, included it in the data. We also used the ThaiCERT Threat Group Cards as an additional source to identify operations, identify alternative actor names, and retrieve the sectors and countries targeted by actors/operations.[8] We also checked two other sources: APT Map; and the MITRE repository. Through this set-up we have added over 35 operations that are not listed in Roth et al.'s dataset.

There were several additional challenges for coding our dataset. First, the private sector is often inconsistent in their use of language about APT activity. There are no clear rules when a set of activity is called an 'operation,' 'campaign,' 'attack,' 'incident' or nothing. This is true both for internal consistency of a given threat intelligence company (a company might describe something as an 'operation', but not do so for a different, but seemingly comparable set of activities), and also for reports of different companies (one company might call a set of activities an 'operation', while another one does not).

Second, in many cases, the first observed activity of an APT actor is not referred to with a specific operation/campaign name, but with the name of the APT actor itself. Sometimes, the threat intelligence reports will also use the same name to describe the actor and the activity/operation (for example, 'Operation Ke3chang', which both refers to the actor, and a specific set of activity observed). This becomes problematic when the same actor subsequently conducts other operations, which then (not always, but sometimes) are referred to by a specific operation name that is different from the name of the actor. Also, some operations have received multiple names (like APTs can also receive multiple labels by different companies). We therefore coded both the most used name as well as the alternative name.

Third, sometimes specific pieces of malware are described as 'operations' which is inconsistent with the notion of operation we have in mind for this dataset. In these cases, we have either not included these cases, or have included them but labelled accordingly.

Media coverage

There is no clearly established way of measuring media coverage in the literature. Various studies use a binary measure for whether a certain event was covered in the printed media. For example, Meyer (2021) seeks to explain media coverage of a constitutional court decision in Germany and looks at whether a decision by the court was covered in at least one printed newspaper.[9] This article also focuses on printed newspapers. However, given the pervasiveness of cyber coverage, we count the number of articles per CO. In line with the coding of our dataset described above, to conduct a thorough search for publications, we employed multiple keywords for each operation. Furthermore, some COs are characterized by common names such as 'Cloud Hopper' and 'Operation Hangover.' To exclude irrelevant articles, the keyword 'cyber' was included in all searches conducted. It is important to note that for each operation, searches were restricted to the time-period following its initial public disclosure.

We do this at several points in time – within three months, six months, 12 months, 18 months and 24 months of first disclosure by a cyber threat intelligence company – to capture variation in how coverage of different operations evolves over time. We used the LexisNexis database for newspaper coverage. This database is by far the most used one in academic studies (e.g. Freudenburg et al., 1996; Haider-Markel et al., 2006).[10] LexisNexis covers more than 650 news outlets covering legal, corporate, and governmental issues. The collection encompasses prominent global English-language publications such as the *New York Times*, as well as local newspapers such as the *Baltimore Sun*. Nevertheless, it excludes smaller independent sources or blogs that report on hacking-related news. We focus on general news outlets rather than these more specialized ones for two reasons. First, we are interested in public opinion in general. Second, we are aiming to determine whether reporting on COs exhibits similar biases to those known in reporting on conflict, terrorism, and disasters or not. Research on the latter has also focused on general news outlets, consequently our focus on the same media types allows for a direct comparison. As such, we are pursuing a purposive sampling strategy, namely one that does not aim at completeness, but rather focuses on a particular set of publications guided by the logic of the research project (Riffe et al., 2019: 76).

Our approach to capture media coverage has the advantage of not only providing more identifying variation since we are now working with a continuous variable (and we can apply count estimators, such as Poisson and Negative Binomial models), but also allowing us insight into the intensity of a CO, rather than simple coverage. As Makridis (2021) has found, the average-sized publicly-reported data breach has a positive effect on company reputation, whereas the biggest and most salient breaches have a negative effect.

Independent variables

We analysed the content of commercial threat intelligence reports to capture characteristics of COs relevant for the theoretical propositions developed in the previous section. First, we coded *Operation Type* referring to the type of operation conducted by the APT.

Table I. Descriptive statistics: the proportion of cybersecurity incidents and the SD across time and within different type periods for the major variables examined in the study

	Pooled		*2010–2014*		*2015–2020*	
	Mean	*Standard deviation (SD)*	*Mean*	*SD*	*Mean*	*SD*
Operation type						
Espionage	0.73	0.45	0.69	0.46	0.76	0.43
Attack	0.10	0.30	0.15	0.36	0.07	0.26
Other	0.07	0.26	0.05	0.22	0.08	0.28
Origin						
Global North	0.50	0.50	0.59	0.50	0.47	0.50
United States	0.26	0.44	0.36	0.48	0.22	0.42
G20	0.48	0.50	0.56	0.50	0.45	0.50
Operations from China	0.27	0.45	0.34	0.48	0.25	0.43
Operations from Russia	0.18	0.38	0.17	0.38	0.17	0.37
Operations. from North Korea	0.11	0.32	0.08	0.28	0.13	0.33
Operations from Iran	0.11	0.32	0.12	0.33	0.12	0.33
Operations from United States	0.02	0.15	0.07	0.25	0.00	0.00
Sector hit						
Government hit	0.36	0.48	0.32	0.47	0.35	0.48
Military hit	0.18	0.38	0.22	0.42	0.18	0.38
Finance hit	0.07	0.26	0.07	0.25	0.07	0.26
Energy hit	0.12	0.33	0.17	0.38	0.10	0.30
Health hit	0.04	0.21	0.03	0.18	0.05	0.22
Media hit	0.07	0.26	0.08	0.28	0.06	0.24
Sophistication						
Zero-day	0.15	0.36	0.26	0.44	0.12	0.32
Media coverage						
Total news articles	204.14	763.73	287.07	819.78	165.71	752.10
Articles after six months	177.08	706.08	195.75	555.83	164.68	761.01
Articles after 12 months	148.20	585.35	166.23	469.57	134.17	620.87
Articles after 18 months	174.66	829.50	316.24	1,341.89	116.30	535.04
Observations	228		59		157	

We distinguish between three categories: ‘attack’ refers to those operations with the aim to disrupt, deny, degrade, and destroy (above described as CNA); ‘espionage’ refers to those operations with the aim to collect intelligence; and ‘other’ refers to all other operations. Second, we coded a set of binary variables based on the APT’s country of origin. We coded whether the operation was conducted by American, Chinese, Russian, North Korean, or Iranian based APTs, as well as two more general categories: if the operation came from a country in the Global North;[11] or a country part of the G20. Third, corresponding to Shoemaker and Cohen’s (2006) model of target significance, we track targeting of the following sectors to measure significance of COs: government; military; financial; energy; health; and media. We add a sixth sector as well, namely targeting of the media itself. If distance bias does influence reporting quantity, operations targeting the media should hit ‘closest to home’ for journalists and thus generate higher reporting quantity than against other sectors further ‘away’ from the media. Third, we analysed which CO is known to have used a zero-day. Table I provides an overview with descriptive statistics.

Intercoder reliability test

Finally, we conducted an intercoder reliability test to verify the validity of our coding scheme. Following established practice, we selected a random sample of the operations in our dataset consisting of 70 reports (30% of the entire sample) included in the dataset and calculated Cohen’s Kappa scores for the variables corresponding to our four hypotheses (Lombard et al., 2002; Neuendorf, 2017: 235). We selected reports based on a random identifier (a randomized number), trained a second researcher in our coding scheme, and had the researcher code this random sample to compare their results with our own. The results were encouraging. The analysis revealed complete agreement across all variables, except one: the type of CO. However, even regarding this variable, there existed a high level of agreement among the coders, with an agreement rate of 98.39% (compared to an expected agreement of 68.78%). This high level of concordance resulted in a Kappa score of 0.95.

Statistical specification and main results

To understand the relationship between media coverage and the characteristics of COs, we consider regressions of the form:

$$TM_i = \gamma o_i + \xi North_i + \phi a_i + \zeta z_i + \lambda_t + \varepsilon_i \qquad (1)$$

where TM_i denotes the total number of media articles about a CO i, o denotes an indicator for whether it was a cyber effect operation (i.e. CNA) or other operation, *North* denotes an indicator for whether a country in the Global North was targeted, a denotes a vector of indicators for the sector the attack was directed (government, military, finance, energy, health and media), and z denotes an indicator for whether the attack was a zero-day. We also control for a linear time trend, λ. Standard errors are heteroskedasticity-robust.

Because our outcome variable of the total number of articles is often censored at zero (roughly 25% of the sample), we estimate Equation 1 using a negative binomial model, which handles count variables especially when there are many zeros. We also take the hyperbolic sine of the total number of news articles to accommodate zeros. We sequentially layer on the controls in Equation 1 to assess the different potential factors behind the dispersion in news. By doing so, we can also gauge the potential importance of omitted variables as potential confounding factors that would otherwise bias our results: by studying the change in our coefficients of interest as more controls are added, we can learn about the possible severity of omitted variables bias.

Table II documents these results. Starting with column 1, we group operations and other attacks together, finding that these attacks generate 1759 more news stories than their espionage counterparts. That is, operations that aim to disrupt, deny, degrade and/or destroy have 1759 more stories, but the estimate is not statistically significant. In column 2, we add in indicators for whether the United States or G20 countries were targeted. Now, the coefficient estimate on operations and other attacks declines by roughly half in magnitude, as does its standard error, and we find that United States targeted attacks gain less attention and G20 targeted attacks gain more attention, relative to their non-United States or G20 counterparts. However, again these estimates are not statistically significant. If we omit the G20 indicator, we still find a similar coefficient on the United States.

Column 3 subsequently adds indicators for all the different sectors, which enter insignificantly, but nonetheless reduce the estimate on operations and other activity. Column 4 adds an indicator for the use of one or more zero-days by the actor in the operation. Here, in our preferred specification, we find that the use of one or more zero-day exploits in a CO is associated with 359 more news stories, statistically significant at the 10% level. Now, we also find that the military, finance and media sectors receive slightly fewer stories, relative to their counterparts, whereas healthcare and energy sectors more news stories. Importantly, we also find a statistically significant effect of 515 more stories on operations and other attacks now that zero-days are included, suggesting that failure to control for zero-days creates attenuation bias on our coefficients of interest. Column 5 adds indicators on the origin of the attack, but the results are not altered substantially.

One concern with the results so far is that news stories are highly skewed: some events receive substantial coverage, whereas others receive little. To address this concern, we take the hyperbolic sine of news stories, which is equivalent to taking the logarithm and keeping values of zero. Here, we find qualitatively similar results: operations and other attacks now enter significantly at a 10% level and are associated with upwards of 80% more news stories. We also find that the use of zero-day exploits is associated with 165–169% more news stories.

Next, we examine the persistence of cyberattacks (see Table III). We regress future values of news stories, that is, news stories after 12 or 18 months, on news stories in previous months, that is. six or 12 months. In columns 1 and 2 of Table III, we find little statistically significant evidence of persistence. Specifically, news after six months is negatively associated with news after 18 months, but news after 12 months is strongly associated with news after 18 months. This suggests that there are important intertemporal dynamics in media coverage, but the dynamics are very noisy and must be treated with caution.

Finally, columns 3 and 4 of Table III take the hyperbolic sine of the outcome and right-hand side variables and replicate the analysis, again allowing for events that trigger no news stories. Now, we find substantial economically and statistically significant effects. For example, a 1% increase in news stories after six months

Table II. Baseline determinants of cyber news salience

	Number of news stories					*Log (news stories)*				
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*	*(8)*	*(9)*	*(10)*
Computer network attacks/other operations	1759.35	869.80	705.34	514.92†	408.39†	0.84†	0.89†	0.87†	1.00*	0.85†
	[3768.45]	[1320.24]	[609.88]	[309.52]	[238.93]	[0.48]	[0.48]	[0.49]	[0.47]	[0.50]
United States		−419.07	−128.81	22.81	214.97		0.13	0.11	0.28	0.16
		[692.07]	[214.66]	[140.42]	[187.95]		[0.49]	[0.49]	[0.47]	[0.47]
G20 member		360.62	320.65	117.94	51.46		0.37	0.38	0.22	0.26
		[577.35]	[258.94]	[130.22]	[110.90]		[0.41]	[0.42]	[0.37]	[0.39]
Government hit			360.92	176.49	132.66			0.30	0.28	0.30
			[305.60]	[123.48]	[125.82]			[0.35]	[0.33]	[0.35]
Military hit			−830.07	−503.09*	−521.06*			−0.53	−0.62†	−0.66†
			[618.30]	[219.84]	[241.62]			[0.35]	[0.35]	[0.35]
Finance hit			−756.63	−500.76*	−514.93*			−1.34*	−1.14†	−1.07†
			[603.27]	[228.28]	[243.58]			[0.63]	[0.61]	[0.61]
Energy hit			453.49	420.21	605.65			0.49	0.85	0.91†
			[394.51]	[258.86]	[411.06]			[0.55]	[0.53]	[0.53]
Health hit			320.23	159.53	287.43			1.26	1.23	1.35
			[367.56]	[136.54]	[179.21]			[0.78]	[0.79]	[0.85]
Media hit			−564.59	−227.69	−227.82			−0.12	−0.23	−0.16
			[453.68]	[141.73]	[146.63]			[0.51]	[0.55]	[0.55]
Zero-day				359.18*	398.83*				1.65**	1.69**
				[165.88]	[200.42]				[0.43]	[0.43]
Year	−406.85	−175.79	−82.12	−39.90†	−45.42†	−0.26**	−0.25**	−0.25**	−0.23**	−0.24**
	[929.65]	[304.07]	[86.77]	[23.81]	[26.44]	[0.04]	[0.04]	[0.04]	[0.04]	[0.04]
Attack from China					−50.03					−0.01
					[132.44]					[0.41]
Attack from Russia					388.07					0.47
					[256.85]					[0.54]
Attack from North Korea					73.84					0.42
					[141.87]					[0.51]
Attack from Iran					84.18					0.80
					[164.57]					[0.50]
Attack from USA					374.87					0.15
					[310.24]					[1.53]
R-squared						0.146	0.154	0.187	0.250	0.262
Sample size	224	224	224	221	221	224	224	224	221	221

Sources: Authors. Reported are the coefficients associated with negative binomial regressions of the number of news stories (columns 1–6) and the log (operationalized with the hyperbolic sine to allow for values of zero) of the number of news stories (columns 6–10 to accommodate for zeros) on various characteristics of the cyberattack. Standard errors are heteroskedasticity-robust.

† $p < 0.1$; * $p < 0.05$; ** $p < 0.01$.

Table III. Evaluating the persistence of cyberattacks in the media

	News after 12 months	*News after 18 months*	*Log (news after 12 months)*	*Log (news after 18 months)*
	(1)	*(2)*	*(3)*	*(4)*
News after six months	175.51	-1,769.08		
	[933.53]	[7,863.96]		
News after 12 months		3,500.11		
		[15,112.12]		
Log (news after six months)			0.96**	0.09
			[0.01]	[0.07]
Log (news after 12 months)				0.87**
				[0.07]
Computer network attacks/other operations	33,570.88	229,225.84	-0.05	-0.02
	[154,033.12]	[831,571.62]	[0.8]	[0.07]
United States	52,606.10	452,255.83	-0.00	-0.00
	[253,020.80]	[1,753,519.81]	[0.07]	[0.07]
G20 member	8,629.14	32,942.69	-0.07	0.00
	[42,861.70]	[215,005.67]	[0.08]	[0.06]
Government hit	-29,026.64	-175,931.50	0.02	0.02
	[148,411.80]	[698,967.55]	[0.09]	[0.05]
Military hit	-66,985.91	-454,336.42	-0.18†	0.04
	[331,198.01]	[1,814,496.47]	[0.09]	[0.05]
Finance hit	-115,890.46	-994,803.40	0.06	-0.13
Energy hit	36,108.01	63,204.90	-0.06	-0.11
Health hit	31,388.79	139,687.14	-0.11	0.04
Media hit	-10,175.92	-537.94	-0.18	-0.04
Zero-day	67,914.00	584,175.61	-0.10	0.05
Year	-15,420.26	-115,097.02	-0.02	-0.2*
Attack from China	-12,797.84	-16,6251.62	-0.04	0.04
Attack from Russia	14,282.75	91,768.25	-0.13	0.07
Attack from North Korea	4,965.56	23,753.11	-0.11	0.11
Attack from Iran	42,232.23	306,366.43	-0.14	0.05
Attack from United States	-74,078.29	-689,916.07	-0.18	-0.12
R^2			0.958	0.984
sample size	221	217	221	217

Sources: Authors. Reported are the coefficients associated with negative binomial regressions of the number of news stories (columns 1 and 2) and log (operationalized with the hyperbolic sine to allow for values of zero) of the number of news stories (columns 3 and 4 to accommodate for zeros) on historical numbers of news stories and various characteristics of the cyberattack. Standard errors are heteroskedasticity-robust. † $p < 0.1$; * $p < 0.05$; ** $p < 0.01$.

is associated with a 0.96% increase in news stories after 12 months. Furthermore, a 1% increase in news stories after 12 months is associated with a 0.87% increase in news stories after 18 months. Again, we include a year time trend to purge variation coming from general increase in cyberattacks.

Importantly, we also have run diagnostics where we restrict the sample to events that have non-zero stories in any of the news variables. Doing so for the specification in column 4 produces a coefficient of -0.25 on stories after six months (p-value = 0.119) and 1.27 on stories after 12 months (p-value = 0.00). This is consistent with our expectation that there are important intertemporal dynamics, but that in the short run (i.e. six months), a story might pass and eventually come back and trigger more stories 6–12 months down the road.

In the Online appendix Tables A.I and A.II document additional diagnostics, where we focus on variation from the Global North versus the G20 and where we take a different estimation approach using a Poisson distribution to estimate our count model. Results are robust.

Discussion

Our statistical analysis found that cyber effect operations tend to receive more coverage compared to espionage, but results are not statistically significant. While this could be due to the sample size, we have sufficient balancing between espionage and operations (see Online Figure A.1) and, therefore, can rule out that the absence is from statistical imbalance in the share of attacks. Instead, we find that the use of novel techniques, specifically zero-day exploits, is a statistically and economically significant predictor of coverage quantity. We also found that the sector targeted by a CO correlates with media reporting quantity. A surprising finding was that operations targeting the military or financial sectors generate less coverage. Nonetheless, the predictive models explain limited variation in news coverage – indicating that COs are treated differently in the media than other forms of conflict – although once we add lagged values of media coverage, the R^2 in our models spikes substantially. That suggests that media coverage is persistent and momentum builds behind stories.

Although COs have become part and parcel of international politics, these findings suggest that media reporting still treats them as a curiosity item. The selection bias in favour of novel techniques our study shows corresponds to the 'gee-whiz' bias identified by Moeller (2006), and combined with the evident absence of a bias towards more intense effects, suggests that the media treats cyber threats qualitatively different from other types of threats where intensity is a key variable determining the quantity of coverage. These findings are surprising since if the media does contribute towards threat inflation and drives fears of cyber war, as many have argued, one would expect there to be a bias towards reporting on the most dramatic effects. Consequently, media reporting in response to attacks may not be the main driver of heightened threat perception and the resulting instability, although it remains a transmission mechanism.

However, this 'gee-whiz' bias could still explain prevailing fears among the public because it emphasizes the nature of the cyber threat as complex and uncontrollable – two key characteristics known to increase feeling of dread and perception of heightened risk (Slovic, 1987). Indeed, some scholars have long noted how cyber threat perception is likely intertwined with a growing sentiment of dread due to the perceived vulnerability of modern societies towards unknown and uncontrollable threats (Dunn-Cavelty, 2012; Dunn-Cavelty and Søby, 2020). If this bias does apply systematically, it is likely to distort public perception by promoting operations based on their curiosity value while dismissing or neglecting those with simple tools yet possibly far more significant impact. Consequently, one would expect public perception to equate the danger of a cyber threat with its level of novelty – or, in established jargon, its level of 'sophistication.' This perception may explain why policy-makers called the SolarWinds operation, a cyber espionage operation, an 'act of war' (Williams, 2020): according to Microsoft President Brad Smith, this was 'the largest and most sophisticated attack the world has ever seen' (Reuters, 2021).

Conclusion

Academic research and policy debates on CC widely rely on media reporting as a data source. Journalists in turn get much of their information from commercial threat intelligence reports, which often constitute the only publicly available source of information on CC. Prior research has shown that reporting by these firms is biased by their business interests, prioritizing high-profile threats and neglecting threats to weaker actors. Media reporting thus builds on a data source already subject to significant bias. Meanwhile, research in political communication has identified a set of distinct biases in media reporting on conflict. Building on this research, we hypothesized that cyber threat reporting is subject to a 'double bias,' where only a fraction of a fraction of activity gets reported on, distorting public perception and academic and policy debates. These bias matter because media reporting is likely a key influence shaping persistent cyberwar fears among the population – with significant implications for stability.

To test this theory, this study provided the first systematic analysis of the determinants of media attention to commercial threat reporting. We contributed an original dataset of commercial threat reporting on COs, coding each operation based on a set of characteristics likely to predict coverage quantity based on the media bias literature. Subsequently, we traced media reporting on each of these operations and estimate a series of multivariate regressions to identify the characteristics – motivated by theory – that correlate with reporting quantity.

We document surprising results. Many, if not at all, of the classical biases present in conflict and disaster reporting apply to cyber conflict reporting. Neither effects intensity, nor distance, nor attacker identity (as a perceived adversary) showed any statistically significant correlation with reporting quantity on a given CO. Rather, only the presence of easily observable indicators of sophistication, namely zero-days, correlated significantly with reporting quantity. These results challenge our expectations about biases in cyber threat reporting and their impact on public threat perception being subject to similar biases as reporting on conventional conflict and disaster. Instead, these findings indicate that cyber threat reporting exhibits a distinct 'gee-whiz' bias towards technological novelty. Building on research on risk perception, we showed that this type of bias can still drive heightened threat perceptions by emphasizing the nature of cyber threats as technologically novel, complex, and uncontrollable.

These surprising findings may thus help explain not only the persistently high threat perception among the public, but recent survey results show perception of cyber threats as more existential than any other types of threats. This result is in line with the perception of cyber threats as belonging to a distinct category that we would expect to prevail based on our findings. Consequently, media reporting biased in favour of cyber threats with higher technological sophistication (based on easily observable indicators) likely drives public perception of such threats as exceptional, justifying exceptional responses. Even absent evidence of the damaging effects of COs in practice, such bias would continue to feed associated perceptions of uncontrollable risks – thus explaining why public fears have not only persisted, but increased, despite a lack of tangible evidence of cyber dangers to the public. Since recent research has linked such threat perception to greater willingness to retaliate (Shandler et al., 2022), this type of media bias may thus directly contribute towards instability and escalation risks in CC.

These points indicate several avenues for future research. First, to better assess the hypothesis that news coverage on cyber threats is qualitatively different from coverage of other threats, research comparing news reporting on cyber threats and conventional security threats with similar characteristics will be highly valuable. Specifically, it will be interesting to compare coverage of cyber threats to coverage of other novel technological threats.

Second, more case study research is needed to trace the causal mechanisms behind the correlations between the characteristics of a CO and reporting volume we have identified. For example, the initial findings of this research potentially explain why the COs of the Mask received so little attention but Sandworm's Olympic Destroyer was widely discussed. Olympic Destroyer caused clear visible effects, whereas the Mask was only involved in espionage operations. Also, while the Mask was seen as highly advanced, it did not use any zero-days – that is, it did not rely on any easily observable indicators of sophistication that reporters can easily pick up on. Future case study research can assess these findings more systematically, tracing how reporting from one commercial threat intelligence company ultimately ends up in certain media articles. Specifically, it will be useful to examine in detail how the language used to describe COs in threat reporting shapes media reporting, both in quality and in quantity.

Third, and possibly most importantly, foundational qualitative and quantitative research on how the impact of COs shapes news reporting is urgently needed. Intuitively, one would expect the actual damage caused by COs to be the key determinant of news coverage quantity and quality. Yet there is little data available on the impact of the operations examined, and a lack of systematic framework to assess the impact of COs in general. The reasons are relatively obvious, namely the secrecy that surrounds such operations and the incentives of victims to hide both their identity as well as the true scope and scale of impact. Perhaps the uncertainty resulting from the lack of information about what COs do is what sets them apart from other threats – and why reporting privileges the novelty aspect over other characteristics.

Replication data

The dataset, codebook, and do-files for the empirical analysis in this article, along with an Online appendix, can be found at https://www.prio.org/journals/jpr/replicationdata.

Acknowledgements

The authors thank Mariya Alexandra Green and Tobias Pulver for their excellent work as research assistants.

Funding

The author received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Lennart Maschmeyer https://orcid.org/0000-0003-4666-2387
Max Smeets https://orcid.org/0000-0003-4057-6445

Notes

1. It thus marked the first time a commercial threat company potentially disclosed advanced cyber operations of a state in continental Europe. (Kaspersky, 2022).
2. The first assessment by McAfee that a North-Korean threat actor called Lazarus group was behind the attack turned out to be incorrect (Greenberg, 2018).

3. This is based on a search in Lexis Nexis, including newspaper articles until December 2021.
4. The original source of the term was JP 1-02 (United States Department of Defense dictionary). While the terminology was later replaced by JP 3-12, the terms computer network exploitation and computer network attacks are still widely used.
5. To clarify, we exclusively focus on geographical location. We are not examining attribution.
6. For previous data collection efforts and analyses on cyber operations (COs) see: Akoto, 2021; Council on Foreign Relations, 2023; EuRepoC, 2023; Valeriano and Maness, 2015. We were not able to use these datasets as they draw upon both threat intelligence reports and media articles, which would introduce biases in our data (because it is covered in the media, it is then also more likely listed as an operation). Instead, we draw on two separate sources: the information about COs only comes from threat intelligence reports; and the data about news coverage comes from Lexis Nexis (as discussed below).
7. Sometimes the private companies use the same name for the threat actor and the operation, complicating data collection efforts.
8. Although since the ThaiCERT cards are coded on the actor-level, and because their coding might sometimes be flawed, we still consulted the relevant threat intelligence reports themselves.
9. For similar measurement only looking at the *New York Times* see Kaitlyn, Metzgar and Rouse (2013). Categorical variables are frequently used, for example, Collins and Cooper (2012).
10. A minority of studies uses ProQuest, Factiva or Google News Search.
11. We use the United Nations Finance Center for South–South Cooperation list of 77 countries and China, the most widely used and generally accepted list.

References

Aitel D (2016) CyberSecPolitics: Useful fundamental metrics for cyber power. *CyberSecPolitics*, 22 June. Available at: https://cybersecpolitics.blogspot.com/2016/06/useful-fundamental-metrics-for-cyber.html (accessed 12 April 2022).

Akoto W (2021) International trade and cyber conflict: Decomposing the effect of trade on state-sponsored cyber-attacks. *Journal of Peace Research* 58(5): 1083–1097.

Akoto W (2024) Who spies on whom? Unraveling the puzzle of state sponsored cyber economic espionage. *Journal of Peace Research* 61(1): XX–XX.

Alterman E (2003) *What Liberal Media?: The Truth About Bias and the News*. New York: Basic Books.

Arsenault A, Kreps S, Snider K, et al. (forthcoming) Cyber scares and prophylactic policies: Cross-national evidence on the effect of cyberattacks on public support for surveillance. *Journal of Peace Research*.

Baron DP (2006) Persistent media bias. *Journal of Public Economics* 90(1): 1–36.

Bejtlich R (2020) Greg Rattray invented the term advanced persistent threat. *TaoSecurity Blog*, 10 October. Available at: https://taosecurity.blogspot.com/2020/10/greg-rattray-invented-term-advanced.html (accessed 8 July 2022).

Berlemann M and Thomas T (2019) The distance bias in natural disaster reporting – Empirical evidence for the United States. *Applied Economics Letters* 26(12): 1026–1032.

Buchanan B (2017) The legend of sophistication in cyber operations. *Belfer Center for Science and International Affairs*. Available at: https://www.belfercenter.org/publication/legend-sophistication-cyber-operations (accessed 18 November 2019).

Buchanan B (2020) *The Hacker and the State: Cyber Attacks and the New Normal of Geopolitics*. Cambridge, MA: Harvard University Press.

Buchanan B and Cunningham FS (2020) Preparing the cyber battlefield: Assessing a novel escalation risk in a Sino-American crisis. *Texas National Security Review* 3(4): 54–81.

Clarke RA and Knake RK (2010) *Cyber War: The Next Threat to National Security and What to Do about It*. 1st edition. New York: Ecco.

Collins TA and Cooper CA (2012) Case salience and media coverage of supreme court decisions: toward a new measure. *Political Research Quarterly* 65(2): 396–407.

Council on Foreign Relations (2023) Tracking state-sponsored cyberattacks around the world. Available at: https://www.cfr.org/cyber-operations (accessed 11 May 2023).

Donohue B (2014) The mask – Unveiling the world's most sophisticated APT campaign. *Kaspersky Daily*, 11 February. Available at: https://www.kaspersky.co.uk/blog/the-mask-unveiling-the-worlds-most-sophisticated-apt-campaign/3114/ (accessed December 2023).

Dunn-Cavelty M (2008) Cyber-terror—Looming threat or phantom menace? The framing of the US cyber-threat debate. *Journal of Information Technology & Politics* 4(1): 19–36.

Dunn-Cavelty M (2012) The militarisation of cyberspace: Why less may be better. In: *2012 4th International Conference on Cyber Conflict (CYCON 2012)*. Tallinn: NATO CCD COE, pp. 1–13.

Dunn-Cavelty M (2013) From cyber-bombs to political fallout: Threat representations with an impact in the cyber-security discourse. *International Studies Review* 15(1): 105–122.

Dunn-Cavelty M and Søby KK (2020) *Securing 'the Homeland': Critical Infrastructure, Risk and (in)Security*. London: Routledge.

Egloff FJ (2020) Contested public attributions of cyber incidents and the role of academia. *Contemporary Security Policy* 41(1): 55–81.

Ellman M and Germano F (2009) What do the papers sell? A model of advertising and media bias. *Economic Journal* 119(537): 680–704.

EuRepoC (2023) European Repository on Cyber Incidents. Available at: https://eurepoc.eu/databases (accessed 11 May 2023).

Freudenburg WR, Coleman C-L, Gonzales J and Helgeland C (1996) Media coverage of hazard events: Analyzing the assumptions. *Risk Analysis* 16(1): 31–42.

Gomez MA (2019) Past behavior and future judgements: Seizing and freezing in response to cyber operations. *Journal of Cybersecurity* 5(1): tyz012.

Gomez MA and Villar EB (2018) Fear, uncertainty, and dread: Cognitive heuristics and cyber threats. *Politics and Governance* 6(2): 61–72.

Greenberg A (2018) North Korea's Olympic diplomacy hasn't stopped its hacking. *Wired*, 15 February. Available at: https://www.wired.com/story/north-korea-olympics-diplomacy-south-korea-hacking/#:∼:text=He%20argues%20that%20the%20Kim,it%20safe%20from%20Western%20invasion (accessed December 2023).

Groseclose T and Milyo J (2005) A measure of media bias. *Quarterly Journal of Economics* 120(4): 1191–1237.

Gross ML, Canetti D, and Vashdi DR (2017) Cyberterrorism: Its effects on psychological well-being, public confidence and political attitudes. *Journal of Cybersecurity* 3(1): 49–58.

Haider-Markel DP, Allen MD and Johansen M (2006) Understanding variations in media coverage of US Supreme Court decisions: Comparing media outlets in their coverage of Lawrence v. Texas. *Harvard International Journal of Press/Politics* 11(2): 64–85.

Harknett RJ and Smeets M (2022) Cyber campaigns and strategic outcomes. *Journal of Strategic Studies* 45(4): 534–567.

Healey J (2016) The US government and zero-day vulnerabilities: From pre-heartbleed to shadow brokers. *Journal of International Affairs* 1: 1–20.

Innis HA (1951) *The Bias of Communication*. Toronto: University of Toronto Press.

Jardine E, Porter N and Shandler R (2024) Cyberattacks and public opinion – The effect of uncertainty in guiding preferences. *Journal of Peace Research* 61(1): XX–XX.

Jervis R (2017) *Perception and Misperception in International Politics: New Edition*. Princeton, NJ: Princeton University Press.

Joseph MF and Poznansky M (2018) Media technology, covert action, and the politics of exposure. *Journal of Peace Research* 55(3): 320–335.

Joyce R (2016) Disrupting nation state hackers. In: *USENIX*, San Francisco, CA, 27 January 2016. Available at: https://www.usenix.org/conference/enigma2016/conference-program/presentation/joyce (accessed 11 May 2023).

Kaspersky (2022) Animal Farm is an advanced threat actor with its own toolkit. Available at: https://apt.securelist.com/apt/animal-farm%20 (accessed 8 July 2022).

Kearns EM, Betus AE and Lemieux AF (2019) Why do some terrorist attacks receive more media attention than others? *Justice Quarterly* 36(6): 985–1022.

Kello L (2013) The meaning of the cyber revolution: Perils to theory and statecraft. *International Security* 38(2): 7–40.

Lawson S (2013) Beyond cyber-doom: Assessing the limits of hypothetical scenarios in the framing of cyber-threats. *Journal of Information Technology & Politics* 10(1): 86–103.

Lawson S and Middleton MK (2019) Cyber pearl harbor: Analogy, fear, and the framing of cyber security threats in the United States, 1991–2016. *First Monday* 24(3–4): DOI: https://doi.org/10.5210/fm.v24i3.9623.

Leyden J (2014) Kaspersky rips the mask from sneaky Spanish spy campaign. *Register*, 11 February. Available at: https://www.theregister.com/2014/02/11/mask_cyberspy_campaign/ (accessed December 2023).

Lindsay JR (2013) Stuxnet and the limits of cyber warfare. *Security Studies* 22(3): 365–404.

Lockheed Martin (2015) Gaining the advantage applying cyber kill chain® methodology to network defense. Available at: https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/Gaining_the_Advantage_Cyber_Kill_Chain.pdf (accessed December 2023).

Lombard M, Snyder-Duch J, and Bracken CC (2002) Content analysis in mass communication: Assessment and reporting of intercoder reliability. *Human Communication Research* 28(4): 587–604.

Makridis CA (2021) Do data breaches damage reputation? Evidence from 45 companies between 2002 and 2018. *Journal of Cybersecurity* 7(1): tyab021.

Maschmeyer L (2021) The subversive trilemma: Why cyber operations fall short of expectations. *International Security* 46(2): 51–90.

Maschmeyer L, Deibert RJ and Lindsay JR (2021) A tale of two cybers – How threat reporting by cybersecurity firms systematically underrepresents threats to civil society. *Journal of Information Technology & Politics* 18(1): 1–20.

McCarthy JD, McPhail C and Smith J (1996) Images of protest: Dimensions of selection bias in media coverage of Washington demonstrations, 1982 and 1991. *American Sociological Review* 61(3): 478–499.

Menn J and Finkle J (2014) Researchers uncover cyber spying campaign dubbed 'The Mask'. *Reuters*, 10 February. Available at: https://www.reuters.com/article/cybersecurity-espionage-mask/researchers-uncover-cyber-spying-campaign-dubbed-the-mask-idUKL2N0LF0KM20140210/ (accessed December 2023).

Meyer P (2021) Explaining media coverage of constitutional court decisions in Germany: The role of case characteristics. *Political Communication* 38(4): 426–446.

Miller RA and Albert K (2015) If it leads, it bleeds (and if it bleeds, it leads): Media coverage and fatalities in militarized interstate disputes. *Political Communication* 32(1): 61–82.

Moeller SD (2006) 'Regarding the pain of others': Media bias and the coverage of international disasters. *Journal of International Affairs* 59(2): 173–196.

Mueller C (1997) International press coverage of East German protest events, 1989. *American Sociological Review* 62(5): 820–832.

Nacos BL (2016) *Mass-Mediated Terrorism: Mainstream and Digital Media in Terrorism and Counterterrorism*. Third edition. Lanham: Rowman & Littlefield.

Neuendorf KA (2017) *The Content Analysis Guidebook*. Second edition. Los Angeles, CA: Sage.

Niven D (1999) Partisan bias in the media? A new test. *Social Science Quarterly* 80(4): 847–857.

Reuters (2021) SolarWinds hack was 'largest and most sophisticated attack' ever: Microsoft president. *Reuters*, 15 February. Available at: https://www.reuters.com/business/media-telecom/solarwinds-hack-was-largest-most-sophisticated-attack-ever-microsoft-president-2021-02-16/#:∼:text=WASHINGTON%2C%20Feb%2014%20(Reuters),O)%20President%20Brad%20Smith%20said (accessed December 2023).

Riffe D, Lacy S and Watson BR (2019) *Analyzing Media Messages: Using Quantitative Content Analysis in Research*. Fourth edition. Routledge communication series. New York, London: Routledge.

Roth F and Stirparo P, Bizeul D, et al. (2015) APT groups and operations. Available at: https://docs.google.com/spreadsheets/d/1H9_xaxQHpWaa4O_Son4Gx0YOIzlcBWMsdvePFX68EKU/edit?usp=embed_facebook (accessed December 2023).

Shandler R and Canetti D (2024) Special issue introduction: Cyber-conflict - Moving from speculation to investigation. *Journal of Peace Research* 61(1): XX–XX.

Shandler R, Gross ML, Backhaus S, et al. (2022) Cyber terrorism and public support for retaliation – A multi-country survey experiment. *British Journal of Political Science* 52(2): 850–868.

Shandler R, Gross ML and Canetti D (2023) Cyberattacks, psychological distress, and military escalation: An internal meta-analysis. *Journal of Global Security Studies* 8(1): ogac042.

Shoemaker PJ and Cohen AA (2006) *News around the World: Content, Practitioners, and the Public*. New York: Routledge.

Slovic P (1987) Perception of risk. *Science* 236(4799): 280–285.

Smeets M (2022) *No Shortcuts: Why States Struggle to Develop a Military Cyber-Force*. London: Hurst.

Snider KLG, Shandler R, Zandani S, et al. (2021) Cyberattacks, cyber threats, and attitudes toward cybersecurity policies. *Journal of Cybersecurity* 7(1): tyab019.

Snyder D and Kelly WR (1977) Conflict intensity, media sensitivity and the validity of newspaper data. *American Sociological Review* 42(1): 105–123.

Spring T (2018) Olympic destroyer: A false flag confusion bomb. *Threat Post*, 8 March. Available at: https://threatpost.com/olympic-destroyer-a-false-flag-confusion-bomb/130262/ (accessed December 2023).

Valeriano B and Maness R (2015) *Cyber War versus Cyber Realities: Cyber Conflict in the International System*. Oxford, New York: Oxford University Press.

Vicic J and Gartzke E (2024) Cyber-enabled influence operations as a 'center of gravity' in cyberconflict: The example of Russian foreign interference in the 2016 US federal election. *Journal of Peace Research* 61(1): XX–XX.

Visé D de (2023) Americans now fear cyberattack more than nuclear attack. *The Hill*, 7 April. Available at: https://thehill.com/policy/cybersecurity/3938210-americans-now-fear-cyberattack-more-than-nuclear-attack/ (accessed December 2023).

Williams J (2020) Durbin says alleged Russian hack 'virtually a declaration of war'. *The Hill*, 16 December. Available at: https://thehill.com/policy/cybersecurity/530461-durbin-says-alleged-russian-hack-virtually-a-declaration-of-war/ (accessed December 2023).

Work JD (2020) Evaluating commercial cyber intelligence activity. *International Journal of Intelligence and Counter-Intelligence* 33(2): 278–308.

Zetter K (2014) Hacker lexicon: What is a zero day? *Wired*, 11 November. Available at: https://www.wired.com/2014/11/what-is-a-zero-day/ (accessed December 2023).

Zetter K (2016) Hacker lexicon: What are CNE and CNA? *Wired*, 6 July. Available at: https://www.wired.com/2016/07/hacker-lexicon-cne-cna/ (accessed December 2023).

CHRISTOS A MAKRIDIS, b. 1991, PhD in Management Science & Engineering, and Economics (Stanford University, 2018); Managing Director, Center for Digital Finance and Technologies, Columbia University; Digital Fellow, Stanford University; Associate Research Professor, Arizona State University; Adjunct Associate Professor, University of Nicosia. Publications in a wide array of economics, finance and social science journals.

LENNART MASCHMEYER, b. 1984, PhD in Political Science (University of Toronto, 2020); Senior Researcher at Center for Security Studies, ETH Zürich (2019); co-chair of the Threat Intel Coalition SIG at the Forum for Incidence Response and Security Teams; forthcoming book: *Subversion – From Covert Action to Cyber Conflict* (Oxford University Press, 2024).

MAX SMEETS, b. 1990, DPhil in International Relations (University of Oxford, 2017); Senior Researcher, Center for Security Studies, ETH Zurich; and co-director, European Cyber Conflict Research Initiative. Most recent book: *No Shortcuts: Why States Struggle to Develop a Military Cyber-Force* (Oxford University Press & Hurst, 2022). Co-editor, *Deter, Disrupt or Deceive? Assessing Cyber Conflict as an Intelligence Contest* (Georgetown University Press, 2023) and *Cyberspace and Instability* (Edinburgh University Press, 2023).