Can IOCs Impose Cost? The Effects of Publishing Threat Intelligence on Adversary Behavior

Bouwman, X.B.; Ethembabaoglu, A.M.; Hermans, B.; Hernandez Ganan, C.; van Eeten, M.J.G.

DOI
10.1145/3719027.3765026

Publication date
2025

Document Version
Accepted author manuscript

Published in
CCS 2025 - Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security

Can IOCs Impose Cost? The Effects of Publishing Threat Intelligence on Adversary Behavior

Xander Bouwman
Delft University of Technology

Aksel Ethembabaoglu
Delft University of Technology

Bart Hermans
Delft University of Technology

Carlos Gañán
Delft University of Technology

Michel van Eeten
Delft University of Technology

Abstract

Exposing intrusion campaigns has become a geopolitical tool, with governments and commercial firms publishing threat intelligence reports about hacking attempts and modus operandi. U.S. government officials have explained this as not just a defensive practice but also as a way to 'impose cost' on attackers by forcing them to develop new infrastructure, tools, and techniques, consuming their scarce resources. We empirically examine this claim by analyzing attacker behavior before and after the publication of indicators of compromise (IOCs). Using IOC feeds from two leading commercial providers – deemed to best enable detection of sophisticated threats – we matched IOCs against a large dataset of real-world network traffic metadata. This enabled us to generate sightings retroactively, capturing malicious activity up to 150 days before and after publication. Unlike prior work focused on post-publication malicious activity, our method provides a more complete view over time. Our results show that most IOCs point to resources that attackers had already abandoned by the time of IOC publication, limiting their utility for detecting ongoing attacks and undermining the idea of 'imposing costs'. Statistical modeling further reveals that publication status has low explanatory power for sightings, suggesting that confounding variables exist. We also observed a 30-day delay between the peak of threat actor activity and IOC publication for one provider. This study is the first empirical assessment linking threat intelligence publication to attacker behavior, bridging computer science and international relations.

CCS Concepts

• **Security and privacy → Economics of security and privacy**; *Intrusion detection systems*; *Social aspects of security and privacy.*

Keywords

Threat Intelligence; Intrusion Detection; Cyber Conflict Studies

ACM Reference Format:
Xander Bouwman, Aksel Ethembabaoglu, Bart Hermans, Carlos Gañán, and Michel van Eeten. 2025. Can IOCs Impose Cost? The Effects of Publishing Threat Intelligence on Adversary Behavior. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25), October 13–17, 2025, Taipei, Taiwan.* ACM, New York, NY, USA, 21 pages. https://doi.org/10.1145/3719027.3765026

CCS '25, Taipei, Taiwan

ACM ISBN 979-8-4007-1525-9/2025/10
https://doi.org/10.1145/3719027.3765026

1 Introduction

The US Department of Defense (DoD) strategy on countering state-sponsored espionage has been heavily influenced by cyber persistence theory, representing a shift towards a more active stance on cybersecurity [17, 20]. A key element of this more active stance is to frustrate adversaries' operations, for example, by disrupting their infrastructure, publicizing modus operandi, or indicting operators [4, 26, 33, 34]. The concept of 'imposing costs' on threat actors is now central to the US theory of victory for cyber operations: the term occurs 27 times in the 2020 report of the Cyberspace Solarium Committee on cybersecurity policy reform [50] and a 2023 White House strategy document called on the US to "execute disruption campaigns at scale" [21].

Publishing threat intelligence is one method of imposing costs. US Cyber Command has been declassifying intelligence about threat actors on VirusTotal since 2018 and later directly to its partners, aiming "to impose costs on adversary nation-state [malicious cyber actors] and increase the resiliency of vulnerable networks" [24, 32]. Its former commander, General Paul Nakasone, has voiced the ambition to provide "insights to domestic and foreign partners to mitigate and respond to malign activity" [37]. The idea that better intelligence leads to better detection and incurs 'costs' on attackers is also embodied in the foundational 'pyramid of pain' model for threat intelligence [6]. And indirectly, this idea is reflected by the value of the commercial threat intelligence industry, now worth hundreds of billions of dollars. One firm alone, Crowdstrike, had a market capitalization of over $90B in April 2025 [14]. In short, having access to threat intelligence is thought to improve intrusion detection, thereby forcing the adversary to adapt tactics. Because this occupies their scarce resources, high-quality threat intelligence is believed to lead to a strategic advantage for defenders and perhaps even be able to ultimately reduce the willingness of the adversary to attempt future operations [20]. We refer to this as the 'impose costs' mechanism. If this mechanism works as implied, it should, in principle, be observable in network traffic that adversaries abandon resources after those have been published as threat intelligence.

In this paper, we present the first attempt to measure an observable relationship between publishing IOCs and adversary behavior. We employed IOCs from two market-leading threat intelligence vendors, which represent the best view on the resources of advanced persistent threats (APTs). Although US Cyber Command also published IOCs, the size of this set was too small to enable large-scale analysis. We matched the IOCs against a unique dataset of 12 months of historical network traffic metadata from government enterprise networks, containing IP addresses, domains, and hashes for all traffic that passed through the intrusion detection systems (IDS) in the networks. The resulting dataset of 1 billion matches – which we call 'hits' or 'sightings' of an IOC – allowed us to analyze when threat actors were using the resources in the period *before and after* the IOCs were published.

Earlier research with a similar methodology has focused on malicious activity *after* IOCs publication and could therefore not provide systematic insights into attacker behavior over time, nor how discovery and publication of the IOC affects attacker behavior (section 2). With our approach, we could test if publication is related to sightings – where fewer sightings would indicate that attackers were moving away from the resource once it was published, in line with the 'impose cost' mechanism.

Since our telemetry is generated by large government networks and high-end threat intelligence, our findings are indicative of APT behavior, but it cannot capture the full range of APTs. We assume these networks are high-value targets for espionage, rather than for financially-motivated attacks. Hence, we would expect state-affiliated threat actors to dominate the attack patterns we discovered. Large organizations in other sectors, like the financial industry, might be more attractive to cybercriminal actors – as would be small and medium-sized businesses. Thus, the patterns we observe might be different in those contexts.

We aim to answer the following main research question: *Do threat actors abandon resources after these have been published as indicators of compromise in commercial threat intelligence?* Our analysis bridges the empirical gap that exists between intrusion detection and the strategic notion of 'imposing costs' on threat actors. We make the following contributions:

- We provide the first systematic study of threat actor activity on IOCs before and after their publication. For IOCs to 'impose cost' they must point defenders to resources that attackers are actively using at publication time. We find that for 81% of IOCs that caused sightings, attackers had already stopped using the resources by the time that IOCs had been published to the customer. This challenges the assumption that IOCs could 'impose cost' by enabling intrusion detection (section 4).
- Through statistical modeling with panel regression analysis, we explore factors that determined when IOCs caused sightings. We found a credible negative relationship with our primary variable of interest: whether the IOC had been published or not. For the same IOC, the likelihood of a sighting occurring on a given day decreased by 3.2 percentage points after publication, holding all other factors constant using fixed effects. This statistically significant drop in sightings provides empirical support for the 'impose costs' hypothesis, although the effect is modest, with the model explaining just 8.2% of observed variance. We analyzed various covariates but did not find other strong explanatory factors, pointing to the existence of confounding variables (section 5).
- We provide two interpretations of our data. Most threat intelligence IOCs we studied captured attacker behavior that had already ended, making it unlikely that they imposed direct costs on adversaries. For 81% of the IOCs that led to sightings, the resources they pointed to were seemingly no longer being used by attackers when customers received the IOC. Another way to read our findings is that second-order effects occurred. For example, it is possible that the broader presence of threat intelligence industry imposes indirect or anticipatory costs – e.g., by pushing threat actors to adopt stronger operational security practices (section 6).
- We highlight a potential issue with timeliness of commercial threat intelligence. For one of the market-leading providers we observed a consistent 30-day lag between peak attacker activity and IOC publication. This delay suggests that even certain high-end providers feeds may not be able to enable real-time detection and raises questions about their effectiveness as instruments for imposing costs on adversaries. However, this finding varied between vendors, with our second vendor showing different timing patterns (subsection 4.2).

Our findings establish new data about the relationship between publishing information and threat actor behavior. While our results have practical implications for network defenders, they also raise important questions for policymakers relying on threat intelligence sharing as a mechanism to impose costs on adversaries (section 6). This article is structured as follows. After we explain our data and method (§3), we explore IOCs and sightings with descriptive statistics and measurements (§4). We then perform statistical modeling of the relationship between IOC publication and sightings (§5) and answer our research question.

2 Related work

Our measurement study is positioned between two fields: international relations and intrusion detection. To the former, we contribute knowledge about the hypothesized 'impose cost mechanism', while to the latter, we contribute insights into the ability of commercial threat intelligence IOCs to enable detection.

2.1 Strategic studies

From a close reading of policy documents, speeches of key decision-makers, and other primary sources, Healey [20] reconstructed how the persistent engagement strategy is thought to lead to stability, i.e., more overall security. One step in the implied chain of causality entails for "the USA to observe adversary behavior and warn targets of the details of coming (or ongoing) attacks, improving US defense". Doing so repeatedly is thought to reduce the ability of adversaries to attack and, eventually, also their willingness to attack [20]. But what evidence is this based on?

International relations scholars have studied the effects of public exposure of cyber operations on threat actors. However, this research has been primarily qualitative, e.g., based on case studies of indictments. Buchanan summarizes that "in general, hackers from democratic governments seem to fear exposure the most." [11] A Columbia University student report – not peer-reviewed – investigated "the impact of leaks and information disclosures on adversary operations" with interviews and a literature review. Based on this qualitative work, the students concluded that public disclosures did not cause any of the studied threat actors to cease operations altogether, although they may have compelled some actors to become more sophisticated [3].

Strategic studies often take a broader qualitative approach than our measurement-based study. Authors have pointed out that infrastructure is just one resource that an adversary can be denied use of, with other examples being software vulnerabilities and competent staff, which is why the US has pursued criminal charges against foreign operators since 2013 [34]. Even just signaling capabilities and credibility have been said to be able to impose costs [50], but whether or not such steps, or even reciprocal attacks, could actually deter attackers from consecutive operations has been hotly debated over the years [44]. Our study provides new empirical data for future theoretical work on such questions.

2.2 Threat intelligence

Signature-based intrusion detection rests on the idea that once an adversary has been found in one network, he can be found in all networks [22]. Threat intelligence uses practices from traditional counterintelligence to capture adversary behavior in order to build such signatures and to inform other security controls [42]. In the technical cybersecurity community, 'imposing costs' appears primarily when discussing security controls, e.g., controls that introduce deception in defensive environments [5, 16].

Existing research has described threat intelligence IOCs from open sources and commercial sources and measured them along various quality dimensions [8, 19, 28, 52], also investigating the ability to enable intrusion detection and to do so in a timely fashion [2, 25, 49]. Some authors also measured the *sightings* that are generated from performing intrusion detection using threat intelligence. Vermeer et al. used data from a managed security service provider (MSSP) and found that its detection rules triggered most of the alerts in the first week after they were created [51].

Our study is unique in that it systematically measures threat intelligence sightings prior to IOC publication, comparing measurements before and after publication to examine the 'impose cost' mechanism. As far as we are aware, there exist only two other studies that included any threat intelligence sightings prior to IOC publication in their measurements, but neither attempted to measure those systematically, nor did they compare measurements before and after publication. Griffioen et al. measured the timeliness of open source blocklists [19], which capture a different kind of activity than commercial IOCs [8]. Tostes et al. performed survival analysis to optimize the 'shelf life' of IOCs, i.e., determine how long they generated relevant hits after they had been published [49]. These last authors found just 10% of their IOCs to have led to sightings ahead of publication, which is inconsistent with our findings. The authors do not go into specifics about the logs used to generate sightings and their retention period, although the focus of their work was on the period after publication.

No study provides empirical data about the relationship between publication of threat intelligence and behavior of targeted attackers. The most approximate work to ours can be found in studies of cybercrime, i.e., about threat actors who are primarily financially motivated [10, 48]. Relevant to our study is the recurring finding from this work that criminals adapted to takedown actions against their infrastructure – in this context, primarily IP addresses and domains – by rapidly replacing these resources. A comprehensive study into the spam ecosystem described how the constant blocklisting of spammy servers led criminals to adopt a tactic of rapidly cycling through infrastructure, using 66% of the abusive IPs for just a single day [47]. These findings about cybercrime actors are consistent with the 'pyramid of pain' model popular in threat intelligence [6], which characterizes denying infrastructure to attackers as inflicting little 'pain', as it is relatively simple to replace. Although the takedowns did not prevent actors from consecutive criminal activities, they did incur 'costs' on them in the sense of 'friction' as understood in persistent engagement.

Some measurements have been performed on the exploitation of certain software vulnerabilities prior to publication. Symantec researchers found in a 2012 study that threat actors exploited software vulnerabilities on average 8 months before they were publicly disclosed [7]. Such zero-day attacks are still the exception, however, and most attack campaigns described in threat intelligence reports make use of publicized software vulnerabilities [15].

3 Methodology

We take a quantitative approach to a theme that is normally discussed primarily using qualitative data, such as legal indictments. Our dataset consists of two high-end feeds with indicators of compromise from two firms leading in the threat intelligence industry. A partner organization matched these IOCs against their unique dataset consisting of a year of real-world network traffic metadata from multiple government enterprise networks. The resulting hits form our dataset of 'sightings' of the respective IOCs.

With this approach, we could retroactively identify if and when an IOC pointed to resources that threat actors were actively using, also before that IOC was published by the threat intelligence vendor. The impose cost mechanism implies that the publication of an IOC forces attackers to abandon the resource. We should be able to observe a reduction in sightings of the IOC post-publication. We fit a panel regression model to estimate whether this effect is visible in the data.

3.1 Positionality and ethics

The dataset analyzed in this study consists of sightings of commercial indicators of compromise (IOCs), generated by a partner organization using logs from multiple government enterprise networks. These networks are used by government personnel in a professional capacity. The partner organization requested us to not name them nor the threat intelligence vendors they use, so as to no publicly expose their security practices. The fact that logs are derived from government enterprise networks, rather than e.g. corporate, leads to a limitation in our ability to generalize our finding. We discuss this and other limitations in section 7.

At the time of the study, all authors were employed at a public university. They have expertise in cybersecurity measurement and policy. The partnership enabled access to unique data while maintaining the academic independence of the study. The partner performed the matching of threat intelligence indicators to historical network metadata. The research team did not access raw traffic content or personally identifiable information. The study design, analysis, and interpretation were led exclusively by the authors. No financial or editorial input was provided by the threat intelligence vendors whose feeds were analyzed.

The research team did not have access to any traffic content or personally identifiable information (PII). Each sighting consisted only of a timestamp, the triggering IOC, and a coarse log category (e.g., inbound or outbound connection). As outlined in section 3, this minimized data exposure and reduced risk to individuals. The study design and data handling protocol were reviewed and approved by the authors' Institutional Review Board (IRB).

In line with our data-sharing agreement with the partner organization, and to minimize reputational or economic harm, we deliberately avoided naming or identifying these firms. All references to providers are made in general descriptive terms, and any metadata or field names that could lead to their identification have been excluded.

There are theoretical risks associated with this study. Threat actors could use our findings to estimate typical delays between attacker activity and IOC publication, potentially improving their operational security and dwell time. However, this information is already indirectly accessible to them by monitoring CTI feeds against their infrastructure. We judge the overall risk to be limited. Conversely, the potential benefits of this work are substantial: it helps CTI consumers, researchers, and policymakers understand the empirical limits of current threat intelligence and evaluate the efficacy of the 'impose cost' doctrine in cyber operations.

The CTI providers included in this study were selected based on publicly available rankings of market leaders in the industry [18]. We have made good-faith efforts to hypothesize plausible internal explanations for our findings based on publicly available information about these firms. We have also disclosed our findings to the CTI providers involved and invited them to respond to a draft of this article. One firm responded but did not offer feedback, nor did they answer our questions about their services and internal processes. No changes were therefore made to the article based on vendor response.

The dataset did not contain personal data as defined by applicable privacy regulations. We believe this work serves the public interest by contributing evidence-based insights into a widely adopted cybersecurity strategy and the performance of a multi-billion-dollar threat intelligence industry. In line with open science principles, we recognize the importance of sharing research artifacts. However, due to the sensitivity of the dataset and the risk of re-identifying CTI providers through IOC publication dates or metadata fields, we are unable to release the dataset or associated code. This decision was made in accordance with our confidentiality agreement and ethical obligations to our research subjects.

3.2 Network traffic metadata

The partner organization runs a passive network monitoring system using well-known open-source tools on multiple government enterprise networks. This system generates logs of network traffic metadata that can be matched against IOCs, e.g., incoming and outgoing connections, URLs in HTTP sessions, and domains on TLS certificates. These logs spanned the period from 13-09-2023 to 02-09-2024 (356 days) which we refer to as our measurement period in the rest of this paper. The retention period of the logs was one year. This meant that the partner organization could match IOCs against the logs *retroactively* for one year. The resulting dataset consisted of the timestamp of each sighting, an identifier for the IOC that caused it, and log category that the sighting was found in (e.g., whether corresponding to an inbound or outbound connection). In other words, the IP addresses and other metadata from the enterprise network were not included. This data minimization was our way of reducing the potential for harm to users of the network.

3.3 Threat intelligence IOCs from two market leaders

Indicators of compromise (IOCs) enable intrusion detection by capturing artifacts of adversary behavior [22]. The IOCs in this study were produced by two firms described by Gartner as leading the threat intelligence industry [18]. The vendors provide feeds for automated ingestion of the IOCs. Some IOCs are provided alongside reports that provide more context on the associated attacks and advanced persistent threat (APT) actor, but we used all IOCs in the feed, regardless of whether they had been mentioned in a report. We had access to multiple years of the IOCs published by the two vendors, but as we will describe next, only IOCs published in our measurement period are used for collecting sightings.

Besides the date when a sighting took place, the publication date of the respective IOC was a key variable in our study. We found that the vendors sometimes re-published IOCs in their feeds. We took as publication date the moment an IOC was *first* included in one of the two vendor feeds. This is best aligned with our research question, as the moment of feed inclusion is when customers were first able to use the IOC for intrusion detection and also when the adversary was first able to learn that the resource had been detected.

In our measurement period of nearly a year, the two commercial vendors of cyber threat intelligence published a substantial number of IOCs in their feeds: 12 million and 224 million for Vendor 1 and Vendor 2, respectively. Where Vendor 1 provided all types of IOCs roughly evenly – including domain names, IP addresses, and malware binary hashes – the IOCs that Vendor 2 provided were in majority hashes. The vendors provided certain metadata with all IOCs, like the type of IOC (IP, domain, etc.). A fraction of the IOCs were also tagged by the vendors with metadata about the actor involved, which we use in subsection 4.3 and subsection 5.2.

3.4 Measuring sightings

We speak of a *hit* or *sighting* when an IOC appears in the network traffic metadata. This could be because of, e.g., an outgoing connection to the IP address of a phishing panel (IP IOC) or a malware binary being transferred unencrypted (file hash IOC). In a security monitoring context, such signals would be investigated in the Security Operations Center. This triage process is labor intensive and, even then, often remains unresolved [13]. In our research project, we had no insight into whether such investigations took place or what their outcomes were. This means we cannot verify whether a sighting actually represents an intrusion attempt or that it might be a false positive alert. Perhaps the IOC was incorrectly tied to a campaign by the threat intelligence vendor. Perhaps the IOC was correctly attributed, but at the moment when the sighting took place, it was no longer associated with the campaign. There is a

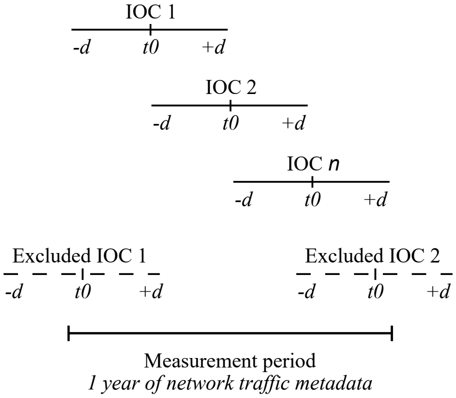

Figure 1: Illustration of our methodology. In order to reduce noise, we assumed sightings within a window of $d = 50$ before or after publication of the respective IOC to be relevant, so a total of 101 days (subsection 3.4). We performed sensitivity analysis for different values of d.

lively debate over how long threat intelligence IOCs lead to relevant sightings [49]. This period is likely variable over time and over different attackers.

Rather than attempting to triage sightings to determine if they were true positives, we can be sure that we are measuring signal because false positives would be uncorrelated with IOC publications dates. In other words, the false positives add random noise to the overall dataset. As long as a substantial portion of IOCs are true positives, then we should be able to see a difference between before and after IOC publication. This enables us to make meaningful assessments about the relationship between IOC publication and attacker behavior at a macro level. If we assume that the IOCs of the market-leading vendors are of high quality, then the bulk of them should be true positives. In this case, the overall pattern for the true positives would outweigh those for the false positives. We would still be able to see the effects of the publication on attacker behavior.

We took three further steps to reduce the noise in our dataset. First, to keep those sightings most likely to represent malicious activity, we used only IOCs labeled by the vendors themselves as being malicious with high confidence, eliminating roughly one-third of IOCs with lower confidence.

Second, we dropped sightings corresponding to 194 IOCs that had been labeled by GreyNoise as being 'benign scanners' in the time period in which the sighting occurred.

Third, we kept only the sightings closest to IOC publication date. We defined a parameter d that is the number of days before and after the publication date of an IOC – a kind of 'time-to-live' value for the IOC, except also *before* publication (see Figure 1). We only included sightings that fall inside that window. So, for a window of $d = 50$ days, we would include the sightings that occurred in the period of 101 days. We conduct sensitivity analysis to assess the effects of different values of d on the results. Where vendors had reuploaded IOCs multiple times in their feed (subsection 3.3) or the two vendors had uploaded the same IOC, we used the first publication date for this filtering process. This time-to-live window also means that we exclude IOCs that were published within $d = 50$ days of the start or end of our measurement period, since for those IOCs the network logs do not cover the full window, thus rendering them no longer comparable to the other IOCs. To illustrate: the first and last IOCs in Figure 1 are excluded. Simply put, if our experiment was repeated in a SOC-setting, and researchers had – over the course of nearly a year – generated sightings on logs of enterprise networks with a retention period of 50 days and also stopped generating sightings on IOCs 50 days after they were published, those researchers would get a dataset similar to ours. Except that we were able to experiment with optimal 'retention period' (value for d).

There is a lively debate over how long threat intelligence IOCs lead to relevant sightings [49]. This is the reason for including the parameter d. It allows us to conduct sensitivity analysis on if the size of the window has any impact on our findings, which we discuss in Section 4.2. The window you choose impacts what sightings you observe. On our networks we found a peak of activity around 30 days before publication, which window $d = 50$ properly captured, and let us answer our research question. A shorter window would have missed that. A longer window would come with more noise and with more IOCs being excluded because their window would fall partially outside of the measurement period. We performed sensitivity analysis and fitted our models for the four values of d, which did not lead to substantively different outcomes. We report the regression tables in Appendix B.

The steps outlined above generated the final dataset. To attempt to better capture signal with our measurements, we defined two cohorts of IOCs that consisted of domains: *new domains* and *established domains*, where we expected the first cohort to lead to more true positive sightings than the second. As figure Figure 3 shows, the cohort of new domains does indeed to higher numbers of sightings. Most domains fell somewhere in between and thus were not included in either cohort. The cohort of *established domains* were those domains: *i)* registered more than 50 days before publication of the IOC, as asserted by DomainTools, and at the same time *ii)* included in the Tranco 1M of popular domains at the start of our measurement period[1]. This led to a cohort of 57 domain IOCs that led to sightings. In our most commonly used window of $d = 50$ days, 29 of these remained. Vice-versa, we defined the cohort of *new domains* as those registered less than 50 days before IOC publication and that were not included in the Tranco 1M list. This cohort contained 156 domains, of which 92 remained in a window of $d = 50$. Cohort sighting counts are listed in Appendix A.

3.5 Statistical modeling

In section 5, we investigate the relationship between publication of IOCs and the corresponding sightings by fitting two types of panel data models. Panel data, also known as longitudinal data, consists of repeated observations of the same entities (in our case, IOCs) over time. This rich data structure allows us to control for unobserved heterogeneity. Both models take as the dependent variable if sightings occurred for an IOC on a given day, i.e., a binary variable. We did not use a threshold because for our research question, our goal was to measure if the resource was still being used by

[1] Tranco is a research-oriented top site ranking hardened against manipulation [27]. The list we used is available at https://tranco-list.eu/list/3VPNL/1000000

the threat actor at all, no matter to what degree. This choice for a binary dependent variable also addresses the outsize number of sightings that around 20% of IOCs lead to (see Figure 3). As we report in Appendix B, we also experimented with using another dependent variable based on normalized sighting counts as part of the sensitivity analysis. This did not lead to a different overall outcome.

First, we fitted a fixed effects model to address our research question directly. Fixed effects models are commonly used in panel data analysis to control for omitted variable bias. This let us control for time, meaning to eliminate bias arising from unobserved variables that were constant across IOCs but which evolved over time, such as global trends in the threat landscape or defender awareness. They allowed us to eliminate bias arising from unobserved variables that, conversely, varied across IOCs but did not change over time, like an IOC's type and detectability [45]. We operationalized this with the Python library `linearmodels.panel.PanelOLS`. We confirmed the that a fixed effects model was more suitable for our situation than random effects – which assume that the entity-specific intercepts are uncorrelated with the regressors – using the Wu-Hausman specification test for the different window sizes of d, finding that the unique error was uncorrelated with our regressors.

Second, simple ordinary least squares (OLS) regression in order to also shed light on between-group variation (which fixed effects does not allow). To shed light on between-group variation (which fixed effects does not allow), we also fitted ordinary least squares (OLS) models. These models allow us to investigate the impact of time-invariant covariates, such as IOC type and vendor, on sightings.

We used linear rather than logistic regression models in order to provide readily interpretable coefficients, something like to an effect size. Logistic regression would ensure more internally consistent estimates, at the cost of less intuitive interpretation. To be clear, neither approach would allow us to make causal claims – that would require randomized experiments where some IOCs are published and others are not (section 7).

4 Measurements of sightings

We perform various measurements to understand when IOCs lead to sightings, and discuss the implications for the ability of IOCs to impose cost. Unless otherwise noted, all figures are for sightings recorded using a window of $d = 50$ days before and after publication of the respective IOC.

4.1 Sightings of IOCs

Our approach generated a dataset with a *lot* of sightings: 1.04 billion to be precise. This is despite excluding sightings and IOCs that did not meet specific criteria, as explained in the previous section. Considering the work that SOC analysts perform triaging sightings on a daily basis, we can only acknowledge that alert fatigue must be a real and serious issue [2]. As Figure 2 shows, IP addresses from Vendor 1 were the largest category of IOCs to cause sightings. From the feed of 12 million IOCs, 248,000 led to sightings (2% of the total). This stands in sharp contrast to Vendor 2, where the feed contained 224 million IOCs, so an order of magnitude larger, yet only 893 of these led to sightings (0.0004%). Part of the explanation of this huge difference is that Vendor 1 provided many network-based IOCs, and that our dataset was generated from logs from a network monitoring system, whereas Vendor 2 provided more hashes, better suited for working with host-based detection. Malware hashes cause sightings in network monitoring logs only when the corresponding files are transferred unencrypted, which seldom happens anymore due to increased TLS usage. Figure 2 also shows the counts of

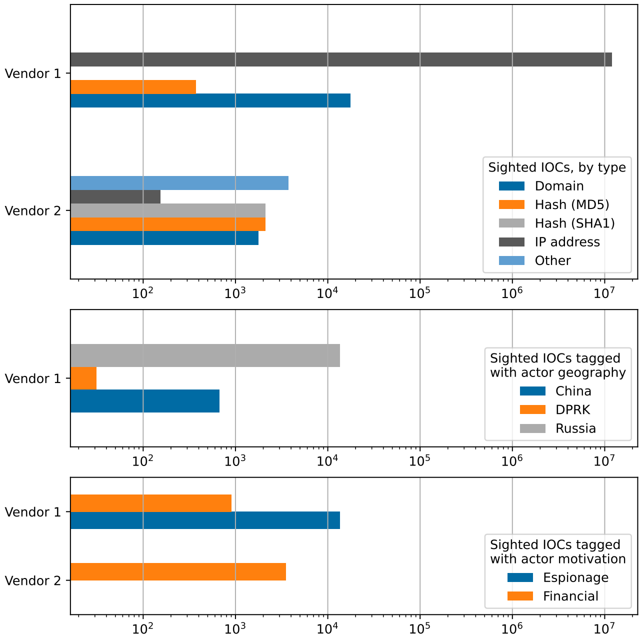

Figure 2: Counts of IOCs in two commercial feeds that led to sightings during our measurement period of about 1 year, *before* filtering for the most relevant sightings (subsection 3.4).

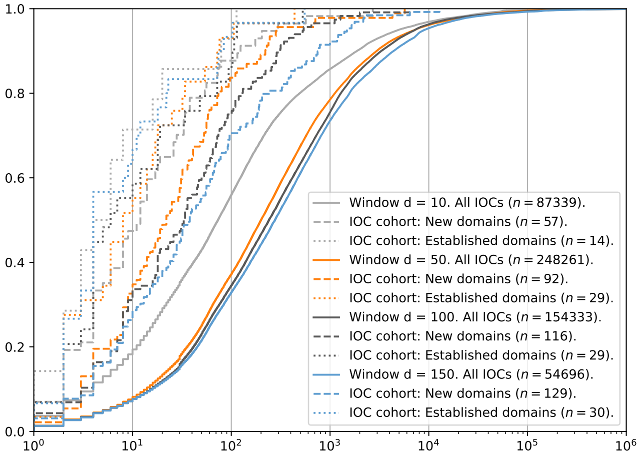

Figure 3: Empirical distribution function of sightings per IOC in our dataset, *after* filtering for the most relevant sightings (subsection 3.4). Around 20% of IOCs lead to more than 1,000 sightings. Our domain IOC cohorts receive fewer sightings per IOC, especially the established domains cohort.

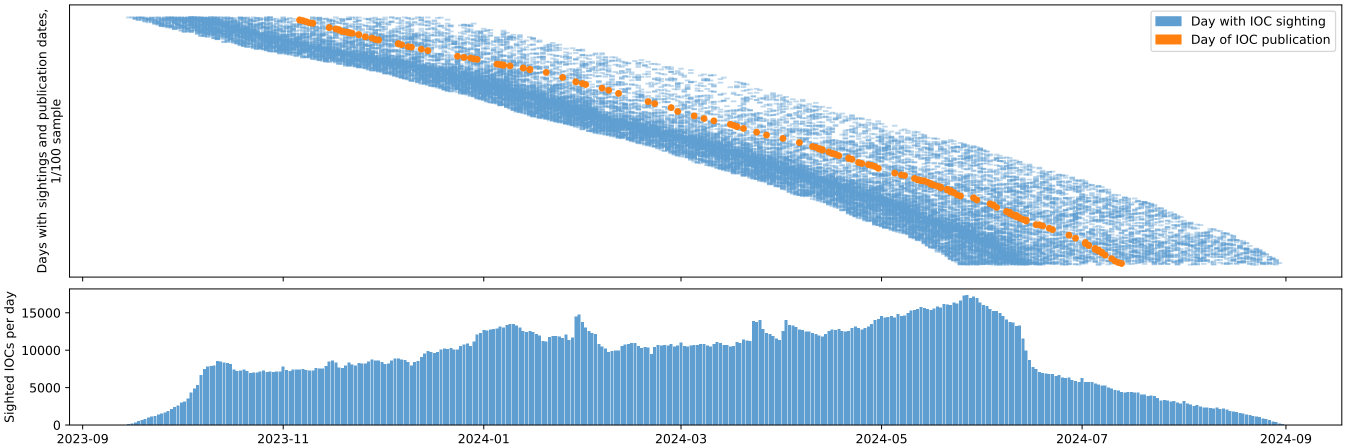

Figure 4: Sightings and IOC publication dates over our measurement period of 356 days, for window $d = 50$.

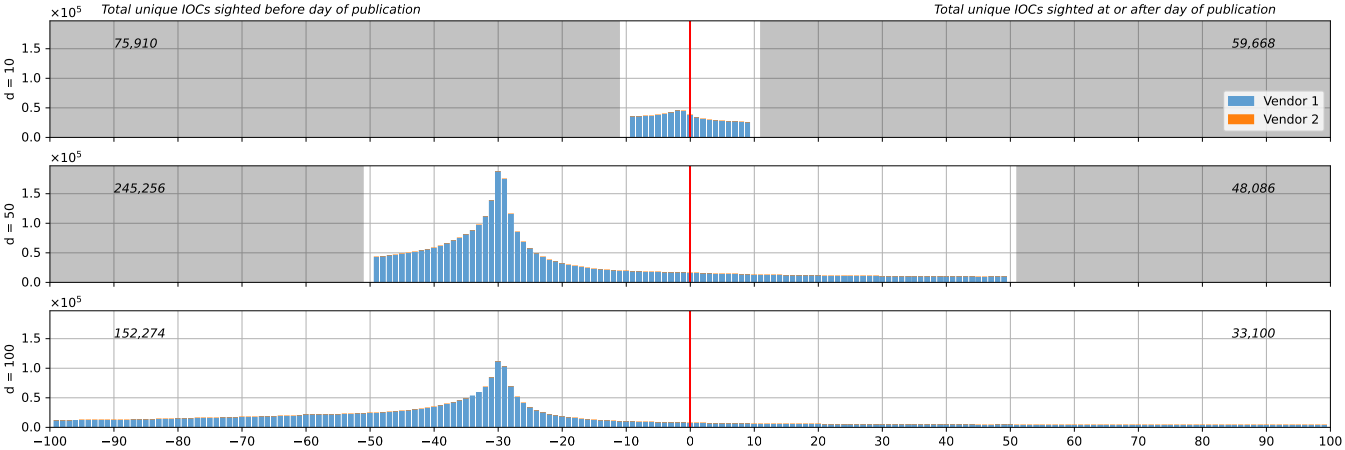

Figure 5: Counts of IOCs sighted daily relative to their respective publication dates, for three windows of d. A peak in sightings can be observed around 30 days before IOC publication for windows $d = 50$ and $d = 100$, as well as for $d = 150$ (not pictured). This pattern is not visible in $d = 10$, demonstrating that defenders need to retain network logs at least 30 days to observe this activity (see subsection 4.2 and section 6). Sightings from Vendor 2’s IOCs did not reach the visualization threshold.

IOCs that were tagged by the vendors with information about the related adversary *and* that we were able to map to a nation-state, making these numbers more of a function of the research process than representative data about threat actor activity. The number of IOCs tagged with these metadata is nevertheless relevant to place coefficients in subsection 5.2 into context.

Around 20% of IOCs caused disproportionately many sightings, as the curve beyond 0.8 in Figure 3 shows. This is not dependent on which window of d days we use to filter. This skewed distribution motivated the decision to binarize sightings – so whether an IOC had any sighting on a specific day, rather than use the absolute number of sightings – for use as a dependent variable in statistical modeling (section 5). Otherwise, the number of sightings of this 20% of IOCs would overwhelm the whole dataset. Also visible in the figure are our (small) domain IOC cohorts: new IOCs and established IOCs, as defined in subsection 3.4. Both cohorts lead on average to fewer sightings per IOC, which reflects that domains lead to fewer overall sightings (Table 3). New domains are noisier than established domains for all windows of d, which aligns with our expectation for these IOCs. In Figure 4, we plotted days with sightings over our measurement period and used a sample so that the variations within the figure would still be discernible. Most sightings occurred *before* IOC publication, as the density of blue dots around the orange dots shows. The figure also provides a histogram with a daily count of unique IOCs that led to sightings, showing normal variations over the measurement period with a slight overall increase over the measurement period. The slopes at the start and end of the histogram represent the sightings corresponding to increasing vs. decreasing numbers of IOCs, as a result of the number of IOCs meeting the criteria to be included in the dataset (see subsection 3.4). The reason that the slope at the end takes a different shape is that most sightings occurred before the IOC publication.

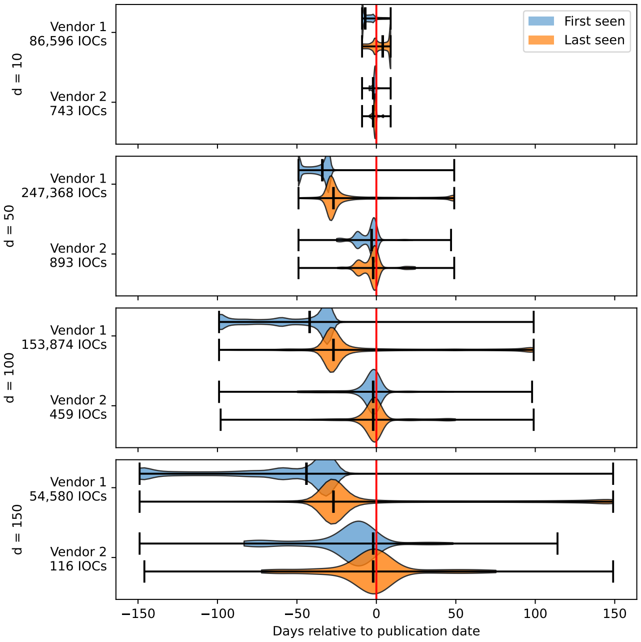

Figure 6: Distributions and medians of first and last sightings of IOCs, as days relative to the IOC's publication, per vendor. Note that plot area size between Vendor 1 and Vendor 2 is *not* proportional to the number of observations.

4.2 Timeliness of commercial IOCs

For IOCs to 'impose cost', they must point defenders to resources that attackers are actively using at publication time (section 1). This was true for just 19.0% of the IOCs in our dataset, i.e., out of the IOCs that led to sightings on the target networks, just 19.0% had any sightings occur after publication (Table 1). Only this portion could enable detection of potentially malicious activity that was ongoing at time of IOC publication. In other words, for 81% of those IOCs that pointed to relevant malicious activity, that activity had already stopped by the that time that IOCs became available to the customer. The inverse was true for a fraction of IOCs, 1.2% (2921 IOCs) which could be considered highly timely because they had *all* activity occur after their publication, and thus enabled consumers of the threat intelligence feeds to detect *all* the activity as it happened. In sum, despite significant timeliness differences between providers shown in Table 1, IOCs from both vendors demonstrated limited ability to detect ongoing threats at time of publication. This challenges their ability to 'impose cost'.

Figure 5 shows sightings on IOCs relative to their publication date, for window $d = 50$. The upper subplot shows a 1/100 sample of all sightings in this window in a 'calendar' heatmap. And the lower subplots provides the total daily sighted IOCs. The window of $d = 50$ days of relevant sightings that we included is visible in these plots as the slope at the start and end in the lower subplot (see subsection 3.4). The density of blue markers suggests a higher density before publication. Recall that our methodology ensures that IOCs have an equal probability of leading to sightings on each day within the windows (section 3).

We further investigate sightings relative to IOC publication date with Figure 6, which shows the first and last sighting of each IOC for the various windows of d. Table 2 provides the corresponding mean and median values. Although the overall timeliness of both providers was low, as we discussed, we observed large differences between the two vendors in how sightings were distributed. For Vendor 1, there is a peak of activity around 30 days ahead of IOC publication, which provided the vast majority of sighted IOCs (86,596 vs. 743 from Vendor 2).

Table 1: Counts of IOCs with sightings *after* publication date

Metric	Overall	Vendor 1	Vendor 2
Total IOCs ($d = 50$)	248 261	247 368	893
With any sightings after publication	47 110 (19.0%)	47 024 (19.0%)	86 (9.6%)
With first sighting after publication	2 921 (1.2%)	2 886 (1.2%)	35 (3.9%)

Table 2: IOC timing statistics. Median and mean first and last sightings, in days relative to IOC publication.

		d = 10	d = 50	d = 100	d = 150
Overall	Median	-7, 4	-34, -27	-41, -27	-44, -27
	Mean	*-5, 3*	*-36, -15*	*-52, -10*	*-63, -1*
Vendor 1	Median	-7, 4	-34, -27	-42, -27	-44, -27
	Mean	*-5, 3*	*-36, -15*	*-53, -10*	*-63, -1*
Vendor 2	Median	-2, -2	-3, -2	-2, -2	-2, -2
	Mean	*-3, -2*	*-10, -5*	*-7, 0*	*-22, -7*

The implication of these findings is that customers of Vendor 1 will miss the bulk of the adversarial activities associated with the IOCs as they occur. They will detect most incidents only in historical logs, provided they store them for at least 30 days. This effect was less pronounced for Vendor 2 (Table 1 although the small number of observations prevents us from drawing definite conclusions.

It is important to note that although Vendor 1's 30-day delay between the peak in potentially malicious activity and IOC publication was surprising to us, Vendor 1 made certain metadata about the timeliness of its IOCs transparently available to its customers, for example, pointing to open source blocklists and the time when IOC had been observed there before inclusion in its feed. Comparing that metadata to IOC publication dates revealed the same delay of around 30 days. This could, for example, be explained by the vendor performing its own research and validation (section 6).

4.3 Factors influencing sighting patterns

To understand which properties affect how sightings are distributed relative to IOC publication, we qualitatively compare various subsets of the total set of IOCs. This helps us understand potential threat actor considerations related to abandoning resources. Note that some subsets have very small sample sizes, so we do not treat these findings as generalizable results, but rather as informative for future work. This is a result of IOC metadata being quite sparse, where vendor reports are much richer in context [8]. In subsection 5.2 we perform statistical modeling using the same factors.

As Table 3 shows, different types of IOC show distinct sighting patterns: domains persist longer than IP addresses, likely in part because threat actors are known to use the two types of resources jointly to rapidly rotate through IP addresses in a technique known as fast flux [35]. Because a significant portion of domains leads to sightings beyond publication date, they may prove more relevant for intrusion detection and could also 'impose cost' more than other types of IOCs. We attribute the differences in timing of hash-based IOCs to vendor preferences and their sighting timing patterns – vendor 1 only provided MD5 hashes and vendor 2 provided more hashes overall (Figure 2, Figure 6). We note again that our measurements are based on logs of network traffic metadata section 3. This visibility is reflected by network-based IOCs having a higher probability to lead to sightings.

The results by domain IOC cohort in Table 4 let us compare established domains to new domains, as explained in subsection 3.4. Here, we were interested in contrasting the two cohorts, rather than producing generalizable measurements from either category of domains. With that in mind, the median last seen date of new domains is notable – 28 days after publication, compared to 14 days before publication for established domains. On the one hand, this suggests that the vendors had provided relevant threat intelligence about these freshly registered domain names, which still generated sightings after IOC publication – although not especially timely, given the median first seen date 38 days before IOC publication. On the other hand, it shows that actors did not abandon that infrastructure for quite a while after they had been published. Although the number of observations (92 IOCs) is too small to draw conclusions from, as compared to our overall dataset, this finding suggests that newly registered malicious domain names could perhaps be a category of threat intelligence where vendors excel in providing IOCs that are suitable for real-time detection. Furthermore, because of how we defined this cohort, these sightings were relatively likely to reflect actual malicious events, as compared to the overall dataset.

Table 5 shows results for the small subset (0.21% of IOCs that led to sightings) where the vendor had provided some information about what the related threat actors motivation was. We coded IOCs as belonging to threat actors that according to the vendor were either financially motivated or motivated to perform espionage[2] based on metadata the vendors had added, or where we could deduce this using the vendors' APT naming conventions [38, 41].

IOCs linked to espionage are observed much earlier than those linked to financial motives. This is consistent with them being more stealthy, so more time passes before they are discovered and published, compared to the financially motivated attacks. They are also abandoned earlier, well before publication. The median date of abandonment is 28 days before publication. (We do see a small tail of espionage IOCs that last after publication.) Perhaps surprisingly, the attacks tagged as financial have much shorter lifespans: the median first seen date is 16 days before publication, and the median last seen date is 12 days before publication. This might reflect the fact that criminal operations are not as stealthy and anticipate being blocklisted quickly. Thus, criminals have internalized the need to proactively rotate through resources, rather than wait for detection. This is consistent with prior research in cybercrime [1, 53].

Table 3: Distributions of first and last sightings, in days relative to IOC publication, by IOC type and using $d = 50$.

IOC Type	Count	Medians	FS/LS distribution
IP address	246,275	-34, -27	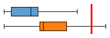
Domain	1,034	-33, -19	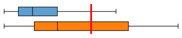
Hash (MD5)	429	-4, -3	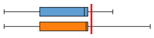
Hash (SHA1)	361	-2, -2	
Other	162	-23, -23	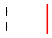

Table 4: Distributions of first and last sightings, in days relative to IOC publication, by domain IOC cohort (as defined in subsection 3.4) and using $d = 50$.

Domain IOC cohorts	Count	Medians	FS/LS distribution
New domains	92	-38, 28	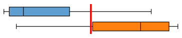
Established domains	29	-26, -14	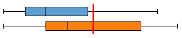
All domain IOCs	1034	-33, -19	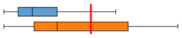

Finally, we look at a very small subset of tags that attribute the IOCs to the geography of the threat actor (Table 6) and that we were also able to map to nation-states based on publicly available information or the vendors' actor naming conventions [38, 41]. Russian actors have been thought to conduct relatively 'noisy' operations, also as a form of signaling [30]. However, our sightings show these IOCs are published later after the attacks than the IOCs associated with Chinese threat actors. Just 0.02% of IOCs had this metadata (using $d = 50$), and this small sample size means that this data only supports hypotheses that should be explored in future work.

In sum, while some patterns suggest reactive abandonment, the majority of resources were abandoned well before publication. A fraction continued to be used despite public disclosure, showing that threat actor responses to IOC publication are more complex than simple abandonment models suggest.

5 Statistical modeling

To characterize the relationship between IOC sightings and potential covariates, we fitted four linear regression models on our dataset, applying window $d = 50$. The regression table provides the resulting coefficients and model diagnostics (Table 7, on page 11). We provide descriptive statistics for our dependent variable and covariates in Appendix A. Although our results are based on a window of $d = 50$ days around IOC publication date, as described

[2] We also classified actors as motivated by ideology (hacktivism) but the corresponding IOCs did not result in sightings during our measurement period.

Table 5: Distributions of first and last sightings, in days relative to IOC publication, by actor motivation (0.21% of IOCs have this property) and using $d = 50$.

Actor motivation	Count	Medians	FS/LS distribution
Financial	516	-16, -12	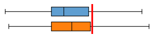
Espionage	16	-48, -28	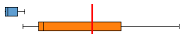
Untagged IOCs	247,729	-34, -27	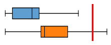

Table 6: Distributions of first and last sightings, in days relative to IOC publication, by actor geography (0.02% of IOCs have this property) and using $d = 50$.

Actor geography	Count	Medians	FS/LS distribution
China	43	-27, -17	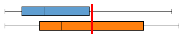
Russia	14	-49, -28	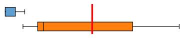
Untagged IOCs	248,203	-34, -27	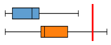

in subsection 3.4, we also fitted our models for the three other values of d. This did not lead to substantively different outcomes. We included these additional regression tables in Appendix B.

5.1 Relationship between IOC publication and sightings, with fixed effects regression

Panel data analysis with fixed effects let us investigate within-group variation, i.e., how sightings varied over time for each individual IOC in our dataset. This approach allowed us to make claims about our research question because it let us control for variations common to all IOCs in the dataset, like the slight increase in unique IOCs observed daily over the measurement period (Figure 4). Statisticians have emphasized that significant levels are not always informative for large datasets, and so we focus on effect size [54].

As expected from the measurements in the earlier sections of this paper, there is a negative coefficient, meaning that IOCs had a lower probability of leading to sightings after their publication: on average 3.2 percentage points, within the same IOC, holding all time-invariant factors constant. The effect size is modest but meaningful: publication explains 8.1% of the within-group variance. (To compare, in policy analysis, the effects of a policy on an outcome – say, the impact of minimum wage increases on employment rates – in the range of 5-20% are seen as meaningful, given that there are many other unobserved factors at play.) This is consistent with the hypothesis that threat intelligence can impose costs.

However, as we also learned from the measurements, the peak of sightings occurred around 30 days before publication. This suggests that while the fixed effects model shows a statistically significant negative effect on the occurrence of sightings, the exact publication day might not capture the precise number of days from which the decline in sightings begins. In fact, a downward trend in sightings seems to have already been occurring well before the publication. This raises the possibility that the publication date itself may not be the primary driver of the observed decline. Instead, the decrease in sightings could reflect an ongoing trend, with the publication potentially coinciding with this natural decline rather than directly causing it[3]. We explore some other possible covariates in the next section, as the fixed effects model does not allow for time-invariant covariates, like properties of IOCs (which are the same for all sightings of that IOC). Note that we added as a covariate in the fixed effects model whether the day was a weekday to allow for comparison of coefficients, but that this covariate was absorbed by the model.

5.2 Relationship between other factors and sightings, with OLS regression

We used simple linear regression to explore *between-group* variation, including properties of the respective IOCs. The main finding from fitting three models (2, 3, 4) is that, despite including a host of covariates, none of the models provides a higher explained variance (Adjusted R^2) than the fixed effects model (1) did for within-group variation. In other words, there are many confounding factors that all of these models don't capture. We go deeper into limitations of this study in section 7. The only properties of IOCs that stand out are covariates indicating that the IOC was of the type IP address or domain (the latter being the reference category). These were included in model 2. As discussed in section 4, hash-type IOCs leading to fewer sightings is a result of our dataset being generated on logs from network monitoring software.

As covariates, we included in models 2-4 the factors that we explored in subsection 4.3, based on our cohorts, vendor metadata and naming conventions. We warn again that the sample size is small for these factors. We therefore do not draw any conclusions from the resulting coefficients or attempt to generalize. We observe that in our dataset, IOCs tagged by the vendor to be related to Russia-nexus actors were more likely to lead to sightings than those related to China or DPRK – or other states that either weren't tagged as threat actors by the vendors, that we couldn't map, or that didn't lead to any sightings. We speculate the small negative coefficient in model 4 for IOCs belonging to financially motivated adversaries, as compared to those we classified as being motivated by espionage, to be a result of the shorter lifespan of the former category of IOCs (subsection 4.3).

6 Discussion

We found modest evidence of a publication effect that would be in line with the 'impose cost' mechanism: fewer days with sightings occurred after IOC publication than before. Our panel regression model also reflected this, indicating a negative relationship between IOC publication and sighting probability, capturing a small but meaningful 8.1% of within-group variance in whether or not sightings occurred on a given day.

On the other hand, the IOC feeds from two market-leading CTI providers largely captured threats that occurred in the past. For the IOCs that generated sightings, only 19% provided sightings after

[3] We do not evaluate R^2 values against benchmarks because we are doing empirical research and not building a predictive model.

Table 7: Regression models for IOCs that led to sightings in a window of $d = 50$ days around publication date. In subsection 5.1, we address our research question directly, finding a small but meaningful effect size using the fixed effects model (model 1), in which publication explains 8.1% of the within-group variance of sightings. In subsection 5.2, we analyze other covariates using simple linear regression models (models 2, 3, 4). None of combinations of included covariates lead to a higher explained variance than model 1.

	Dependent variable: IOC led to sightings (binary, daily)			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Day was after IOC publication	-0.032***	-0.161***	-0.161***	-0.161***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant		0.207***	0.128***	0.129***
		(0.000)	(0.006)	(0.006)
Day was a weekday (ref. Weekend)		0.002***	0.002***	0.002***
		(0.000)	(0.000)	(0.000)
Vendor 2 (ref. Vendor 1)			-0.004*	0.003
			(0.002)	(0.003)
IOC type: IP adress (ref. Domain)			0.087***	0.087***
			(0.001)	(0.001)
IOC type: MD5 hash (ref. Domain)			-0.022***	-0.021***
			(0.002)	(0.002)
IOC type: SHA1 hash (ref. Domain)			-0.021***	-0.022***
			(0.003)	(0.003)
IOC type: Other, e.g. regex (ref. Domain)			0.004	0.013***
			(0.004)	(0.004)
Domain IOC cohort: None (ref. Established domains)			-0.007	-0.008
			(0.006)	(0.006)
Domain IOC cohort: New domains (ref. Established domains)			0.029***	0.027***
			(0.007)	(0.007)
Actor geography: Russia (ref. China)				0.133***
				(0.009)
Actor geography: DPRK (ref. China)				-0.078**
				(0.032)
Actor motivation: Financial (ref. Espionage)				-0.018***
				(0.002)
Included effects:	Entity, Time	-	-	-
Absorbed covariates:	Weekday	-	-	-
Observations	25079007	25079007	25079007	25079007
N. of groups	248261			
R^2	0.001	0.059	0.059	0.059
Within R^2	0.081			
Adjusted R^2		0.059	0.059	0.059
Residual Std. Error	0.008	0.323	0.323	0.323
F Statistic	21042.850***	781826.089***	175906.742***	131958.455***

Note: Standard errors in parentheses.
* p < 0.1; ** p < 0.05; *** p < 0.01.

their publication. The other 81% were only seen before publication. Thus, their publication was unable to 'impose cost' by enabling intrusion detection of ongoing threats. Furthermore, the peak in the number of IOCs sighted was 30 days before their respective publication dates, with a median last sighting 27 days before publication. Not to mention that over 98% of the total set of published IOCs were never sighted at all. In the simplest terms, the answer to our research question would therefore be: threat actors largely abandon resources *before* these have been published as IOCs in commercial threat intelligence.

Since we see that attacker resources are being abandoned before the publication of the associated IOC, does this mean that the impact of the publication is at best modest and perhaps negligible? Not necessarily. There are second-order effects of publication to consider. If attackers know that their resources are going to be detected and published at some point, and this point is earlier rather than later, they have an incentive to invest in making their infrastructure resilient to rapid detection and publication. Attackers could choose to abandon infrastructure proactively rather than wait for defenders to detect it. This also saves resources in not having to do counter-detection – i.e., detecting the detection by defenders.

In other words, the pattern we are seeing is consistent with threat actors having internalized the cost of publication into their operations with improved operational security. Vendors have themselves described behaviors of certain threat actors that indeed seem to point in this direction – raising the spectre of "IOC Extinction" [40]. Such a dynamic is similar to what the security community has previously observed in regular cybercrime. For example, spammers have played a cat-and-mouse game with email service operators since the '90s, in which constant blocklisting of spammy servers led criminals to adopt a tactic of rapidly cycling through infrastructure, using 66% of abusive IPs for just a single day [47]. Other examples are the short life span and rapid cycling of phishing sites [36] and responses of cybercrime groups to takedowns [10, 48].

In strategic terms, 'friction' or costs are being imposed on attackers at a macro level, rather than by the specific instances of detection and publication. If the cost of publication is internalized by attackers, then that has interesting implications. Everyone benefits from higher costs for attackers since it raises the barriers for conducting attacks. These benefits are not restricted to the paying customers of high-end threat intelligence products. In economic terms: these products have positive security externalities that are basically a public good (non-excludable and non-rivalrous).

Our findings do raise three important implications for enterprise customers of high-end threat intelligence. First, the value of threat intelligence might be lower than what they are currently willing to pay for it, if detection is a relevant use case for them. Some portion of enterprises are spending six figure numbers on licenses for commercial threat intelligence. In reality, the costs are likely even higher, since enterprises report using 7 threat intelligence sources on average [39]. These customers might want to re-evaluate how much they are willing to spend on threat intelligence feeds, if they intend to use them for real-time intrusion detection.

The second implication of our findings is that threat hunting, i.e., retrospective-first detection models, will become even more essential. Effective use of commercial CTI now depends far more on retroactive hunting in stored telemetry than on real-time blocking. This increases the value of good logging and might increasingly collide with data retention and privacy laws.

A third implication is for all customers of commercial threat intelligence. Our findings suggest the need for more comprehensive benchmarking of IOC timeliness based on network logs. We found that Vendor 1's IOCs peak approximately 30 days before publication, whereas Vendor 2 shows far fewer sightings but with a shorter lag (subsection 4.2). If enterprises combine forces, for example via an ISAC or CSIRT, they could generate a more comprehensive evaluation of the vendors relevant for their sector.

Before we turn to the implications for the government strategy to impose cost, we should ask: would analyzing government-published IOCs, rather than through from commercial vendors, have led to different results? The original aim of this research project was to do just that, but the number of government-published IOCs was nowhere near high enough to allow a similar quantitative approach. We can, however, discuss similarities and differences. Most importantly, the threat actors being tracked are APTs, and therefore, the sightings that stem from either government-published or industry-published IOCs would likely be similar.

One difference is that governments can draw on the use of special investigative powers by, for example, the intelligence agencies. This might enable them to deliver more timely IOCs than industry can. On the other hand, to publish these indicators typically requires a declassification process, which is likely to be time-consuming. So it is not clear government would be able to move faster. Another difference is that governments needn't worry about a business model and are free to release IOCs to the public at large, which might cause adversaries to respond differently than releasing them only to the customer base of threat intelligence vendors.

However, these considerations are made moot by what is happening in practice. Public releases by US Cyber Command on its VirusTotal profile have faltered. According to a spokesperson, this is because it got "additional [legal] authorities to share information directly with industry partners" [32]. However, according to the same article, operators are reportedly still sharing threat intelligence on the platform, just without "going out loud and proud.".

Does it ultimately matter if there exists a quantitative empirical basis for the 'impose cost' mechanism? Should the policy be changed because of our findings? In international relations, sending signals to opponents is as much a matter of perception as it is of hard, measurable impacts. Some authors have argued that the popularity of the persistent engagement strategy may be driven by institutional as well as strategic imperatives, as part of the bureaucratic struggle for government resources [29] or simply by decision-makers lacking better options [31].

A final implication for public policy has to do with cybersecurity advice, which often encourages organizations to share threat information with each other, for example, in sectoral platforms like ISACs (Information Sharing and Analysis Centers) [23, 43]. If the market leaders in threat intelligence publish IOCs that are in the majority not capturing ongoing threats but past threats, can individual organizations do better? A 2022 study suggests that, yes, peer-to-peer sharing of threat intelligence might, in certain cases, be able to report IOCs faster than corporations [9]. However, given the cat-and-mouse dynamic described above, the shorter-lived malicious resources will become, the more formidable the challenge will be to capture and share relevant resources as threat intelligence.

A question to reflect on is why Vendor 1 shows a roughly 30-day delay between the peak of malicious activity and the publication of the IOCs, as opposed to Vendor 2. The firm is at least somewhat transparent to its customers about the delay in IOC publication by including metadata about the earlier publication of the IOC by third parties, e.g., open-source blocklists (subsection 4.2). Vendors could choose not to immediately publish IOCs because their products are evaluated by metrics other than just timeliness [19]. In a 2020 study [8], when customers were asked what commercial threat intelligence services are good for, more so than timeliness, many of them mentioned 'providing context' and 'curation'. Curation, i.e., providing only highly vetted IOCs, is directly tied to the fact that customers need to avoid overwhelming their security operations analysts with false positive alerts. Curation takes time. As a result, commercial IOCs are not as likely to lead to false-positive alerts, contrary to open-source blocklists. The latter may be faster to report IOCs, but they are often unsuitable for intrusion detection. Furthermore, subsection 4.3, we did identify freshly registered domain names as a category of IOCs provided by the two vendors that were able to lead to relevant sightings after IOC publication date, suggesting that there are differences in IOC quality that our quantitative approach does not fully capture. And lastly, we must remember that threat intelligence firms are not simply IOC factories; rather, they track activity clusters – from world-famous 'bears' and 'pandas' to developing trends – tying observations together and reporting about attribution as well as tools, techniques, and procedures. We are optimistic that commercial threat intelligence firms could be more open about their investigative process and what customers may expect in terms of, e.g., timeliness, accuracy, and coverage of certain threats. Until this happens, we can only speculate about the explanation of the delay between sightings on IOCs and their publication.

7 Limitations

The main limitation of our method is that we measure adversary behavior indirectly, with sightings of IOCs as a proxy for adversary behavior. This approach causes three problems. First, our sightings highly likely include false positives, i.e., observations that were actually legitimate traffic instead of threat actor behavior. It is hard to say how large we should expect the portion of false positives to be – there exists no ground truth of maliciousness of IOCs, as this is bound to time and place of observation. A study on alert fatigue in SOCs illustrates that the number of false positives could well be substantial [2], though SOC alerts are generated by many more detection rules than mere IOCs. In the end, we assume that the IOCs of two market-leading firms contain some portion of false positives but that they do not overwhelm the dataset. An additional benefit of our approach is that it provided new insights into the ability of commercial threat intelligence IOCs to enable detection.

Second, even if a sighting does capture actor behavior (i.e., true positive), different classes of IOCs can lead to different security outcomes: a sighting in our dataset might represent anything from network reconnaissance to ransomware being deployed. In other words, not all IOCs are created equal, yet in this quantitative study, we do treat them as such.

And third, threat intelligence feeds provide only a partial coverage of actor behavior [8]. Threat actors act strategically to evade detection, and even a hypothetical perfect IOC set would only lead to sightings if intrusion detection controls happen to be set up on the right network, and at the right place. In sum, our data reflects the governmental enterprise networks where the logs were collected. It might have failed to capture those attacks directed against, say, financial institutions. Still, the networks where the sightings occurred do form a typical and important use case – i.e., client base – for these kinds of IOCs.

Another limitation exists in the fact that we used observational data. While panel data models are powerful tools for inferring causal relationships, the gold standard for causal inference is randomized controlled trials, where the treatment (in our case, IOC publication) is randomly assigned. The threat intelligence vendors would be uniquely positioned to perform such measurements for scientific benefit – although a true experimental setup would require that certain IOCs not be published, likely causing ethical (and commercial) concerns. For observational data like ours, which the vendors could also look to, future work could benefit from recent developments in econometrics that better allow causal inference [12, 46]. We see further opportunities for future work to explore factors that contribute to sightings like those that we identified in subsection 4.3, but that had too small a sample size to draw any conclusions from, like domain IOC cohorts and actor motivation. Given that it would require a serious number of observations to be able to drill down into categories, threat intelligence vendors would again be best positioned to study these factors. Alternatively, they could label a larger portion of their IOCs with metadata.

Finally, we had to make assumptions in the interpretation of our data. An especially relevant parameter is the number of days of sightings to include before and after IOC publication (d). We performed sensitivity analysis on this variable to alleviate the impact of this assumption. Similarly, we fitted multiple regression models, which did not lead to substantially different outcomes. Nevertheless, we were careful not to assert causality. We hypothesized potential confounding factors that may have contributed to our results.

8 Conclusions

In this study, we found limited empirical evidence for the strategic notion of 'imposing costs' on threat actors. We wanted to bridge the gap between computer science and international relations and asked: *Do threat actors abandon resources after these have been published as indicators of compromise in commercial threat intelligence?* We find that no, they largely abandoned resources earlier.

We described how sightings on commercial threat intelligence are distributed: just 19% of IOCs that led to sightings still did so after publication. In other words, for 81% of those IOCs that pointed to relevant activity for our networks, threat actors had already stopped using the resources by the time customers received the IOCs – which would therefore have been unable to 'impose costs' (section 4).

To quantify the relationship between sightings and IOC publication, we fitted a panel regression model using fixed effects, finding a negative relationship between IOC publication and the probability of future sightings occurring on a given day. The model was able to explain a small but meaningful 8.2% of variance in whether or not an IOC was sighted (section 5). We raised questions about timeliness of commercial threat intelligence, because the IOC feed of one vendor captured malicious activity that occurred largely around 30 days before publication to the customer (section 4).

In sum, the commercial threat intelligence IOCs feeds that we studied largely captured malicious activity that occurred before customers were able to use them for intrusion detection. Does this mean that threat intelligence cannot 'impose cost'? Not necessarily. Clearly, more sightings occurred before IOC publication than after, as both the distribution of the data and statistical modeling show. We provide various interpretations of our findings, one of which is that second-order effects may have occurred because threat actors internalized a higher degree of operational security. This explanation has interesting implications for network defenders as well as policymakers (section 6). We are careful with claims of causality and point to the challenges of trying to measure malicious activity (section 7).

Acknowledgments

The authors are grateful to Jon Lindsay, Rainer Böhme, and the anonymous reviewers for providing feedback on this work. This research has been partially funded by the Ministry of the Interior and Kingdom Relations of the Netherlands.

References

[1] S. Agarwal and M. Vasek. 2025. Examining Newly Registered Phishing Domains at Scale. In *The 24th Workshop on the Economics of Information Security (WEIS)*. Tokyo, Japan. https://discovery.ucl.ac.uk/id/eprint/10209951

[2] Bushra A. Alahmadi, Louise Axon, and Ivan Martinovic. 2022. 99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 2783–2800. https://www.usenix.org/conference/usenixsecurity22/presentation/alahmadi

[3] M. Armelli, S. Caudill, J.P. Dees, M. Eager, J. Keltz, I. Pelekis, J. Sakellariadis, V.V. Singh, and K. von Ofenheim. 2020. *Cyber Threat Intelligence - What is the Impact of Information Disclosures on an Adversary's Operations?* Technical Report. University of Columbia, SIPA. https://www.sipa.columbia.edu/academics/capstone-projects/cyber-threat-intelligence-what-impact-information-disclosures-adversary

[4] J. E. Barnes. 2021. U.S. Military Has Acted Against Ransomware Groups, General Acknowledges. (Dec. 2021). https://www.nytimes.com/2021/12/05/us/politics/us-military-ransomware-cyber-command.html

[5] R. Bevington. 2024. *Turning the tables: Using cyber deception to hunt phishers at scale.* https://www.bleepingcomputer.com/news/security/microsoft-creates-fake-azure-tenants-to-pull-phishers-into-honeypots

[6] Bianco, D. 2013. The Pyramid of Pain. http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html

[7] L. Bilge and T. Dumitraş. 2012. Before we knew it: an empirical study of zero-day attacks in the real world. In *CCS '12: Proceedings of the 2012 ACM conference on Computer and communications security*. Association for Computing Machinery, New York, NY, USA, 833–844. doi:10.1145/2382196.2382284

[8] Xander Bouwman, Harm Griffioen, Jelle Egbers, Christian Doerr, Bram Klievink, and Michel van Eeten. 2020. A different cup of TI? The added value of commercial threat intelligence. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 433–450. https://www.usenix.org/conference/usenixsecurity20/presentation/bouwman

[9] Xander Bouwman, Victor Le Pochat, Pawel Foremski, Tom Van Goethem, Carlos H. Ganan, Giovane C. M. Moura, Samaneh Tajalizadehkhoob, Wouter Joosen, and Michel van Eeten. 2022. Helping hands: Measuring the impact of a large threat intelligence sharing community. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 1149–1165. https://www.usenix.org/conference/usenixsecurity22/presentation/bouwman

[10] Ryan Brunt, Prakhar Pandey, and Damon McCoy. 2017. Booted: An analysis of a payment intervention on a ddos-for-hire service. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*. La Jolla, CA, USA.

[11] Ben Buchanan. 2020. *The Hacker and the State: Cyber Attacks and the New Normal of Geopolitics.* Harvard University Press, Cambridge, MA, USA.

[12] B. Callaway and P. H. C. Sant'Anna. 2021. Difference-in-Differences with multiple time periods. 225, 2 (12 2021), 200–230. doi:10.1016/j.jeconom.2020.12.001

[13] D. Chiba, M. Akiyama, Y. Otsuki, H. Hada, T. Yagi, T. Fiebig, and M. Van Eeten. 2022. DomainPrio: Prioritizing Domain Name Investigations to Improve SOC Efficiency. 10 (March 2022), 34352–34368. doi:10.1109/ACCESS.2022.3161636

[14] CompaniesMarketCap.com. 2025. CrowdStrike (CRWD) - Market capitalization. https://companiesmarketcap.com/usd/crowdstrike/marketcap

[15] G. Di Tizio, M. Armellini, and F. Massacci. 2022. Software Updates Strategies: a Quantitative Evaluation against Advanced Persistent Threats. (May 2022), 1. doi:10.1109/TSE.2022.3176674

[16] J. Dykstra, K. Shortridge, J. Met, and D. Hough. 2022. Sludge for Good: Slowing and Imposing Costs on Cyber Attackers. (Nov. 2022). arXiv:2211.16626 doi:10.48550/arXiv.2211.16626

[17] M. P. Fischerkeller, E. O. Goldman, and R. J. Harknett. 2022. *Cyber Persistence Theory.* Oxford University Press, Oxford, England, UK. https://global.oup.com/academic/product/cyber-persistence-theory-9780197638262

[18] Gartner. 2021. Market Guide for Security Threat Intelligence Products and Services.

[19] Harm Griffioen, Tim M. Booij, and Christian Doerr. 2020. Quality Evaluation of Cyber Threat Intelligence Feeds. In *International Conference on Applied Cryptography and Network Security (ACNS).*

[20] J. Healey. 2019. The implications of persistent (and permanent) engagement in cyberspace. 5, 1 (1 2019), tyz008. doi:10.1093/cybsec/tyz008

[21] U.S. White House. 2023. National Cybersecurity Strategy. (March 2023). https://www.whitehouse.gov/briefing-room/statements-releases/2023/03/02/fact-sheet-biden-harris-administration-announces-national-cybersecurity-strategy

[22] Vincenzo Iozzo. 2019. The Case for Scale in Cyber Security. https://media.ccc.de/v/36c3-11220-the_case_for_scale_in_cyber_security [Conference talk; accessed 4. Feb. 2020].

[23] C. Johnson, L. Badger, D. Waltermire, J. Snyder, and C. Skorupka. 2016. *Guide to Cyber Threat Information Sharing.* Technical Report. National Institute of Standards and Technology. http://dx.doi.org/10.6028/NIST.SP.800-150

[24] Joseph Cox. 2020. Internal Docs Show Why the US Military Publishes North Korean, Russian Malware. Vice.com. https://www.vice.com/en/article/5dmwyx/documents-how-cybercom-publishes-russian-north-korean-malware-virustotal

[25] A. Kodituwakku, C. Xu, D. Rogers, D. K. Ahn, and E. W. Fulp. 2023. Temporal Aspects of Cyber Threat Intelligence. In *2023 IEEE International Conference on Big Data (BigData)*. IEEE, 15–18. doi:10.1109/BigData59044.2023.10386664

[26] Z. Lanz. 2022. Cybersecurity Risk in U.S. Critical Infrastructure: An Analysis of Publicly Available U.S. Government Alerts and Advisories. 5, 1 (2022), 43–70. doi:10.52306/2578-3289.1121

[27] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS 2019)*. doi:10.14722/ndss.2019.23386

[28] Vector Guo Li, Matthew Dunn, Paul Pearce, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. 2019. Reading the Tea leaves: A Comparative Analysis of Threat Intelligence. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 851–867. https://www.usenix.org/conference/usenixsecurity19/presentation/li

[29] J. R. Lindsay. 2021. Cyber conflict vs. Cyber Command: hidden dangers in the American military solution to a large-scale intelligence problem. 36, 2 (Feb. 2021), 260–278. doi:10.1080/02684527.2020.1840746

[30] E. D. Lonergan and S. W. Lonergan. 2022. Cyber Operations, Accommodative Signaling, and the De-Escalation of International Crises. (Jan. 2022). doi:10.1080/09636412.2022.2040584

[31] Lonergan, E. and Poznansky, M. 2023. Are We Asking Too Much of Cyber? https://warontherocks.com/2023/05/are-we-asking-too-much-of-cyber

[32] Martin Mathishak. 2024. As Cyber Command evolves, its novel malware alert system fades away. https://therecord.media/cyber-command-virustotal-twitter-malware-alerts-cnmf

[33] Martin Matishak and Jonathan Greig. 2024. US confirms takedown of China-run botnet targeting home and office routers. https://therecord.media/china-run-botnet-takedown-fbi-doj-routers

[34] Tim Maurer and Garret Hinck. 2020. Persistent Enforcement: Criminal Charges as a Response to Nation-State Malicious Cyber Activity. (1 2020). https://jnslp.com/2020/01/23/persistent-enforcement-criminal-charges-as-a-response-to-nation-state-malicious-cyber-activity

[35] L. B. Metcalf, D. Ruef, and J. M. Spring. 2017. Open-source Measurement of Fast-flux Networks While Considering Domain-name Parking. 13–24 pages. https://www.usenix.org/conference/laser2017/presentation/metcalf

[36] T. Moore and R. Clayton. 2007. Examining the impact of website take-down on phishing. In *ACM Other conferences*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/1299015.1299016

[37] P.M. Nakasone. 2022. Posture statement of Gen. Paul M. Nakasone, commander, U.S. Cyber Command before the 117th Congress. https://www.cybercom.mil/Media/News/Article/2989087/posture-statement-of-gen-paul-m-nakasone-commander-us-cyber-command-before-the

[38] Kevin Poireault. 2023. Understanding Threat Actor Naming Conventions. https://www.infosecurityeurope.com/en-gb/blog/threat-vectors/understanding-threat-actor-naming-conventions.html

[39] Ponemon Institute. 2019. *The Value of Threat Intelligence: The Second Annual Study of North American & United Kingdom Companies.* Technical Report February. Ponemon Institute.

[40] Michael Raggi. 2024. IOC Extinction? China-Nexus Cyber Espionage Actors Use ORB Networks to Raise Cost on Defenders. https://cloud.google.com/blog/topics/threat-intelligence/china-nexus-espionage-orb-networks

[41] Florian Roth. 2018. The Newcomer's Guide to Cyber Threat Actor Naming. https://medium.com/@cyb3rops/the-newcomers-guide-to-cyber-threat-actor-naming-7428e18ee263

[42] J. Sigholm and M. Bang. 2013. Towards Offensive Cyber Counterintelligence: Adopting a Target-Centric View on Advanced Persistent Threats. (2013),

166–171. https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A640955&dswid=6375

[43] F. Skopik, G. Settanni, and R. Fiedler. 2016. A problem shared is a problem halved: A survey on the dimensions of collective cyber defense through security information sharing. *Computers & Security* 60 (7 2016), 154–176. doi:10.1016/j.cose.2016.04.003

[44] S. Soesanto and M. Smeets. 2020. Cyber Deterrence: The Past, Present, and Future. In *NL ARMS Netherlands Annual Review of Military Studies 2020: Deterrence in the 21st Century—Insights from Theory and Practice*. T.M.C. Asser Press, The Hague, The Netherlands, 385–400. doi:10.1007/978-94-6265-419-8_20

[45] James H. Stock and Mark W. Watson. 2020. *Introduction to Econometrics, 4th edition.* Pearson.

[46] L. Sun and S. Abraham. 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. 225, 2 (12 2021), 175–199. doi:10.1016/j.jeconom.2020.09.006

[47] K. Thomas, R. Amira, A. Ben-Yoash, O. Folger, A. Hardon, A. Berger, E. Bursztein, and M. Bailey. 2016. The Abuse Sharing Economy: Understanding the Limits of Threat Exchanges. In *Research in Attacks, Intrusions, and Defenses*. Springer, Cham, Switzerland, 143–164. doi:10.1007/978-3-319-45719-2_7

[48] Kurt Thomas, Danny Yuxing Huang, David Wang, Elie Bursztein, Chris Grier, Thomas J Holt, Christopher Kruegel, Damon McCoy, Stefan Savage, and Giovanni Vigna. 2015. Framing Dependencies Introduced by Underground Commoditization. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*. Delft, The Netherlands.

[49] B. Tostes, L. Ventura, E. Lovat, M. Martins, and D. Menasché. 2023. Learning When to Say Goodbye: What Should be the Shelf Life of an Indicator of Compromise? In *2023 IEEE International Conference on Cyber Security and Resilience (CSR)*. IEEE, 2023–02. doi:10.1109/CSR57506.2023.10224937

[50] Brandon Valeriano. 2020. Cost Imposition Is the Point: Understanding U.S. Cyber Operations and the Strategy Behind Achieving Effects. https://www.lawfareblog.com/cost-imposition-point-understanding-us-cyber-operations-and-strategy-behind-achieving-effects

[51] Mathew Vermeer, Carlos Gañán, and Michel van Eeten. 2022. Ruling the Rules: Quantifying the Evolution of Rulesets, Alerts and Incidents in Network Intrusion Detection. In *ASIA CCS '22: ACM Asia Conference on Computer and Communications Security, Nagasaki, May 30-June 3, 2022.* ACM.

[52] M. Vermeer, N. Kadenko, M. van Eeten, C. Gañán, and S. Parkin. 2023. Alert Alchemy: SOC Workflows and Decisions in the Management of NIDS Rules. In *ACM Conferences*. Association for Computing Machinery, New York, NY, USA, 2770–2784. doi:10.1145/3576915.3616581

[53] T. Vissers, J. Spooren, P. Agten, D. Jumpertz, P. Janssen, M. Van Wesemael, F. Piessens, W. Joosen, and L. Desmet. 2017. Exploring the Ecosystem of Malicious Domain Registrations in the .eu TLD. In *Research in Attacks, Intrusions, and Defenses*. Springer, Cham, Switzerland, 472–493. doi:10.1007/978-3-319-66332-6_21

[54] R. L. Wasserstein, A. L. Schirm, and N. A. Lazar. 2019. Moving to a World Beyond "$p < 0.05$". (3 2019). doi:10.1080/00031305.2019.1583913

A Covariate descriptive statistics

Table 8: Descriptive statistics of dependent variable and covariates in section 5.

Binary Variables

Variable	Mean	Frequency of 1s	% of 1s	% Units Ever 1	% Units Change
IOC seen	0.127	3,177,614	12.7	100.0	100.0
IOC published	0.505	12,663,657	50.5	100.0	100.0
Weekday	0.715	17,919,227	71.5	100.0	100.0

Categorical Variables

IOC Source

Category	Frequency	%	% Units Ever in Category	% Units Change
Vendor 1	24,984,168	99.6	99.6	0.0
Vendor 2	94,839	0.4	0.4	0.0

IOC Type

Category	Frequency	%	% Units Ever in Category	% Units Change
IP address	24,873,775	99.2	99.2	0.0
Domain	106,656	0.4	0.4	0.0
Hash (MD5)	44,541	0.2	0.2	0.0
Hash (SHA1)	37,673	0.2	0.1	0.0
Other	16,362	0.1	0.1	0.0

Domain IOC cohorts

Category	Frequency	%	% Units Ever in Category	% Units Change
New domains	713	0.0	0.0	0.0
Established domains	147	0.0	0.0	0.0
None	25,068,301	100.0	100.0	0.0

Actor Geography

Category	Frequency	%	% Units Ever in Category	% Units Change
China	4,343	0.0	0.0	0.0
Russia	1,414	0.0	0.0	0.0
DPRK	101	0.0	0.0	0.0

Actor Motivation

Category	Frequency	%	% Units Ever in Category	% Units Change
Financial	54,742	0.2	0.2	0.0
Espionage	1,616	0.0	0.0	0.0

Note: Number of observations: 25,079,007; Number of unique units: 248,261

B Regression tables for sensitivity analysis

Analysis in the body of this article was instrumented using a window of sightings around $d = 50$ days around IOC publication date, as explained in the method section (see subsection 3.4).

As part of our sensitivity analysis, we also fitted regression models for other values of d. Tables 9, 10 and 11 show regression models for $d = 10$, $d = 100$, and $d = 150$, respectively.

We furthermore ran our analysis for another unit of analysis (using the standard window of $d = 50$), and report the results in Table 12. All regression models in this article so far had as dependent variable the binarized daily sightings – i.e., 1 if any sightings occurred for an IOC on a given day and 0 if there did not.

The dependent variable for Table 12 is the count of daily sightings, normalized using the sum of sightings for the respective IOC throughout the measurement period (only for that window, i.e., in this case, $d = 50$).

$$DV = \frac{daily\ sightings\ of\ IOC}{total\ sightings\ of\ IOC\ (for\ this\ window)}.$$

Table 9: Regression models for IOCs that led to sightings in a window of $d = 10$ days around publication date.

	Dependent variable: IOC led to sightings (binary, daily)			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Day was after IOC publication	-0.103***	-0.099***	-0.099***	-0.099***
	(0.001)	(0.001)	(0.001)	(0.001)
Constant		0.397***	0.158***	0.160***
		(0.001)	(0.027)	(0.027)
Day was a weekday (ref. Weekend)		0.018***	0.018***	0.018***
		(0.001)	(0.001)	(0.001)
Vendor 2 (ref. Vendor 1)			-0.026*	-0.018
			(0.015)	(0.015)
IOC type: IP adress (ref. Domain)			0.251***	0.251***
			(0.006)	(0.006)
IOC type: MD5 hash (ref. Domain)			-0.023	-0.024
			(0.015)	(0.015)
IOC type: SHA1 hash (ref. Domain)			-0.021	-0.023
			(0.016)	(0.016)
IOC type: Other, e.g. regex (ref. Domain)			0.417***	0.417***
			(0.029)	(0.029)
Domain IOC cohort: None (ref. Established domains)			-0.009	-0.011
			(0.028)	(0.028)
Domain IOC cohort: New domains (ref. Established domains)			0.032	0.030
			(0.030)	(0.030)
Actor geography: Russia (ref. China)				0.027
				(0.017)
Actor geography: DPRK (ref. China)				-0.290***
				(0.073)
Actor motivation: Financial (ref. Espionage)				-0.018**
				(0.008)
Included effects:	Entity, Time	-	-	-
Absorbed covariates:	Weekday	-	-	-
Observations	1834455	1834455	1834455	1834455
N. of groups	87339			
R^2	0.006	0.011	0.015	0.015
Within R^2	0.018			
Adjusted R^2		0.011	0.015	0.015
Residual Std. Error	0.027	0.477	0.476	0.476
F Statistic	9709.255***	10248.144***	3191.585***	2395.690***

Note: Standard errors in parentheses.
* p < 0.1; ** p < 0.05; *** p < 0.01.

Table 10: Regression models for IOCs that led to sightings in a window of $d = 100$ days around publication date.

	Dependent variable: IOC led to sightings (binary, daily)			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Day was after IOC publication	-0.154***	-0.121***	-0.121***	-0.121***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant		0.151***	0.091***	0.091***
		(0.000)	(0.004)	(0.004)
Day was a weekday (ref. Weekend)		0.003***	0.003***	0.003***
		(0.000)	(0.000)	(0.000)
Vendor 2 (ref. Vendor 1)			0.002	0.006***
			(0.002)	(0.002)
IOC type: IP adress (ref. Domain)			0.071***	0.071***
			(0.001)	(0.001)
IOC type: MD5 hash (ref. Domain)			-0.013***	-0.014***
			(0.002)	(0.002)
IOC type: SHA1 hash (ref. Domain)			-0.013***	-0.014***
			(0.002)	(0.002)
IOC type: Other, e.g. regex (ref. Domain)			-0.019***	-0.017**
			(0.007)	(0.007)
Domain IOC cohort: None (ref. Established domains)			-0.010***	-0.010***
			(0.004)	(0.004)
Domain IOC cohort: New domains (ref. Established domains)			0.028***	0.028***
			(0.004)	(0.004)
Actor geography: Russia (ref. China)				0.068***
				(0.005)
Actor motivation: Financial (ref. Espionage)				-0.009***
				(0.002)
Included effects:	Entity, Time	-	-	-
Absorbed covariates:	Weekday	-	-	-
Observations	31022139	31022139	31022139	31022139
N. of groups	154333			
R^2	0.023	0.044	0.044	0.044
Within R^2	0.058			
Adjusted R^2		0.044	0.044	0.044
Residual Std. Error	0.037	0.284	0.283	0.283
F Statistic	731733.953***	707022.086***	159628.265***	130622.606***

Note: Standard errors in parentheses.
* p < 0.1; ** p < 0.05; *** p < 0.01.

Table 11: Regression models for IOCs that led to sightings in a window of $d = 150$ days around publication date.

	Dependent variable: IOC led to sightings (binary, daily)			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Day was after IOC publication	-0.096***	-0.096***	-0.096***	-0.096***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant		0.119***	0.062***	0.062***
		(0.000)	(0.003)	(0.003)
Day was a weekday (ref. Weekend)		0.003***	0.003***	0.003***
		(0.000)	(0.000)	(0.000)
Vendor 2 (ref. Vendor 1)			0.019***	0.057***
			(0.004)	(0.004)
IOC type: IP adress (ref. Domain)			0.057***	0.057***
			(0.001)	(0.001)
IOC type: MD5 hash (ref. Domain)			-0.012***	-0.009***
			(0.003)	(0.003)
IOC type: SHA1 hash (ref. Domain)			-0.012***	-0.010**
			(0.004)	(0.004)
IOC type: Other, e.g. regex (ref. Domain)			-0.030***	-0.034***
			(0.006)	(0.006)
Domain IOC cohort: None (ref. Established domains)			0.000	0.000
			(0.003)	(0.003)
Domain IOC cohort: New domains (ref. Established domains)			0.044***	0.044***
			(0.003)	(0.003)
Actor geography: Russia (ref. China)				0.027***
				(0.008)
Actor motivation: Financial (ref. Espionage)				-0.053***
				(0.003)
Included effects:	Entity, Time	-	-	-
Absorbed covariates:	Weekday	-	-	-
Observations	16463496	16463496	16463496	16463496
N. of groups	54696			
R^2	0.005	0.034	0.035	0.035
Within R^2	0.043			
Adjusted R^2		0.034	0.035	0.035
Residual Std. Error	0.016	0.255	0.255	0.255
F Statistic	80953.181***	292456.007***	66518.343***	54455.516***

Note: Standard errors in parentheses.
* p < 0.1; ** p < 0.05; *** p < 0.01.

Table 12: Regression models for IOCs that led to sightings in a window of $d = 50$ days around publication date, using another dependent variable: normalized daily sightings.

	Dependent variable: Normalized count of daily sighting			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Day was after IOC publication	-0.005***	-0.017***	-0.017***	-0.017***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant		0.019***	0.018***	0.018***
		(0.000)	(0.001)	(0.001)
Day was a weekday (ref. Weekend)		0.000***	0.000***	0.000***
		(0.000)	(0.000)	(0.000)
Vendor 2 (ref. Vendor 1)			-0.001***	-0.001***
			(0.000)	(0.000)
IOC type: IP adress (ref. Domain)			0.000	0.000
			(0.000)	(0.000)
IOC type: MD5 hash (ref. Domain)			0.001*	0.001*
			(0.000)	(0.000)
IOC type: SHA1 hash (ref. Domain)			0.001**	0.001**
			(0.001)	(0.001)
IOC type: Other, e.g. regex (ref. Domain)			0.001**	0.002**
			(0.001)	(0.001)
Domain IOC cohort: None (ref. Established domains)			0.000	0.000
			(0.001)	(0.001)
Domain IOC cohort: New domains (ref. Established domains)			0.000	0.000
			(0.001)	(0.001)
Actor geography: Russia (ref. China)				-0.000
				(0.002)
Actor geography: DPRK (ref. China)				0.000
				(0.006)
Actor motivation: Financial (ref. Espionage)				-0.000
				(0.000)
Included effects:	Entity, Time	-	-	-
Absorbed covariates:	Weekday	-	-	-
Observations	25079007	25079007	25079007	25079007
N. of groups	248261			
R^2	0.000	0.021	0.021	0.021
Within R^2	0.021			
Adjusted R^2		0.021	0.021	0.021
Residual Std. Error	0.001	0.059	0.059	0.059
F Statistic	11183.947***	272792.168***	60621.943***	45466.484***

Note: Standard errors in parentheses.
* p < 0.1; ** p < 0.05; *** p < 0.01.