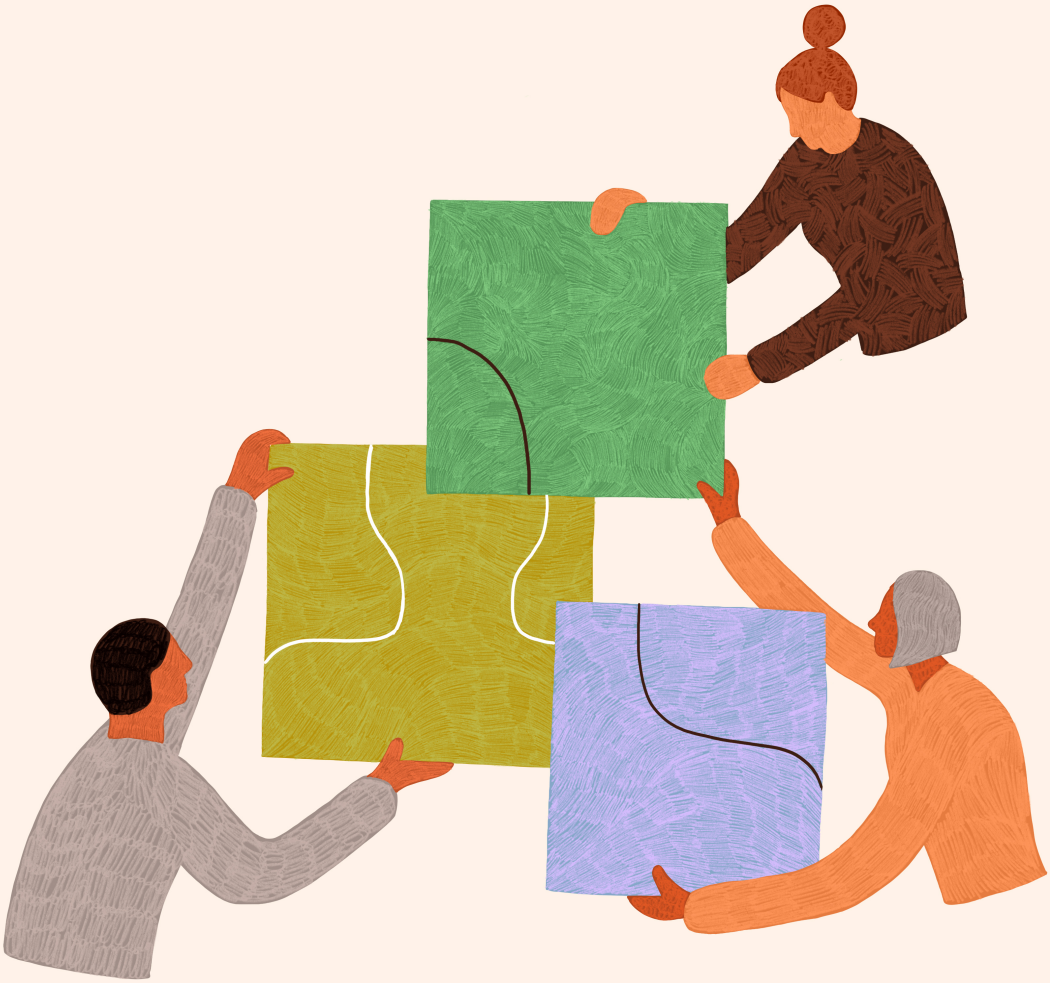
The Economics of Cybersecurity Threat Intelligence



X.B. Bouwman

THE ECONOMICS OF CYBERSECURITY THREAT INTELLIGENCE

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus,
Prof.dr.ir. H. Bijl,
chair of the Board for Doctorates
to be defended publicly on
Tuesday, 15 December 2026, 15:00

by

Xander Bastiaan BOUWMAN

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus	Chairperson
Prof.dr. M.J.G. van Eeten	Delft University of Technology, promotor
Prof.dr.ing. A.J. Klievink	Leiden University, promotor
Independent members:	
Prof.dr. G. Smaragdakis	Delft University of Technology
Prof.dr. D.W.J. Broeders	Leiden University
Prof.dr. P.A.L. Ducheine	University of Amsterdam
Prof.dr. R.G. Holz	University of Münster
Prof.dr.mr. J.A. de Bruijn	Delft University of Technology, reserve member

The work presented in this dissertation was supported by funding from the Ministry of the Interior and Kingdom Relations of the Netherlands.



Keywords: Threat intelligence, Information sharing, Internet measurements

Printed by: Ridderprint

Cover image: Sjoerd van Leeuwen

Copyright © 2026 by X.B. Bouwman

An electronic version of this dissertation is available at
http://repository.tudelft.nl

There was never a plan. There was just a series of mistakes.

— Robert A. Caro [1]

CONTENTS

Summary ix

Samenvatting xi

1 Introduction 1
- 1.1 Background . . . 1
- 1.2 Knowledge gaps . . . 5
- 1.3 Approach . . . 6

2 Market Incentives 11
- 2.1 Introduction . . . 12
- 2.2 Background . . . 13
- 2.3 Ethics . . . 15
- 2.4 Methodology . . . 15
- 2.5 Description of paid TI . . . 19
- 2.6 Comparison with open TI . . . 23
- 2.7 Uses and value of TI . . . 25
- 2.8 Discussion . . . 29
- 2.9 Related work . . . 30
- 2.10 Limitations . . . 31
- 2.11 Conclusions . . . 32

3 Peer Production 39
- 3.1 Introduction . . . 39
- 3.2 The CTC Community . . . 41
- 3.3 Production of the blocklist . . . 43
- 3.4 Evaluation of the blocklist . . . 45
- 3.5 Impact of the CTC blocklist . . . 52
- 3.6 Ethics . . . 56
- 3.7 Discussion . . . 57
- 3.8 Limitations . . . 59
- 3.9 Related work . . . 60
- 3.10 Conclusions . . . 61

4 Security Outcomes 63
- 4.1 Introduction . . . 64
- 4.2 Related work . . . 66
- 4.3 Methodology . . . 68
- 4.4 Measurements of sightings . . . 73
- 4.5 Statistical modeling . . . 79

4.6 Discussion . . . 81
4.7 Limitations . . . 84
4.8 Conclusions . . . 85

5 Conclusion **87**
5.1 Consolidated findings . . . 87
5.2 Central research question . . . 88
5.3 Implications for actors in the market for threat intelligence . . . 90
5.4 Future work . . . 97

Acknowledgments **99**

Authorship Contributions **101**

A Appendices Chapter 2 **105**
A.1 Interview protocol . . . 105
A.2 Codebook . . . 105

B Appendices Chapter 3 **109**
B.1 COVID-19 keyword list . . . 109

C Appendices Chapter 4 **111**

Bibliography **125**

SUMMARY

This dissertation investigates how economic incentives in the market for cybersecurity threat intelligence (CTI) shape security outcomes. It analyzes CTI across three production models: commercial firms, open-source, and sharing communities. Drawing on computer science, security economics, and international relations, it considers CTI not only as a technical input to intrusion detection, but also as a product whose production and use are shaped by human behavior and market incentives.

Chapter 2 finds that high-end commercial CTI vendors provide differentiated, non-overlapping data sets rather than comprehensive coverage. This suggests that not only customers but also vendors may be unable to properly assess the quality of threat intelligence. Interviews with security professionals show that, in the absence of clear measures of success, consumers prioritize workflow optimization over security outcomes, for example, by preferring sources with low false-positive rates.

Chapter 3 evaluates the largest CTI sharing community to date, consisting of thousands of volunteers aiming to aggregate data. The results show that community models can generate unique data, but struggle to maintain this uniqueness when they rely on existing verification mechanisms such as VirusTotal. In the case study, automated vetting helped keep false-positive rates low, but also filtered out novel indicators and shifted the community's output toward generic internet abuse already tracked by other sources.

Chapter 4 empirically tests high-end commercial sources against the 'impose cost' doctrine. It finds that 81% of commercial Indicators of Compromise (IOCs) that led to sightings on the studied networks were published only after threat actors had abandoned the associated infrastructure, challenging the efficacy of CTI for real-time detection. Although statistical modeling indicates a modest reduction in sightings following publication, this decline often began before publication and cannot be interpreted as evidence of a direct effect. In other words, the CTI studied in this case captured mostly historic threats, for which attackers had already cycled through to new infrastructure.

Taken together, the findings suggest that CTI sources are valuable complements rather than substitutes. We assert that because neither buyers nor sellers of CTI have good measures of quality, the market is more characterized by mutual uncertainty than vendor-exploited information asymmetry. This raises broader questions: when is delayed, curated intelligence preferable to fast but noisy information? Are organizations evaluating CTI's contribution to security, or merely its fit with existing workflows? And does public reporting disrupt adversaries, or principally document campaigns that have already ended? In the conclusion, we call for greater market transparency and improved methods for evaluating organizational outcomes to answer these questions.

SAMENVATTING

Dit proefschrift onderzoekt hoe economische prikkels in de markt voor *cybersecurity threat intelligence* (CTI), ofwel cybersecurity dreigingsinformatie, de uitkomsten van informatiebeveiliging beïnvloeden. Het analyseert dreigingsinformatie uit drie productiemodellen: commerciële dienstverleners, open bronnen en *communities* die dreigingsinformatie delen. Vanuit de informatica, *security economics* en internationale betrekkingen beziet het CTI niet alleen als technisch middel voor detectie, maar ook als product waarvan productie en gebruik worden gevormd door menselijk gedrag en marktprikkels.

Hoofdstuk 2 stelt vast dat hoogwaardige commerciële dienstverleners elk gedifferentieerde, niet-overlappende data leveren in plaats van een allesomvattend product dat het gehele dreigingslandschap dekt. Dit suggereert dat niet alleen de klanten, maar ook de commerciële aanbieders zelf mogelijk niet in staat zijn de kwaliteit van *threat intelligence* te beoordelen. Interviews met beveiligingsprofessionals laten zien dat afnemers, bij gebrek aan duidelijke maatstaven voor succes, hun eigen werkprocessen boven beveiligingsresultaten optimaliseren, bijvoorbeeld door bronnen met weinig fout-positieve meldingen te verkiezen.

Hoofdstuk 3 beoordeelt de grootste *threat-sharing community* tot nu toe, bestaande uit duizenden vrijwilligers die als doel hadden om dreigingsinformatie te aggregeren. De resultaten laten zien dat zulke gemeenschappen unieke data kunnen genereren, maar moeite hebben om die toegevoegde waarde te behouden wanneer zij afhankelijk zijn van bestaande verificatiemechanismen, zoals VirusTotal. In deze case-study bleek dat geautomatiseerde validatie het aantal fout-positieve meldingen laag kon houden, maar dat dit ook nieuwe indicatoren wegfilterde. Hierdoor verschoof de output van de gemeenschap naar meer generieke vormen van internet-misbruik die andere bronnen al volgden.

Hoofdstuk 4 toetst hoogwaardige commerciële bronnen empirisch aan de '*impose cost*'-doctrine. Hieruit blijkt dat 81% van de commerciële *Indicators of Compromise* (IOCs) die leidden tot alerteringen in de bestudeerde netwerken pas werden gepubliceerd nadat actoren de bijbehorende infrastructuur al hadden verlaten. Dit plaatst vraagtekens bij het nut van CTI voor *real-time* detectie van cybersecuritydreigingen. Onze berekeningen wijzen weliswaar op een bescheiden afname van alerteringen na publicatie van kenmerken, maar deze afname begon vaak al vóór publicatie en kan niet als bewijs voor een direct effect worden uitgelegd. Met andere woorden, de CTI die dit hoofdstuk beschrijft, ving vooral historische dreigingen waarvoor de aanvallers al waren doorgewisseld naar nieuwe aanvalsinfrastructuur.

Gezamenlijk wijzen de bevindingen erop dat CTI-bronnen elkaar op waardevolle wijze kunnen aanvullen. Wij stellen dat, omdat noch kopers noch verkopers van CTI beschikken over goede manieren om kwaliteit te meten, de markt eerder wordt gekenmerkt door wederzijdse onzekerheid dan door informatie-asymmetrie die door vendors wordt uitgebuit. Dit roept bredere vragen op: wanneer heeft vertraagde, zorgvuldig samengestelde dreig-

ingsinformatie de voorkeur boven snelle maar ruisrijke informatie? Meten organisaties de bijdrage van CTI aan cybersecurity, of zoeken zij vooral CTI die aansluit bij hun bestaande werkprocessen? En verstoren publieke CTI-rapporten tegenstanders daadwerkelijk, of documenteren zij vooral campagnes die al zijn beëindigd? In de conclusie roepen wij op tot meer transparantie in de markt en betere manieren om de organisationele uitkomsten te evalueren om deze vragen te beantwoorden.

1

INTRODUCTION

One day after the Democratic National Committee announced that it had been hacked in the lead-up to the 2016 U.S. presidential election, 'lone hacker' Guccifer 2.0 took credit for the intrusion, later leaking emails and documents of DNC staffers online [2, 3]. Although the mysterious hacker claimed to be of Romanian descent, cybersecurity threat intelligence firm CrowdStrike, which had investigated the DNC's systems, stood by its original assessment that Russian intelligence operatives were responsible for the operation based on similarities it had found to earlier Russian operations. The U.S. intelligence community later declassified part of a report confirming the involvement of the Russian military intelligence service, the GRU [4, 5].

The significance of these and other stolen emails for the 2016 presidential elections is material for another dissertation, but the role of CrowdStrike in understanding the DNC hack illustrates that Cybersecurity Threat Intelligence (CTI) is more than a category of products and services; it is the lens through which society understands major cybersecurity events. Threat intelligence firms attribute intrusions to actors, shaping political discourse of these events. Furthermore, by investigating certain events and not others, threat intelligence firms also influence which events are worthy of public attention [6, 7]. In addition to this public-facing role, threat intelligence is essential to daily defensive cybersecurity operations. It is integrated into environments to achieve security outcomes – such as informing automated intrusion detection and guiding resource allocation in Security Operations Centers (SOCs) – and to provide decision advantage for strategic business choices, fueling a significant and growing global market, now worth billions of dollars. Threat intelligence is interdisciplinary, and practitioners are found not only in commercial firms. Sharing in the threat intelligence ecosystem is driven by a combination of altruism and commercial agreements [8], and so various types of economic agents are involved.

1.1. BACKGROUND

Threat intelligence is practitioner-led and although there exists no rigorous definition, others have described it as "information about attacks and groups of perpetrators" and

1

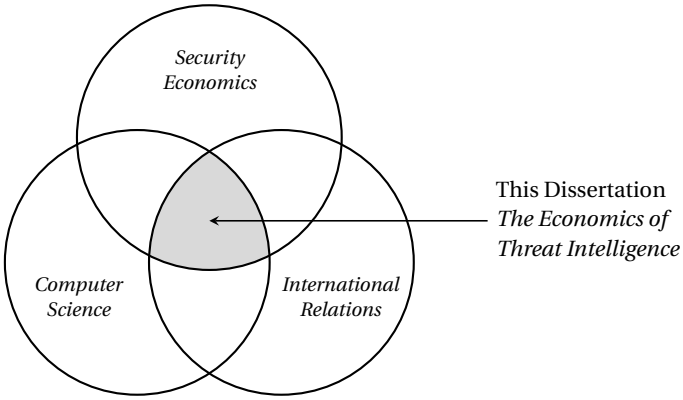


Figure 1.1: This dissertation focuses on the interplay between three bodies of literature, each of which provides a distinct perspective on cybersecurity threat intelligence.

"knowledge of attacker's activities [ranging] from a simple narrative around a threat actor's motivations, through to in-depth technical descriptions of an attacker's tactics, techniques and procedures" [9, 10]. Two popular conceptual models – the Cyber Kill Chain and the Diamond Model of Intrusion Analysis – emphasize, respectively, the stages of a cyber operation and the socio-technical (threat actors, victims) and technical factors (capabilities, infrastructure) of an intrusion [11, 12].

This dissertation defines cybersecurity threat intelligence simply as **the outcome of an analytical process investigating cyber threat actors**. Because these outputs – be it blogs, reports, indicators of compromise or otherwise – have economic value, they are sold commercially as well as shared on the basis of quid pro quo. Therein lies the focus of this dissertation. What determines the value of threat intelligence? And to which economic incentives do the economic actors involved in the 'market' for threat intelligence respond?

To answer these economic questions, one must look at the actors involved. This dissertation might have been called 'the ecosystem of threat intelligence'. But it is not a how-to guide, nor an anthropological study of threat intelligence researchers. The focus of this work is on the market dynamics that bring about threat intelligence and how they shape its outcomes. Like the subject being studied, this research is interdisciplinary. It primarily focuses on the interplay among three bodies of literature, each with a different perspective on threat intelligence: computer science, security economics, and international relations.

COMPUTER SCIENCE

Within computer science, threat intelligence exists as an applied subfield of cybersecurity. It is not a traditional academic discipline but rather an industry practice that academia is retrospectively analyzing and attempting to formalize. Its central premise is that sharing threat information between organizations improves overall security. The dominant intuition frames threat intelligence primarily as an input to intrusion detection systems: to identify adversaries on networks, defenders must know their signatures, tactics, and indicators of compromise (IOCs). This hierarchy of indicators is often conceptualized as the "Pyramid of Pain" [13], which distinguishes between simple indicators like hash

values and Internet Protocol (IP) addresses – which are easy to detect but also easy for attackers to change – and complex Tactics, Techniques, and Procedures (TTPs), which are harder to detect but impose more cost on the adversary when blocked, establishing a link between technical indicators and the strategic objective of cost imposition.

Computer science has studied threat intelligence mainly from the perspective of quantifying and measuring such threat intelligence inputs. Scientific articles have described at least 32 quality dimensions of threat intelligence data, but most researchers agree on three main metrics: timeliness, coverage (or: completeness), and accuracy [14, 15, 16]. Coverage is often found to be low, with various studies showing that threat intelligence feeds have very little overlap [17, 15, 18]. This fragmentation implies that each individual feed provides only a partial view of the threat landscape. Much effort has also been devoted to standardizing the automated exchange of this data, resulting in dedicated threat intelligence platforms that facilitate sharing, often in the de facto STIX 2.0 format [19, 20]. However, the automated ingestion of these indicators sometimes leads to "alert fatigue" in Security Operations Centers (SOCs), where analysts are overwhelmed by the volume of low-fidelity alerts [21].

But while computer science excels at objective performance metrics, CTI resists them. The three main quality dimensions should be understood relative to the requirements of the defensive environment. Moreover, standard metrics like precision and recall are difficult to apply because there is no "ground truth" of the total universe of threats – defenders do not know what they do not know. Unlike traditional security engineering, which models adversaries as stochastic processes, threat intelligence views adversaries as agents with preferences and capabilities that make certain targets more attractive than others. This conceptual shift moves the problem beyond the traditional bounds of computer science and into the social sciences.

SECURITY ECONOMICS

Threat intelligence has a cost to produce, but it also has a value, yet this value is notoriously difficult for a buyer to ascertain pre-purchase due to the measurement problem described above. This creates a market rife with information asymmetries. Because the quality of threat intelligence is hard to evaluate even *after* purchase, this market mirrors Akerlof's "Market for Lemons" [22], where information asymmetry degrades market quality. Extending Akerlof's argument, Grigg argued that markets for security products and services exist where *neither* the vendor nor the buyer truly knows the objective quality. Rather than information asymmetry, such a market exhibits information insufficiency, leading Grigg to term it "Market for Silver Bullets" [23]. Markets characterized by either information asymmetry or information insufficiency can cause perverse incentives and are grounds for market failures. These give rise to institutions, in the sense of formal and informal rules, such as certifications and insurance [22, 23].

Economic incentives also shape how threat intelligence is shared. While broad information sharing is theoretically optimal for defense [24], in practice, the "free-rider problem" often discourages firms from contributing to public pools of intelligence [25]. Instead, sharing often occurs in closed groups or "trusted circles," transforming intelligence from a public good into a "club good" [26]. Furthermore, there is a fundamental misalignment of incentives: vendors are incentivized to sell subscriptions by demonstrat-

ing activity (e.g., volume of indicators), while defenders are incentivized to stop attacks, which requires high-fidelity, actionable intelligence rather than volume.

The first two chapters of this dissertation consider three types of sources of threat intelligence production and dissemination: commercial threat intelligence, open-source threat intelligence, and threat intelligence shared in peer groups. Each is characterized by different incentive structures. This is relevant because the first mechanism is often missing from research in security economics and public policy on cybersecurity information sharing, which focuses on the latter two. While individual organizations bear costs for producing threat intelligence, the benefits of sharing diffuse across the community of defenders [25, 24]. The third study in the dissertation analyzes the effects on security outcomes and finds evidence of such positive externalities for commercial threat intelligence (Chapter 4).

INTERNATIONAL RELATIONS

Threat intelligence markets exist within a broader geopolitical context that shapes both threat production and consumption. As the first paragraph of this chapter has argued, threat intelligence has become the lens through which political events are understood, and as such, threat intelligence itself can be seen as a component of modern statecraft. Many significant threat actors are not financially motivated cybercrime groups but state-sponsored or state-affiliated groups, using CTI-tracked activities as instruments of national power [27]. Although tying cyber threat actors to responsible real-world organizations was first believed to be unfeasible due to the 'attribution problem', between 2010 and 2020, states became more willing to name and shame their adversaries as a political tool [28, 6]. This period in time corresponds roughly to the emergence of the commercial threat intelligence industry.

Threat intelligence practitioners furthermore use concepts and theories from international relations to understand threat actor behavior, such as coercion, signaling, and cost imposition. Unlike traditional intelligence, which derives value from secrecy, CTI is increasingly used for public signaling. Strategies like Persistent Engagement rely on the public exposure of adversary operations to impose costs and friction [29, 30]. In this ecosystem, private CTI firms play a unique role. By attributing state actions and shaping narratives around cyber conflict, they effectively function as "private intelligence agencies" [7], wielding significant influence over how geopolitical events are interpreted by the public and policymakers.

This dissertation empirically tests the validity of these strategic assumptions, specifically whether the public dissemination of threat intelligence actually imposes measurable costs on adversaries.

CONCEPTUAL MODEL

Figure 1.2 provides a generalized overview of a market for threat intelligence, integrating supply and demand into the traditional intelligence cycle [31].

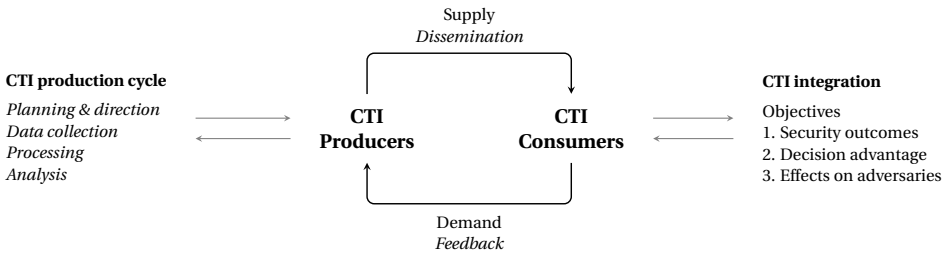


Figure 1.2: Conceptual framework of the market for cyber threat intelligence (CTI)

1.2. KNOWLEDGE GAPS

This dissertation analyzes the market structure, production processes, and operational impact of cybersecurity threat intelligence. It seeks to answer the central research question: *How do the economic incentives in the market for cybersecurity threat intelligence shape security outcomes?* This question is decomposed into three knowledge gaps, which the dissertation addresses with three corresponding empirical chapters. Together, these three studies provide a comprehensive analysis of the CTI ecosystem, moving from its production incentives (Chapter 2 & 3) to its ultimate, measurable impact (Chapter 4).

The scope of this dissertation is twofold: first, to analyze the economic and socio-technical incentives that shape the CTI production cycle, and second, to empirically measure the resulting impact when CTI is integrated by consumers, specifically observing defensive security outcomes and the ability to 'impose costs' on adversaries.

GAP 1: COMMERCIAL PRODUCTION

Previous efforts to formalize and measure the quality of threat intelligence have focused almost exclusively on open threat intelligence, such as public abuse feeds and blocklists. High-end commercial threat intelligence has remained largely absent from the state of the art due to its closed nature and the high costs of subscriptions. While commercial threat intelligence feeds are costly, their actual composition and value proposition relative to open sources are empirically unverified. It is unknown to what extent their data meets quality dimensions, and how the market's incentive structures influence their product. Furthermore, as explored in the conceptual framework, it is unclear how consumers perceive value in a market characterized by information asymmetry. Without understanding what is being sold, we cannot assess whether the market is efficient or if it is a "market for lemons".

Knowledge Gap 1: *What are the measurable characteristics of commercial threat intelligence compared to open sources, and how do consumers perceive its value?*

GAP 2: COMMUNITY PRODUCTION

Prior work on threat intelligence sources has consistently demonstrated the problem of fragmentation and low coverage among individual feeds. While data sharing is a frequently proposed solution to overcome this limitation, empirical investigations have

typically evaluated closed, trusted-tier exchanges or targeted crowdsourcing platforms. As a counterpoint to the commercial market, it remains an open question whether large-scale, open, community-driven models can overcome the fragmentation and quality challenges seen in threat intelligence and defensive security controls produced under market incentives. The socio-technical processes of community-based vetting and quality control are poorly understood. If community models can outperform commercial ones, it would suggest a failure of the market mechanism to provide optimal security outcomes.

Knowledge Gap 2: *Can community-driven CTI models improve coverage and actionability compared to existing defenses?*

GAP 3: TIMELINESS AND SECURITY OUTCOMES

Much of the scientific literature around threat intelligence measures properties of the sources of CTI itself – such as accuracy, coverage, and timeliness – or evaluates sightings only *after* the publication of indicators.

A significant gap in the literature is the disconnect between these measured properties and their ability to influence strategic security outcomes. Specifically, this raises questions regarding an unproven assumption in cyber conflict studies that holds that attributing and publishing threat intelligence can disrupt active campaigns and force threat actors to adapt their behavior by abandoning resources, taking up their scarce resources and thus contributing to strategic objectives. If CTI does not lead to measurable changes in adversary behavior, the strategic value of "naming and shaming" and public reporting must be reconsidered. To empirically test these assumptions, it is necessary to study high-end commercial CTI. Because commercial providers are generally believed to have the most comprehensive coverage of advanced persistent threats and sophisticated campaigns, their intelligence products are best positioned to shed light on whether public reporting actually forces adversaries to abandon infrastructure.

Knowledge Gap 3: *Can security outcomes of threat intelligence be empirically validated? Specifically, is commercial threat intelligence sufficiently timely to enable real-time intrusion detection, and does this cause threat actors to abandon their infrastructure?*

1.3. APPROACH

1.3.1. STUDY 1: COMMERCIAL PRODUCTION (CHAPTER 2)

Chapter 2 addresses this gap by answering the following questions: (i) What do paid threat intelligence (TI) services consist of? (ii) How is paid TI different from open TI? And (iii) how do customers use CTI and perceive value? It provides the first empirical measurements of high-end commercial threat intelligence sources and compares them to open source threat intelligence. The findings suggest that even market leaders struggle to provide substantial coverage of cyber threats. Instead, the firms provide largely unique IOCs, and so point customers to non-overlapping slices of the threat landscape. This empirical analysis supports the "market for silver bullets" hypothesis.

In addition to investigating PTI and OTI as sources of threat intelligence, the chapter provides insights into demands put on them by the consumers of threat intelligence through interviews, e.g. with SOC analysts. A core demand that consumers mention is the need for CTI to "provide context" to their work. Interestingly, the limited coverage

of paid sources was not of much concern to consumers. Instead, the chapter finds that consumers optimize for low false-positive rates, optimizing their workflow, rather than demand comprehensive coverage.

1.3.2. STUDY 2: COMMUNITY PRODUCTION (CHAPTER 3)

Chapter 3 investigates the effectiveness of non-commercial production models for CTI. It addresses two questions: (i) By pooling data from its community, did the CTC improve coverage of COVID-19-related threats over existing defenses? And (ii) Did publishing threat data in a freely available blocklist improve the ability of defenders to act against threats, compared to the existing abuse mitigation infrastructure? This chapter describes the results of a large-scale measurement study performed during the period of COVID lockdowns, where the opportunity arose to perform a 'natural experiment' because CTI specialists came together in the largest sharing community to date in order to collectively combat threat actors that were targeting healthcare institutions.

Measurements performed on the threat intelligence created by the community found that community models succeeded in adding unique data about novel threats, as compared to existing sources of threat intelligence, but that they face their own socio-technical challenges. This chapter demonstrates a similar prioritization of workflow optimization over coverage, showing the challenge is not exclusive to commercial markets. The community scaled up operations using an automated vetting workflow, based on submitted IOCs meeting a threshold on VirusTotal, in order to ensure accuracy (i.e., low false-positives). Inadvertently, this automation filtered out novel indicators, causing the focus of the created threat intelligence to drift from unique new CTI toward threats like regular phishing that other sources were already tracking.

1.3.3. STUDY 3: SECURITY OUTCOMES (CHAPTER 4)

Chapter 4 describes the results of a study that measured threat actor activity on IOCs in commercial threat intelligence feeds, asking: Do threat actors abandon resources after these have been published as indicators of compromise in commercial threat intelligence? To answer this, the study empirically tests the strategic assumption that disseminating CTI can disrupt active campaigns and force adversaries to adapt their behavior. Crucially, the methodology captured activity both before and after publication of threat intelligence to commercial feeds, allowing claims about the timeliness of this CTI and validating its operational capability for real-time intrusion detection.

Findings in this chapter challenge the efficacy of CTI for real-time intrusion detection and its systemic influence on real-world threat actor behavior. The study reveals that the vast majority of threats captured by market-leading CTI providers had already occurred in the past; for the small subset of published IOCs that generated any sightings on monitored networks, over 80% were only seen prior to their publication date. Because most indicators had lost their operational relevance before defenders received them, their ability to enable active defense and thereby impose costs on adversaries was limited. Furthermore, while panel regression modeling demonstrated a statistically significant downward effect on the probability of sightings following publication – suggesting that threat actors may eventually abandon exposed infrastructure – this effect typically mirrored a natural decay that had commenced well before the intelligence was published. These empirical results

cast doubt on the timeliness of commercial CTI as a reliable tool for real-time defense, and prompt a reconsideration of the expectations placed on 'naming and shaming' via intelligence reporting as a mechanism for cost imposition.

LIST OF PUBLICATIONS

Table 1.1 provides a list of publications of the author. Three of these publications form the chapters of this dissertation. Contributions of the co-authors are specified in the appendix on Authorship Contributions (Appendix 5.4).

Table 1.1: List of publications of the author and corresponding chapters.

Chapter	Publication
Ch. 2	Xander Bouwman, Harm Griffioen, Jelle Egbers, Christian Doerr, Bram Klievink, and Michel van Eeten. "A different cup of TI? The added value of commercial threat intelligence". In: *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 433–450. ISBN: 978-1-939133-17-5. URL: https://www.usenix.org/conference/usenixsecurity20/presentation/bouwman
Ch. 3	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. "Helping hands: Measuring the impact of a large threat intelligence sharing community". In: *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 1149–1165. ISBN: 978-1-939133-31-1. URL: https://www.usenix.org/conference/usenixsecurity22/presentation/bouwman
Ch. 4	Xander Bouwman, Aksel Ethembabaoglu, Bart Hermans, Carlos Hernandez Ganan, and Michel van Eeten. "Can IOCs Impose Cost? The Effects of Publishing Threat Intelligence on Adversary Behavior". In: *32nd ACM Conference on Computer and Communications Security (CCS)*. ACM CCS 2025. United States: ACM, 2025. DOI: 10.1145/3719027.3765026. URL: https://www.sigsac.org/ccs/CCS2025/
–	Stephanie de Smale, Rik van Dijk, Xander Bouwman, Jeroen van der Ham, and Michel van Eeten. "No One Drinks From the Firehose: How Organizations Filter and Prioritize Vulnerability Information". In: *2023 IEEE Symposium on Security and Privacy (SP)*. 2023, pp. 1980–1996. DOI: 10.1109/SP46215.2023.10179447

2

Market Incentives

Commercial threat intelligence is thought to provide unmatched coverage on attacker behavior, but it is out of reach for many organizations due to its hefty price tag. This paper presents the first empirical assessment of the services of commercial threat intelligence providers. For two leading vendors, we describe what these services consist of and compare their indicators with each other. There is almost no overlap between them, nor with four large open threat intelligence feeds. Even for 22 specific threat actors – which both vendors claim to track – we find an average overlap of only 2.5% to 4.0% between the indicator feeds. The small number of overlapping indicators show up in the feed of the other vendor with a delay of, on average, a month. These findings raise questions on the coverage and timeliness of paid threat intelligence.

We also conducted 14 interviews with security professionals that use paid threat intelligence. We find that value in this market is understood differently than prior work on quality metrics has assumed. Poor coverage and small volume appear less of a problem to customers. They seem to be optimizing for the workflow of their scarce resource – analyst time – rather than for the detection of threats. Respondents evaluate TI mostly through informal processes and heuristics, rather than the quantitative metrics that research has proposed.

This chapter was previously published as a research paper: Xander Bouwman, Harm Griffioen, Jelle Egbers, Christian Doerr, Bram Klievink, and Michel van Eeten. "A different cup of TI? The added value of commercial threat intelligence". In: *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 433–450. ISBN: 978-1-939133-17-5. URL: https://www.usenix.org/conference/usenixsecurity20/presentation/bouwman

2.1. INTRODUCTION

Cyber threat intelligence (TI) has acquired a strong presence in the market for security services. TI is, simply put, information on attacker behavior that can be used to adapt one's defenses to the threat landscape. A well-known form of TI are indicators of compromise (IOCs): machine-readable data feeds with resources – typically IP addresses, domains, or file hashes – that have been observed in malicious behavior. Commercial vendors also regularly release analyst-focused reports which go beyond indicators and paint a picture of the tactics, techniques and procedures (TTPs) of specific actors.

While organizations can generate limited forms of TI from their internal systems, they increasingly turn to procuring TI from external sources. Roughly speaking, there are three sources of TI: open, shared and paid. Open TI (OTI) typically consists of public lists of indicators, such as Abuse.ch [35], AlienVault [36], and Malwaredomains.com [37]. Here, threat intelligence is often another name for abuse feeds and blocklists. Shared TI (STI) is sourced via trusted communities where members exchange their own threat information, and where there is no payment associated with this exchange. The community can be formalized, such as the membership in an Information Sharing and Analysis Center (ISAC), or it can be informal, where membership is based on personal trust relationships. The third source is paid TI (PTI). A 2019 survey among 1,908 IT and security professionals in North America and the U.K. found that 44% of respondents say that the primary source of threat intelligence in their organization is purchased [38]. The commercial market for TI products and services was estimated at over USD 10 billion globally in 2025 and was projected to continue rapid growth in the coming years [39].

Despite its importance, the commercial market for TI is largely uncharted territory in academic research. In light of the high fees and license restrictions associated with PTI, this is understandable. The nearest work, conducted by Li et al. [40], was not focused on PTI specifically, but did include two edge cases of paid services. These services did not provide original TI sources, but helped curate and aggregate otherwise free or low-end indicator feeds. They were priced in the range of USD 1-10k per year. In this paper, we focus on high-end, original TI sources from market leaders, which charge around USD 100-600k per year (Section 2.5.2). Except for the edge cases, prior research has analyzed only open sources in the form of abuse feeds and blocklists. Recent work [40, 41] confirmed the results of earlier studies [42, 43, 44]: the feeds are highly heterogeneous, overlap among any two feeds is often very small and never more than 10%, and the feeds are dominated by 'singletons'; entries that do not appear in any other feeds. This also explains why, in order to achieve better coverage of their threats, network defenders combine on average 7.7 TI sources [38].

No prior work has analyzed what TI services the high-end vendors in the commercial market offer, how their data compares to OTI and how customers evaluate PTI. This lack of insight also affects firms facing security investment make decisions. The subscription fees of the market leaders usually range in the hundreds of thousands of dollars per year [45]. Buyers have to decide if these substantial costs are worth it compared to open or shared TI. According to research firms Gartner and Forrester, purchasers are struggling to compare services [45, 46]. Prior work [40, 47, 48, 49, 50, 41] assumed that buyers are in need of quantitative metrics on criteria like coverage, volume, accuracy, and timeliness. This assumption, however, has not been empirically validated via a user study.

This paper reports on a mixed-method study that sheds light on the market for paid TI and reduces the information asymmetry currently confronting buyers. We present the first qualitative and quantitative analysis of commercial TI services. Using grounded theory, we analyze 14 interviews with PTI customers to understand what sources they are buying, how they use them and how they evaluate added value in the absence of independent analyses. We complement this user study with a high-level quantitative analysis of services of two market leaders, comparing these to several open TI feeds.

With this approach, we aim to answer the following questions: *(i)* What do paid TI services consist of? *(ii)* How is paid TI different from open TI? *(iii)* How do customers use TI and perceive value? We make the following contributions:

- We present the first empirical analysis of paid TI from market leaders, comparing the data of two leading vendors to each other and to OTI.

- We demonstrate that there is almost no overlap between paid and open TI sources, signaling that they capture a different part of the threat landscape. Surprisingly, there is also little overlap among the two paid feeds, although they focus on the same topic areas. Even when tracking the same 22 threat actors, only 2.5 % to 4.0% of indicators are found by both vendors, depending on the type of indicator. Timeliness, as measured on the small overlap among sources, shows delays of more than a month between sources. These findings suggest serious issues with coverage and timeliness of commercial TI.

- We find that customers use TI less exclusively for network detection than is often assumed. Other use cases include understanding of the threat landscape, informing business decisions, awareness programs, and threat hunting.

- Where prior work has assumed that customers work with metrics like coverage in order to evaluate TI, we find a different logic in practice. Customers value the better curated and more selective paid TI sources over other sources – especially the larger and potentially more noisy open sources – because they consume less analyst time. Surprisingly, the fact that smaller sets may also imply more false negatives, i.e., limited coverage, is much less a concern. Costs hardly play a role at all.

2.2. BACKGROUND

A major challenge in an organization's risk management process is to identify and understand all relevant threats. For many cybersecurity threats, their existence, likelihood and impact are not known to the organization. Threat intelligence services claim to address this by providing the necessary information to identify risks, aid in their quantification, guide the selection of controls, provide indicators to detect adversaries, and show possible courses of action.

To a limited extent, TI can be extracted from an organization's own security controls. Think of firewall logs that observe external IP addresses involved in brute-forcing Secure Shell (SSH) passwords or of spam filters that contain emails with phishing uniform resource locators (URLs). Such IP addresses and URLs are typically referred to as indicators

of compromise (IOCs). The downside of producing in-house TI is that any single organization will only observe a small fraction of the threat landscape. Another challenge is that extracting the most relevant signals, rather than the most obvious ones, requires resources and expertise. For this reason, many organizations acquire external sources of threat intelligence, like open sources or sharing communities.

Compared to OTI and STI, paid threat intelligence makes a different value proposition. It does not only contain indicators and information observed from an ongoing threat somewhere else (e.g., an IP address that has brute-forced or phished a different organization), but insights based on active research, proprietary vantage points, and potentially insider information by a specialized provider. PTI is often perceived as being of higher quality and providing better and earlier warning. To protect value and exclusivity, vendors typically vet their customers, so that adversaries cannot readily see that their activities were detected. Vendors also typically provide integration into products, like a malware detection middlebox. This provides them with unique visibility across the networks of their clients, where they have probes for monitoring and middleboxes for protection. The vendors in this market typically claim that with these vantage points and cloud-based aggregation and analysis of data, they can track advanced attacks and threat actors. Vincenzo Iozzo, Senior Director at CrowdStrike, a key player in this market, articulated this advantage as follows: "If you [the attacker] get detected on one machine, all of your offensive infrastructure has to be scrapped" [51]. Later in this paper, we explore the extent to which this advantage allows PTI vendors to uncover offensive infrastructure.

In general, assessing quality of TI data through metrics is very hard, as there is no ground truth on global maliciousness [43]. Below we briefly describe some quality characteristics and metrics that were developed in earlier work, as these concepts return at various points in the paper.

Coverage pertains to the proportion in which the TI actually observes the attacks it promises to observe – i.e., the proportion to which the intended indicators are actually contained in the feed [40, 48]. The opposite is how much relevant information it fails to provide, which is the rate of false negatives [49, 50, 43].

Accuracy is the proportion of indicators in a feed that actually belong in the feed. This pertains to the degree of true positives. Its opposite is the degree of false positives [40, 49, 50] and this is a major factor in the value of a data-source. Depending on the organization, even a relatively low number of false positives may lead to notification fatigue and thereby reduce the feed value.

Timeliness of information in a feed pertains to the time gap between an attack vector occurring and its associated indicators being included in the feed [52]. Some authors refer to this as latency [40], or speed [49]. Timeliness may be essential for the value of a feed used in active defense mechanisms, e.g., intrusion detection middle boxes, but for forensics purposes this is less critical.

Ingestibility relates to the structure and consistency of structure in a feed [52], i.e., how well it can be automatically processed.

Relevance of a TI feed [52], also referred to as fitness [48], describes how well the indicators and contents of a feed fit an organization's use case. A TI feed rarely attempts to cover *all* malicious activity but often focuses on a certain type of threat. A feed may hence be of low value if it is not relevant to an organization.

In summary, assessing the quality of TI feeds is a difficult subject for OTI, and–due to limited availability–even more so for PTI. In the remainder of the paper, we will produce some of these metrics for PTI, replicating the approach of [40] (Sections 2.5 & 2.6), while also using interviews with users to see how organizations assess the value of PTI with or without such metrics (Section 2.6).

2.3. ETHICS

Research on PTI data is hindered by high fees and license restrictions. Firms who buy a subscription are not allowed to share the data with third parties, including researchers. We were able to overcome this barrier when one of the authors was set up as an intern in an organization that has a subscription to the TI of two market leaders. These two vendors are included in Gartner's market overview [45] and they are positioned among the most expensive suppliers. The analysis of the data was conducted on the organization's premises and within the conditions of their vendor license agreements. Only the aggregate results of the analysis were shared with the rest of the author team. The organization was willing to collaborate with our study on condition that we would not name them nor the vendors included in the study, and that we would not include characteristics of the feeds that would make the vendors easily identifiable. Hence, some numbers are reported as ranges rather than exact counts.

Our second data source consists of interviews with 14 security professionals who work with paid threat intelligence. We received approval from our Institutional Review Board for this human-subjects research. All respondents explicitly gave their consent to have their interview transcribed and used in this study. To enable our respondents to talk about their use and evaluation of PTI without risking reputational repercussions for themselves or their organizations, we have anonymized their identities. To provide context for specific quotes, we describe a respondent's role and sector. Respondents were provided with information on the research objectives and the interview protocol before the interview. Afterwards they could check and correct quotes attributed to them.

2.4. METHODOLOGY

We use a mixed-methods approach, combining a qualitative user study with a quantitative analysis of the TI data. To answer the first question – what does paid TI consist of? – we report on the answers from our respondents, rather than impose our own definition of TI. We complement these answers with a high-level description of the feeds and reports that were provided by two market leaders from 2013-2018. For the second question – how PTI compares to OTI – we analyzed indicator feeds of the two market leaders and four open feeds. The third question – how do customers use TI and perceive its value? – we answered based on our interviews. We answer the questions consecutively in Sections 2.5, 2.6, and 2.7. Here, we describe the data collection and analysis. A high-level overview is presented in Table 2.1.

2.4.1. THREAT INTELLIGENCE DATA

As described in Section 2.3, an internship of one of the authors allowed for access to the TI services of two market leaders, both included in Gartner's market overview [45].

Type	Data	Source	Contents	Period
PTI	Paid TI services	Two leading providers	7,308 reports; and 420,173 indicators (IPs, domains, MD5)	2013/01/01 to 2018/12/31
OTI	Alienvault OTX	Community aggregator	59,290 IPs	2018/10/01 to 2018/10/31
OTI	Blocklist.de	Independent	121,540 IPs	2018/10/01 to 2018/10/31
OTI	CINSscore	Security firm	55,906 IPs	2018/10/01 to 2018/10/31
OTI	Emergingthreats	Security firm	876 IPs	2018/10/01 to 2018/10/31
TI	14 interviews	Professionals using PTI	Qualitative findings	2019/08/27 to 2019/12/23

Table 2.1: Data sources for this mixed-methods study.

The organization which provided us with the access chose these vendors because they are among the market leaders and were deemed to have the most relevant TI. It did not conduct any analysis of the overlap in indicators among the two vendors before we started our research.

The offerings of the vendors consisted of 5-10 subsets around specific topic areas, e.g., 'financial industry', 'cyberespionage' or 'cybercrime'. Customers typically subscribe to the subsets most relevant for them, rather than to all. As explained in Section 2.3, we cannot identify the vendors, nor can we list the exact topic areas the customer organization subscribed to. We can only say that we had access to 3-5 subsets for each of the vendors and that these subsets focus on the same topic areas. The selection of these topic areas likely influences how the indicators are distributed over target industries, as visualized in Figure 2.4. The degree of indicator overlap between vendors might also vary across focus areas.

We assessed only the indicators that were packaged with the TI reports that the vendors release. These reports analyze the developments in the threat landscape and actor groups. One vendor also had a bulk feed of indicators that were not associated with reports. We did not include this in our comparison, because the other vendor provided no such feed.

We assessed the overlap between the paid TI sources as follows. Vendors label their intelligence products with metadata, referring to a specific threat actor for 35% of their reports and 60% of all indicators. Vendors use their own naming schemes for threat actors. The same actor may thus appear as Deep Panda, APT19, or KungFu Kittens across different vendor reports. We mapped the names used by the vendors to a common set of threat actors using an overview maintained by well-known security researcher Florian Roth [53]. About 30% of all indicators could be mapped to a common threat actor listed in the overview. These indicators form the basis for the analysis of the overlap visualized in Figure 2.3. In sum, we measure overlap specifically where the vendors claim that they are tracking the same actor groups. The results are reported in Section 2.5.1. In Section 2.6.2,

we also report on the overall overlap across the feeds.

To map the distribution of indicators across targeted industries, we mapped 179 labels from both vendors to a common set of 16 categories. These are listed in Figure 2.4.

To represent OTI, we collected four freely available sources. Three indicator feeds, for the high degree to which they were reused by other open sources [41] – which we expected might lead to reuse in paid TI– and one community-based aggregator that enables its users to extract indicators from blogs and reports, which vice versa we expected might lead to reuse of PTI indicators in OTI. This way, we selected for an upper bound of overlap between the two types. We compare one month of data from these collected OTI feeds (October 2018) against five months of indicators from the two PTI vendors (July to December 2018), in order to compensate for the higher churn in the OTI feeds and, again, to find an upper bound in overlap. The preparation of the OTI data consisted of deleting all duplicate IPs except for the first occurrences, then removing all IPs present on the first measurement day of each set, as it was unsure when those were first received. To remove inconsistencies from the threat intelligence, we normalized them as follows. For URLs and domains, we removed `http(s)://`, but kept prefixes that are part of the domain – including the `www` domain which may point to another location. For file hashes in indicators, we compared only MD5 hashes for each file to prevent duplicate results. For IPs, no normalization was necessary, as the format was consistent between sources. For the reports, we removed all punctuation and casing from labels.

When matching indicators, we had to assume a time-to-live period during which indicators could be reasonably assumed to remain valid. Prior work chose periods of 30 days [40] or does not make such assumptions explicit. Since overlap among sources is an important indicator of coverage, we proceeded conservatively and chose a time-to-live period of 360 days. This again to provide us with an upper bound, i.e. over- rather than under-estimating the amount of overlap.

2.4.2. INTERVIEWS

As the field of PTI is marked by ambiguous terms and complex practices, we opted for a data collection method that can help us unpack the ambiguity, namely semi-structured interviews in a grounded-theory approach – similar to [54]. Grounded theory means that the researchers draw conclusions through a reflexive process of inductive reasoning [55, 56]. This approach means that our findings can help form an understanding about how the market for PTI functions, but not about how the reported views are distributed across the global population of security professionals who use PTI.

Between August and December of 2019, we conducted 14 interviews with professionals who work with PTI. Participants were selected from different sectors in business and government (Table 2.2). A requirement was that their organization purchased commercial threat intelligence. We contacted participants via their personal networks, as well as via LinkedIn. Geographically, the participants were located in The Netherlands (11) and Japan (3). They fulfilled positions from analyst to management in both security operations and threat intelligence teams. Two of the respondents worked at Managed Security Service Providers (MSSPs) which use TI to protect the networks of their clients. In the sample, the financial industry supplies the most interviewees. While this might simply be an artifact of our recruitment effort, it is consistent with the fact that financial sector firms have

INDUSTRY	Respondents *n=14*
Finance	4
R&D	3
Government	2
Managed Security Service Provider	2
Infrastructures	2
Oil and gas	1

Table 2.2: Interviewees by industry, all experienced with PTI.

the highest investment levels in cybersecurity [57] and are thus more likely to acquire expensive PTI subscriptions. One of our respondents, a teamlead TI at an MSSP, remarked: "I would say that the financial industry is one of the most mature sectors and can do more with threat intel."

The interviews we conducted with security professionals were semi-structured, meaning that respondents' views were central and conversations were open-ended [58]. The researchers did not ask loaded questions to avoid steering the conversation [59, 56]. A simple interview protocol was used in which participants were asked about: *(i)* Their definition of threat intelligence; *(ii)* What commercial sources their organization pays for and, for each of the sources, discuss costs, source properties, use-cases, and valuable organizational outcomes; *(iii)* Their experiences with use of non-commercial sources; *(iv)* Their reasons for having discontinued a source, if ever. The full interview protocol is included in Appendix A.1.

This approach was chosen deliberately to minimize the influence of pre-conceived ideas about what TI is and how it should be evaluated. Due to the nature of open questions it also means that participants' answers cannot be seen as an exhaustive description of their opinion. Rather, outcomes of the interviews are the perceptions that participants have prioritized, that were 'top of mind' for them at the time.

We transcribed the interviews and coded them using the ATLAS.ti software. Analytic codes were drawn from the interviews and used as labels to identify recurring answers. The codebook was iteratively saturated over the course of 9 interviews [56]; it reached the point where no new codes could be observed from the interviews (see Figure 2.1). One researcher carried out initial coding. We developed the codebook through meetings with co-authors and two other researchers, each time independently coding and then refining codes as needed. According to [60], this is a suitable way to ensure reliability of findings for our purposes. This eventually led to analytic codes for TI service types, sources, use cases, and value perceptions. The codebook is reflected in Tables 2.3 and 2.4, and is included in Appendix A.2.

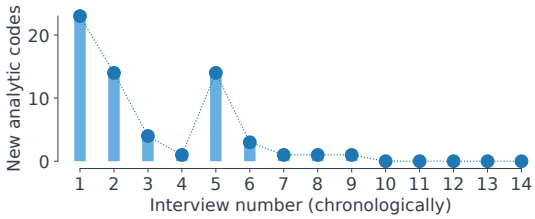


Figure 2.1: We reached analytical saturation after 9 interviews, as no more new codes were uncovered.

TI SOURCE TYPES	Respondents *n=14*
Paid threat intel providers (PTI)	100%
Open sources (OTI)	79%
Shared sources (STI)	64%
Government	50%
With product or service	50%
Collective procurement	36%
Own research	36%

Table 2.3: Prominent TI source types among our respondents are paid TI, open TI, and shared TI. We selected for the first.

2.5. DESCRIPTION OF PAID TI

We now address the first sub-question: What do paid TI services consist of? We coded the different answers to the interview question of what TI meant to the respondents. While prior work has focused on the indicator feeds, respondents also mentioned other PTI services, such as reports, requests for information, portals and custom alerts. For some, these services were more important than the indicators. We describe these different forms of PTI. For the two main services, indicators and reports, we also take a look at the offerings of the two PTI vendors in our study. We end with a brief exploration of price levels in the commercial market for threat intelligence.

2.5.1. TI SERVICES

Customers might subscribe to multiple TI services [38]. The variety of services reflects different needs in the market. As stated by a Team Lead TI at a bank: "Intelligence requirements differ per department. The SOC [Security Operations Center] would like to see indicators of compromise and to know TTPs [tactics, techniques and procedures], in order to understand what criminals are targeting, while the Risk department wants to know if these criminals have the capability and intent to disrupt our business, and if we are in control of those risks."

Vendors of paid TI attach metadata to their intelligence products. It describes what

INTELLIGENCE PRODUCTS	Respondents *n=14*
Indicators	71%
Reports	71%
Requests for information	57%
Portal	50%
Data mining and aggregation	29%
Custom alerts	14%

Table 2.4: Most respondents name indicators and reports as intelligence products that they receive.

industry a report relates to, as well as what threat actors the provider believes are involved. If the provider attaches indicators to the report, they will place in the metadata the degree of confidence they have that the indicator is malicious. This information is used to interpret TI and determine how it can be used.

Indicators are signals of attacker presence on a network. They are also referred to as indicators of compromise (IOCs). Examples are an IP address of known attacker infrastructure, the hash of a piece of malware or a domain associated with a phishing campaign. Indicators are provided in proprietary formats via an application programming interface (API), making them ingestable by detection systems such as a security information and event management (SIEM) system or an intrusion detection system (IDS). Indicators might be used for network-based or host-based detection, but also in different business processes, from security engineering to various business decisions.

The two PTI vendors in our study attach indicators to their reports. We find that the volumes of new indicators and reports of both vendors have steadily increased over the course of five years (Figure 2.2). The publication of indicators is unevenly distributed over time. For example, two hikes of the indicator volume can be explained by a vendor's introduction of new report types, leading them to publish many of the corresponding indicators at once. Indicators are also very unevenly distributed over the reports they are published with.

Coverage is the extent to which a TI source actually includes indicators for all the threats that the source intends to capture. In the absence of ground truth on all ongoing threats, we look at the overlap among the two vendors – a similar approach to Li et al. [40]. Less overlap means that each vendor is observing unique indicators that are missed by the other providers, suggesting limitations in the coverage of threats.

We analyzed the overlap for the same 22 actors that both vendors claim to track. The results are summarized in Figure 2.3. We find very low overlap – on average between 2.5 and 4.0%, depending on the indicator type. The overlap is unevenly distributed and mostly concentrated on a handful of actors. The highest overlap is 21.0% for the IP addresses of threat actor 21. This overlap is low, considering that these indicators are supposed to provide coverage of the activity of the same threat actors. These results beg the question

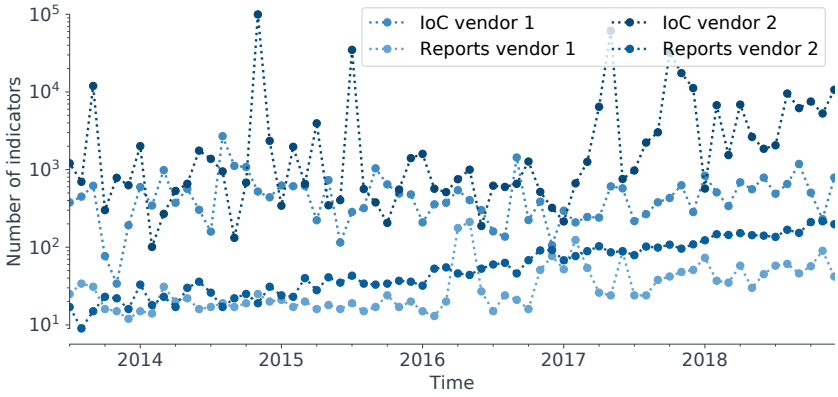


Figure 2.2: An upward trend is visible in the indicators and reports published by two leading paid TI providers between 2013 and 2019. Note that the Y-axis uses a log scale. Over 2018, a customer of one of the vendors might receive some 100 reports and 2500 indicators per month. The distribution of indicators over reports (and thus over months) is highly irregular.

what coverage these indicators actually provide of all malicious activity, by these actors and otherwise. We extend the analysis of overlap in Section 2.6 with OTI.

Reports come in various flavors. Malware reports describe the technical results of reverse-engineering captured binaries. Threat reports describe the goals and modus operandi of threat actors. Advisories and alerts describe current events, such as a software vulnerability being targeted or a threat group expanding their focus to another geographic region.

As can be seen in Figure 2.4, paid TI reports primarily consider threats to government, military, and financial services. These sectors likely form an important part of the customer base of PTI providers. A surprising third in subject matter was civil society, including non-governmental organizations (NGOs) and international organizations. Possibly this prominence could be explained by the political significance of the targeting of civil society, or by the relatively large number of such campaigns being observed.

In our interviews, respondents describe that reports of PTI providers are helpful on "all decks of the organization", from SOC analyst to Chief Information Security Officer (CISO). A manager at a managed service provider (MSP) described reports as the most important form of PTI, as they allowed his analysts to provide context for the clients of why certain alerts happened and what they mean: "We need to inform the customer if we get a critical alert. It may be 2 AM and we call the customer's mobile phone, only to find out that it's a false positive – that happens very often. Then we need to explain why this [security appliance] product says this, so we need some kind of reasoning. Then the customer asks us to filter out the event, or keep monitoring it."

Requests for information (RFIs) are inquiries from customers to the vendor's analysts. These were described by eight of our respondents as an important form of PTI. A request might work as follows. From reading a threat report, an analyst could be wondering about the relevance of the described threat for their organization. They might inquire with

the report's authors if campaigns in some specific sector or geographic area had been observed. The vendor would then search their own data and report back with information.

Our respondents described that requests for information are budgeted as part of the contract at around 10 inquiries per year. One respondent explained that, in practice, the PTI providers were willing to share information if they could, even without a formal RFI. The quota for inquiries were mostly a formality. Another respondent described that as part of their contract, they had been assigned an analyst at the vendor for 0.5 full-time equivalent (FTE). They would be always in touch with the same person, who over time came to understand their information needs.

Portals provide access to information that a PTI vendor has delivered over time. They consist of websites with historic data on threat actors and their campaigns, overviews of reports by target sector, as well as other data that the vendor may provide, such as indicators or malware samples. Our respondents describe that portals contain most of the information a vendor has, sometimes more than what can be requested from the API.

Data mining platforms and aggregators are effectively OTI as-a-service. These are subscription-based platforms on which customers can run queries, sometimes coming pre-loaded with open source security data. Some TI aggregators focus on the analysis tool (TI platform) as a product, and have a curated OTI feed as an additional service. Customers can plug in their own PTI data sources in these platforms. The categorization of such trade tools as 'intelligence products' is up for debate, but our respondents did indeed name them as part of their paid threat intelligence.

Custom alerts notify customers of specific risks to their organizations. An example is when a domain is registered that is similar to the customer domain, which could be used for typosquatting. Another example is when compromised credentials of the customer organization occur in credential dumps. We note that the two leading providers that we have analyzed do not offer this service. Providers that offer it may be targeting a different audience with their services, possibly customers that are in need of a managed TI capability, rather than external data sources as input to their existing TI team.

2.5.2. PRICING OF PTI

Public information on pricing of paid threat intelligence information is sparse. We did not identify a single instance of a PTI provider transparently providing pricing information on its website. In general, a recent Gartner report lists the services of the market leaders as upwards of USD 100,000 [45].

We collected 38 price points for 6 popular PTI providers [46] as well as 2 smaller providers. These are displayed in Figure 2.5. The data points were derived from publicly available schedule price lists [61, 62], as well as by requesting quotes from these vendors. Note that services offered by vendors are not directly comparable and that therefore this figure gives only a rough indication of pricing in the market as a whole.

On the right side ($100,000-$600,000 per year), we find high-end vendors which sell their own TI, while on the left ($30,000-$100,000 per year) we find paid aggregators, whose services primarily consist of providing a platform to integrate TI from other sources and to support analysts with analytics.

The wide bandwidth may be explained by pricing models and negotiations. Our respondents describe that pricing models are sometimes based on per-user licensing,

where costs increase as the number of analysts that have access to the provider's portal grows. Furthermore, pricing can be negotiated. A TI analyst at a major bank, said: "Vendor pricing is arbitrary. It's based on the size of the customer organization in most cases." Another respondent described a negotiation with a TI provider in which the asking price was lowered by a factor of ten.

Although there is leeway and room for negotiation, PTI currently seems restricted to enterprise organizations with large budgets for information security. A Team Lead TI at a bank shared: "The global costs of a CTI team is around €1-1.5 million per year, including 6-7 staff members and tooling costs. [...] Purchasing two or three feeds at €115,000 to €135,000 per year each may sound like a lot, but is actually not so bad for a bank."

2.6. COMPARISON WITH OPEN TI

We now address the second question: *How is paid TI different from open TI?* We define sources of OTI by the simple fact that they are freely available, but they are very diverse in nature. A news article may be seen as open source intelligence, as may a thread on a message board on which criminal activity is discussed. According to one analyst we spoke to, open TI included the cybersecurity podcast he listened to on his commute to work.

We then compare indicators from two paid providers and four prominent open source feeds, and establish that there is almost no overlap between the two types of sources.

2.6.1. REPORTS

Security researchers, also those at PTI vendors, post blogs, write-ups, and tweets that are accessible to everyone. These are basically the OTI analogue to the reports from PTI vendors. In fact, most vendors make a small fraction of their reports publicly available for marketing, moving them into the OTI domain.[1] These are the same reports as their customers receive, but without the associated indicator sets.

Open and paid sources are by no means decoupled in terms of the attackers they describe. Vendor reports often refer to work posted by OTI researchers, in some cases copying in screenshots of tweets that first observed certain attacker behavior. We find that paid providers often draw on such open sources, but their reports are much more complete in their description of the context, implications, and possible mitigation options. A manager at a research and development (R&D) institute in Japan had the view that this property makes paid TI useful especially for larger organizations: "[Paid] threat intelligence is useful for organizations wanting to know much more details about an attack, but using OSINT should be enough for the purposes of a SOC in a medium-sized company."

2.6.2. INDICATORS

Respondents collect indicators from open, shared, and paid sources, often loading these in their TI platform or SIEM system. In these systems, indicators are labeled with their origin, allowing analysts to interpret them, e.g. based on their source. Three respondents stated that they did not distinguish between paid and open sources in their detection

[1] Freely available vendor reports are indexed by the APTNotes project and can be referenced and searched on https://threatminer.org.

processes – indicators are indicators.

Indicators from open source feeds are more commonly referred to as blocklists or abuse feeds. Respondents discussed the confidence they place in TI sources. In this context, PTI was thought to be more 'accurate' than those from OTI. When pressed, it seems that respondents actually meant: more curated and smaller feeds, rather than more accurate. Smaller sets produce, by definition, fewer false positives. Respondents were emphasizing the element of accuracy that impacts the analyst's workflow, namely minimizing the number of false positives. The other side of accuracy, the rate of false negatives, was not mentioned, even though this rate might actually be higher for smaller and more curated feeds. Because of this perceived accuracy, PTI indicators are used with more confidence in detection and other use cases.

Overlap among paid and open sources is negligible, even though the OTI lists are vastly more voluminous than the lists of PTI indicators, as can be seen in Figure 2.6. This is relevant, because overlap helps us understand the coverage of a source – the level to which it captures the intended threats – as described in Section 2.5. The overlap between the individual sources is shown in Figure 2.7 as a fraction of the total volume of that source for that period. Vendor 1 and vendor 2 (PTI) share some indicators with each other – 1.3% and 13.0% respectively. This seems low, considering that the feeds are focused on the same topic areas in the overall landscape. It appears that PTI has the same pattern as OTI: a lack of overlap among feeds and the dominance of 'singletons'. Less than 1% of the PTI indicators overlap with any of the OTI sources. Vice versa, the OTI sources share indicators with each other, quite substantially in some cases, but there is basically no overlap with PTI: 0.0% of all OTI indicators are also observed in the PTI sources. As we describe in the Methodology, the uncommonly large amount of overlap between the OTI sources [40] is explained by the fact that we selected specifically for OTI that is often re-used [41] in hopes that this would increase overlap with PTI sources – to no such effect.

Timeliness of indicators means that the information is available to the customer early enough to actually detect and stop an attempt at compromise. While there is no ground truth as to when a particular resource (domain, IP, binary) was first used by a threat actor, we can assess timeliness via pairwise comparisons of the different feeds and measuring the delays in which indicators are made available to customers. For this analysis, we work with the small number of indicators that occur in more than one set.

We first compare the two PTI vendors (n=16 and n=28, respectively). In Figure 2.8, we can see drastic delays in when PTI vendors observe indicators. On average, it takes more than a month before an indicator observed by one vendor is also observed by the other vendor. In terms of defending against sophisticated threat actors, this is a very long delay. We find almost no instances where threat intelligence is distributed by other PTI within the same week. It seems that vendors do not use the TI from their competitors, or at least not successfully, to find the same indicators in their own telemetry.

We conducted the same analysis for the indicators of the two vendors versus the four OTI sources. The latter comparison is only based on the tiny overlap of 2 to 7 cases, which means that we cannot draw any strong conclusions. One might expect that paid providers would be faster in all cases, but PTI sources were faster in only half of the cases we have analyzed. So it seems PTI is not faster in finding indicators than OTI. The same appears to hold vis-à-vis STI, in the words of a team lead TI at an oil and gas company: "We also have

our own networks with companies in [our sector] and I must honestly say, we quite often get the information earlier there than from our [paid] intelligence providers."

While we have no ground truth as to when a particular resource (domain, IP, binary) was first used by an actor, the fact that we see major delays in all pairwise comparisons of our PTI and OTI sources suggests timeliness is a major problem.

2.7. USES AND VALUE OF TI

This section addresses our final question: *How do customers use TI and how do they perceive its value?* Given that customer organizations of PTI are paying a substantial amount of money compared to OTI and STI, they apparently value PTI to be worth the asking price. We tried to understand their perceptions of this value in two ways: by asking them how they used PTI (use cases) and by asking about what they see as strengths and weaknesses of their sources. These two ways are aligned with the economic distinction between 'stated preferences' versus 'revealed preferences'. The former infers preferences from what people explicitly state as preference, the latter from their actual choices and behaviors.

2.7.1. USE CASES OF TI

Based on the analysis of the interviews, we found 9 use cases for TI (Table A.1). The percentages refer to what percentage of the respondents mentioned this use case.

The top three use cases are central to SOC operations. **Network detection** (93% of respondents mentioned this) is still the main use of threat intelligence. This includes all instances in which TI is used to reduce attacker dwell time in an automated fashion, including correlating TI to logs, ingesting it in a SIEM or IDS, or using it in host-based detection controls.

Situational awareness was mentioned in two out of three interviews (64%) as a use case. This is the ability of TI or SOC analysts to have a general understanding of their organization's threat environment and risk profile. Situational awareness is broader than detection. It is relevant in the planning and direction phases of the intelligence cycle [63]. A Team Lead TI at a bank said it is not just about cybersecurity: "I believe that is too limited. It's about understanding who is a threat to my organization, not just the technical channel used for the attack."

SOC prioritization (50%) is a more practical use of TI, e.g. to assess how critical alerts are or to direct threat hunting efforts. This way, resources – especially attention of analysts – can be allocated toward most relevant threats.

Informing business decisions (36%) concerns uses of TI to improve organizational decision-making. For example, a CISO used TI to evaluate the return on various options to invest in security controls. But there are also organizations that use paid TI to assess the risks associated with a potential acquisition of international competitors, to gain a 'business decision advantage'. An analyst at a bank said: "Threat intelligence means engaging with various business units in order to understand their information needs, and then developing a way to answer those in a timely way. As a consultative practice within [our organization] we could provide information on geopolitical affairs, and intelligence on physical risks in a country that [we are] operating in. It doesn't necessarily have to

relate to cyber." This use case, like situational awareness, also underlines the value of reports compared to raw indicators.

Enrichment of own threat intelligence (36%) is where PTI is processed with the aim to improve the organization's own services – e.g., by managed SOC providers and government computer emergency response teams (CERTs), as well as by TI teams to internal stakeholders. As an MSSP manager stated: "We can give our customer the reason why [their security appliance] generates this alert. We can get background information to understand why [the PTI vendor] detects this as suspicious."

Improving end user awareness (29%) is about using TI to educate the wider employee population of the organization, e.g., security-awareness based on reports about recent phishing campaigns.

Threat hunting (29%) is active investigation using TI. This is the type of research which requires human creativity and is currently hard to automate. Combining TI and other data can generate insights for an analyst on where and how to search for attacker activity in systems and networks.

Informing security engineering (21%) includes using TI to organize vulnerability management as part of maintaining the organization's own systems. It also includes the prioritization of developer tasks, e.g., on a customer-facing app, based on observed attacker tactics.

Reducing financial fraud (14%) is a specific use case for banks. Their PTI vendors are supplying them with lists of compromised credit cards. Based on this data, they can decide to block cards or investigate accounts for money laundering. A Team Lead TI at a bank shared: "We are not in a position to buy something from dark web criminals, but [the PTI vendors] are."

2.7.2. VALUE PERCEPTION OF TI

Prior work has assumed that users would like to evaluate TI based on quality criteria such as volume, overlap, timeliness, accuracy, and coverage [40, 47], as discussed in Section 2.2. The assumption is that those criteria capture what users value. Rather than work from this assumption, we have followed an inductive approach. When our respondents made evaluating comments about TI, we labeled them. This resulted in a wider set of 16 separate codes of the properties that made TI perceived as more valuable. We describe these labels, printed here in italics, and list them in Table A.2 in Appendix A.2.

First of all, only three respondents (21%) mentioned anything about price or affordability. If price is not a key factor, this begs the question: what is? We distinguished three clusters of values: confidence, relevance and actionability.

Confidence relates to how much the user trusts the TI to provide useful results. This was primarily interpreted by respondents as not wasting the time of analysts. A head of SOC at an Internet service provider (ISP) described: "One of our commercial sources actually even has a negative value for us, because it costs us time to look into alarms that it generates, which turn out to be mostly false positives." Most of our respondents desire low 'noise' in the TI they receive, which we labeled as to *automatability* (79% of respondents named this property). Loaded into systems such as firewalls or IDSes, low-quality information will immediately have drastic operational impacts. Closely related to automating TI is trust in the vendor. A head of SOC at an ISP described wanting

to be able to verify the origin of intelligence they receive, or what one might call the vendor's transparency: "For us, a provider's ability to answer questions about their intel is an indication of the confidence we can place in them. Of course I understand if they sometimes cannot name their sources. But we need some understanding of the process that led to an indicator being placed on a list in order to use it." Respondents perceive a source as more valuable if it provides an *original contribution* (50%), as demonstrated by this quote by a head of SOC: "We notice a lot of re-use between providers and people in the community. As a rule, we prefer original intelligence over curated or aggregated intelligence because you can cut out the middleman and directly ask questions about the assessment or provide feedback on the intelligence." Confidence is evidently related to *accuracy* (43%) of TI, but respondents also mentioned *selectiveness* (29%). As an analyst at a bank told us: "From [this trusted community] we get emails so often that we filter and tend to ignore them. Whereas [another source] only emails us twice per year. In that case we are likely to look into it." Selectiveness indicates a preference for a low volume, which is seen as an indication of accuracy.

These perceptions of value reflect an intriguing implicit trade-off that users are making; smaller, more curated TI sources are valued higher, as these require fewer organizational resources and may prevent information overload for analysts [64]. But they also imply more false negatives – something our analysis of the overlap of sources (Section 2.5) has confirmed. The risk of having high false negative rates was, remarkably, much less of a worry.

The second cluster of properties perceived as valuable is on **relevance**. Here, we find properties that value the degree to which the TI is tuned to the specific situation of the organization. Two respondents mentioned ending a contract with a PTI vendor because their intelligence mostly covered a *sector* (64%) not relevant to them. *Geographic focus* (50%) is a valued property because our respondents seem to understand attacker groups to choose targets based on earlier successful campaigns in a given country. Furthermore, a certain bias may be related to the geographic focus, as a respondent in Japan described the market for PTI as too US-centric, with certain information not being usable for their organization. A Team Lead TI at an oil and gas company warned: "Be aware of biases of your intelligence providers. For example, a US-based provider will never report on US spying activities." Hence, coverage of one source may have the *ability to correct bias* (14%) in another. *Coverage of relevant threats* (50%) was mentioned by respondents, but not in terms of if the indicator feeds exhaustively contained all the relevant infrastructure of the relevant threat groups. Rather, respondents seemed to interpret this in terms of coverage of their threat landscape, i.e. providing information about threats relevant for their organization. A Team Lead TI at a bank explained how at one point they were confronted with an advanced threat actor. They then separately asked four PTI vendors to tell them what they knew about this threat actor and validated this with their own observations. They noted: "One vendor had nothing, the other three came with a theory. Based on the data that we could observe ourselves, we saw that [vendor X] was totally wrong. The other two vendors were right. It might have been a coincidence, but we did this a few more times and then decided to work with those two vendors." Again, these value perceptions reflect a way of thinking that aims to reduce impact of TI on analysts, in this case by valuing TI that reduces the inputs into the workflow to what is considered the

most relevant.

Finally, there is a cluster of values around **actionability**. This was defined by a Team Lead TI at a bank as: "intelligence which you can use to influence your business." The *capability to provide context* (100%) means that TI helps the user to understand and explain events and alerts. Paid sources are seen by respondents as better at providing context than open sources. One analyst from the Netherlands said: "Intelligence is about context, about putting threats into perspective for your organization." *Timeliness* of TI (50%) is a valued property because indicators lose their relevance rapidly. Once TI is outdated, it is no longer actionable. In Section 2.6, we compared the timeliness of OTI and PTI and did not find a significant difference. This was consistent with the remarks of our respondents. Some said that they do see a difference with STI: they receive certain TI earlier from trusted communities. Yet, even timely TI is only actionable if it is *comprehensive* (50%) enough to be able to base decisions on it, such as suggesting possible mitigation strategies. Some respondents spoke about 'rich' information and about the difference between 'raw' vs. 'polished' intelligence, where PTI is deemed of higher quality because it is more polished toward use. An analyst at a bank stated: "A [colleague] at another bank is just going to post some IOCs to you, or it will be a small write-up, because their time is limited. It will not be of the same quality [as that of PTI providers]. That's the key difference: you're paying for polished intel rather than what we would call raw intel." *Interpretability* (50%) refers to the property that the analyst can make sense of the information, e.g., it has good meta-data. *Data visualization* (14%) is a related property that aids in putting the TI to use by making it accessible.

In sum, we find that TI is evaluated on a much broader set of criteria than prior work assumed. Furthermore, an underlying logic in the properties that respondents value is that they are optimizing the workflow of their organization – most notably their analysts – rather than the detection of threats. This is one of the key reasons why they value the smaller and more curated PTI sources. The fact that these smaller sources might have limited coverage and uncertain timeliness (Section 2.5 & 2.6) is not described as a major problem.

2.7.3. EVALUATING TI

Customers of TI found it difficult to compare sources, which corroborates findings in market research [45, 46]. Evaluation happens mostly in informal processes and based on tacit criteria and heuristics. One research manager described: "So far, we don't have any kind of scientific evaluation process or method. Just a feeling of the analysts. They are using the threat intel daily, and they can feel if they are comfortable with it." Six out of fourteen respondents did define some criteria or intelligence requirements in order to evaluate TI sources, often in the form of information gaps in the organization – i.e., what questions the TI team needed to answer.

One analyst described using metrics within the network detection use case: "You can [demonstrate the effectiveness] by generating a metric on IOC feeds. For example, how many times does this commercial IOC feed purchased from [vendor X] create security events within our organization? And then, what is the outcome of those security events? Is it a false positive? In which case, that means that IOCs sent by that vendor are inaccurate. We can feed that back to the vendor when it comes that negotiation about the contract."

On calculating metrics for the use case of informing business decisions, he added: "That's slightly more difficult to develop metrics around and quantify. But really, what we're looking for from stakeholders [in our organization] is very simple feedback: Was this useful? Did this aid your decision? [...] That is good enough to say if [my team's] reporting is having an impact."

2.8. DISCUSSION

In this paper, we have attempted to lift the veil of paid TI services. We confirm that, indeed, paid TI seems to be a different cup of tea, with distinct intelligence products and low overlap with open TI sources. The interviews we conducted display an apparent contradiction in the practice of TI use: professionals discuss at length the properties that they believe make TI valuable to them, yet hardly attempt to measure or validate these beliefs.

This contradiction questions if threat intelligence metrics, as proposed by [40] and others, can actually capture the right value properties. Research has focused on developing metrics that could be used to understand the coverage, accuracy and timeliness that PTI providers can provide. In our interviews we found, however, that customers are much more pragmatic in how they evaluate the added value of TI, namely through the impact it has on their analysts and security operations. To optimize the analysts' workflow, poor coverage is not necessarily a big problem, while the number of alerts is. This drives customers to smaller, curated sets – the opposite direction of where a coverage metric would point them. In detection parlance: one might expect customers of PTI to select sources for low false negatives, while actually they seem to be selecting for low false positives. These are two distinct goals that are both part of the concept of accuracy. Another limitation of these metrics is that although quantification makes sense for network detection – events can be measured – it makes less sense for the other uses of PTI services that customers described. Thoughtfully composed threat actor reports do not lend themselves to quantitative analysis. Further, just counting network events does not tell much about organizational outcomes: an event may occur without it having much relevance or impact. Analyst skill and experience therefore remain essential for triage in the SOC [64]. Carefully prepared analysis reports could contribute to answering strategic questions in organizations [65], yet TI is currently used mostly in operational processes. That being said, metrics could help to optimize the selection of TI sources for event detection and to understand the potential for false negatives by looking at coverage and overlap. Metrics for TI are useful in this, more narrow, context.

Currently we lack a good understanding of the coverage of PTI vendors due to secrecy around their methods. While that is understandable in order to maintain operational security in the face of advanced attackers, it does make it harder for customers to evaluate what they are actually buying. This paper seeks to address this by describing overlap and timeliness through the comparison of indicators. We find that even when looking at the same actor groups, two of the leading PTI providers have diverging information with very small overlap. The secrecy around their methods to observe threat actors also benefits vendors economically: as long as their methods remain opaque, myths will live on about how TI providers may offer some special degree of TI coverage, possibly through an exclusive skillset, 'hacking back', or by means of access to restricted information. As

described by Shires [66], vendors use "cyber noir" symbols that portray their work as deploying unconventional tactics in mythical battles between good and evil, often aligned with national security. Such stories and symbols give rise to an understanding of detection and attribution capabilities of PTI vendors that currently cannot be substantiated nor vetted.

As a consequence of the low transparency, the market for paid TI shows signs of asymmetric information, in which the vendors know what they are selling, but customers don't know what they are buying. Consumers in the market for TI therefore find it hard to compare services [45, 46]. As Metcalf concluded already in 2015 for blocklists: "secrecy does not benefit the operational analyst who must decide which lists to apply" [43]. And indeed, five years later, our respondents say it is still "mostly guesswork" to understand the visibility and methods of paid TI providers, and with that the value of the services they offer.

Under conditions of information asymmetry, buyers rely on signals. One such signal is whether the firm is seen as a market leader, which is partially signaled via a high price for its services. In this sense, the phrase 'nobody ever got fired for buying IBM' also rings true for threat intelligence. Customers are incentivized to purchase from leading providers – the safe choice under uncertainty. This way, economic value is linked to vendor reputation. In the longer run though, structural information asymmetry holds the risk for vendors that customers may lose trust in the value of PTI services, which would decrease the willingness to pay. This effect is known as a 'market for lemons' [22]. Grigg [23] went one step further and argued that even vendors might lack reliable information on the quality of their products. Providers of PTI might know what data they collect and how, but they do not know – and can't know, Grigg would argue – how effective their product is in improving the security of their clients. Our analysis suggests that, in light of lacking ground truth and low overlap in indicators, vendors themselves may not even know how well they are able to track specific threat actors. When both seller and the buyer lack reliable information on the quality of a product, this creates – in Grigg's analysis – a market for 'silver bullets', where herding behavior and arbitrary best practices triumph over rational purchasing decisions.

Finally, we note that through their forensic work, TI vendors have profound influence on how the general public and the political leadership understands security incidents. Reporting on such incidents is not just neutral technical analysis but also requires interpretation and 'sense-making', as Stevens (2019) showed for the analysis of Stuxnet by Symantec [67]. Indeed, public understanding of such incidents is shaped by the political and economic prisms of the experts who carry out the analysis [68, 69]. Information asymmetry in the market for paid threat intelligence is therefore not only of economic, but also of political significance.

2.9. RELATED WORK

There is a rich line of research that has studied the properties of open threat intelligence, also known as abuse feeds and blocklists – e.g., [42, 43, 44]. Problems in coverage, timeliness and accuracy have consistently been observed in these studies.

In recent years, proposals have been put forward to formalize and measure the quality of TI [40, 47, 41, 48, 70, 50]. This includes metrics on features such as coverage, accuracy,

timeliness, relevance, overlap, latency, and volume. [71] has investigated how to present TI quality to analysts. Applications of these approaches to measure quality of TI have been limited to OTI, also in the recent studies by Li et al. [40] and Griffioen et al. [41].

Aside from the availability of high quality information, it is essential how this information is used. [38] identifies that organizations have issues interpreting threat intelligence, triaging large volumes of threat information or dealing with large numbers of false positives. In this sense, TI has similar operational issues as blocklists of IP addresses and domain names, which have an established history in computer security. While TI as contextualized, high-level information has the potential to remediate these issues [72], a 2019 SANS survey nonetheless found (low-level) indicators of compromise to be valued higher by respondents than information about (high-level) tactics, techniques and procedures (TTPs) [73]. The authors attribute this to the fact that most of their respondents were security operations analysts, who might view the value proposition of TI primarily as enriching alerts with technical details. Our study provides a detailed analysis of how threat intelligence is actually being used within organizations, and how the value is perceived by those directly affected by it.

We go beyond the related work in two key ways. First, we present the first empirical study of the PTI of market leaders. The nearest study is [40], which was not focused explicitly on PTI, but did include two edge cases of paid services. These services were not providing original high-end TI sources, but helped curate and aggregate otherwise free or low-end indicator feeds, and were priced in the range of USD 1-10k per year, as kindly confirmed to us by one of the paper's co-authors. We followed the measurement approach developed by [40] but provide the first application to 'real' PTI: services of commercial threat intelligence providers which operate their own detection network and perform forensic analysis to generate original data about threats. With this value proposition, vendors justify pricing between USD 100-600k per year. A common sentiment in the TI industry is that 'real' high-quality threat intelligence may only be obtained from these exclusive closed-source commercial providers, and [38] finds PTI sources are used twice as often as OTI in industry.

Second, we contextualize these quantitative approaches to measuring quality by conducting a user study of PTI customers and identify their perceptions of value. This has enabled us to find that users use and evaluate TI differently than the measurement approaches developed by researchers assume. In reality, users hardly calculate the proposed metrics. Their perception of value is determined by various use cases in which this quantification is not only missing, but sometimes points in conflicting directions – as around the issues of accuracy and coverage.

2.10. LIMITATIONS

Our mixed-methods approach introduces several limitations. First, we only analyzed the services of two PTI vendors. As they are among the market leaders at the high end of the market, we assume that our findings are representative for that market, but future work is needed to corroborate this.

Second, our analysis was based on data of a single customer of these two vendors. This customer acquired 3-5 subsets of indicators of each vendor in the same topic areas, of a total offering of 5-10 subsets that each vendor offers. Other subsets might show

somewhat different results for the target industries (Figure 2.4) or the overlap between vendors. Given that the available selection of subsets form a significant portion of all subsets, we expect that they provide a valid basis for comparison. The exact numbers, however, are likely to vary across other subsets.

Third, our analysis of PTI has to contend with a lack of ground truth. We followed the approach from prior work on OTI [40] and conducted a comparative analysis among different feeds. For the analysis of indicators on different threat actors, we relied on the well-known mapping developed by Florian Roth across the different threat actor naming schemes used by PTI vendors. We cannot ascertain how reliable this mapping is, other than the fact that Roth is an expert in the field, his mapping is well known, and he is collaborating on it with other industry insiders – so mistakes would presumably be corrected.

Fourth, the comparison with OTI was limited to a single month of four feeds. While these feeds were chosen because they are actually re-used by many other feeds in the OTI landscape [41], a broader set of feeds will provide a more reliable result. That said, the lack of overlap with PTI was quite stark and unlikely to change when analyzing other feeds. In OTI research, the low overlap among any two feeds has been a consistent finding for years.

Fifth, regarding our user study, our main limitations stem from a small sample size (n=14). Our sample contains a variety of organizations, but it may contain selection bias as respondents are geographically located in the Netherlands and Japan only, were working with TI (rather than choosing not to), and willing to talk about this in an interview. All of this means that we clearly do not claim that our findings are generalizable for all organizations using TI. Given that no prior work existed, neither on the PTI feeds nor on users of PTI, we chose to do an in-depth exploration of the views of such users using the grounded theory method. For more generalizable results, a survey could be designed based on our findings.

2.11. CONCLUSIONS

This study explored services in the market of commercial threat intelligence. We analyzed the indicators of two paid TI vendors and found 13.0% of vendor 1's indicators appear in vendor 2's set and – vice versa – a mere 1.3% of vendor 2's indicators in vendor 1's set. If we drill down to the 22 threat actors for which both vendors have indicators, we find an average overlap of these indicators of no more than 2.5 to 4.0% per group, depending on the type of indicator. Further, this overlap occurs primarily with a handful of actors. The fact that the indicators of two vendors are largely separate sets, even when assessed for specific threat actors that they both track, raises questions on the coverage that services of these vendors actually provide.

Reports produced by paid TI providers describe the tactics of threat actors, the results of malware reverse-engineering, or give advisories for current events, among other things. The reports concern primarily government, military, and financial institutions – important customers of TI vendors – but surprisingly also pay a lot of attention to campaigns targeting civil society, possibly due to their political significance.

Besides indicators and reports, paid TI services also consist of requests for information from analysts, portals with historic information, data mining platforms, and custom alerts.

These services are expensive, with subscription costs of major vendors often upwards of $100,000 per year.

Whereas paid sources offer 'polished' TI, open sources contain 'raw TI' as one respondent described it. We find that this statement holds for the two types of paid intelligence products that we have compared with open TI, namely indicators and reports. In terms of substance, paid reports are for example similar to open source blog posts and tweets of security researchers that are freely available online, but the paid reports are more comprehensive in their descriptions of context and recommendations. Further, paid reports are packaged with machine-readable indicators. We compare these to open TI indicators feeds (which are much larger in volume), and find less than 1% overlap between them, suggesting that PTI providers successfully differentiate themselves and are capturing a different part of the threat landscape. In terms of timeliness, we find no evidence that PTI is faster than OTI, surprisingly enough, although this is based on the small sample of overlapping indicators. There is a delay of around one month before indicators from one set are found in another.

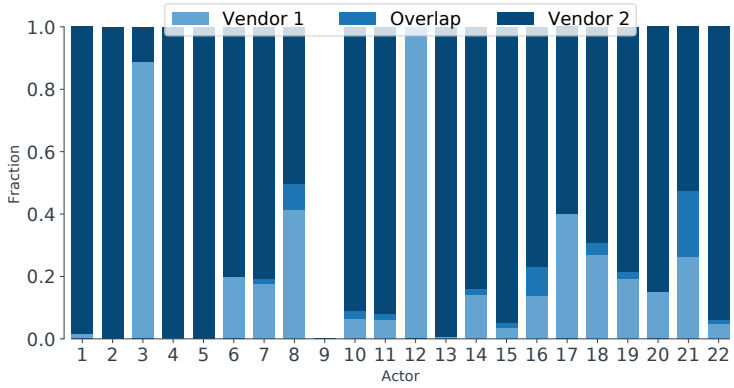
The main use case for TI is network detection, followed by situational awareness – which we understand to mean: informing your threat profile – and prioritization of resources in the SOC. We find that one-third of respondents use threat intelligence to improve organizational decision-making: to inform security engineering, to reduce financial fraud, but in one instance also for risk management around international mergers and acquisitions. Asked what makes TI valuable, respondents name properties related to actionability, relevance, and confidence. All respondents describe valuing the ability for TI to provide context, which suggests that they view TI as a reference. Further, the ability to automate using TI is important for respondents: almost all name valuing a low false-positive ratio and interoperability with their detection systems. Importantly, only half of our respondents discuss coverage as something they value in threat intelligence – reducing false-negatives or misses seems to be much less of a concern. We conclude that TI consumers evaluate threat intelligence mostly through the impact on their organization's detection processes.

Evaluation of TI sources is done mostly through informal processes by our respondents. When a subscription renewal comes up, TI professionals decide whether to continue largely based on implicit criteria and tacit understanding of the value of that source. This is surprising, given that research has focused on developing metrics and heuristics that could enable a quantitative understanding of the value of TI. In practice, metrics or intelligence requirements are used by less than half of the professionals we interviewed.

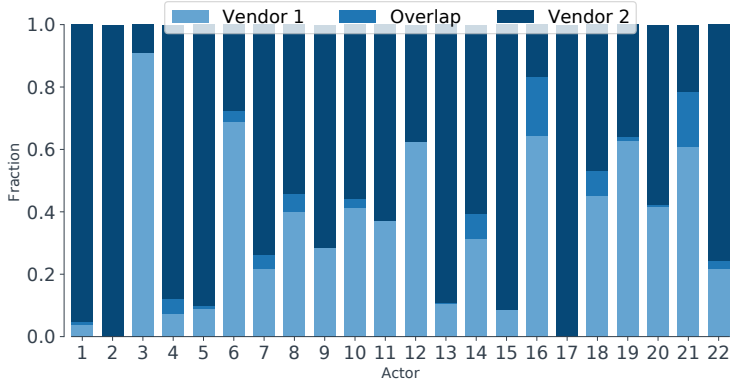
There is this promise that leading paid TI vendors would be able to overcome the persistent problem of sharing threat information among defenders. They can aggregate and analyze threat data at scale from vantage points across different clients and networks – as was also argued by CrowdStrike Senior Director Vincenzo Iozzo (see Section 2.2). We do not dispute that important advances are being made. That being said, our study raises doubts as to the extent to which this promise has been fulfilled today. Even when the vendors claim to track the same threat actor, they each see only a tiny fraction of the associated indicators. The fact that almost all PTI indicators are unique to one vendor, is a pattern we know all too well from OTI sources. So the pay-off of aggregating data across clients and networks, as claimed by Iozzo, is not very clear in terms of detection

capability, to say the least. Even when a client would be willing to pay the steep price of simultaneously acquiring the feeds of all market leaders – a proposition that would cost them millions each year – it is likely that this strategy would reproduce the pattern that we know from OTI: feeds contain mostly singletons and more feeds still get nowhere close to comprehensive coverage.

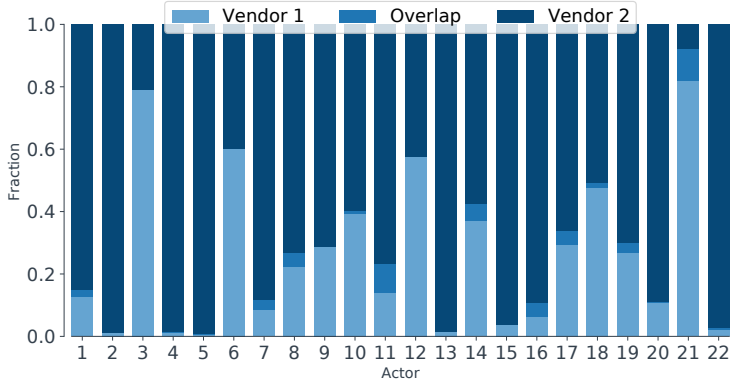
The sharing of indicators across vendors would still be a first step to improve coverage and the detection of attackers. The current state of affairs in paid TI resembles the market for anti-phishing services about a decade ago. The lack of data sharing meant that each anti-phishing company thought it had strong coverage and could protect its client brands well, while the truth was that they missed most of the attacks they were hired to detect [74]. In that market, like in the market of malware detection, sharing across vendors was eventually recognized as a superior security strategy. Until that happens for commercial threat intelligence, the problem of information sharing persists.



(a) IP indicators, avg. overlap 2.8%



(b) Domain indicators, avg. overlap 4.0%



(c) MD5 indicators, avg. overlap 2.5%

Figure 2.3: Indicators for threat actors tracked by both vendors from 2013-2018. Overlap is tiny and concentrated on a handful of actors.

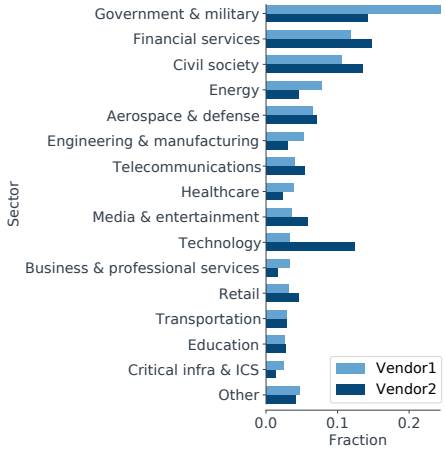


Figure 2.4: Government and financials represent the main industries as reported on by two leading paid TI providers from 2013-2018. Civil society is the third most targeted sector observed, which is surprising if we assume that these organizations are not really in a position to pay for and use high-end TI. We mapped vendor labels to a common structure.

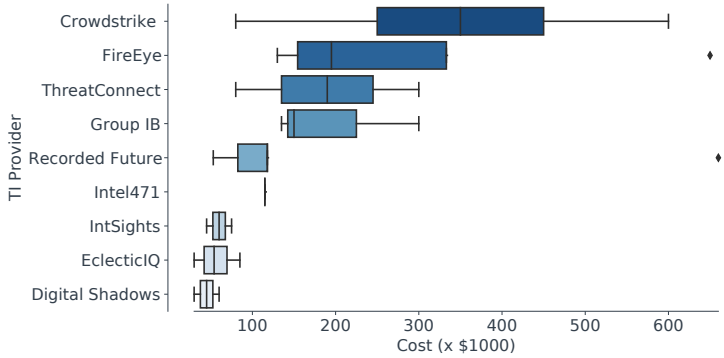


Figure 2.5: Subscription costs of various PTI providers. Pricing is often based on the size of the customer organization, with major vendors demanding upwards of $100,000 per year.

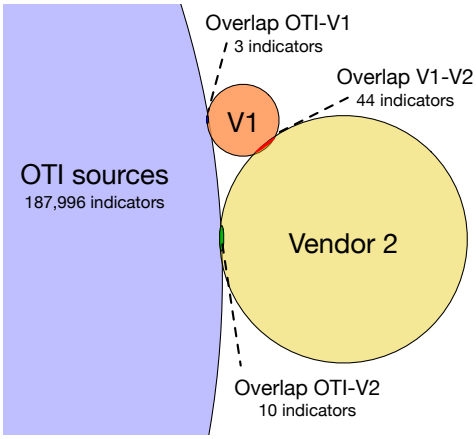


Figure 2.6: Overlap between OTI indicators published in one month (October 2018) and PTI indicator sets published in the enclosing five months (Aug-Dec). The latter time frame allows for variations in timing when the indicators are first reported and thus provides us with an upper-bound estimate of overlap. The overlap turns out to be negligible, suggesting that these types of intelligence are different in kind, as well as in volume. Areas in the diagram are proportional to the number of indicators.



Figure 2.7: Indicator overlap as a percentage of the row's total volume. From Vendor 1's indicators, 13% is also listed by Vendor 2, and 0.3% is listed by AlienVault. Overall, PTI and OTI sources hardly share any indicators, at most 0.9% relative to the PTI set, and at most 0.02% relative to the OTI set. The same subsets are used as for Figure 2.6.

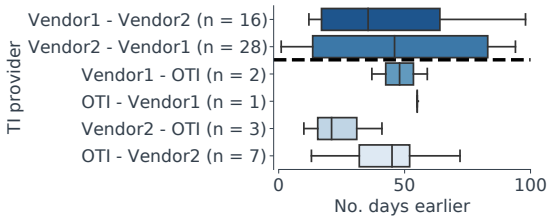


Figure 2.8: Timeliness comparison between sources. There is often a delay of over a month before an indicator is listed by a second source. We show three pairwise comparisons. To illustrate: we found 16 instances where Vendor 1 was earlier than Vendor 2, by an average of 45 days. Note that the number of indicators (n) is low, because it is based on the small overlap between sources. We aggregate the 4 OTI sources into a single set. For the comparison between the two vendors, we use the period of 2013-2018. For the comparison between vendors and OTI, we used the same subset as for Figure 2.6.

3

PEER PRODUCTION

We tracked the largest volunteer security information sharing community known to date: the COVID-19 Cyber Threat Coalition (CTC), with over 4,000 members. This enabled us to address long-standing questions on threat information sharing. First, does collaboration at scale lead to better coverage? And second, does making threat data freely available improve the ability of defenders to act? We found that the CTC mostly aggregated existing industry sources of threat information. User-submitted domains often did not make it to the CTC's blocklist as a result of the high threshold posed by its automated quality assurance using VirusTotal. Although this ensured a low false positive rate, it also caused the focus of the blocklist to drift away from domains related to COVID-19 (1.4%-3.6%) to more generic abuse, such as phishing, for which established mitigation mechanisms already exist. However, in the slice of data that was related to COVID-19, we found promising evidence of the added value of a community like the CTC: just 25.1% of these domains were known to existing abuse detection infrastructures at time of listing, as compared to 58.4% of domains on the overall blocklist. From the unique experiment that the CTC represented, we draw three lessons for future threat data sharing initiatives.

3.1. INTRODUCTION

For years now, research has consistently found that threat data feeds each cover just a fraction of the landscape. Numerous comparisons have been made among different threat intelligence sources and they all find very little overlap: feeds are dominated by data points that appear only in a single source and in no other one. This holds across the spectrum of threat data, from freely available blocklists and abuse feeds [15, 16] to closed industry sources [17], all the way up to the most expensive feeds at the high end of the

This chapter was previously published as a research paper: 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. "Helping hands: Measuring the impact of a large threat intelligence sharing community". In: *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 1149–1165. ISBN: 978-1-939133-31-1. URL: https://www.usenix.org/conference/usenixsecurity22/presentation/bouwman

market [18]. It points to poor coverage of the threat landscape, a problem for the industry and – more importantly – for its customers. To illustrate: anti-phishing companies missed a large portion of the phishing sites targeting their customers, while those sites were being discovered by their competitors [74].

While prior work on threat intelligence sources demonstrated the problem of low coverage, it does not discuss how it could be overcome. One obvious and often proposed solution is more data sharing. This typically takes place in informal trusted communities of specialists or in formalized sharing agreements among firms, such as the Cyber Threat Alliance [75] and the Anti-Phishing Working Group [76], which often demand reciprocity, or 'quid pro quo', in order to avoid free-riding behavior. This forms a high entry barrier for access to the shared data, as not everyone has enough to contribute in order to gain access, and so benefits are typically limited to corporate entities. Non-participants can only get access to the separate services of these firms. In economic terms, such sharing arrangements create club goods, not public goods.

A potentially more effective form of information sharing would go beyond these boundaries: open to any contributor, with free access for anyone to the pooled data. Under normal conditions, market incentives would prevent such a public good from emerging. Generating quality threat information costs money and firms have to recoup their investments. However, the 2020 global pandemic gave rise to a real-world experiment that temporarily suspended normal economics in data sharing: the COVID-19 Cyber Threat Coalition (CTC) [77]. The CTC was a volunteer-based response where firms and individuals shared threat intelligence on pandemic-related threats posed by cybercriminals as well as nation states [78]. The coalition's mission was "to operate the largest professional-quality threat lab in the history of cybersecurity" and reduce gaps in availability and coverage of existing defense mechanisms [77, 79]. After 3 months, over 4,000 individuals and organizations had signed up, with companies like Symantec, Microsoft and Cofense contributing data [80, 81]. Contrary to sharing based on quid pro quo, the resulting blocklist was freely available to anyone.

Does such large-scale open data sharing actually improve our defenses against threats? To the best of our knowledge, research on the effectiveness of large data-sharing arrangements is extremely sparse. More than a decade ago, Moore and Clayton studied the inner workings of crowdsourcing at PhishTank, but not its impact on mitigation [82]. Thomas et al. found that a Google-based threat exchange could have more impact than existing standalone anti-abuse pipelines [83] but access to the pooled data in the exchange was restricted. A recent study on VirusTotal did not analyze the information-sharing aspects of the service nor its impact, but evaluated the aggregation of detection results of the participating vendors [84].

In this paper, we set out to extend the literature on data-sharing arrangements by learning from the CTC experiment. We investigate two questions: *(i)* By pooling data from its community, did the CTC improve coverage of COVID-19-related threats over existing defenses? And *(ii)* Did publishing threat data in a freely available blocklist improve the ability of defenders to act against threats, compared to the existing abuse mitigation infrastructure? To answer these questions, we first describe the organizational setup of the CTC and observe how the community pooled data and conducted quality assurance. We then evaluate the blocklist through manual classification of a sample and by identifying false

negatives for COVID-19-related domains. Next, we conduct longitudinal measurements to infer *who* acted *when* against domains on the CTC blocklist: registry, registrar, browser vendor, or security provider? We end by identifying key lessons for improving the impact of open large-scale data sharing mechanisms.

In sum, our contributions are:

- We present the first empirical evaluation of a large and open threat sharing community, the COVID-19 Cyber Threat Coalition, and describe its mechanisms for producing and vetting threat intelligence.
- We find that user contributions were heavily skewed, with just 10 users making 90% of contributions. Further, the high threshold posed by the indicator vetting process, which relied on VirusTotal, resulted in a mere 5.14% of user-contributed indicators actually being propagated to the blocklist. Instead, most data came from commercial firms, and at least part of this was vetted against a lower threshold.
- We show that the CTC list went well beyond the scope of just COVID-19-related threats. Generic phishing made up a large portion of the blocklist: domains containing the word *whatsapp* (2.8%) outnumbered those containing keywords related to COVID-19 (2.6%).
- Based on longitudinal measurements, we demonstrate that for 58.4% of the domains on the CTC blocklist, existing abuse-mitigation mechanisms were faster: domain-level or client-side interventions had already taken place before the domains appeared on the blocklist. For COVID-19-related domains, this share is smaller: 25.1%, which is evidence of added value of the community. The remaining portion was intervened against later or not at all. This means that the blocklist improved the ability of defenders to protect themselves. Its impact could have been larger, if it had not depended as much on VirusTotal for vetting indicators.

3.2. THE CTC COMMUNITY

The CTC started as a Slack community on March 19 – a week after the World Health Organization (WHO) declared COVID-19 a pandemic. Joshua Saxe, a security specialist at Sophos, founded the community out of a "personal sense of alarm", conceiving it as a "crisis commons model" where "traditional competition and grievance [were] set aside in a moment of exceptional need" [85]. The CTC founders articulated objectives for the community and then set up community support services for volunteers to join and participate. We also joined and revealed ourselves as researchers interested in learning more about how the community was functioning. The CTC's mission was threefold: fostering collaboration across organizations to uncover otherwise missed threats, producing professional-grade output that the community can rely on, and prioritizing the public good over the interests of individual actors [77].

3.2.1. INFORMATION PRODUCTS

The CTC set out to publish an open high-quality blocklist containing COVID-19-themed abuse to "supplement the existing defensive structure" by setting up a "separate threat

intelligence platform dedicated just to pandemic-related cyber activity" [79, 86]. Blocklist data was published on the CTC's website from March 29, 2020 onwards [87]. In May, it reported that over 60,000 distinct IP addresses had been consuming the blocklist [88]. Initially, they offered two types of lists: one with vetted indicators of compromise (IOCs) and a larger list with unvetted IOCs submitted by the community. The latter was eventually discontinued to "produce the highest quality feeds with the least amount of false positives possible" [89]. Within the vetted category, four blocklist files were available: for domains, URLs, IP addresses and file hashes. The latter two, however, have remained empty at all times. Our analysis focuses on what is arguably the main list, i.e., the domain list.

The technical indicators were supplemented by threat advisories and community meetings. The advisories featured themes like phishing and ransomware related to COVID-19, trends in pandemic-related domain registrations, and the security challenges of remote work. They were published on the CTC website[1] and sent out via a mailing list every week from April 6 until May 26. From April 16 onward, community leaders also organized online community meetings or 'town halls' hosted on Zoom and archived on YouTube [90]. These webinars were intended to disseminate findings from the reports, to update the community on changes to the CTC procedures and infrastructure, to interview representatives of security companies, and to allow for Q&A with the community leaders. For the first two months, town halls were organized weekly; after June 11, the CTC planned to organize them every two weeks, but no town halls took place since.

3.2.2. COMMUNITY STRUCTURE

The primary means of collaboration in the Cyber Threat Coalition was a Slack workspace. The aim of this Slack workspace, which anybody could join, was for members of the security community to share information and make the necessary contacts for interventions [91]. After a brief period of unstructured posts, the workspace was organized into themed channels for posting community updates, IOCs per topic (e.g., web, email, malware), and region-based networking. Membership grew to 1,500 in one week [92] and to over 4,000 after 3 months [93]. Members were requested to add their affiliation and job title to their screen name, although this information was not validated. Of the users who indicated a country in their username, around half came from North America, around a third from Europe (mainly the United Kingdom and the Netherlands), 3% from Australia, with the rest scattered around the world. Based on stated affiliations, the CTC counted representatives of security vendors, healthcare providers, CERTs, financial institutions, domain name infrastructure providers, major technology firms and law enforcement agencies among its members. Activity peaked at the end of March and then slowly decreased (Figure 3.1). By July, the public channels saw much less activity, with daily posts in the single digits, and no official announcements were made except for a single update in February 2021 stating that the CTC was in "low battery mode" but "evolving".

The community had a simple organizational structure [88]: a 'steering committee' functioned as administrators for the community, with its members committed to "sacrificial" time contributions [91], leading teams for various tasks like vetting process development, advisory writing, and media outreach. Within the larger community, over 100 'vetted volunteers' [94] underwent identity verification via their social media profiles to

[1] https://www.cyberthreatcoalition.org/advisories

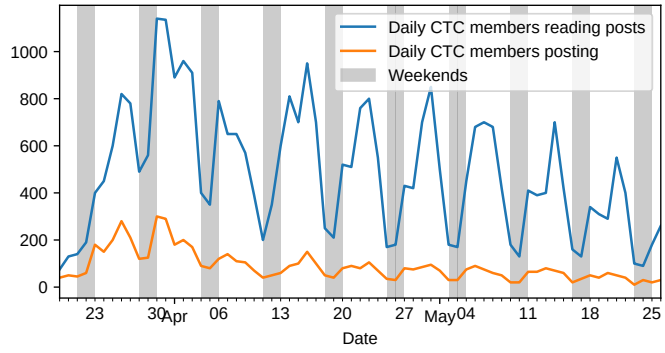


Figure 3.1: Member activity on the CTC Slack workspace peaked 1.5 weeks after the community was founded [88].

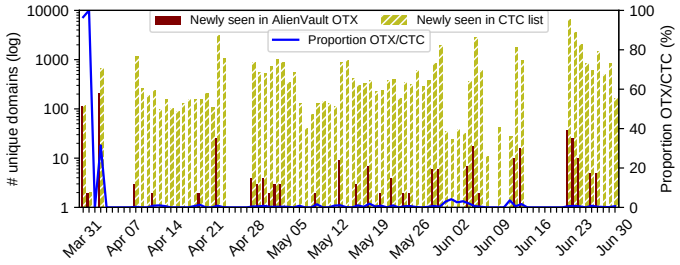


Figure 3.2: Counts of unique domains newly seen in the CTC AlienVault OTX group and on the CTC blocklist (logarithmic scale), and the proportion of AlienVault OTX domains that were propagated to the CTC blocklist.

obtain access to private channels, where more sensitive data could be shared [91].

3.3. PRODUCTION OF THE BLOCKLIST

In this section, we describe how the CTC's prime information product was sourced and vetted. Four sources fed the CTC blocklist: *(i)* user contributions on Alienvault OTX; *(ii)* user contributions via a Slackbot; and industry lists from both *(iii)* named and *(iv)* unnamed vendors.

Initially, TI sharing was ad hoc, consisting of community members posting free-form IOCs, first in the general Slack channel and soon after in field-specific channels (e.g., for email and domains). After one week, the CTC set up groups for machine-readable TI sharing on AlienVault OTX [92]. Users could submit IOCs to the CTC Slack workspace. After they had been checked for maliciousness , indicators were published on the 'vetted' list [91]. In June, the CTC introduced a Slackbot that let users contribute indicators as well as vet indicators supplied by other users [80]. We describe these two vetting mechanisms in Section 3.3.2. Further, commercial security providers contributed "hundreds of millions of indicators per day" [91] outside of the AlienVault group, although these were also subject to the same vetting procedure. These indicators appear to represent the vast majority of domains on the final vetted blocklist, as demonstrated in Figure 3.2. Known industry

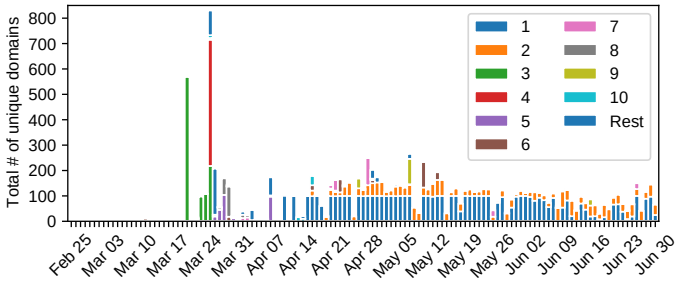


Figure 3.3: Daily count of unique domain indicators contributed to the CTC AlienVault OTX group, with the top 10 and other contributors ranked on the total number of contributions over time.

contributors to the CTC are Symantec [81], Microsoft and Cofense [80], but other vendors have asked to remain anonymous [91].

3.3.1. USER CONTRIBUTIONS ON ALIENVAULT OTX

We analyze the indicators contributed to the CTC group on AlienVault OTX, using the API to gather all submissions over time. By July 2020, 738 users had been accepted into the CTC's closed AlienVault OTX group. Only 47 users actually contributed IOCs – 10 of whom made 90% of all contributions (Figure 3.3). The two heaviest contributors did so on a fixed schedule, and described drawing on newly registered domains and certificate transparency for their lists. Others added more opportunistically, with a downward trend over time. In general, recipients were often left guessing as the source or method behind the contributed indicators.

In other words, even though the community has over 4000 members, only 1–2% of them contributed indicators, with 10 users contributing the bulk. This skewed distribution resembles the pattern found in the PhishTank community project, where volunteers collect and classify phishing sites [82]: there, the top 10 contributors made 69.9% of submissions and 57.4% of votes [95]. It also resembles free and open source software (FOSS) development, where a tiny fraction of contributors supplies most code and most others only contribute once, typically in the form of a bug report [96]. As we discuss in the next section, most user contributions to the CTC never made it past the vetting stage (see also Section 3.3.2). At the outset, a large proportion of AlienVault submissions made it to the CTC blocklist, but after the list was reset on April 12 (Section 3.4.2), domains added to the CTC blocklist came mostly from named and unnamed industry sources. Over the entire period of our data collection, just 1.23% of domains on the blocklist came from user contributions on OTX.

3.3.2. QUALITY ASSURANCE

The community admins instituted vetting mechanisms for submitted indicators, in order to "provide reasonable assurance that what we re-share with the public are examples of truly malicious artifacts" [85]. Initially, vetting consisted of a pool of volunteers manually verifying maliciousness of submitted indicators. The workflow was eventually automated using a security orchestration service that integrated VirusTotal: if a domain received 10

or more hits, this would lead to a domain being marked as malicious and propagated to the blocklist. A domain with between 4 and 10 hits would require manual review, while for fewer than 4 hits, it would be marked as "clean" [80] and dropped.

Although automation allowed for higher volumes of indicators to be processed, this workflow left the community with three problems. First, only indicators that were already known to be malicious by many VirusTotal scanners could be added to the blocklist. Earlier research has consistently found very low overlap among TI sources [15, 18], so requiring that indicators are found by 10 scanners imposes a high threshold to propagation. It meant that just 5.14% of domains from the AlienVault OTX group made it onto the blocklist. Moreover, VirusTotal aggregates labels from 84 established automated scanners, negating the community contribution aspect of the CTC. The second problem: some of the contributed industry sources, in particular Cofense [80], were seen as reliable, yet their indicators did not appear on the vetted blocklist as too few VirusTotal scanners flagged them. As a solution, the CTC admins lowered the thresholds for these trusted sources [80].

The third problem was how to evaluate the indicators that fell in-between the thresholds for 'clean' or 'malicious'. At first, this was done manually by a small pool of volunteers. It is unlikely that they could keep up with the overall volume, so probably these indicators did not make it onto the vetted list. On June 11, the CTC admins announced a new Slackbot that would allow all members to contribute to vetting [80]. Upon request, the bot would serve a domain or URL to an individual user for evaluation. It resembled crowdsourcing mechanisms like PhishTank, albeit without the built-in validation of multiple users checking the same indicator. The bot potentially increased the scalability of manual vetting and would allow indicators to be included long before VirusTotal would provide enough 'hits'. The downside was that vetting might be done by members with unverified expertise or even potentially adversarial motives. Unfortunately, when the crowdsourcing functionality arrived by June, the peak of user activity had already passed, so it came too late to actually change the vetting process. Despite this process, some false positives slipped through, leading domain owners to join the CTC Slack workspace and request removal from the blocklist.

3.4. EVALUATION OF THE BLOCKLIST

In this section we address our first research question: *By pooling data from its community, did coverage of COVID-19-related threats improve over existing defenses?* First, we manually label a sample of domains to evaluate the nature of new IOCs appearing on the CTC blocklist. Second, we track the evolution of the blocklist composition, and measure its focus on COVID-19-related abuse. Third, we measure the role and impact of VirusTotal on the vetting process. Fourth, we estimate the coverage of COVID-19-related domains through a comparison with external sources.

3.4.1. MANUAL CLASSIFICATION OF DOMAINS

Our first assessment of the blocklist content is to manually inspect a sample of domains. Over the course of 5 days (17-22 May 2020), we took a daily sample of 50 domains that were newly added to the CTC blocklist. We visited those within 4 hours of their appearance on the list, in order to minimize the chance that the domain would already be affected

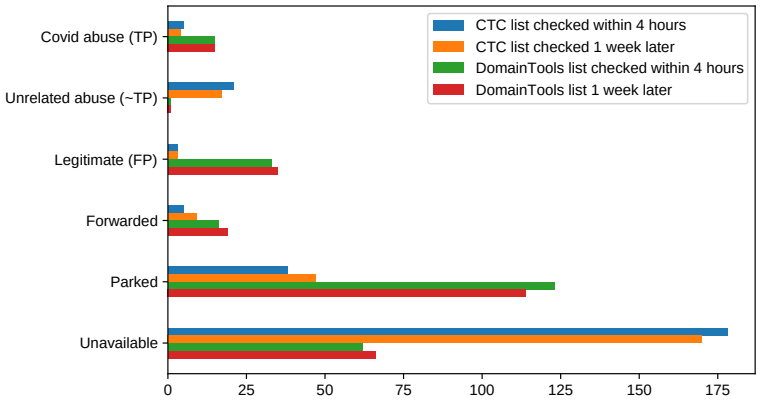


Figure 3.4: Manual classification of 500 domains, visited at two moments in time.

by a countermeasure. We navigated to the site with Microsoft Internet Explorer 11 in a virtual machine on a computer located in the Netherlands. If necessary, we translated page contents using Google Translate. After one week, we visited each site again to see if anything had changed. As a reference, we carried out this process not only for the CTC blocklist but also for a COVID-19 blocklist published by DomainTools [97].

We based our label taxonomy on existing classifications [98, 99]. In many cases, the labeling could be applied in a straightforward manner, such as with parked domains. Some categories contain more ambiguity. When we were uncertain about what the appropriate label was, we took a conservative approach and counted them as true positives, giving the blocklist publishers the benefit of the doubt. As shown in Figure 3.4, out of 250 domains visited from the CTC list, we encountered just 5 cases of COVID-19 related abuse (2%), including phishing sites themed with the pandemic and sites selling dubious protective materials. The CTC list also included 21 examples (8%) of internet abuse that was not visibly tied to COVID-19, such as generic phishing sites, counterfeit products and pharmaceuticals – although it is conceivable that these domains were in fact used in a campaign that somehow played upon COVID-19, for example in a spam run. We deemed 3 websites as legitimate – and therefore false positives on the list (1%). The majority of domains were unavailable (71%).

The DomainTools list contained more examples of actual COVID-19 related abuse (6%), but at the cost of more false positives (13%), many of which had in common that they contained keywords related to the pandemic, such as a Wuhan-based welding hardware supplier.

When revisiting the domains after one week, we saw minor changes (4% overall), most of which were previously unavailable domains becoming reachable as parking pages. We saw seven examples (1.5% overall) of generic abuse not related to COVID-19 becoming active in this interval of one week.

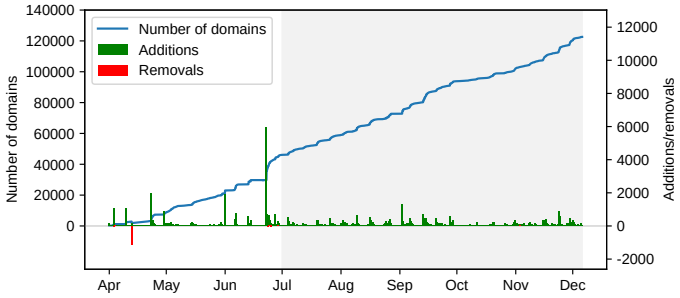


Figure 3.5: Evolution of the composition of the CTC vetted domain blocklist.

3.4.2. COMPOSITION OF THE BLOCKLIST

Where Section 3.4.1 analyzed a sample, further analysis in this paper is based on the full CTC blocklist of vetted domains from March 31 up to July 1, 2020, during the peak of community activity. We retrieved this list once a day starting March 31, every hour starting April 4, and every five minutes from April 15 onward. Over time, the vetted domain blocklist steadily grew (Figure 3.5), reaching 46,103 domains on July 1, 2020. Cumulatively, 46,832 fully qualified domain names (FQDNs) on 27,096 unique second-level domain names (SLDs; i.e., domains that can be bought from a registrar) appeared on the vetted list by July 1, 2020. Domain removals were rare, except when the list was reset on April 12 due to the shift toward automated vetting [93] which suggests that stale entries were not purged from the list, potentially generating false positives after cleanup or takedown of a malicious domain[2]. Further, we suspect that back-end changes caused a temporary eight-day blockage in updates, followed by the addition of 5,918 domains on June 22, as users reported download issues around the same time [93]. The list continued to be updated until December 6, 2020, albeit without much community input, after which the list remained available, but no longer changed, indicative of dwindling efforts in the CTC.

Given the CTC's objective to track COVID-19-related abuse, we examine the coverage of the blocklist in this area. For our evaluation of coverage of COVID-19-related malicious domains (here and in Section 3.4.4), we generated 370 COVID-19-related keywords in 15 languages (see Section B.1). At its peak, over 5,000 domains containing such keywords were registered every day [100]. However, over our measurement period, only 1,229 (2.6%) out of the 46,832 domains seen on the CTC blocklist contain at least one keyword. The share of COVID-19-related domains was stable between 1.4% and 3.6% throughout time, except before the list reset on April 12, when up to 73% of the (much shorter) list was COVID-19-related[3].

Overall, generic phishing was much more frequent: 17.3% of domains matched a brand tracked by PhishTank[4]: the keyword *whatsapp* occurred in 2.8% of domains, making it more prevalent than all pandemic-related keywords combined. Likewise, Google Safe Browsing and VirusTotal classify 96.8% and 85.7% of flagged domains respectively as

[2] 1,140 domains were removed; of these domains, 38.8% re-emerged at some later point in time.
[3] Out of the 1,140 domains removed then, 826 were COVID-19-related, of which 474 would never return.
[4] `https://www.phishtank.com/target_search.php`; the list was pruned to reduce the likelihood of false positives, and no lookalike terms were added.

Provider	Phishing (%)	Malware (%)	Other (%)	Total flagged
GSB*	27 462 (96.8)	1 017 (3.6)	364 (1.3)	28 383
VirusTotal**	40 132 (85.7)	6 680 (14.3)	20 (0.0)	46 832

* Domains may be assigned to multiple abuse types.
** By majority vote over engines with a specific abuse classification.

Table 3.1: Abuse types for domains on the CTC list as identified by Google Safe Browsing (GSB) and VirusTotal.

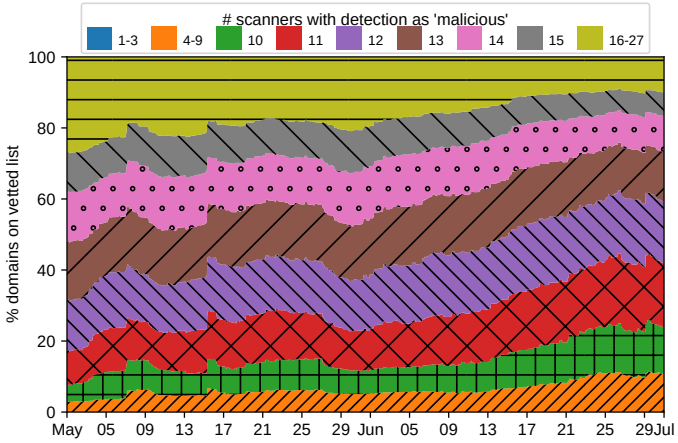


Figure 3.6: Evolution of the proportion of domains with a given detection count by VirusTotal domain scanners at the time of presence in the CTC blocklist.

engaging in phishing (Table 3.1).

Our findings contrast with the CTC's stated goal of sharing "high quality threat intelligence related to the COVID-19 pandemic" [101]. The small share of COVID-19-related domains suggests that the collected TI goes beyond COVID-19-specific abuse, and instead captures *any* abuse observed during the pandemic. This could be the result from automated submission processes, with generic TI sources being redirected to the CTC instead of curated and targeted feeds. It also seems directly related to the decision to move to vetting based on VirusTotal. Before that shift, the proportion of COVID-19-related domains was much higher. Afterwards, the list relied on the ability of existing scanners to detect the abuse and therefore it converged on conventional forms of abuse, potentially discarding highly relevant IOCs on new threats. We revisit this issue in Section 3.4.4.

3.4.3. EFFECT OF VIRUSTOTAL SCANNERS ON VETTING

We independently replicated the CTC's vetting procedure based on VirusTotal, in order to gain insight into the effects on the outcomes. We requested VirusTotal data once a day for all domains that up to that date had appeared on the CTC blocklist, from April 28 until July 1, 2020.

Any domain with more than 10 detections in VirusTotal was automatically considered to be vetted as malicious. We indeed observe that domains meeting that criterion make up the large majority (88.7–97.2%) of domains throughout our measurement period

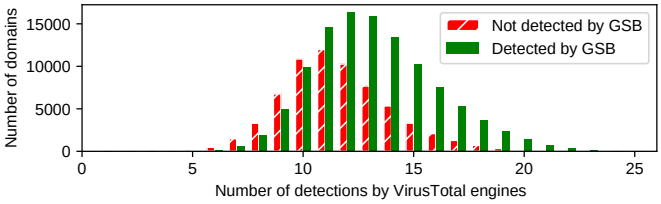


Figure 3.7: Distribution of domains that were (not) detected by Google Safe Browsing (GSB) over the number of detections by VirusTotal engines.

(Figure 3.6). We also see a smaller set (2.8%–11.3%) of domains with between 4 and 10 detections, which suggests that they either went through the manual review process or they came from trusted industry feeds that were vetted against a lower threshold (Section 3.3.2). The share of domains with fewer than 4 detections is negligible, which indicates that the industry sources contributed only domains that were already known to multiple security vendors in VirusTotal.

No domain was marked as malicious by more than 27 of VirusTotal's 84 domain scanning engines; this low ratio is consistent with that found by Peng et al. [84]. Among the 84 scanners, only 23 detected over 5% of vetted domains, and 21 scanners detected over 5% of the domains with COVID-19-related keywords in them. In other words, the CTC threshold for vetting is quite high: 10 detections means that a domain is already known to nearly half of all engines that contribute a non-trivial amount of detections.

As an external corroboration, we determine whether domains on the CTC list were also flagged by Google Safe Browsing (GSB). We find that the more VirusTotal engines detect a domain, the more likely it is to also be flagged by GSB (Figure 3.7). We confirm with a χ^2 test that the distributions of VirusTotal engine counts for domains that are and are not detected by GSB significantly differ ($\chi^2 = 14595$, critical value at $\alpha = 0.05$: 40.113, $p < 0.0001$). The CTC's threshold of 10 detections is close to the crossover where more domains are flagged by GSB than are not (11 detections). While this supports a low false positive rate, the threshold also makes it very hard for the CTC to contribute new threat intelligence beyond that of existing anti-abuse infrastructure.

Overall, slightly fewer scanners tend to detect malicious domains as our measurement period progresses (Figure 3.6). Among the top 10 engines (Figure 3.8), most have a consistent detection rate and therefore contribution over time. AegisLab WebGuard achieved a high detection rate only after mid May: this is a possible indicator that they may have then started ingesting the CTC blocklist, instead of proactively flagging domains and therefore contributing to the vetting process. Decreasing detection rates for Avira and ESET may translate into a decreased contribution to the vetting process over time. Once a domain was included on the vetted blocklist, few additional engines marked it as malicious (Figure 3.9), even as its listing duration increased; a larger increase was only visible for the earliest listed domains. This suggests that the vetting decision by the CTC was based on stable detections by the scanners, so domains warrant their near-indefinite presence on the vetted list (Section 3.4.2).

In summary, the CTC vetted list comprised (only) domains where a relatively large

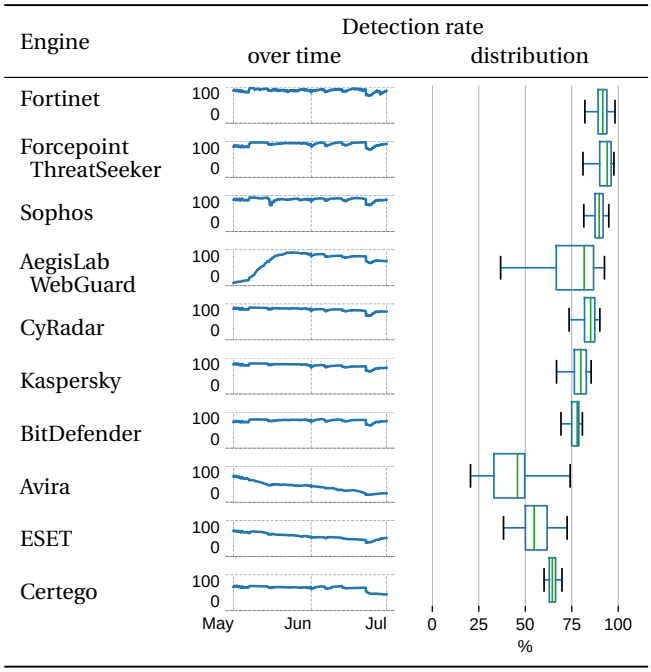


Figure 3.8: Detection rates for the top 10 engines in VirusTotal sorted by maximum detection rate.

proportion of security scanners agree on their maliciousness, making false positives unlikely. However, it is also *required* that at least 10 VirusTotal scanners flag a domain. Unless these scanners successfully adapt to the novel COVID-19-related abuse, this strict threshold causes many false negatives. This led to a problem for the CTC: their whole rationale was to "supplement the existing defensive structure". If the existing scanners were to adapt, then the CTC no longer supplemented them. If they did not adapt, then using them for vetting while maintaining a high threshold also meant that the CTC no longer supplemented them. This reaffirms our observation in Section 3.4.2 that instead of capturing COVID-19-related abuse not seen in other TI sources, as was the CTC's primary target, the CTC list rather contained primarily 'generic' TI on which many existing sources (here VirusTotal scanners) agreed. While the TI may therefore conform to the CTC's goal of being of "high quality" (few false positives), the trade-off is that coverage of the threats for which it was set up is undermined (many false negatives). To empirically evaluate this trade-off, we explore the coverage of COVID-19-related domains in the next subsection.

3.4.4. COVERAGE OF COVID-19 RELATED MALICIOUS DOMAINS

What COVID-19-related malicious domains were missed by the CTC? For this question, we use passive Domain Name System (DNS) data from DNSDB by Farsight Security, produced through passive, real-time collection and aggregation of DNS query-response traffic between authoritative servers and recursive resolvers around the world [102]. In particular, we expect higher coverage through DNSDB where discovery of malicious domains may

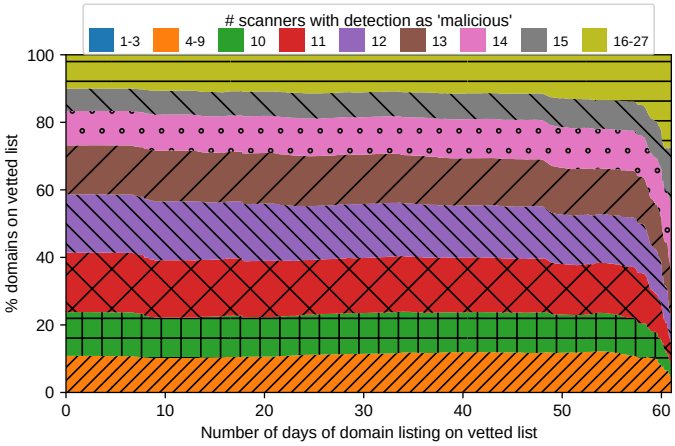


Figure 3.9: Proportion of domains with a given detection count by VirusTotal domain scanners given the duration a domain has been present on the vetted list.

be harder: subdomains as well as top-level domains (TLDs) without available registry zone files (usually country-code TLDs, ccTLDs).

In order to evaluate the CTC blocklist quality, for each day between March 1 and June 30, 2020, we extracted from DNSDB all FQDNs that match our COVID-19-related keywords from Section 3.4.2. In total, we obtain 3,011,717 COVID-19-related domains that contain at least one keyword. We want to know which of these domains are flagged by security vendors as being malicious. In total, 188,305 domains were flagged by at least 1 VirusTotal scanner; however, 162,406 of these are flagged by only one scanner (Fortinet). As presented in Figure 3.10, 5,767 domains were flagged by 4 to 10 scanners, while only 610 domains were flagged by the CTC's threshold of 10 or more scanners. Of these, 535 domains were on the CTC blocklist. According to Google Safe Browsing, the majority of these domains were related to social engineering (562 – 92.13%) followed by malware (33 – 5.40%) and unwanted software (15 – 0.24%). Recent work reported similar patterns in keyword-matched COVID-19 domains [103, 104].

We found 75 COVID-19-related domains in DNSDB that had more than 10 detections on VirusTotal and therefore met the inclusion criterion of the CTC's vetting process, yet were not included on the CTC vetted list. These missed domains are therefore false negatives of the CTC vetted list. In Table 3.2, we compare how false negatives and domains on the CTC vetted list are distributed across TLDs. While there are no significant differences per TLD type, we can observe differences in how ccTLDs are represented. For instance, .gg is the most prominent ccTLD among the false negatives, whereas it does not appear in the top 10 ccTLDs for domains on the CTC list.

In sum: the CTC blocklist contained false negatives on COVID-19-related abuse, possibly because of the CTC setting a high threshold of 10 detections on VirusTotal. Furthermore, that threshold caused an extreme reduction in how many COVID-19-related domains made it onto the blocklist, resulting in more false negatives compared to using lower thresholds. Of course, all of this underlines again the more fundamental problem

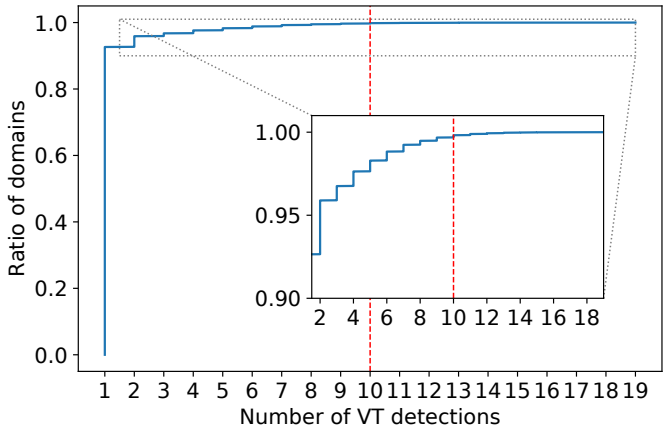


Figure 3.10: Proportion of COVID-19 keyword domains of Section 3.4.2 detected by a given number of VirusTotal domain scanners. The CTC vetting process used a threshold of 10.

TLD Type	# False Negatives		# CTC domains	
gTLD	99	(80.49%)	1,042	(84.78%)
grTLD	3	(2.44%)	3	(0.24%)
ccTLD	21	(17.07%)	183	(14.89%)
	gg 3	(14.29%)	ru 41	(22.40%)
	ga 2	(9.52%)	cl 20	(10.93%)
	eu 2	(9.52%)	cc 18	(9.84%)
	ru 2	(9.52%)	tk 14	(7.65%)
	pl 2	(9.52%)	br 9	(4.92%)
	tk 2	(9.52%)	gd 7	(3.83%)
	de 2	(9.52%)	ml 6	(3.28%)
	su 1	(4.76%)	gq 6	(3.26%)
	co 1	(4.76%)	ca 5	(2.73%)
	us 1	(4.76%)	in 4	(2.19%)

Table 3.2: Number of false negatives versus CTC vetted domains per TLD.

that the CTC's reliance on VirusTotal undermines its goal to cover threats that are not well covered by existing anti-abuse infrastructure.

3.5. IMPACT OF THE CTC BLOCKLIST

In this section we address our second research question: *Did publishing threat data in a freely available blocklist improve the ability of network defenders to act against threats, compared to the existing abuse mitigation infrastructure?* We combine various data sources to understand which, how and when actors intervene to take down domains on the blocklist, and produce a longitudinal measurement of these countermeasures. We look specifically at *domain-level* interventions and *client-side* interventions. Registrars, registries and hosting providers are responsible for domain-level interventions. This method protects all users, as it prevents them from accessing the domain, but is relatively invasive, as the intervention cannot be circumvented; it may therefore be applied more cautiously [105]. Client-side interventions inherently only protect those clients who enable the interven-

		Intervention by			
TLD type	# domains	Registry	Registrar	p-value	Any ▽
New gTLDs	4 714	45.2%	20.6%	< 0.0001	56.3%
Legacy gTLDs	14 604	2.7%	23.9%	< 0.0001	26.0%
ccTLDs	2 206	3.9%	4.7%	0.181	8.3%
Total	21 524	12.1%	21.2%	< 0.0001	30.8%

Table 3.3: Registry and registrar interventions on second-level domains from the CTC blocklist, grouped by TLD type, with p-value given for χ^2 test.

tion, but may be able to respond more quickly and aggressively to threats. For the blocklist of the CTC to improve the ability of defenders to act against COVID-19-related threats, it should flag domains that existing defenses do not act against or it should flag domains more quickly than existing defenses.

3.5.1. DOMAIN-LEVEL INTERVENTIONS

A core countermeasure is the takedown of a website. This is requested by law enforcement agencies, by targeted organizations, or specialized – e.g., brand-protection – companies acting on their behalf [106, 107, 108]. The takedown can subsequently be implemented at (sub)domain level by registries, registrars and/or hosting providers. We measure takedowns by registrars and registries primarily through 'Extensible Provisioning Protocol' (EPP) status codes [109, 110] within the WHOIS domain registration data: we consider a domain as taken down when the EPP status codes CLIENTHOLD for registrars and SERVERHOLD for registries respectively are set, which indicates that the domain is not delegated, i.e., activated in the DNS [109, 111] [5]. We retrieve historical WHOIS domain registration data for all vetted domains through VirusTotal, on July 6 and 7.

In total, 6,635 (30.8%) out of the 21,524 distinct second-level domains where we could obtain WHOIS data saw either a registrar or a registry intervention[6]. The coverage and actors of domain-level interventions depends on the type of TLD (Table 3.3). The intervention rate was the highest among new generic TLDs (gTLDs) at 56.3%, with most enacted by registries as they must comply with the most stringent requirements [112]. Meanwhile, on legacy gTLDs (com/net/org), intervention is less prevalent (26.0%) and rather fell to registrars, as this has historically been their responsibility [113, 114]. Finally, ccTLDs saw relatively little intervention (8.3%), potentially owing to their independence in setting (abuse) policies [115, 116, 117]. Consequently, assuming the CTC list contained useful TI, it could have acted as a complementary defense where other interventions were less common, in this case especially for ccTLDs.

Next, we measure if appearance on the CTC list is more timely than these interventions, as this would provide those using the blocklist with an advance warning of a live threat before it is intervened upon. We find that if an intervention did take place, it was usually faster than the CTC: 61.3% of registry and 77.1% of registrar interventions already

[5] As noted by Alowaisheq et al. [107] and confirmed by our own observations, other status codes (in particular *PROHIBITED) are at best unreliable indicators of takedown, and often reflect registry- or registrar-specific behavior (e.g., default configurations).

[6] For 35 domains, we ignore the intervention as it occurred before January 1, 2020 and is therefore unlikely to be COVID-19- or CTC-related.

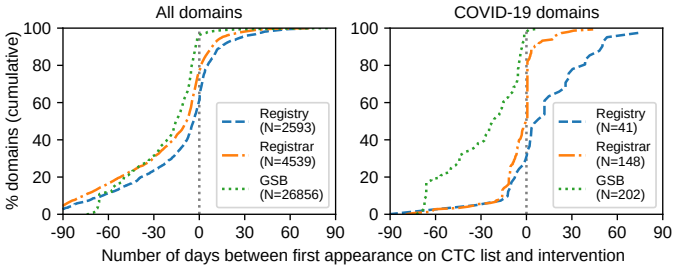


Figure 3.11: Delay between the first appearance of a domain on the CTC blocklist and interventions by registries, registrars or Google Safe Browsing (GSB).

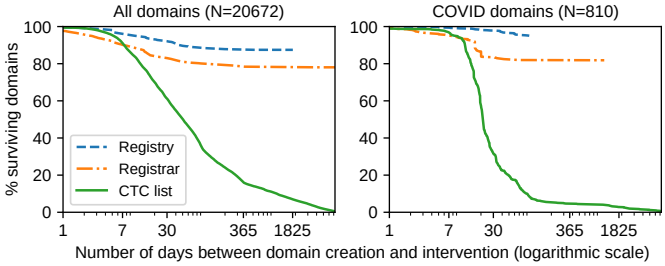


Figure 3.12: Empirical survival function for domains on the CTC blocklist where WHOIS registration data is available. We measure lifetime from domain creation until registry or registrar intervention or first appearance on the CTC list.

occurred before the domain appeared on the vetted CTC list (Figure 3.11). In cases where the CTC blocklist does predate the intervention, the delay tends to be small.This therefore suggests that for the share of domains that registrars and registries do intervene upon, their malicious contents were mostly already unavailable by the time the domains appeared on the CTC vetted list, so end users would not need the CTC's blocklisting to be protected from those domains.

Another way to see this effect is that for newly registered domains, interventions by registries and in particular registrars tend to occur closer to the registration time than CTC blocklisting (Figure 3.12). This may indicate close monitoring of new and suspicious domains by registrars and registries that results in more immediate action than the publication of the malicious domain on the CTC blocklist. However, the CTC list covers many more domains, suggesting that registrars and registries may be more careful in taking action against domains, especially once they are older, while the CTC captures TI more broadly. Moreover, the CTC list may also include more novel threats that registrars and registries are not well able to detect and take down. This is further supported by interventions across domains on the blocklist that contain any of the COVID-19 keywords from Section 3.4.2: we see registry or registrar interventions for 185 out of 821 COVID-19-related second-level domains (22.5%), lower than the 30.8% seen across the whole list ($\chi^2 = 32.591$, $p < 0.0001$) even though multiple registries and registrars had subjected COVID-19-related domains to additional verification [118, 119, 120]. The

CTC list is therefore even more comprehensive for this novel abuse type: 68.3% and 49.3% of domains appear on the CTC blocklist before a registry or registrar intervention respectively, meaning the CTC blocklist is also more proactive and therefore more useful in flagging COVID-19 domains than domains overall.

3.5.2. CLIENT-SIDE INTERVENTIONS

Client-side solutions such as domain scanning engines, firewalls, DNS-based filters and browser interstitials provide a complementary countermeasure by blocking access to malicious content, although only for their users. These solutions typically generate or ingest threat intelligence – such as the CTC blocklist – in order to determine if a resource should be considered as malicious [15, 18]. Note that VirusTotal, which the CTC used to confirm maliciousness, and its constituent scanners also serve as client-side solutions. We discussed in Section 3.4.3 how only 23 of its 84 engines succeed at detecting at least 5% of domains on the CTC blocklist, reinforcing the blocklist's utility, as it can be ingested separately to complement the other engines' protection.

Browser-based intervention Major browsers check every URL that a user navigates to against the Google Safe Browsing (GSB) service, and display a warning interstitial when the URL is known to be malicious [121], which therefore has the potential to protect a large user base. We use the Google Safe Browsing API[7] to receive hash prefixes of detected malicious URLs[8]. We retrieved an initial state on April 17, 2020, and afterwards collected updates every half hour until July 15, which allowed us to determine when a domain was first flagged. Throughout our measurement period, 28,383 domains on the CTC list (60.6%) were flagged at some point, meaning that the CTC list provides some complementary protection. However, as discussed in Section 3.4.3, the domains that GSB fails to flag tend to have fewer detections by VirusTotal engines (Figure 3.7), suggesting that their malicious status is less agreed upon. By the end, 20,962 domains (44.8%) remained flagged: it is unclear whether GSB removes domains because they are no longer considered malicious, or automatically after a certain delay. Importantly, GSB performs worse on COVID-19-related domains, at some point flagging 302 out of 1,229 domains (24.6%, $\chi^2 = 713.858$, $p < 0.0001$). This might reflect that these COVID-19-related domains contain more scams and forms of abuse outside of the normal scope of GSB, which is focused on the conventional abuse categories of phishing, malware and spam. The CTC list may therefore provide greater benefits for these domains. While GSB does not achieve full coverage, if it flags domains, it does so before the CTC in 96.3% of general cases, and 98.0% of COVID-19 domains (Figure 3.11), with remaining domains being added quickly to GSB after their first appearance on the CTC list. The CTC list therefore leaves users vulnerable for longer across the domains that GSB detects, and would only be useful if its additional domains consisted of novel TI.

[7] https://developers.google.com/safe-browsing/v4

[8] Discrepancies exist between the output of the Google Safe Browsing API and actual browser interventions, among others due to "data sharing restrictions" [122]. Our data are therefore an approximation of the latter interventions.

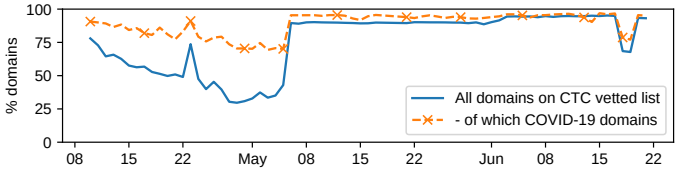


Figure 3.13: NXDOMAIN responses (i.e., blocking interventions) by the Quad9 DNS resolver, indicating ingestion of the CTC blocklist on May 6.

DNS-based intervention Quad9 is a public DNS resolver that blocks malicious domains by responding to queries with NXDOMAIN [123]. We retrieved DNS records for all vetted domains from Quad9 [124] once a day from April 10 until June 21. Quad9 relies on threat intelligence from at least 18 providers [124] and was reported to include the CTC blocklist from May onward [88]. This inclusion proved beneficial: before May 6, Quad9's detection rate was 30% at its lowest point (70% for COVID-19-related domains), but from then onward Quad9 included almost all of the blocklist (Figure 3.13). This shows that the CTC blocklist managed to incorporate threat intelligence that was unknown to at least some security service providers.

In summary, did the CTC improve the ability of defenders to act against threats, compared to existing abuse mitigation infrastructures? We find that for 58.4% of the FQDNs on the CTC blocklist, existing abuse mitigation pipelines at the domain level or in the browser were all faster than the CTC at intervening. For these domains, the pooling and sharing of data in a public blocklist then provided little additional value. For the remaining 41.6%, defenders – such as public and private organizations or managed security service providers – who ingested the open CTC blocklist did however improve their ability to defend themselves, compared to relying on existing anti-abuse pipelines. This advantage was even more sizable for COVID-19-related domains, at 74.9% additional coverage, showing once again that the CTC was more effective when focusing on its original goal of collecting COVID-19-related abuse. One additional indication of the CTC's utility occurred in May 2021, when the DNS provider Quad9 started ingesting the CTC list. This created an almost complete overlap between the CTC blocklist and Quad9's client-side interventions, leading to better protection for its users. To put it briefly: if actors in abuse mitigation intervened, they tended to do so more quickly than the CTC, but often they didn't intervene, so the list improved the ability of defenders to act, especially on COVID-19-specific threats.

3.6. ETHICS

Our study describes a threat information sharing community that anyone could join [101] but was nevertheless not public: information was TLP:GREEN, i.e., restricted to the community [125]. Throughout this work, we adhere to the CTC's Code of Conduct [126]. In accordance with the cited "Chatham House Rule", we take great care to prevent identifying any community members because this might impact them professionally, except for the founder, who has publicly stated his role. Further, we do not cite from observed conversations, as this might impede trust in the confidentiality of this and future threat

information sharing initiatives. To still provide a rich image of the community, we rely on public sources such as the CTC's webinars [90] and interviews with steering committee members [127, 128, 79, 85, 86]. The blocklist that we analyze in-depth has been made fully public by the CTC [87]. The lead author's institutional review board (IRB) has approved this study design.

3.7. DISCUSSION

We have described collaboration in the Cyber Threat Coalition community, which removed many of the entry barriers normally present on threat information sharing arrangements. The CTC relied on Slack and AlienVault OTX groups to pool threat information, then vetting it through VirusTotal before making the resulting blocklist freely available. To learn whether this unique arrangement led to improved threat defenses, we investigated two questions.

First, *by pooling data from its community, did the CTC improve coverage of COVID-19-related threats over existing defenses?* We find that the CTC primarily consolidated existing sources, rather than produce or propagate new threat intelligence. Just 10 users contributed most of the domains in the CTC's AlienVault group and over time, user contributions on the blocklist were outstripped by data from named and unnamed security firms. The community drifted away from its initial goal of tracking pandemic-related abuse: on the CTC blocklist, domains with COVID-19-related keywords (2.6%) were overshadowed even by those with simply the keyword *whatsapp* (2.8%). General phishing domains eventually made up most of the CTC blocklist. This drift was caused by the reliance of the CTC on VirusTotal for vetting domains to be propagated on its blocklist: it required 10 detections by VirusTotal engines. This high threshold meant that the resulting list tended to reproduce the detection of conventional abuse by antivirus engines, rather than contribute new threat intelligence focused on COVID-19. In terms of coverage of these latter threats, we found 75 false negatives – domains that should have been on the list according to the standards of the CTC itself, yet were missing. Thousands more COVID-19-related domains might have been missed by the CTC, with the exact number depending on what threshold one chooses in terms of VirusTotal detections.

Second, *did publishing threat data in a freely available blocklist improve the ability of network defenders to act against threats, compared to the existing abuse mitigation infrastructure?* Our analysis presents a dual view on who acted on domains listed by the CTC and when. On the one hand, no actor achieved full coverage in their interventions – be it at domain-level or client-side – meaning that the CTC vetting process succeeded in delivering an aggregated, more complete set of malicious domains warranting action. This is perhaps best demonstrated by its inclusion in the Quad9 DNS service. On the other hand, where actors in abuse mitigation did intervene, they were usually faster. For 58.4% of the FQDNs on the blocklist, the CTC lagged in incorporating the indicators in its list. Here, they provided little added value – in particular because other interventions such as domain takedowns or browser interstitials typically have a much wider reach. For the small fraction of COVID-19-related domains, the CTC blocklist was more effective in terms of coverage and improving existing defenses. Here, the apparent lack of focus of the CTC on COVID-19-related abuse impaired the overall utility of its blocklist.

Based on our findings, we draw three lessons for future open source threat information

sharing initiatives. Our first lesson is that *scaling up the community does not automatically lead to better pooling of threat information.* In just a few weeks' time, the CTC managed to set up a pipeline for collecting, vetting, and disseminating threat intelligence. Throughout its progression, CTC admins repeatedly pressed members to "signal boost our social media posts", inviting more people to join the community [91]. Possibly, the admins had assumed that network effects would only increase as the community grew, and did not anticipate the dynamics of a volunteer organization. The CTC did not capitalize as much as it could have on its pool of 4,000 volunteers: contributions on AlienVault were made by just a fraction of community members. This is in line with earlier research on open source software development (Section 3.3.1). However, that research also found that open-source communities have a long-tail of members with small contributions, like a bug report. Along the same lines, the CTC could have benefited from indicator vetting by its members (our second lesson). Scaling up the CTC community may also have disincentivized threat information sharing, because it exacerbated the free-rider problem: having a large number of untrusted participants may have discouraged some contributors from sharing indicators that were sensitive or that they feared would be used commercially by other firms.

The second lesson is that *openness of the community requires a scalable quality assurance process for the contributed indicators.* The CTC chose to fulfill that need via VirusTotal, but thereby undid some of the benefits of pooling new threat information in the first place. As described in Section 3.3.2, indicators had to meet a threshold level of 10 scanners in VirusTotal before they were propagated to the CTC blocklist. Only 21 scanners were able to detect more than 5% of the COVID-19-related domains on the blocklist, so a threshold of 10 means that half of the dominant scanners need to already detect the indicator before it was shared and published. Although this workflow is scalable and produced a blocklist with a low number of false positives, it also meant that valuable indicators that were not yet known to the dominant VirusTotal engines were discarded because they did not meet the threshold. In light of the CTC's mission, it was ironic that this particularly affected indicators related to COVID-19. Had the CTC's solution for manual vetting of indicators, the Slackbot, come earlier, then it might have prevented the impact of relying on VirusTotal. Crowdsourced vetting of indicators can be successful, as PhishTank has shown [82]. User participation had already tapered off, however, when the Slackbot was introduced, so the transition away from VirusTotal was not successful.

The CTC was founded on the premise that existing abuse and threat information sharing mechanisms were unsuitable or unprepared for the risks posed by COVID-19-related abuse [79] and that a new response was needed. Our third lesson is that *existing threat intelligence and abuse mitigation structures are actually quite resilient and able to adapt to 'new' types of threats.* Where the existing anti-abuse pipelines intervened, they did so faster than the CTC could detect the domain and include it on its blocklist. For example, 96% of the client-side interventions through Google Safe Browsing occurred before the domains were featured on the CTC list (Section 3.5.2). Another way to read these findings is that due to its loss of focus on COVID-19-related domains, the CTC was forced to 'compete' with general-purpose abuse sharing mechanisms, a battle it was unlikely to win.

3.8. LIMITATIONS

The fast-paced evolution of the CTC and its operational processes introduces inherent limitations for our study. Changes to the CTC vetting process and our own data collection, as well as blocklist hosting issues, caused temporal gaps in our data. Moreover, where we do not have access to historical data, in particular for VirusTotal detections and active DNS records, we only collect data starting from the moment when a domain was first included on the blocklist. Finally, certain improvements to the vetting process, such as the Slackbot or the preferential treatment of curated third-party sources, were only introduced by the time community participation had dwindled. Given the low number of contributions, we could therefore not investigate if these would have had a significant impact on the blocklist.

As VirusTotal is the main driver of the CTC vetting process, our analyses are inherently biased by its classifications. Figures 3.6-3.10 provide some insight into how these classifications are distributed for our data. However, researchers have questioned the reliability of VirusTotal [84, 129] and other (phishing) blocklists [130]. These shortcomings may be exacerbated by the novelty of COVID-19-related abuse, as well as the semantic discussion on whether tactics such as scams or price gouging constitute 'maliciousness' at all. Nonetheless, VirusTotal provides us with objective and independent detection metrics across a large set of domain scanning engines, serving as a strong signal for maliciousness, meaning that domain-level and client-side interventions can be expected. We attribute these interventions based on indicators that carry a level of uncertainty, as also observed in previous work [107]. In particular, the availability of sufficiently detailed WHOIS data is skewed toward gTLDs [108], and we assume correct parsing of its non-standard format by VirusTotal. Our results in Section 3.5 therefore serve as a lower bound to actual interventions. Similarly, the count of COVID-19-specific abusive domains is a lower bound, as we assume the presence of pandemic-related keywords. However, other domains in Section 3.4 may have only carried COVID-19-related content on their web page, or have been propagated within a COVID-19-related context (e.g., a spam email). Where we quantify potential false negatives, we equally did not intend to do so exhaustively.

Our research focused on the CTC blocklist. We evaluated it through its fit for purpose to inform real-time enforcement actions, in line with the CTC's mission, but the list of COVID-19 related abuse material could conceivably also be used in retrospect to evaluate security controls or even as training data for machine learning purposes. Finally, although the blocklist was an important and visible outcome of the community's efforts, it was not the only one. The CTC produced threat advisories and community meetings, and facilitated communication between members of the security community. We refrain from analyzing the posts in the CTC Slack workspace due to our adherence to the CTC code of conduct [126]. These conversations likely continued in private messages between members, where we would not have visibility on their outcomes, such as takedowns or law enforcement intervention. Indeed, an important added value of the CTC may be not in the data that they output, but in the network of peers that they managed to bring together at short notice, and as a showcase for what open source threat sharing might look like.

3.9. RELATED WORK

The performance of blocklists as sources of threat intelligence has been the topic of previous studies, which raised questions about the coverage of relevant threats by open sources [15, 16, 44, 131] as well as closed, commercial threat intelligence sources [18]. Peng et al. [84] found that even the best engines on VirusTotal missed 30% of submitted phishing sites and Oest et al. [132] managed to evade being blocklisted for 55% of phishing domains using simple cloaking techniques. Metcalf and Spring [17] hypothesized such problematic coverage to be an artifact of the collection method of abuse infrastructures, with each using disparate methods to detect threats from their specific vantage points. Notably in the blocklist literature, Li et al. [15] proposed metrics by which to understand threat intelligence and calculated those metrics for a set of open source blocklists. We draw on its coverage and timeliness metrics and our study also describes a blocklist, that of the CTC, but we go beyond descriptive metrics and measure the blocklist's ability to inform countermeasures. More generally, where Li et al. and other studies take blocklists as-is for their analysis, we conduct measurements on the CTC blocklist to shed light on the open source threat sharing process by which it was produced.

The open source threat sharing model is not entirely new: security information has always propagated through "informal networks of trusted security professionals that exist across [organizations]" [133] and earlier research has described initiatives that share the CTC's objective to pool threat information, but do not match it in terms of access and scale. Thomas et al. described aggregating data from within Google services, in a lab environment and with access restricted to the authors [83]. Another aggregator, Facebook ThreatExchange, was included in the set of feeds analyzed by Li et al. [15], who described it as a "closed-community platform" of "hundreds of companies and organizations". Here, access is based on being involved in software development for the Facebook platform. Outside of academic research, we see the Cyber Threat Alliance, an industry partnership of 33 firms that share reports and indicators with each other ahead of publication on the basis of quid pro quo, and therefore not available to the public [75]. VirusTotal combines scan engine logic to classify files and URLs [129], and ingests external data feeds as inputs [84]. Despite the fact that users can also contribute binaries and URLs as inputs on a limited free plan, VirusTotal is not a community but a commercial service. Probably most approximate to the CTC, PhishTank is an aggregator of domains and URLs suspected of phishing that provides a practical, real-world example of the promise of collaboration to identify threats. It has a crowdsourcing capability to let users validate maliciousness and like the CTC blocklist, its feed is shared free of charge. It is an older initiative, with the last study of it from 2008, which did not evaluate the ability to inform countermeasures [74]. New user registration on PhishTank has been closed since 2020. Three properties set the experiment of the CTC in threat information sharing apart from these initiatives. First, the low barriers to entry. Formalizing a threat information sharing community has not been attempted before at the scale of thousands of volunteers [92]. Second, the community's efforts were documented in open sources such as webinars and interviews [127, 128, 79, 85, 86, 90]. And third, it produced information goods that were made freely available – potentially magnifying the impact of the threat information [87]. Because of these properties, the CTC approached what Benkler [134] called peer production: open creation and sharing performed by online groups. In this case, peer production of security

information. We conclude that the scale at which threat information sharing occurred in the CTC in response to the pandemic offered a unique opportunity to investigate the collaborative model.

More generally, authors have recently described internet abuse related to COVID-19, which the CTC also tracked. They have pointed to signs of coordinated campaigns [135, 136] and drew attention to the risk of overzealous filtering of COVID-19-related material [137, 103]. Bitaab et al. investigated examples of phishing related to COVID-19 and concluded that the existing anti-phishing ecosystem fell short, based on the sheer volume of COVID-19 related fraud reported on by the U.S. Federal Trade Commission (FTC) [138]. Our measurements support their finding that for the niche of COVID-19 related material, an organization like the CTC could play a valuable supplementary role.

3.10. CONCLUSIONS

The Cyber Threat Coalition had the aim to "break down traditional barriers to intelligence sharing [and] produce a professional-quality threat feed that the broad IT security public [could] rely upon" [77]. We found that by pooling data from its community, the CTC managed to improve coverage of threats related specifically to COVID-19 over that of existing defenses, and we found evidence that the CTC was faster than other defenses to list such domains. Therefore the community improved the ability of network defenders to take action by publishing its threat data in a freely available blocklist. Over time the CTC lost focus, until it aggregated mostly generic abuse information. We described how this can be traced back to choices that the community made in scaling up its quality assurance processes using VirusTotal. This prevented the CTC from delivering some of its value, as it relied on a threshold number of scanners to recognize a domain as malicious, dropping valuable new indicators in the process and causing it to lag behind established defense mechanisms.

Looking back on his experiences, founder Joshua Saxe said: "We had a lot of [volunteering] energy, but we didn't have the right organizational machinery to funnel that energy." [88]. Given that no open source threat sharing community of this kind previously existed, it is unsurprising that the Cyber Threat Coalition went through growing pains. Although we have examined its impact critically, it is only as a result of the hard work of the members of this community that we have been able to investigate the principles of open source threat information sharing at all – and such a community may have impact in subtle ways that are not easily observed or quantified, such as building trust and facilitating contact. We dedicate this paper to the volunteers of the CTC and hope that the lessons discussed in this work may contribute to future threat information sharing.

4

SECURITY OUTCOMES

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. We empirically examine this claim by analyzing attacker behavior before and after publication of indicators of compromise (IOCs). Using IOC feeds from two leading commercial providers, 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 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.

This chapter was previously published as a research paper: Xander Bouwman, Aksel Ethembabaoglu, Bart Hermans, Carlos Hernandez Ganan, and Michel van Eeten. "Can IOCs Impose Cost? The Effects of Publishing Threat Intelligence on Adversary Behavior". In: *32nd ACM Conference on Computer and Communications Security (CCS)*. ACM CCS 2025. United States: ACM, 2025. DOI: 10.1145/3719027.3765026. URL: https://www.sigsac.org/ccs/CCS2025/

4.1. INTRODUCTION

The US Department of Defense (DoD) strategy on countering state-sponsored espionage has been heavily influenced by cyber persistence theory, representing a shift toward a more active stance on cybersecurity [139, 29]. 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 [140, 141, 142, 143]. 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 Commission on cybersecurity policy reform [30] and a 2023 White House strategy document called on the US to "execute disruption campaigns at scale" [144].

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" [145, 146]. 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" [147]. 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 [13]. And indirectly, this idea is reflected by the consolidated value of the commercial cybersecurity industry, now worth hundreds of billions of dollars. One firm alone, CrowdStrike, had a market capitalization that exceeded $113B in early 2026 [148]. 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 [29]. 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 4.2). 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 they 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 4.4).

- 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 8.1% of observed variance. We analyzed various covariates but did not find other strong explanatory factors, pointing to the existence of confounding variables (Section 4.5).

- 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 4.6).

- 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 (Section 4.4.2).

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 4.6). This article is structured as follows. After we explain our data and method (§4.3), we explore IOCs and sightings with descriptive statistics and measurements (§4.4). We then perform statistical modeling of the relationship between IOC publication and sightings (§4.5) and answer our research question.

4.2. 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.

4.2.1. STRATEGIC STUDIES

From a close reading of policy documents, speeches of key decision-makers, and other primary sources, Healey [29] 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 [29]. 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." [149] 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 [150].

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 [143]. Even just signaling capabilities and credibility have been said to be able to impose costs [30], but whether or not such steps, or even reciprocal attacks, could actually deter attackers from consecutive operations has been hotly debated over the years [151]. Our study provides new empirical data for future theoretical work on such questions.

4.2.2. 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 [51]. Threat intelligence uses practices from traditional counterintelligence to capture adversary behavior in order to build such signatures and to inform other security controls [152]. In the technical cybersecurity community, 'imposing costs' appears primarily when discussing security controls, e.g., controls that introduce deception in defensive environments [153, 154].

Existing research has described threat intelligence IOCs from open sources and commercial sources and measured them along various quality dimensions [16, 15, 18, 155], also investigating the ability to enable intrusion detection and to do so in a timely fashion [21, 156, 157]. Some authors also measured the *sightings* that are generated from performing intrusion detection using threat intelligence. Vermeer, Gañán, and Eeten 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 [158].

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, Booij, and Doerr measured the timeliness of open source blocklists [16], which capture a different kind of activity than commercial IOCs [18]. 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 [156]. 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 [159, 160]. 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 [161]. These findings about cybercrime actors are consistent with the 'pyramid of pain' model popular in threat intelligence [13], 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 [162]. Such zero-day attacks are still the exception, however, and most

attack campaigns described in threat intelligence reports make use of publicized software vulnerabilities [163].

4.3. 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.

4.3.1. 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 4.7.

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 4.3, 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 [164]. 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.

4.3.2. 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 Hypertext Transfer Protocol (HTTP) sessions, and domains on Transport Layer Security (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.

4.3.3. THREAT INTELLIGENCE IOCS FROM TWO MARKET LEADERS

Indicators of compromise (IOCs) enable intrusion detection by capturing artifacts of adversary behavior [51]. The IOCs in this study were produced by two firms described by Gartner as leading the threat intelligence industry [164]. 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 the 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 Section 4.4.3 and Section 4.5.2.

4.3.4. 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 [165]. 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 lively debate over how long threat intelligence IOCs lead to relevant sightings [156]. 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 publication 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

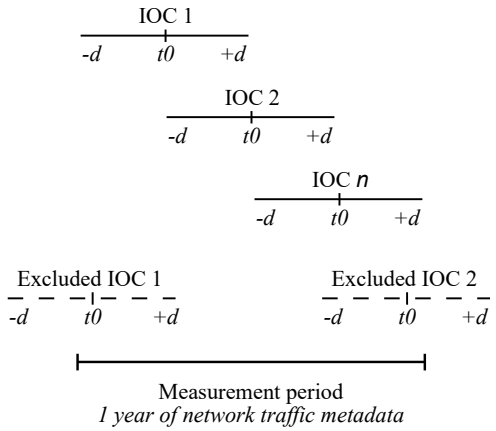


Figure 4.1: Illustration of our methodology. In order to reduce noise, we assumed sightings within a window of $d = 50$ before or after publication of the respective IOC to be relevant, so a total of 101 days (Section 4.3.4). We performed sensitivity analysis for different values of d.

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 4.1). 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 (Section 4.3.3) 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 4.1 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 [156]. 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 4.4.2. 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 C.

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 4.3 shows, the cohort of new domains does indeed lead 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 C.

4.3.5. STATISTICAL MODELING

In Section 4.5, 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 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 4.3). As we report in Appendix C, 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 [167]. We operationalized this with the Python library `linearmodels.panel.PanelOLS`. We confirmed that a fixed effects model was more suitable for our situation than random effects – which assume that the

[1] Tranco is a research-oriented top site ranking hardened against manipulation [166]. The list we used is available at https://tranco-list.eu/list/3VPNL/1000000

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, to shed light on between-group variation (which fixed effects does not allow), we also fitted simple 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 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 4.7).

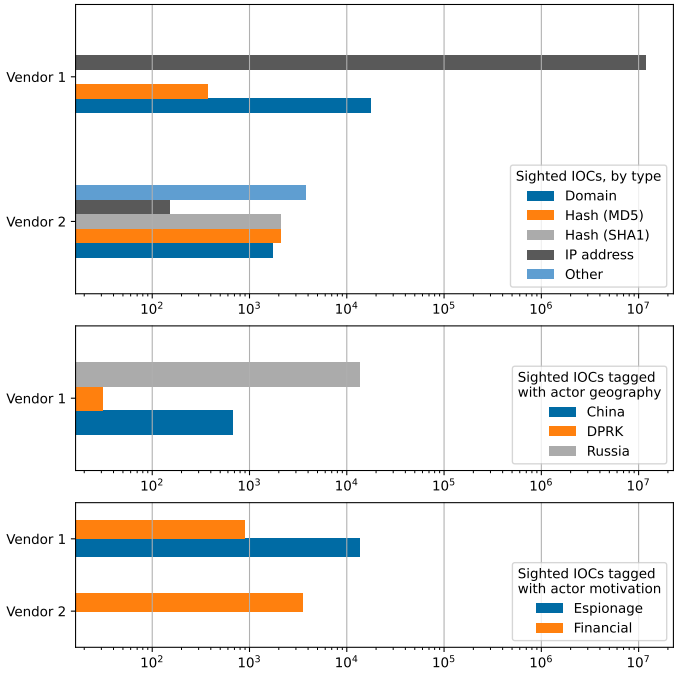


Figure 4.2: Counts of IOCs in two commercial feeds that led to sightings during our measurement period of about 1 year, *before* filtering for the most relevant sightings (Section 4.3.4).

4.4. 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.

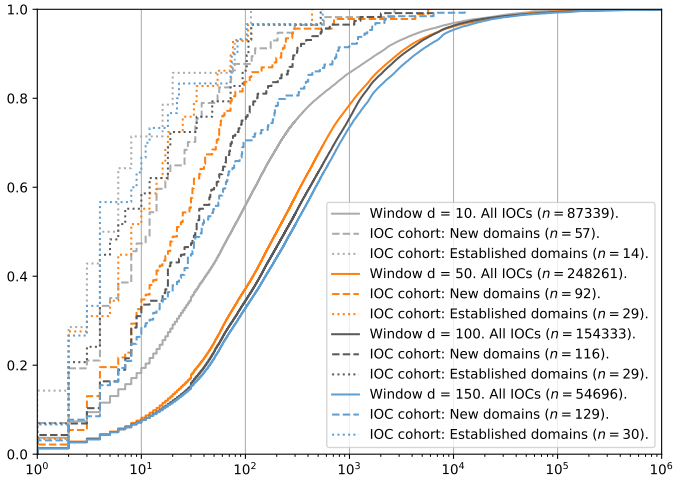


Figure 4.3: Empirical distribution function of sightings per IOC in our dataset, *after* filtering for the most relevant sightings (Section 4.3.4). Around 20% of IOCs lead to more than 1,000 sightings. Our domain IOC cohorts receive fewer sightings per IOC, especially the established domains cohort.

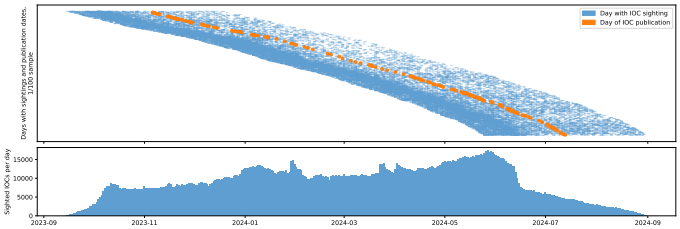


Figure 4.4: Sightings and IOC publication dates over our measurement period of 356 days, for window $d = 50$.

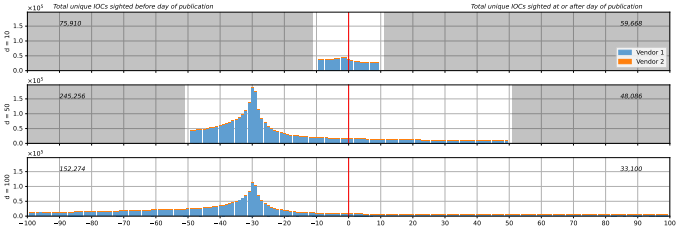


Figure 4.5: Counts of IOCs sighted daily relative to their respective publication dates, for three windows of d. A peak in sightings can be observed around 30 days before IOC publication for windows $d = 50$ and $d = 100$, as well as for $d = 150$ (not pictured). This pattern is not visible in $d = 10$, demonstrating that defenders need to retain network logs at least 30 days to observe this activity (see Section 4.4.2 and Section 4.6). Sightings from Vendor 2's IOCs did not reach the visualization threshold.

4.4.1. 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

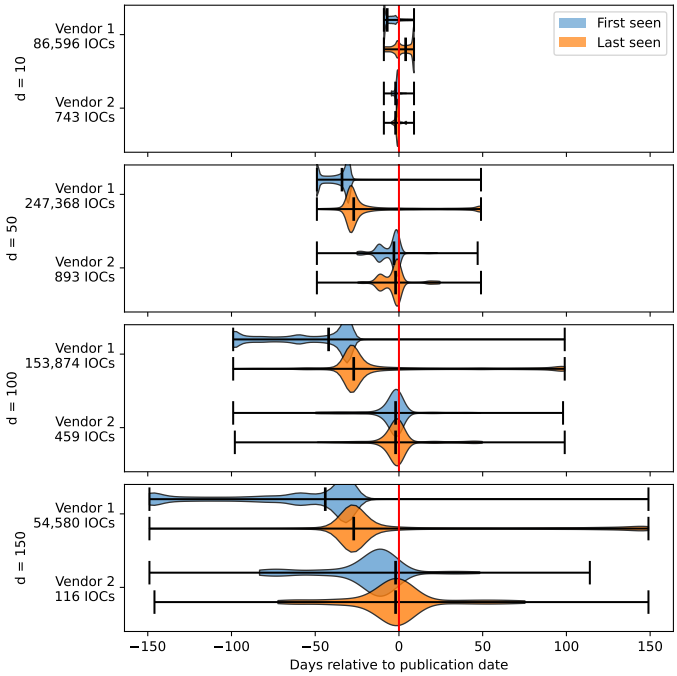


Figure 4.6: Distributions and medians of first and last sightings of IOCs, as days relative to the IOC's publication, per vendor. Note that plot area size between Vendor 1 and Vendor 2 is *not* proportional to the number of observations.

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 [21]. As Figure 4.2 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, 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 4.2 also shows the counts of 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 Section 4.5.2 into context.

Around 20% of IOCs caused disproportionately many sightings, as the curve beyond 0.8 in Figure 4.3 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 4.5). 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 Section 4.3.4. Both cohorts lead on average to fewer sightings per IOC, which reflects that domains lead to fewer overall sightings (Table 4.3). New domains are noisier than established domains for all windows of d, which aligns with our expectation for these IOCs. In Figure 4.4, 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 Section 4.3.4). The reason that the slope at the end takes a different shape is that most sightings occurred before the IOC publication.

4.4.2. TIMELINESS OF COMMERCIAL IOCS

For IOCs to 'impose cost', they must point defenders to resources that attackers are actively using at publication time (Section 4.1). 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 4.1). 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 time that IOCs became available to the customer. The inverse was true for a fraction of IOCs, 1.2% (2,921 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 4.1, IOCs from both vendors demonstrated limited ability to detect ongoing threats at time of publication. This challenges their ability to 'impose cost'.

Figure 4.5 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 Section 4.3.4). 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 4.3).

We further investigate sightings relative to IOC publication date with Figure 4.6, which shows the first and last sighting of each IOC for the various windows of d. Table 4.2 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 4.1: 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 4.2: 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 4.1), 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 the 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 4.6).

4.4.3. 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 [18]. In

Section 4.5.2 we perform statistical modeling using the same factors.

As Table 4.3 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 [168]. 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 4.2, Figure 4.6). We note again that our measurements are based on logs of network traffic metadata (Section 4.3). This visibility is reflected by network-based IOCs having a higher probability to lead to sightings.

The results by domain IOC cohort in Table 4.4 let us compare established domains to new domains, as explained in Section 4.3.4. 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 4.5 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 [53, 169].

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

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

Table 4.3: Distributions of first and last sightings, in days relative to IOC publication, by IOC type and using $d = 50$.

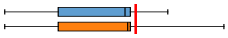
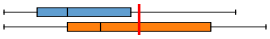
IOC Type	Count	Medians	FS/LS distribution
IP address	246,275	-34, -27	
Domain	1,034	-33, -19	
Hash (MD5)	429	-4, -3	
Hash (SHA1)	361	-2, -2	
Other	162	-23, -23	

Table 4.4: Distributions of first and last sightings, in days relative to IOC publication, by domain IOC cohort (as defined in Section 4.3.4) 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	

the need to proactively rotate through resources, rather than wait for detection. This is consistent with prior research in cybercrime [170, 171].

Finally, we look at a very small subset of tags that attribute the IOCs to the geography of the threat actor (Table 4.6) and that we were also able to map to nation-states based on publicly available information or the vendors' actor naming conventions [53, 169]. Russian actors have been thought to conduct relatively 'noisy' operations, also as a form of signaling [172]. 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.

4.5. 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. It is printed in the Appendices for formatting consistency (Table C.2, on page 114). The Appendices also provide descriptive statistics for our dependent variable and covariates. Although our results

Table 4.5: Distributions of first and last sightings, in days relative to IOC publication, by actor motivation (0.21% of IOCs have this property) and using $d = 50$.

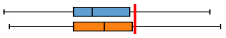
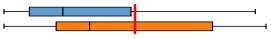
Actor motivation	Count	Medians	FS/LS distribution
Financial	516	-16, -12	
Espionage	16	-48, -28	
Untagged IOCs	247,729	-34, -27	

Table 4.6: 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	

are based on a window of $d = 50$ days around IOC publication date, as described in Section 4.3.4, we also fitted our models for the three other values of d. This did not lead to substantively different outcomes. The Appendices also include these additional regression tables.

4.5.1. 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 4.4). Statisticians have emphasized that significance levels are not always informative for large datasets, and so we focus on effect size [173].

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.

4.5.2. 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 4.7. 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 4.4, 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 Section 4.4.3, 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 the Democratic People's Republic of Korea (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 that 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 (Section 4.4.3).

4.6. 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.

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

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 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" [174]. 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 [161]. Other examples are the short life span and rapid cycling of phishing sites [175] and responses of cybercrime groups to takedowns [159, 160].

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 [176]. 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 (Section 4.4.2). If enterprises combine forces, for example via an ISAC or a Computer Security Incident Response Team (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" [146]. 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 [177] or simply by decision-makers lacking better options [178].

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) [179, 180]. 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 sug-

gests that, yes, peer-to-peer sharing of threat intelligence might, in certain cases, be able to report IOCs faster than corporations [32]. 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 (Section 4.4.2). Vendors could choose not to immediately publish IOCs because their products are evaluated by metrics other than just timeliness [16]. In a 2020 study [18], 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, in Section 4.4.3, 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.

4.7. 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 [21], 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 [18]. 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 [181, 182]. We see further opportunities for future work to explore factors that contribute to sightings like those that we identified in Section 4.4.3, 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.

4.8. 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 4.4).

To quantify the relationship between sightings and IOC publication, we fitted a panel regression model using fixed effects, finding a negative relationship between IOC publica-

tion and the probability of future sightings occurring on a given day. The model was able to explain a small but meaningful 8.1% of variance in whether or not an IOC was sighted (Section 4.5). 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 4.4).

In sum, the commercial threat intelligence IOC 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 4.6). We are careful with claims of causality and point to the challenges of trying to measure malicious activity (Section 4.7).

5

CONCLUSION

The central research question of this dissertation is: *How do the economic incentives in the market for cybersecurity threat intelligence shape security outcomes?* The previous chapters addressed three knowledge gaps in the definition and function of the market for cybersecurity threat intelligence, as per Chapter 1.

5.1. CONSOLIDATED FINDINGS

The first knowledge gap concerned the characteristics of commercial threat intelligence, which had largely been opaque to consumers. Chapter 2 compared the feed composition of two market-leading commercial vendors and open source blocklists, finding that threat intelligence feeds consist mostly of singletons: unique data points. There was almost no overlap between the indicators provided by different vendors, nor with open sources. Even for specific threat actors tracked by multiple vendors, the average overlap was only 2.5% to 4.0%. Interviews with security professionals whose organizations purchased commercial threat intelligence feeds revealed that they did not perceive fragmentation as a failure of the providers. Instead, as consumers, they preferred to conserve scarce analyst attention by selecting feeds with low false-positive rates rather than low false-negative rates, i.e., comprehensive coverage of relevant threats. Effectively, consumers defined threat intelligence primarily in terms of minimizing disruption to their internal workflows rather than its ability to improve security outcomes by detecting threats.

The second knowledge gap regarded community-driven threat intelligence production, and whether this could overcome issues related to fragmentation and quality present in the commercial market. Chapter 3 addressed the second gap by analyzing the COVID-19 Cyber Threat Coalition, the largest volunteer threat data sharing community to date. The study found that while the community succeeded in aggregating existing industry data, it struggled to produce novel intelligence on pandemic-related threats. A key bottleneck was reliance on automated vetting via VirusTotal, which required a high threshold of prior detections. This process inadvertently filtered out the very novel threats the community sought to uncover, causing the blocklist to drift toward generic, already-known

threat data types. However, for the specific subset of COVID-related indicators that were captured, the community model demonstrated clear added value, identifying threats that were not yet known to existing defensive infrastructures.

The third knowledge gap dealt with the extent to which threat intelligence influences security outcomes, specifically in real-time intrusion detection and its effects on threat actor behavior. Chapter 4 examined the strategic assumption that publishing Indicators of Compromise (IOCs) imposes costs on adversaries by forcing them to abandon infrastructure. The study measured attacker activity before and after IOC publication, finding that for 81% of sighted IOCs, threat actors had already abandoned the resources prior to publication. One major vendor consistently showed a lag of approximately 30 days between peak attacker activity and the release of IOC. While statistical modeling indicated a small, meaningful reduction in sightings following publication, the findings largely challenge the efficacy of commercial threat intelligence for real-time intrusion detection and for imposing direct operational costs on sophisticated adversaries.

5.2. CENTRAL RESEARCH QUESTION

The central research question of this dissertation is: *How do the economic incentives in the market for cybersecurity threat intelligence shape security outcomes?* The overall conclusion is that through their different incentive structures, each of the three types of sources of CTI capture distinct parts of the threat landscape, meaning that they are valuable complements rather than substitutes. Incentives are shaped by quality metrics but consumers of threat intelligence struggle to measure quality. Faced with this uncertainty, consumers rationally optimize operational efficiency over security outcomes, inadvertently incentivizing vendors to delay intelligence in favor of curation.

The studies in this dissertation have contributed knowledge about the market structure, production processes, and operational impacts of cybersecurity threat intelligence. Through the findings, recurring questions emerged around three themes: *i)* evaluation of competing production models; *ii)* information asymmetry; *iii)* metrics and bounded rationality. The below sections review findings around these themes, providing a more rich and nuanced answer to the research question.

THREAT INTELLIGENCE PRODUCTION MODELS

In evaluating competing production models, the question arose: do open source and community-based models serve as free, lower-quality substitutes for commercial threat intelligence, or can they effectively replace the need for commercial providers? To address this, we must first recognize that there are broadly three models for producing and sharing threat intelligence: open sources, commercial providers, and community-based models. Open sources (OTI) typically consist of public lists of IOCs, often functioning as abuse feeds or blocklists [15, 16]. Commercial threat intelligence (PTI) is a product sold by vendors who promise high-end, original insights into threat actor behavior. Community-based or shared threat intelligence (STI) is sourced from communities where members exchange information, often on a quid pro quo basis or pro bono, as in the case studied in Chapter 3.

This dissertation has studied these three types of threat intelligence sources and

provided empirical data on them, showing that they are distinct in their content and incentive structures. Chapter 2 established that commercial threat intelligence consists of distinct intelligence products compared to open sources. The study found almost no overlap between the indicators of leading commercial vendors and those of open-source feeds, suggesting they captured distinct parts of the threat landscape. This contradicts the assumption that open source threat intelligence is merely a free, lower-quality version of commercial threat intelligence; instead, they represent fundamentally different views of the threat ecosystem, focusing on different types of activity.

Chapter 3 contrasted community-based models with commercial models, demonstrating that community-based models can approach, and in some cases surpass, commercial and existing models in specific niches. In the 'natural experiment' of the COVID-19 Cyber Threat Coalition, the community successfully pooled resources to track pandemic-related threats. For the specific subset of COVID-related indicators, the community model showed promising evidence of added value, with only 25.1% of domains being known to existing abuse detection infrastructures at the time of listing. This showed that when aligned with a focused mission, community production can identify threats that even established commercial and open industry pipelines miss. In sum, community models are a valuable addition for improving coverage.

INFORMATION ASYMMETRY

A second recurring question evaluates the opacity of the market: does this opacity stem from vendors intentionally exploiting their knowledge advantage, or from a collective inability to measure true quality? As described in Chapter 2, the market for commercial threat intelligence indeed operates with significant opacity. Vendors employ "cyber noir" marketing and mythical symbolism to signal capability in the absence of verifiable metrics. These are indications of information asymmetry. Security economics theory proposes two distinct models for markets where quality is difficult to verify: Akerlof's "Market for Lemons," characterized by information asymmetry where the seller knows more than the buyer [22], and Grigg's "Market for Silver Bullets", characterized by information insufficiency where neither party knows the true quality of the good [23].

Our findings suggest that the market for threat intelligence is more accurately described by Grigg's model of insufficient information. Chapter 2 demonstrated that market leaders lacked a comprehensive view of the threat landscape, with overlap between vendors of 2.5% to 4.0% for specific threat actors tracked by both vendors. Chapter 3 further illustrated this insufficiency in a community model, where the lack of an objective standard for 'novelty' led to reliance on proxy metrics like VirusTotal detections, inadvertently filtering out the very intelligence the community sought to produce. Finally, Chapter 4 challenged the ability of threat intelligence IOCs to enable real-time intrusion detection. The finding that 81% of indicators were actionable only after adversaries had abandoned resources suggests that perhaps neither buyers nor sellers truly understood the efficacy of threat intelligence for real-time intrusion detection. Thus, the market for threat intelligence is best understood as a "Market for Silver Bullets" [23], driven by mutual uncertainty rather than exploited asymmetry. As discussed below, firms have recently demonstrated willingness to collaborate and increase transparency, which we can see as a way to increase both customers' and their own information. In this "Market for Silver Bullets", the

inability to verify quality means that vendors are not fully rewarded for speed or accuracy. This mutual uncertainty prevents the market from optimizing toward better security outcomes.

WORKFLOW OPTIMIZATION OVERRIDES TECHNICAL METRICS

A final question aims to explain market behavior in the face of uncertainty: why do threat intelligence consumers prioritize curated, lower-false-positive feeds despite academic literature emphasizing technical metrics like coverage and timeliness, and how does this demand shape the offered intelligence products? To answer this, we must understand that market failures, such as the information insufficiency in the market for threat intelligence, inevitably lead market participants to develop institutions: formal and informal rules to cope with uncertainty [22, 23]. In the absence of reliable quality metrics, defenders cannot optimize their security posture in the traditional economic sense. Instead, as Herbert Simon posited with the theory of *bounded rationality*, decision-makers faced with cognitive limits and uncertainty 'satisfice', meaning to seek a solution that is good enough to meet their operational constraints [183]. As discussed in the Introduction, the Computer Science approach to threat intelligence has been to quantify various quality dimensions using metrics such as false-positive rate, coverage, and timeliness. This dissertation finds that in practice, local workflow optimization by threat intelligence consumers trumps such technical metrics.

As Chapter 2 showed, security professionals were acutely aware of the scarce resource of analyst attention. Unable to measure the true coverage of their feeds, they instead optimized for the local variable they controlled: false positives. They rationally traded the unknowable risk of missed threats for the tangible benefit of a streamlined workflow, favoring 'curated' feeds over comprehensive ones. Chapter 3 demonstrated how the community-driven sharing effort, in their aim to produce novel intelligence, fell back on the heuristic of a '10-detection threshold' on VirusTotal. This rule guaranteed a low false-positive rate, again satisficing for a workflow requirement while filtering out the novel threats the community intended to find.

Similarly, the 30-day lag identified in Chapter 4 can be understood as a market response to this demand. Vendors prioritized the verification required to prevent false positives over the speed required for real-time detection. The lag between peak attacker activity and IOC publication was not a bug; it was a feature of a market that demanded curation, validation, and 'polished intel' over raw, immediate detection. Ultimately, the consumer's need to optimize workflows creates an incentive for vendors to delay intelligence to ensure a low false-positive rate. In this market, real-time security outcomes are traded away for operational efficiency.

5.3. IMPLICATIONS FOR ACTORS IN THE MARKET FOR THREAT INTELLIGENCE

The findings of this dissertation contribute to more transparency in the market for threat intelligence. They help organizations that procure threat intelligence from commercial providers understand its benefits and limitations, which could translate into more substantive negotiations with commercial providers. Furthermore, the dissertation suggests

that non-commercial sharing communities can offer a substitute for at least some of the offerings of major providers.

The following sections review the structure of the market for threat intelligence, establishing the economic actors involved and the incentive structures that drive their behaviors. Implications and recommendations follow per player in this market.

Before discussing the defensive actors, it is important to note the role of threat actors – after all, the goal of threat intelligence is to provide defenders with intelligence about these actors. Much as extreme weather creates demand for insurance without participating in it, threat actors create demand in this market. And so, although they are the *raison d'être* for its existence, they are not economic actors in the market for the material being produced about them.

Threat actors do influence the outcome of the process. Their behaviors incur economic costs on victims, creating economic opportunities for defenders, including threat intelligence researchers. Furthermore, threat actors employ detection evasion techniques, increasing the 'environmental difficulty' for defenders, and influencing the quality of the resulting products and services.

ORGANIZATIONS CONSUMING THREAT INTELLIGENCE

Consumers of threat intelligence operate under conditions of information asymmetry as boundedly rational agents. As findings in Chapter 2 show, consumers struggle to assess the quality of the products they purchase. In the absence of this information, they cannot act as rational utility maximizers who select the optimal mix of feeds to detect the most threats. Instead, consumers 'satisfice': they optimize for the variables they can observe and control, which in this market is the efficiency of their internal workflows. This leads to a strong preference for 'curated' intelligence with low false-positive rates, even if this comes at the cost of potential false negatives (missed detections). Thus, the consumer's economic behavior in the CTI market is driven not by the maximization of security outcomes, but by the minimization of operational friction. However, under strict resource constraints, this satisficing behavior can be rational. Optimizing for workflow efficiency and accepting a baseline of false negatives allows organizations to reallocate limited security budgets or personnel to other highly effective preventative functions, such as combating phishing through employee awareness programs or improving the secure implementation of IT infrastructure.

The findings of this dissertation show that threat intelligence can, without a doubt, contribute unique data and contextual insights (e.g., actor attribution and TTPs) as compared to traditional sources of information about cybersecurity threats, like blocklists and internet abuse aggregators, which often lack such context (Chapters 2 and 3). Organizations should, however, be clear-eyed about their needs and the extent to which threat intelligence can fulfill them. The research in this dissertation has shown the following caveats that organizations should take into account when planning a threat intelligence program:

1. Organizations should acquire a mix of threat intelligence sources. Not even a high-end commercial provider can enable adequate coverage by itself (Chapter 2). This recommendation, however, primarily applies to organizations wishing to detect cybersecurity incidents (minimize false negatives) and who possess the analytical

resources to integrate disparate feeds. For organizations motivated by compliance or internal signaling (*the IBM effect*, as found in Chapter 2), or those with limited analytical bandwidth, it could be rational to acquire a single source of commercial threat intelligence. Alternatively, they could forego attempting to build an internal capability entirely by outsourcing to a Managed Security Service Provider (MSSP), or relying on security products with implicit threat intelligence integration.

2. Sharing communities can provide unique data and should be part of an organization's threat intelligence portfolio (Chapter 3). While sharing communities do attempt to generate blocklists suitable for a wider audience, active participation and contribution require dedicated personnel capable of vetting unstructured input. For organizations facing standard operational threats or lacking the personnel to filter noisy input, participating in such communities may impose an unjustifiable operational burden. They may realize better outcomes by relying entirely on automated commercial pipelines optimized for low false-positive rates.

3. Threat intelligence may be more useful for retroactive threat hunting rather than real-time intrusion detection. Network logs should be retained for at least 30 days to compensate for delays in IOC propagation (Chapter 4). This suits security operations centers with established threat hunting practices and the infrastructure to retain extensive network telemetry. For organizations constrained by data retention budgets or privacy regulations, extending log retention strictly for delayed IOC matching could yield a negative return on investment. These organizations could instead prioritize preventative controls and zero-trust architectures that reduce reliance on reactive detection.

4. Organizations must recognize that the 'workflow optimization' problem is not just a market failure, but a failure of internal corporate hierarchies. When Chief Information Security Officers (CISOs) establish bureaucratic performance metrics (e.g., tickets closed, low false-positive rates) without aligning them to actual security outcomes, they inadvertently instruct their Security Operations Centers to prefer clean, delayed intelligence over messy, real-time intelligence. Organizations must align their internal reporting hierarchies and incentive structures to tolerate the operational friction necessary to detect novel threats, rather than just demanding efficiency.

This dissertation has been critical of performance metrics. Nevertheless, metrics can be useful for setting up a CTI program. However, they should focus on the program's overall objective – which is likely centered on detecting threats, i.e., minimizing false negatives – rather than depart from the organization's current workflow, which has led practitioners to minimize false positives instead. More practically, organizations could formulate intelligence requirements and track how threat intelligence sources help them meet them. How this works is beyond the scope of this work, but an overall recommendation is for organizations to exchange such best practices with their peers, for example, by attending conferences[1] or joining sectoral or regional Information Sharing and Analysis Centers (ISACs).

[1] An interesting event for practitioners to consider visiting is the yearly Cyber Threat Intelligence conference

As in most of this dissertation, these recommendations have focused on the major use case of threat intelligence: informing intrusion detection. However, as Chapter 2 found, organizations also acquire threat intelligence for more strategic or business-oriented needs, such as risk management and even mergers and acquisitions. For such uses, the ability of threat intelligence to inform outcomes is even harder to measure. Nevertheless, CTI program managers could introduce a simple metric to gauge the program's effectiveness by asking internal consumers of threat intelligence: 'Has this information contributed to your ability to make business decisions?' This question is aligned with the definition of intelligence as information for decision advantage [184].

COMMERCIAL PROVIDERS

Commercial threat intelligence vendors behave as profit-maximizing firms in a market for 'silver bullets' – where consumers purchase the promise of protection rather than easily measurable detection capabilities. In this opaque market, commercial firms are characterized by extreme product differentiation, as demonstrated by the findings in Chapter 2, which found that even market-leading vendors showed almost no overlap in their data. Rather than a deliberate choice, this differentiation is likely a structural necessity driven by the firms' disparate proprietary sensor networks and unique customer bases. The result is that they create exclusive datasets that cannot be easily compared or substituted.

One leading commercial provider's 30-day lag in indicator publication, identified in Chapter 4, could reflect the time required to curate and verify data to the high standards demanded by customers. Perhaps this shows that vendors respond rationally to the consumer's preference for workflow optimization. In this sense, the lack of timeliness is not a failure of production, but a successful alignment with the market's economic incentives. However, we also found that another provider published indicators almost immediately, demonstrating that the market sustains diverse, competing operational models.

While current market incentives strongly favor this opacity and differentiation, vendors who wish to establish objective value and improve overall security outcomes are well-positioned to improve market transparency. They can do this by being more open about their production processes and how their products influence security outcomes for their customers. However, based on evident commercial successes, it would seem that it is in the economic best interests of threat intelligence vendors to maintain their current operational opacity and 'cyber noir' marketing campaigns. Nevertheless, some firms have already demonstrated a willingness to collaborate and improve transparency. Two examples stand out.

The first example is from 2025, when CrowdStrike and Microsoft announced that they would work to harmonize threat actor names. The growing number of players in the threat intelligence space has led to a sharp rise in threat naming schemes, varying from numbered APTs to mythical animals to meteorological phenomena: one firm might refer to a cluster of malicious activities as *Fancy Bear* and another firm as *Forest Blizzard* [185]. This has led to a confusing situation that perhaps can be best compared to many

hosted by the Forum of Incident Response and Security Teams (FIRST), see: https://www.first.org/conference/

biologists discovering a new type of bird and all of them insisting on having the right to name it after themselves. Of course, there also exist situations where uncertainty around the overlap of activity clusters means that analytical rigor requires different names to be maintained. But as the collaboration between CrowdStrike and Microsoft demonstrates, firms can agree to acknowledge cases where they are likely tracking the same activity[2]. It is unclear if this will also translate to the firms' adopting each other's terminology.

The second example is the Cyber Threat Alliance (CTA), a formalized sharing arrangement between largely commercial providers [75] and not to be confused with the individuals sharing COVID-related threat data in Chapter 3. The arrangement was based on quid pro quo, in which the contributions of each member would be tracked and correlated. One way to read the emergence of the Cyber Threat Alliance is that vendors were collaborating for the mutual benefit of each other's customers. Another is that the arrangement allowed participating firms to get a better understanding of what threats their competitors were observing, thus reducing uncertainty about the quality of their own offerings. The collaboration also led to limited joint public reporting.

Going beyond these examples, the threat intelligence industry could demonstrate transparency to its customers and benefit the overall cybersecurity community by considering the following steps. Note, however, that these steps assume a maturing market where consumers will economically reward transparency.

1. Commercial providers could be more open about their production processes by communicating the trade-offs they make. Chapters 2 and 4 suggest that vendors make trade-offs on respectively coverage and timeliness of the data they deliver. Communicating clearly such decisions about the production process would allow competitors to differentiate and improve customers' ability to make informed decisions. However, given the bounded rationality of consumers, vendors might rationally conclude that maintaining operational opacity remains the optimal strategy. If consumers primarily optimize for frictionless workflow integration, providers might achieve better market positioning by investing in polished integrations rather than transparently communicating technical trade-offs.

2. Building on the harmonization of threat actor names, commercial providers could compare the IOCs underlying their respective activity clusters in their private holdings to each other, for example, through a trusted intermediary or by making use of privacy-preserving technologies such as Secure Multi-Party Computation (SMPC). Taking the example above, it would allow CrowdStrike and Microsoft to quantify the overlap of data that makes up their understanding of respectively the *Bear* and *Blizzard*. This relies on vendors having an economic incentive to build industry-wide confidence. Conversely, explicitly revealing overlap (or a lack thereof) could commoditize their offerings or expose coverage gaps. Thus, vendors might choose to avoid such broad comparisons, preferring instead to collaborate only on certain

[2]The result of the effort by CrowdStrike and Microsoft to harmonize threat actor names can be viewed here: https://learn.microsoft.com/en-us/unified-secops/microsoft-threat-actor-naming The studies in this dissertation were published before this list existed and had to rely on the work of independent researchers with the same objective [53].

high-profile public cases while deliberately protecting the opacity of their tactical data feeds to maintain their unique selling propositions.

SHARING COMMUNITIES

Sharing communities, such as the Cyber Threat Coalition analyzed in Chapter 3, are driven by a mix of altruism, professional reputation, and the desire to contribute to a public good. However, their economic behavior is heavily constrained by the collective action problem known as free-riding. We found that contributions followed a steep power-law distribution: a tiny minority of users (less than 1%) contributed the vast majority of indicators, while thousands of others simply consumed the data. While such communities can overcome market failures in specific niches – such as the rapid mobilization against COVID-19 threats – they struggle to sustain high-quality production at scale. Interestingly, many of the participants in the studied community were employed at the commercial firms in the same market, but COVID-19 created an opportunity for them to contribute to a greater cause. The fact that the community eventually faltered likely indicates that large institutionalized sharing communities such as the CTC are not serious alternatives for consumers in this market.

It is important to note that the sharing community studied in this dissertation was exceptionally large. In practice, many smaller communities exist, varying in their degree of formalization. Whether formed between friends, colleagues, or even individuals at competing firms, these communities represent a significant component of the threat intelligence ecosystem that operates largely outside of view.

For these communities to maximize their impact and overcome free-riding, participants should recognize that such communities have the potential to deliver unique threat intelligence compared to open sources and commercial providers. Based on the study about the 'natural experiment' with a large sharing community in Chapter 3, the most important recommendations are:

1. Sharing communities should focus on their comparative advantage: specific, novel, or local threats that slip through the cracks of commercial and automated systems. Chapter 3 found that existing defenses are resilient and fast at mitigating general abuse, and so sharing communities do not attempt to compete on coverage of broad threat categories. This applies to communities organized to generate operational intelligence. However, many communities are formed primarily for educational, networking, or trust-building purposes. Attempting to optimize such groups for unique technical data production could detract from their primary social function; these groups might better serve their members by focusing on sharing contextual cases and strategic trends.

2. To produce threat intelligence for mass consumption, communities should develop scalable vetting processes that still have humans in the loop. Fully automating the quality assurance process through existing scanners (like VirusTotal) risks reproducing known data and filtering out the specific, novel threats that the community has unearthed. This is essential for large-scale communities aiming to produce operational data feeds. In smaller, highly trusted peer-to-peer communities, however, formal vetting processes might introduce unnecessary friction. In these settings,

raw, unvetted intelligence shared quickly among peers may be more valuable than delayed, verified intelligence.

FOR PUBLIC POLICY

Based on our analysis, finding information asymmetries and other imperfections, policymakers might be tempted to determine that a market failure occurs and introduce sweeping policies to address it. However, recall that the market for threat intelligence has only developed since the late 2010s and is still highly dynamic. Recall also that Chapter 4 pointed to signs that commercial providers might be contributing to a public good by investigating threat actors on behalf of their customers, increasing operational security on threat actors and thus incurring costs on them that are non-excludable and thus to the wider public's benefit. Therefore, we suggest policymakers should limit themselves to interventions that mostly address their own role in the market for threat intelligence:

1. Policy initiatives for information sharing have become the go-to response to cyber incidents, to the point where experts have jokingly called them the 'thoughts and prayers' of cybersecurity [186]. Although fostering cooperation is most likely a good idea, a recurring theme in this dissertation has been the limitations of technical metrics for evaluating the quality of threat intelligence. When it comes to threat intelligence sharing initiatives, policymakers should be aware that more is not always better. Rather than sharing arrangements geared toward specific outputs, policy initiatives could focus on reducing barriers to peer-to-peer sharing, as organizations are best positioned to determine which information has value for their needs – such as exchanging contextual cases, trends, and "war stories" that perfectly complement commercial feeds. This is the preferred approach for organically fostering high-trust collaborations via *network* coordination. However, recognizing the limitations of both opaque markets and free-riding networks, policymakers must also consider when *hierarchical* coordination is necessary. In extreme scenarios where systemic vulnerabilities require rapid, comprehensive visibility (e.g., protecting critical infrastructure during a national crisis, or establishing minimum baselines for cyber incident reporting), relying on voluntary sharing networks may be insufficient. In these cases, policymakers are justified in utilizing their hierarchical authority to mandate centralized reporting and enforce transparency.

2. Policymakers should re-evaluate the 'impose cost' strategy. Sharing arrangements and publications that are based on distributing technical indicators of compromise (IOCs) alone are unlikely to effectively impose costs on threat actors, as we found that threat actors have already internalized the discovery of their infrastructure to a degree. Drawing on the 'pyramid of pain' model [72], publicizing both technical identifiers and the modus operandi of threat actors – including with detection rules – is more likely to fulfill decisionmakers' aim to give adversaries a 'bad day' [187]. This applies when the objective is to cause sustained disruption to sophisticated adversaries. If the objective is instead to quickly sanitize domestic networks from a widespread, automated, and unsophisticated campaign, distributing technical IOCs at scale remains highly effective and pragmatic.

5.4. FUTURE WORK

The studies in this dissertation point to several directions for future research, many of which hinge on a better collaboration between the academic and commercial spheres of threat intelligence.

First, we need to move beyond descriptive metrics and develop generalizable models of value in threat intelligence. Chapter 2 showed that current metrics often fail to capture what users actually value – such as context, curation, and workflow optimization. Future comparisons of threat intelligence sources should therefore look beyond simple overlap or count metrics and attempt to measure organizational and security outcomes, for example, based on quantifying decision advantage ([184]). Understanding the drivers of value is not only important for commercial selection but can also inform the design of more effective threat-sharing platforms (Chapter 3).

Second, there is a need to move from observational studies to causal inference. Chapter 4 highlighted the limitations of using sightings as a proxy for adversary behavior. While we identified correlations, the gold standard for understanding the impact of interventions – such as the publication of IOCs – remains randomized controlled trials (RCTs). For instance, an experiment where a subset of IOCs is withheld from publication could definitively isolate the 'impose cost' effect from regular operational security measures.

This leads to our final recommendation: commercial providers are uniquely positioned to drive such research. They possess what academic researchers largely lack, namely, comprehensive telemetry from diverse customer networks and the ability to conduct large-scale experiments. Currently, however, this potential remains untapped due to commercial incentives for opacity. We therefore urge commercial providers to collaborate with academics by sharing their data for independent analysis or by performing in-house research that is published for the benefit of the academic community.

ACKNOWLEDGMENTS

I always tell people who consider to pursue a PhD: make sure your supervisor is someone you look up to and whose way of thinking you admire. For me, this was definitely true, and I am grateful to have watched my supervisor, Michel, at work up close. Like a blacksmith at the anvil, he crafts data into a storyline that makes research not only valuable but enticing. It means a great deal to me to have been his apprentice. Thank you, Michel.

Nevertheless, I still remember my shock as a first-year PhD candidate upon leaving a meeting with him and other colleagues, during which they had torn apart the draft of my first article, now a chapter in this dissertation. After completing the painful process of putting research results to paper for the first time, I had expected a laudation. Instead, I was met with sharp critiques. I left the room thinking: "This isn't what I signed up for!"

But actually, that was exactly what I signed up for. Only later did I realize that my colleagues had tried to shoot holes in the fruit of my labor because they knew it would improve the research and its chances at a top conference. By channeling the rigorous, sometimes downright impossible demands of academic peer review, they were helping me raise my own standards. Later, I would learn to do exactly the same for others.

This anecdote is just one example of the learning path that this dissertation represents. A path that did not always follow a straight line and took a few extra detours. I could not have completed this journey without the support, guidance, and encouragement of many wonderful people. First, I would like to acknowledge my peers in the Economics of Cybersecurity group. The seminars, joint projects, and stimulating conversations we shared have shaped my academic work, and the yearly dinner and Sinterklaas event at Michel's were a reminder that we were all in this together. To the colleagues from the Organization and Governance section at the TPM faculty: I highly appreciate our many lunches and discussions. You have demonstrated what it means to be an academic.

I am immensely grateful to my co-promotor, Bram, and the other co-authors on the papers in this dissertation. I have truly enjoyed diving into the rabbit holes of our data together and collaboratively making sense of it all; your insights and shared dedication were invaluable, sometimes right up until the late-night hours before a deadline. One co-author deserves extra appreciation. Aksel was always a joy to work with and has over the years become a true friend. Without his participation and solidarity, I might have thrown in the academic towel long ago.

A special word of thanks goes out to my family and friends for their support and, above all, for their perseverance and loyalty in asking about the progress of this dissertation year after year. And finally, my deepest love and appreciation go to my wife – for her patience with my occasional absent-mindedness and for bringing the optimism and joy into my life that motivated me to take this dissertation across the finish line.

Xander Bouwman
Delft, October 2026

AUTHORSHIP CONTRIBUTIONS

The author is indebted to his co-authors for their collaboration on the scientific articles that served as the basis for this dissertation. The three tables below show their contributions to the articles on which Chapter 2, Chapter 3, and Chapter 4 were respectively based. Contributions of the co-authors are specified using the CRediT role taxonomy[3].

Table 1: Contributors to Chapter 2, using the CRediT role taxonomy

Name	**Roles**
Xander Bouwman	Conceptualization, Data curation, Investigation, Methodology, Project administration, Software, Visualization, Writing: original draft, Writing: review & editing
Michel van Eeten	Conceptualization, Funding acquisition, Methodology, Project administration, Validation, Writing: review & editing
Harm Griffioen	Conceptualization, Data curation, Investigation, Methodology, Software, Writing: original draft, Writing: review & editing
Bram Klievink	Funding acquisition, Supervision, Writing: review & editing
Christian Doerr	Funding acquisition, Visualization, Writing: original draft, Writing: review & editing
Jelle Egbers	Data curation, Investigation, Methodology, Software, Writing: review & editing

[3]The CRediT role taxonomy: https://credit.niso.org/

Table 2: Contributors to Chapter 3, using the CRediT role taxonomy

Name	Roles
Xander Bouwman	Conceptualization, Data curation, Investigation, Methodology, Project administration, Software, Visualization, Writing: original draft, Writing: review & editing
Michel van Eeten	Conceptualization, Funding acquisition, Methodology, Project administration, Validation, Writing: review & editing
Carlos H. Gañán	Conceptualization, Data curation, Funding acquisition, Investigation, Methodology, Software, Visualization, Writing: original draft, Writing: review & editing
Bram Klievink	Funding acquisition, Supervision, Writing: review & editing
Pawel Foremski	Data curation, Investigation, Methodology, Software, Writing: original draft, Writing: review & editing
Victor Le Pochat	Conceptualization, Data curation, Investigation, Methodology, Software, Writing: original draft, Writing: review & editing
Giovane C. M. Moura	Data curation, Investigation, Methodology, Software, Writing: original draft, Writing: review & editing
Samaneh Tajalizadehkhoob	Conceptualization, Data curation, Investigation, Methodology, Software, Writing: original draft, Writing: review & editing
Tom Van Goethem	Conceptualization, Data curation, Investigation, Methodology, Software, Writing: original draft, Writing: review & editing
Wouter Joosen	Funding acquisition, Writing: review & editing

Table 3: Contributors to Chapter 4, using the CRediT role taxonomy

Name	Roles
Xander Bouwman	Conceptualization, Data curation, Investigation, Methodology, Project administration, Software, Visualization, Writing: original draft, Writing: review & editing
Michel van Eeten	Conceptualization, Funding acquisition, Validation, Writing: review & editing
Carlos H. Gañán	Conceptualization, Data curation, Investigation, Methodology, Software, Visualization, Writing: original draft, Writing: review & editing
Aksel Ethembabaoglu	Conceptualization, Data curation, Investigation, Methodology, Software, Writing: review & editing

ACRONYMS

API Application Programming Interface

APT Advanced Persistent Threat

CERT Computer Emergency Response Team

CISO Chief Information Security Officer

CSIRT Computer Security Incident Response Team

CTA Cyber Threat Alliance

CTI Cyber Threat Intelligence

DNS Domain Name System

DPRK Democratic People's Republic of Korea (North Korea)

EPP Extensible Provisioning Protocol

FIRST Forum of Incident Response and Security Teams

FQDN Fully Qualified Domain Name

FTC U.S. Federal Trade Commission

FTE Full-Time Equivalent

GRU Main Intelligence Directorate (Russian military intelligence)

GSB Google Safe Browsing

HTTP Hypertext Transfer Protocol

IDS Intrusion Detection System

IOC Indicator of Compromise

IP Internet Protocol

IRB Institutional Review Board

ISAC Information Sharing and Analysis Center

ISP Internet Service Provider

MSP Managed Service Provider

MSSP Managed Security Service Provider

NGO Non-Governmental Organization

OLS Ordinary Least Squares

OTI Open source Threat Intelligence

PII Personally Identifiable Information

PTI Paid Threat Intelligence

R&D Research and Development

RCT Randomized Controlled Trial

RFI Request for Information

SIEM Security Information and Event Management

SLD Second-Level Domain

SMPC Secure Multi-Party Computation

SOC Security Operations Center

SSH Secure Shell

STI Shared Threat Intelligence

STIX Structured Threat Information eXpression

TAXII Trusted Automated eXchange of Intelligence Information

TI Threat Intelligence

TLD Top-Level Domain

TLS Transport Layer Security

TLP Traffic Light Protocol

TTP Tactics, Techniques, and Procedures

URL Uniform Resource Locator

WHO World Health Organization

ccTLD Country-Code Top-Level Domain

gTLD Generic Top-Level Domain

A

APPENDICES CHAPTER 2

A.1. INTERVIEW PROTOCOL

We conducted our interviews along the following questions:

- What does 'cybersecurity threat intelligence' mean to you?
- Which threat intelligence sources does your organization pay for?
- Which non-commercial sources are important for your organization? Are these as important as the paid sources?
- Did you ever discontinue a source?
- For each paid TI source, discuss:
  - What are the costs of this source to your organization?
  - What does this source consist of? How do you receive it?
  - How is this source used in your organization? Which systems or processes rely on it?
  - Which results of the use of this source are most valuable? How often does this occur? What are strengths and weaknesses of this source?

A.2. CODEBOOK

The codebook is composed of four tables. Tables 2.3 & 2.4 in Section 2.5, and the two tables below that reflect Section 2.7.

USE CASES OF TI	Respondents *n=14*
Network detection	93%
Situational awareness	64%
SOC prioritization	50%
Informing business decisions	36%
Enrichment of more intel	36%
Improving end user awareness	29%
Threat hunting	29%
Informing security engineering	21%
Reducing financial fraud	14%

Table A.1: Network detection is the most prominent use case of TI, though more than half of respondents also describe situational awareness and SOC prioritization.

TI VALUE PERCEPTIONS	Respondents *n=14*
Actionability	
Providing context	100%
Timeliness	50%
Comprehensiveness	50%
Suitable abstraction level	36%
Interpretability	21%
Visualized well	14%
Relevance	
Sectoral focus	64%
Geographic focus	50%
Coverage of relevant threats	50%
Ability to correct bias	14%
Confidence	
Automatability	79%
Confidence in vendor	71%
Original contribution	50%
Accuracy	43%
Selectiveness	29%
Affordability	21%

Table A.2: We identify actionability, relevance, and confidence as three aspects of value of a TI source. All respondents share the view that TI should provide context.

B

APPENDICES CHAPTER 3

B.1. COVID-19 KEYWORD LIST

We selected COVID-19-related keywords starting from a set of English keywords, which we then translated into 14 major languages: Arabic, Bengali, Mandarin Chinese, Dutch, French, German, Hindi, Italian, Japanese, Malay, Portuguese, Russian, Spanish and Turkish. Finally, we generated lookalike terms using techniques from typosquatting [188] and homoglyphs [189].

B

self-isolation	eovid	covic1	cocvid
corona	co-vid	covld	st1mulus
pandemic	fovid	pandemid	pandernic
mask	eovib	timulus	pandemiic
ncov	vovid	pondemic	pawndemic
vaccine	pademic	pandemec	panndemic
virus	covix	cov8d	pandemmic
hydroxychloroquine	landemic	vccine	panderic
quinacrine	covoid	cov1d	vacvine
chloroquine	ocvid	tsimulus	chloraquine
remdesivir	pandmic	stimlus	cokvid
plaquenil	xovid	vaccione	eovicl
azithromycin	stimulis	covjd	coviud
metformin	cuvid	copvid	vadcine
favipiravir	stimulas	vaccien	pandemci
interferon	stimuls	vaccne	pandem1c
lopinavir	covdi	cpovid	stiumulus
ritonavir	stemulus	covbid	stimullus
arbitol	stimulu	pandimic	stimul8s
stimulus	clovid	vacclne	stumulus
infection	cov-id	pundemic	covid1
n95	stimulos	stimul-us	cov9d
respirator	coivd	c-ovid	covuid
testkit	cpvid	pandeic	covjid
distance	covicl	pandemif	cov8id
quarantine	c0vid	vaecinc	cobvid
lockdown	baccine	vaccone	cov9id
covis	vaccirc	chloroquini	cofvid
covic	clvid	caccine	coivid
cvid	ciovid	stimulous	eovld
coved	covvid	cloroquine	c0vld
covir	vaccime	stimilus	covcid
dovid	covkid	pansemic	pandermic
cevid	coviid	vaxcine	cov1b
covib	coovid	colvid	covid
cavid	chloroquin	panemic	sars-cov
covd	pandemoc	pandernie	sarscov
coviel	mandemic	stimulys	ivomec
cobid	covud	stikulus	ivermectin
civid	covi-d	covibl	mectizan
covici	stimulux	pandemc	iver-dt
cvoid	vacccine	cogvid	ivexterm
accine	pandemix	pandepic	scabo-6
vacine	covilb	stimuus	sklice
andemic	stimuluz	voccine	stromectol
covod	pamdemic	stipulus	soolantra
pandemi	vaccie	pandomic	mk-933
ffp2	ffp3	c19	

Table B.1: List of English keywords and homoglyphs

C

APPENDICES CHAPTER 4

Analysis in the body of Chapter 4 was instrumented using a window of sightings around $d = 50$ days around IOC publication date, as explained in the method section (see Section 4.3.4). The main regression table is printed on page 114.

As part of our sensitivity analysis, we also fitted regression models for other values of d. Regression tables for $d = 10$, $d = 100$, and $d = 150$ are printed on pages 116, 118, and 120, respectively.

We furthermore ran our analysis for another unit of analysis (using the standard window of $d = 50$), and report the results on page 122. 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 C.6 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, $d = 50$).

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

In sum, the following tables are included:

- Covariate descriptive statistics (Table C.1)
- Main regression table (Table C.2)
- Sensitivity analysis tables (Table C.3, Table C.4, Table C.5, and Table C.6)

Table C.1: Descriptive statistics of dependent variable and covariates in Section 4.5.

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	

Table C.1 – *Continued from previous page*

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.

Table C.2: Regression models for IOCs that led to sightings in a window of $d = 50$ days around publication date. In Section 4.5.1, 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 Section 4.5.2, 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)	(0.006)
Day was a weekday (ref. Weekend)		0.002***	0.002***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)
Vendor 2 (ref. Vendor 1)			-0.004*	0.003
	(0.002)	(0.003)	(0.003)	(0.003)
IOC type: IP address (ref. Domain)			0.087***	0.087***
	(0.001)	(0.001)	(0.001)	(0.001)
IOC type: MD5 hash (ref. Domain)			-0.022***	-0.021***
	(0.002)	(0.002)	(0.002)	(0.002)
IOC type: SHA1 hash (ref. Domain)			-0.021***	-0.022***
	(0.003)	(0.003)	(0.003)	(0.003)
IOC type: Other, e.g. regex (ref. Domain)			0.004	0.013***
	(0.004)	(0.004)	(0.004)	(0.004)
Domain IOC cohort: None (ref. Established domains)			-0.007	-0.008
	(0.006)	(0.006)	(0.006)	(0.006)
Domain IOC cohort: New domains (ref. Established domains)			0.029***	0.027***
	(0.007)	(0.007)	(0.007)	(0.007)
Actor geography: Russia (ref. China)				0.133***
	(0.009)	(0.009)	(0.009)	(0.009)

Table C.2 – *Continued from previous page*

	Dependent variable: IOC led to sightings (binary, daily)			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Actor geography: DPRK (ref. China)				-0.078**
	(0.032)	(0.032)	(0.032)	(0.032)
Actor motivation: Financial (ref. Espionage)				-0.018***
	(0.002)	(0.002)	(0.002)	(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.

Table C.3: 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**

Table C.3 – *Continued from previous page*

	Dependent variable: IOC led to sightings (binary, daily)			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
				(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 C.4: 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	-	-	-

Table C.4 – *Continued from previous page*

	Dependent variable: IOC led to sightings (binary, daily)			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
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 C.5: 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	-	-	-

Table C.5 – *Continued from previous page*

	Dependent variable: IOC led to sightings (binary, daily)			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
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 C.6: 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)

Table C.6 – *Continued from previous page*

	Dependent variable: Normalized count of daily sighting			
	FE	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
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.

BIBLIOGRAPHY

[1] C. McGrath. "Robert Caro's Big Dig". In: *New York Times* (May 2012). ISSN: 0362-4331. URL: https://www.nytimes.com/2012/04/15/magazine/robert-caros-big-dig.html.

[2] Ellen Nakashima. *'Guccifer 2.0' claims credit for DNC hack*. June 2016. URL: https://web.archive.org/web/20180417013555/https://www.washingtonpost.com/world/national-security/guccifer-20-claims-credit-for-dnc-hack/2016/06/15/abdcdf48-3366-11e6-8ff7-7b6c1998b7a0_story.html.

[3] Reuters. *'Lone hacker' claims responsibility for cyber attack on Democrats*. June 2016. URL: https://www.reuters.com/article/world/lone-hacker-claims-responsibility-for-cyber-attack-on-democrats-idUSKCN0Z209W.

[4] *CrowdStrike's work with the Democratic National Committee: Setting the record straight*. 2020. URL: https://www.crowdstrike.com/en-us/blog/bears-midst-intrusion-democratic-national-committee.

[5] U.S. Director of National Intelligence. *Background to "Assessing Russian Activities and Intentions in Recent US Elections": The Analytic Process and Cyber Incident Attribution*. Jan. 2017. URL: https://web.archive.org/web/20170106210244/https://www.dni.gov/files/documents/ICA_2017_01.pdf.

[6] F. J. Egloff and M. Dunn Cavelty. "Attribution and Knowledge Creation Assemblages in Cybersecurity Politics". In: *Journal of Cybersecurity* 7.1 (Jan. 2021), tyab002. ISSN: 2057-2085. DOI: 10.1093/cybsec/tyab002.

[7] L. Maschmeyer, R. J. Deibert, and J. R. Lindsay. "A tale of two cybers - how threat reporting by cybersecurity firms systematically underrepresents threats to civil society". In: *Journal of Information Technology & Politics* (Jan. 2021). ISSN: 1933-1681. DOI: 10.1080/19331681.2020.1776658.

[8] Tillson Galloway et al. "Actively Understanding the Dynamics and Risks of the Threat Intelligence Ecosystem". In: *33nd Annual Network and Distributed System Security Symposium, NDSS*. Feb. 2026. DOI: https://dx.doi.org/10.14722/ndss.2026.240102.

[9] T. Steffens. *Attribution of Advanced Persistent Threats*. Berlin, Germany: Springer. ISBN: 978-3-662-61313-9. URL: https://link.springer.com/book/10.1007/978-3-662-61313-9.

[10] NCSC-UK. *Building a Security Operations Centre (SOC), Threat Intelligence*. May 2022. URL: https://www.ncsc.gov.uk/collection/building-a-security-operations-centre/threat-intelligence.

[11] Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin. "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains". In: (2011). URL: https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf.

[12] Sergio Caltagirone, Andrew Pendergast, and Christopher Betz. "The Diamond Model of Intrusion Analysis". In: (2013). URL: https://apps.dtic.mil/sti/tr/pdf/ADA586960.pdf.

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

[14] A. Zibak, C. Sauerwein, and A. C. Simpson. "Threat Intelligence Quality Dimensions for Research and Practice". In: *Digital Threats: Research and Practice* 3.4 (Dec. 2022), pp. 1–22. DOI: 10.1145/3484202.

[15] Vector Guo Li et al. "Reading the Tea leaves: A Comparative Analysis of Threat Intelligence". In: *28th USENIX Security Symposium*. USENIX Security '19. 2019, pp. 851–867. URL: https://www.usenix.org/conference/usenixsecurity19/presentation/li.

[16] Harm Griffioen, Tim M. Booij, and Christian Doerr. "Quality Evaluation of Cyber Threat Intelligence Feeds". In: *18th International Conference on Applied Cryptography and Network Security*. ACNS '20. 2020, pp. 277–296. DOI: 10.1007/978-3-030-57878-7_14.

[17] Leigh Metcalf and Jonathan M. Spring. "Blacklist Ecosystem Analysis". In: *2nd ACM Workshop on Information Sharing and Collaborative Security*. WISCS '15. 2015, pp. 13–22. DOI: 10.1145/2808128.2808129.

[18] Xander Bouwman et al. "A different cup of TI? The added value of commercial threat intelligence". In: *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 433–450. ISBN: 978-1-939133-17-5. URL: https://www.usenix.org/conference/usenixsecurity20/presentation/bouwman.

[19] Daniel Fischer et al. "The Role of Threat Intelligence Sharing Platforms in Companies, Public Authorities, and Universities: Insights from an Exploratory Global Survey". In: *2024 IEEE International Conference on Big Data (BigData)*. 2024, pp. 2616–2623. DOI: 10.1109/BigData62323.2024.10825242.

[20] Sean Barnum. *Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX™)*. Feb. 2014. URL: https://stixproject.github.io/about/STIX_Whitepaper_v1.1.pdf.

[21] Bushra A. Alahmadi, Louise Axon, and Ivan Martinovic. "99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms". In: *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 2783–2800. ISBN: 978-1-939133-31-1. URL: https://www.usenix.org/conference/usenixsecurity22/presentation/alahmadi.

[22] George A. Akerlof. "The Market for "Lemons": Quality Uncertainty and the Market Mechanism". In: *Quarterly Journal of Economics* 84.3 (Aug. 1970).

[23] *The Market for Silver Bullets*. [Online; accessed 7. Jun. 2020]. Mar. 2008. URL: https://iang.org/papers/market_for_silver_bullets.html.

[24] R. Anderson. "Why information security is hard - an economic perspective". In: *Seventeenth Annual Computer Security Applications Conference*. IEEE, Dec. 2001. ISBN: 978-0-7695-1405. DOI: 10.1109/ACSAC.2001.991552.

[25] L. A. Gordon, M. P. Loeb, and W. Lucyshyn. "Sharing information on computer systems security: An economic analysis". In: *Journal of Accounting and Public Policy* 22.6 (Nov. 2003), pp. 461–485. ISSN: 0278-4254. DOI: 10.1016/j.jaccpubpol.2003.09.001.

[26] James M. Buchanan. "An Economic Theory of Clubs". In: *Economica* 32.125 (1965), pp. 1–14. ISSN: 00130427, 14680335. URL: http://www.jstor.org/stable/2552442 (visited on 12/19/2025).

[27] Jon R. Lindsay. *Age of Deception*. Cornell University Press, Oct. 2025. URL: https://www.jonrlindsay.com/age-of-deception.

[28] Thomas Rid and Ben Buchanan. "Attributing Cyber Attacks". In: *Journal of Strategic Studies* 38.January 2015 (2015), pp. 4–37. ISSN: 1743937X. DOI: 10.1080/01402390.2014.977382.

[29] J. Healey. "The implications of persistent (and permanent) engagement in cyberspace". In: *Journal of Cybersecurity* 5.1 (Jan. 2019), tyz008. ISSN: 2057-2085. DOI: 10.1093/cybsec/tyz008.

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

[31] M. Phythian. *Understanding the Intelligence Cycle*. Andover, England, UK: Taylor & Francis, July 2013. ISBN: 978-0-20355847-8. URL: https://www.taylorfrancis.com/books/edit/10.4324/9780203558478/understanding-intelligence-cycle-mark-phythian.

[32] Xander Bouwman et al. "Helping hands: Measuring the impact of a large threat intelligence sharing community". In: *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 1149–1165. ISBN: 978-1-939133-31-1. URL: https://www.usenix.org/conference/usenixsecurity22/presentation/bouwman.

[33] Xander Bouwman et al. "Can IOCs Impose Cost? The Effects of Publishing Threat Intelligence on Adversary Behavior". In: *32nd ACM Conference on Computer and Communications Security (CCS)*. ACM CCS 2025. United States: ACM, 2025. DOI: 10.1145/3719027.3765026. URL: https://www.sigsac.org/ccs/CCS2025/.

[34] Stephanie de Smale et al. "No One Drinks From the Firehose: How Organizations Filter and Prioritize Vulnerability Information". In: *2023 IEEE Symposium on Security and Privacy (SP)*. 2023, pp. 1980–1996. DOI: 10.1109/SP46215.2023.10179447.

[35] *Abuse.ch*. [Online; accessed 3. Feb. 2020]. Feb. 2020. URL: https://abuse.ch.

[36] *AlienVault Open Threat Exchange (OTX)*. [Online; accessed 3. Feb. 2020]. Feb. 2020. URL: https://cybersecurity.att.com/open-threat-exchange.

[37] *DNS-BH Malware Domain Blocklist*. [Online; accessed 3. Feb. 2020]. Feb. 2020. URL: https://www.malwaredomains.com.

[38] Ponemon Institute. *The Value of Threat Intelligence: The Second Annual Study of North American & United Kingdom Companies*. Tech. rep. February. Ponemon Institute, 2019.

[39] *Threat Intelligence Market Size, Share & Growth Forecast 2030*. [Online; accessed 13. Mar. 2026]. Aug. 2025. URL: https://www.marketsandmarkets.com/Market-Reports/threat-intelligence-security-market-150715995.html.

[40] Vector Guo Li et al. "Reading the Tea Leaves : A Comparative Analysis of Threat Intelligence". In: *Proceedings of the USENIX Security Symposium*. 2019.

[41] Harm Griffioen, Tim M. Booij, and Christian Doerr. "Quality Evaluation of Cyber Threat Intelligence Feeds". In: *ACNS*. 2020.

[42] Kurt Thomas et al. "The Abuse Sharing Economy: Understanding the Limits of Threat Exchanges". In: vol. 7462. 2016, pp. 143–164.

[43] Leigh Metcalf and Jonathan M. Spring. "Blacklist Ecosystem Analysis". In: *ACM Workshop on Information Sharing and Collaborative Security*. 2015.

[44] Marc Kührer, Christian Rossow, and Thorsten Holz. "Paint It Black: Evaluating the Effectiveness of Malware Blacklists". In: *RAID* (2014).

[45] Craig Lawson, Ruggero Contu, and Ryan Benson. *Market Guide for Security Threat Intelligence Products and Services*. Tech. rep. Febuary. Gartner Research, 2019. URL: https://www.gartner.com/en/documents/3902168/market-guide-for-security-threat-intelligence-products-a.

[46] Forrester Research. *External Threat Intelligence Services, Q3 2018*. Tech. rep. 2018. URL: https://www.fireeye.com/content/dam/fireeye-www/products/pdfs/pf/intel/rpt-forrester-threat-intel-services.pdf.

[47] Thomas Schaberreiter et al. "A Quantitative Evaluation of Trust in the Quality of Cyber Threat Intelligence Sources". In: *ARES*. 2019.

[48] Alex Pinto. "Determining the Fit and Impact of CTI Indicators on Your Monitoring Pipeline (#tiqtest2)". In: *FIRST Conference*. Kuala Lumpur, 2018. URL: https://www.first.org/conference/2018/program#pdetermining-the-fit-and-impact-of-cti-indicators-on-your-monitoring-pipeline-tiq-test-2-0.

[49] Roland Meier et al. "FeedRank: A Tamper- resistant Method for the Ranking of Cyber Threat Intelligence Feeds". In: *2018 10th International Conference on Cyber Conflict*. Tallinn: NATO CCD COE Publications, 2018, pp. 321–344.

[50] Paweł Pawliński and Andrew Kompanek. "Evaluating Threat Intelligence Feeds FIRST Technical Colloquium for Threat Intelligence". In: *FIRST Technical Colloquium for Threat Intelligence*. Munich, 2016. URL: https://www.first.org/resources/papers/2016#munich2016.

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

[52] Paweł Pawliński et al. *Actionable Information for Security Incident Response*. Tech. rep. January. ENISA, 2015, pp. 1–79. URL: https://www.enisa.europa.eu/publications/actionable-information-for-security.

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

[54] Daniel Votipka et al. "Hackers vs. Testers: A Comparison of Software Vulnerability Discovery Processes". In: *IEEE Symposium on Security and Privacy*. 2018.

[55] John W. Cresswell. *Qualitative Inquiry and Research Design: Choosing Among Five Approaches*. Second edi. 2007.

[56] Kathy Charmaz. *Constructing Grounded Theory: A Practical Guide Through Qualitative Analysis*. London: SagePublication Ltd, 2006.

[57] *UK businesses: average investment in cyber security 2019*. [Online; accessed 1. Jun. 2020]. July 2019. URL: https://www.statista.com/statistics/586587/iinvestment-in-cyber-security-by-businesses-in-the-uk-by-sector.

[58] Dean Hammer and Aaron Wildavsky. "The Open-Ended, Semistructured Interview: An (Almost) Operational Guide". In: *Craftways*. Routledge, Feb. 2018, pp. 57–101.

[59] Beth L Beech. "Asking questions: Techniques for semi-structured interviews". In: *Political Science* 35.4 (2009), pp. 665–668.

[60] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. "Reliability and Inter-rater Reliability in Qualitative Research". In: *ACM on Human-Computer Interaction* 3 (2019).

[61] *DLT NCPA Pricelists*. [Online; accessed 13. Feb. 2020] https://www.dlt.com/sites/default/files/contract-attachments/NCPA%20Pricelist%20-%20PDF_0.pdf. Feb. 2020.

[62] *Caharasoft GSA Pricelists*. [Online; accessed 13. Feb. 2020] https://web.archive.org/web/20200213230809/https://static.carahsoft.com/concrete/files/1414/3223/9520/CMASPricelist.pdf, https://web.archive.org/web/20200213231518/https://static.carahsoft.com/concrete/files/4114/5029/6739/Compiled_new_DIR_Pricelist_-_Revised_12-10-2015.pdf, . Feb. 2020.

[63] Joint Chiefs of Staff of the United States. *JP 2-01 Joint and National Intelligence Support to Military Operations*. Tech. rep. Washington, D.C.: Department of Defense, 2017. URL: http://www.jcs.mil/Doctrine/Joint-Doctrine-Pubs/2-0-Intelligence-Series/.

[64] Tiffany Bao and Gail-joon Ahn. "Matched and Mismatched SOCs: A Qualitative Study on Security Operations Center Issues". In: *ACM CCS*. 2019.

[65] Andreas Sfakianakis. *Let's make CTI great (again): a 5-year lookback in CTI*. [Online; accessed 14. Feb. 2020]. Nov. 2018. URL: https://www.enisa.europa.eu/events/2018-cti-eu-event/cti-eu-2018-presentations.

[66] James Shires. "Cyber-noir: Cybersecurity and popular culture". In: *Contemporary Security Policy* 0.0 (Sept. 2019), pp. 1–26.

[67] Clare Stevens. "Assembling cybersecurity: The politics and materiality of technical malware reports and the case of Stuxnet". In: *Contemporary Security Policy* 0.0 (Oct. 2019), pp. 1–24.

[68] Florian J Egloff. "Contested public attributions of cyber incidents and the role of academia". In: *Contemporary Security Policy* 0 (Oct. 2019), pp. 1–27.

[69] JD Work. "Evaluating Commercial Cyber Intelligence Activity". In: *International Journal of Intelligence and CounterIntelligence* (Jan. 2020), pp. 1–31.

[70] Alex Pinto and Kyle Maxwell. "Measuring the IQ of your Threat Intelligence Feeds". In: *Defcon 22*. 2015.

[71] Daniel Schlette et al. "Measuring and visualizing cyber threat intelligence quality". In: *International Journal of Information Security* (2020).

[72] David Bianco. *The Pyramid of Pain*. [Online; accessed 3. Feb. 2020]. 2014. URL: http://detect-respond.blogspot.nl/2013/03/the-pyramid-of-pain.html.

[73] Rebekah Brown and Robert M. Lee. *The Evolution of Cyber Threat Intelligence (CTI): 2019 SANS CTI Survey*. Tech. rep. SANS Institute, 2019.

[74] Tyler Moore and Richard Clayton. "The Consequence of Non-Cooperation in the Fight Against Phishing". In: *APWG eCrime Researchers Summit*. 2008.

[75] Cyber Threat Alliance. 2020. URL: https://www.cyberthreatalliance.org.

[76] *Anti-Phishing Working Group (APGW)*. URL: https://apwg.org/membership.

[77] COVID-19 Cyber Threat Coalition. *Our mission*. May 2020. URL: https://www.cyberthreatcoalition.org/about-us/our-mission.

[78] Microsoft. *Microsoft report shows increasing sophistication of cyber threats*. Sept. 2020. URL: https://blogs.microsoft.com/on-the-issues/2020/09/29/microsoft-digital-defense-report-cyber-threats.

[79] Jeremy Kirk. "New Platform Collects COVID-19 Threat Intelligence". In: *BankInfoSecurity* (Apr. 22, 2020). URL: https://www.bankinfosecurity.com/collecting-covid-19-threat-intelligence-a-14154.

[80] COVID-19 Cyber Threat Coalition. *CCTC Town Hall 008 June 11, 2020*. YouTube. June 11, 2020. URL: https://www.youtube.com/watch?v=0aBcwNqJWDA.

[81] Chris Larsen. *Sample Results From Processing a Large Feed of Shady Covid-Type Domains*. Symantec. Apr. 22, 2020. URL: https://symantec-enterprise-blogs.security.com/blogs/expert-perspectives/sample-results-processing-large-feed-shady-covid-type-domains.

[82] Tyler Moore and Richard Clayton. "Evaluating the Wisdom of Crowds in Assessing Phishing Websites". In: *12th International Conference on Financial Cryptography and Data Security*. 2008, pp. 16–30. DOI: 10.1007/978-3-540-85230-8_2.

[83] Kurt Thomas et al. “The Abuse Sharing Economy: Understanding the Limits of Threat Exchanges”. In: *19th International Symposium on Research in Attacks, Intrusions, and Defenses*. 2016, pp. 143–164. DOI: 10.1007/978-3-319-45719-2_7.

[84] Peng Peng et al. “Opening the Blackbox of VirusTotal: Analyzing Online Phishing Scan Engines”. In: *2019 Internet Measurement Conference*. IMC ’19. 2019, pp. 478–485. DOI: 10.1145/3355369.3355585.

[85] Lennart Maschmeyer. *Why the Covid-19 Crisis Offers an Opportunity for Lasting Change in Cybersecurity*. The Center for Security Studies at ETH Zurich, Apr. 28, 2020. URL: https://isnblog.ethz.ch/cyber/why-the-covid-19-crisis-offers-an-opportunity-for-lasting-change-in-cybersecurity.

[86] Arielle Waldman. “Volunteers join forces to tackle COVID-19 security threats”. In: *SearchSecurity* (May 11, 2020). URL: https://searchsecurity.techtarget.com/news/252482930/Volunteers-join-forces-to-tackle-COVID-19-security-threats.

[87] @ThreatCoalition on Twitter. Mar. 29, 2020. URL: https://twitter.com/ThreatCoalition/status/1244340555007614976.

[88] COVID-19 Cyber Threat Coalition. *CCTC Town Hall 007 May 28, 2020*. YouTube. May 28, 2020. URL: https://www.youtube.com/watch?v=cMeDAiG1O_4.

[89] @ThreatCoalition on Twitter. Apr. 8, 2020. URL: https://twitter.com/ThreatCoalition/status/1247958809856937985.

[90] COVID-19 Cyber Threat Coalition. YouTube channel. Apr. 19, 2020. URL: https://www.youtube.com/channel/UCHfhxcqhQADRV2h5hFgqAww.

[91] COVID-19 Cyber Threat Coalition. *CCTC Town Hall 001 April 16, 2020*. YouTube. Apr. 20, 2020. URL: https://www.youtube.com/watch?v=2euXCNUMhmw.

[92] @ThreatCoalition on Twitter. Mar. 28, 2020. URL: https://twitter.com/ThreatCoalition/status/1243711825927118853.

[93] COVID-19 Cyber Threat Coalition. *Slack workspace of the Cyber Threat Coalition*. Access upon request through https://www.cyberthreatcoalition.org/join-us. 2020.

[94] COVID-19 Cyber Threat Coalition. *CCTC Town Hall 003 April 30, 2020*. YouTube. Apr. 20, 2020. URL: https://www.youtube.com/watch?v=3EdzxE1tbzk.

[95] *Statistics about phishing activity and PhishTank usage*. PhishTank. 2021. URL: https://www.phishtank.com/stats.php.

[96] Josh Lerner and Jean Tirole. “Some Simple Economics of Open Source”. In: *The Journal of Industrial Economics* 50.2 (2002), pp. 197–234. ISSN: 1467-6451. DOI: 10.1111/1467-6451.00174. URL: http://dx.doi.org/10.1111/1467-6451.00174.

[97] *DomainTools COVID-19 Threat List*. Mar. 2020. URL: https://covid-19-threat-list.domaintools.com.

[98] SANS Internet Storm Center. *Covid 19 Domain Classifier*. 2020. URL: https://isc.sans.edu/covidclassifier.html.

[99] Giovane C. M. Muora, Thymen Wabeke, and Christian Hesselman. *Coronavirus and DNS: view from the .nl ccTLD*. Tech. rep. TR-2020-01. SIDN Labs, Mar. 26, 2020. URL: https://www.sidnlabs.nl/downloads/SzbDWlJkFgwEv9K8pwSCy/9443d3e686ef9a64f7674b0f05b9fd8d/SIDN-LABS-TR-2020-01.pdf.

[100] COVID-19 Cyber Threat Coalition. *2020-05-26 Weekly Threat Advisory*. May 26, 2020. URL: https://www.cyberthreatcoalition.org/advisories/2020-05-26-weekly-threat-advisory.

[101] COVID-19 Cyber Threat Coalition. *Join us*. 2020. URL: https://www.cyberthreatcoalition.org/join-us.

[102] Pawel Foremski, Oliver Gasser, and Giovane Moura. "DNS Observatory: The Big Picture of the DNS". In: *2019 Internet Measurement Conference*. 2019, pp. 87–100. DOI: 10.1145/3355369.3355566.

[103] Ryo Kawaoka et al. "A First Look at COVID-19 Domain Names: Origin and Implications". In: *2021 Passive and Active Measurement Conference*. PAM '21. 2021. DOI: 10.1007/978-3-030-72582-2_3.

[104] Sean McGrath. *COVID-19 Malicious Domain Report: Mapping malicious activity throughout the pandemic*. ProPrivacy. June 16, 2020. URL: https://proprivacy.com/privacy-news/covid-19-malicious-domain-report.

[105] Maciej Korczyński et al. "Reputation Metrics Design to Improve Intermediary Incentives for Security of TLDs". In: *2017 IEEE European Symposium on Security and Privacy*. EuroS&P '17. 2017, pp. 579–594. DOI: 10.1109/EuroSP.2017.15.

[106] Alice Hutchings, Richard Clayton, and Ross Anderson. "Taking down websites to prevent crime". In: *2016 APWG Symposium on Electronic Crime Research*. eCrime '16. 2016. DOI: 10.1109/ECRIME.2016.7487947.

[107] Eihal Alowaisheq et al. "Cracking the Wall of Confinement: Understanding and Analyzing Malicious Domain Take-downs". In: *26th Annual Network and Distributed System Security Symposium*. NDSS '19. 2019. DOI: 10.14722/ndss.2019.23243.

[108] Victor Le Pochat et al. "A Practical Approach for Taking Down Avalanche Botnets Under Real-World Constraints". In: *27th Annual Network and Distributed System Security Symposium*. NDSS '20. 2020. DOI: 10.14722/ndss.2020.24161.

[109] Scott Hollenbeck. *Extensible Provisioning Protocol (EPP)*. STD 69. IETF, Aug. 2009. URL: https://tools.ietf.org/rfc/std/std69.txt.

[110] S. Hollenbeck. *Domain Registry Grace Period Mapping for the Extensible Provisioning Protocol (EPP)*. RFC 3915. IETF, Sept. 2004. URL: http://tools.ietf.org/rfc/rfc3915.txt.

[111] Internet Corporation for Assigned Names and Numbers. *EPP Status Codes | What Do They Mean, and Why Should I Know?* June 16, 2014. URL: https://www.icann.org/resources/pages/epp-status-codes-2014-06-16-en.

[112] Maciej Korczyński et al. "Cybercrime After the Sunrise: A Statistical Analysis of DNS Abuse in New gTLDs". In: *13th ACM Asia Conference on Computer and Communications Security*. ASIACCS '18. 2018, pp. 609–623. DOI: 10.1145/3196494.3196548.

[113] Natasha Tusikov. “Access Chokepoints”. In: *Chokepoints: Global Private Regulation on the Internet*. 1st ed. University of California Press, 2017. Chap. 4, pp. 116–155. ISBN: 9780520291218.

[114] Kevin Murphy. “VeriSign demands website takedown powers”. In: *The Register* (Oct. 11, 2011). URL: https://www.theregister.com/2011/10/11/verisign_asks_for_web_takedown_powers/.

[115] Kim G. von Arx and Gregory R. Hagan. “SOVEREIGN DOMAINS: A Declaration of Independence of ccTLDs from Foreign Control”. In: *Richmond Journal of Law & Technology* 9.1 (Fall 2002), p. 4.

[116] Peter K. Yu. “The Never-Ending ccTLD Story”. In: *Addressing the World: National Identity and Internet Country Code Domains*. Ed. by Erica Schlesinger Wass. Rowman & Littlefield Publishers, Inc., 2003. Chap. 1, pp. 1–16. ISBN: 9780742528109.

[117] National Research Council. *Signposts in Cyberspace: The Domain Name System and Internet Navigation*. The National Academies Press, 2005. ISBN: 9780309096409. DOI: 10.17226/11258.

[118] EURid. *COVID-19 and extraordinary measures on .eu domain names*. Apr. 8, 2020. URL: https://eurid.eu/en/news/doteu-covid19-measures/.

[119] Eleanor Bradley. *Keeping a close watch on coronavirus domains*. Nominet. Mar. 26, 2020. URL: https://www.nominet.uk/keeping-a-close-watch-on-coronavirus-domains/.

[120] Frank Bajak. “Internet Firm Restricts Virus-Themed Website Registrations”. In: *ABC News* (Mar. 26, 2020). URL: https://abcnews.go.com/Technology/wireStory/internet-firm-restricts-virus-themed-website-registrations-69825166.

[121] Devdatta Akhawe and Adrienne Porter Felt. “Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness”. In: *22nd USENIX Security Symposium*. USENIX Security ’13. 2013, pp. 257–272. URL: https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/akhawe.

[122] Sergio Villegas et al. *sbserver differs from online/browser lookup?* GitHub repository `google/safebrowsing`. Issue 30. Aug. 31, 2016. URL: https://github.com/google/safebrowsing/issues/30.

[123] Sean Gallagher. “New “Quad9” DNS service blocks malicious domains for everyone”. In: *Ars Technica* (Nov. 16, 2017). URL: https://arstechnica.com/information-technology/2017/11/new-quad9-dns-service-blocks-malicious-domains-for-everyone/.

[124] Quad9. *Quad9 DNS: Internet Security and Privacy in a Few Easy Steps*. 2019. URL: https://www.quad9.net/.

[125] Forum of Incident Response and Security Teams. *Traffic Light Protocol (TLP)*. Aug. 2017. URL: https://www.first.org/tlp.

[126] COVID-19 Cyber Threat Coalition. *Code of Conduct*. URL: https://www.cyberthreatcoalition.org/about-us/code-of-conduct.

[127] Tammy Xu. *Volunteer Devs Around the World Fight Pandemic-Inspired Cybercrime*. Built In, June 9, 2020. URL: https://builtin.com/cybersecurity/coronavirus-cyber-threat-coalition.

[128] Brian Krebs. "COVID-19 Has United Cybersecurity Experts, But Will That Unity Survive the Pandemic?" In: *Krebs on Security* (Apr. 15, 2020). URL: https://krebsonsecurity.com/2020/04/covid-19-has-united-cybersecurity-experts-but-will-that-unity-survive-the-pandemic.

[129] Ravindu De Silva et al. "Compromised or Attacker-Owned: A Large Scale Classification and Study of Hosting Domains of Malicious URLs". In: *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021. URL: https://www.usenix.org/conference/usenixsecurity21/presentation/desilva.

[130] Adam Oest et al. "PhishFarm: A Scalable Framework for Measuring the Effectiveness of Evasion Techniques against Browser Phishing Blacklists". In: *2019 IEEE Symposium on Security and Privacy*. 2019, pp. 1344–1361. DOI: 10.1109/SP.2019.00049.

[131] Thomas Vissers et al. "Assessing the Effectiveness of Domain Blacklisting Against Malicious DNS Registrations". In: *2019 IEEE Security and Privacy Workshops*. 2019, pp. 199–204. DOI: 10.1109/SPW.2019.00045.

[132] Adam Oest et al. "Sunrise to Sunset: Analyzing the End-to-end Life Cycle and Effectiveness of Phishing Attacks at Scale". In: *29th USENIX Security Symposium*. USENIX Security '20. 2020. URL: https://www.usenix.org/conference/usenixsecurity20/presentation/oest-sunrise.

[133] Michel J. G. van Eeten and Johannes M. Bauer. "Economics of Malware: Security Decisions, Incentives and Externalities". In: *OECD Science, Technology and Industry Working Papers* 2008/01 (May 2008). DOI: 10.1787/241440230621.

[134] Yochai Benkler. "Peer Production and Cooperation". In: *Handbook on the Economics of the Internet*. Ed. by J. M. Bauer and M. Latzer. Edward Elgar, 2013.

[135] Sameera Horawalavithana et al. *Malicious and Low Credibility URLs on Twitter during COVID-19*. 2021. arXiv: 2102.12223 [cs.SI].

[136] Pengcheng Xia et al. "Identifying and Characterizing COVID-19 Themed Malicious Domain Campaigns". In: *11th ACM Conference on Data and Application Security and Privacy*. CODASPY '21. 2021, pp. 209–220. DOI: 10.1145/3422337.3447840.

[137] Anjali Vyas et al. "Lost in Transmission: Investigating Filtering of COVID-19 Websites". In: *Financial Cryptography and Data Security*. FC '21. 2021.

[138] Marzieh Bitaab et al. "Scam Pandemic: How Attackers Exploit Public Fear through Phishing". In: *2020 APWG Symposium on Electronic Crime Research (eCrime)*. eCrime 2020. 2020.

[139] M. P. Fischerkeller, E. O. Goldman, and R. J. Harknett. *Cyber Persistence Theory*. Oxford, England, UK: Oxford University Press, May 2022. ISBN: 978-0-19763826-2. URL: https://global.oup.com/academic/product/cyber-persistence-theory-9780197638262.

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

[141] J. E. Barnes. "U.S. Military Has Acted Against Ransomware Groups, General Acknowledges". In: *New York Times* (Dec. 2021). ISSN: 0362-4331. URL: https://www.nytimes.com/2021/12/05/us/politics/us-military-ransomware-cyber-command.html.

[142] Z. Lanz. "Cybersecurity Risk in U.S. Critical Infrastructure: An Analysis of Publicly Available U.S. Government Alerts and Advisories". In: *Virtual Commons - Bridgewater State University* 5.1 (2022), pp. 43–70. DOI: 10.52306/2578-3289.1121.

[143] Tim Maurer and Garret Hinck. "Persistent Enforcement: Criminal Charges as a Response to Nation-State Malicious Cyber Activity". In: *Journal of National Security Law & Policy* (Jan. 2020). URL: https://jnslp.com/2020/01/23/persistent-enforcement-criminal-charges-as-a-response-to-nation-state-malicious-cyber-activity.

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

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

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

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

[148] CompaniesMarketCap.com. *CrowdStrike (CRWD) - Market capitalization.* Mar. 2026. URL: https://companiesmarketcap.com/usd/crowdstrike/marketcap.

[149] Ben Buchanan. *The Hacker and the State: Cyber Attacks and the New Normal of Geopolitics.* Cambridge, MA, USA: Harvard University Press, Feb. 2020. ISBN: 978-0-67498755-5.

[150] M. Armelli et al. *Cyber Threat Intelligence - What is the Impact of Information Disclosures on an Adversary's Operations?* Tech. rep. University of Columbia, SIPA, 2020. URL: https://www.sipa.columbia.edu/academics/capstone-projects/cyber-threat-intelligence-what-impact-information-disclosures-adversary.

[151] S. Soesanto and M. Smeets. "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*. The Hague, The Netherlands: T.M.C. Asser Press, Dec. 2020, pp. 385–400. ISBN: 978-94-6265-419-8. DOI: 10.1007/978-94-6265-419-8_20.

[152] J. Sigholm and M. Bang. "Towards Offensive Cyber Counterintelligence: Adopting a Target-Centric View on Advanced Persistent Threats". In: *DIVA* (2013), pp. 166–171. URL: https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A640955&dswid=6375.

[153] J. Dykstra et al. "Sludge for Good: Slowing and Imposing Costs on Cyber Attackers". In: *ArXiv e-prints* (Nov. 2022). DOI: 10.48550/arXiv.2211.16626. eprint: 2211.16626.

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

[155] M. Vermeer et al. "Alert Alchemy: SOC Workflows and Decisions in the Management of NIDS Rules". In: *ACM Conferences*. New York, NY, USA: Association for Computing Machinery, Nov. 2023, pp. 2770–2784. DOI: 10.1145/3576915.3616581.

[156] B. Tostes et al. "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, pp. 2023–02. DOI: 10.1109/CSR57506.2023.10224937.

[157] A. Kodituwakku et al. "Temporal Aspects of Cyber Threat Intelligence". In: *2023 IEEE International Conference on Big Data (BigData)*. IEEE, 2023, pp. 15–18. DOI: 10.1109/BigData59044.2023.10386664.

[158] Mathew Vermeer, Carlos Gañán, and Michel van Eeten. "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, 2022.

[159] Kurt Thomas et al. "Framing Dependencies Introduced by Underground Commoditization". In: *Proceedings of the Workshop on the Economics of Information Security (WEIS)*. Delft, The Netherlands, June 2015.

[160] Ryan Brunt, Prakhar Pandey, and Damon McCoy. "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, June 2017.

[161] K. Thomas et al. "The Abuse Sharing Economy: Understanding the Limits of Threat Exchanges". In: *Research in Attacks, Intrusions, and Defenses*. Cham, Switzerland: Springer, Sept. 2016, pp. 143–164. ISBN: 978-3-319-45719-2. DOI: 10.1007/978-3-319-45719-2_7.

[162] L. Bilge and T. Dumitraş. “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*. New York, NY, USA: Association for Computing Machinery, Oct. 2012, pp. 833–844. ISBN: 978-1-45031651-4. DOI: 10.1145/2382196.2382284.

[163] G. Di Tizio, M. Armellini, and F. Massacci. “Software Updates Strategies: a Quantitative Evaluation against Advanced Persistent Threats”. In: *IEEE Transactions on Software Engineering* (May 2022), p. 1. ISSN: 1939-3520. DOI: 10.1109/TSE.2022.3176674.

[164] Gartner. *Market Guide for Security Threat Intelligence Products and Services*. Dec. 2021.

[165] D. Chiba et al. “DomainPrio: Prioritizing Domain Name Investigations to Improve SOC Efficiency”. In: *IEEE Access* 10 (Mar. 2022), pp. 34352–34368. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2022.3161636.

[166] Victor Le Pochat et al. “Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation”. In: *Proceedings of the 26th Annual Network and Distributed System Security Symposium*. NDSS 2019. Feb. 2019. DOI: 10.14722/ndss.2019.23386.

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

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

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

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

[171] T. Vissers et al. “Exploring the Ecosystem of Malicious Domain Registrations in the .eu TLD”. In: *Research in Attacks, Intrusions, and Defenses*. Cham, Switzerland: Springer, Oct. 2017, pp. 472–493. ISBN: 978-3-319-66332-6. DOI: 10.1007/978-3-319-66332-6_21.

[172] E. D. Lonergan and S. W. Lonergan. “Cyber Operations, Accommodative Signaling, and the De-Escalation of International Crises”. In: *Security Studies* (Jan. 2022). ISSN: 0963-6412. DOI: 10.1080/09636412.2022.2040584.

[173] R. L. Wasserstein, A. L. Schirm, and N. A. Lazar. “Moving to a World Beyond “p < 0.05””. In: *American Statistician* (Mar. 2019). ISSN: 0003-1305. DOI: 10.1080/00031305.2019.1583913.

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

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

[176] Ponemon Institute. *The Value of Threat Intelligence: The Second Annual Study of North American & United Kingdom Companies*. Tech. rep. February. Ponemon Institute, 2019.

[177] J. R. Lindsay. "Cyber conflict vs. Cyber Command: hidden dangers in the American military solution to a large-scale intelligence problem". In: *Intelligence and National Security* 36.2 (Feb. 2021), pp. 260–278. ISSN: 0268-4527. DOI: 10.1080/02684527.2020.1840746.

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

[179] C. Johnson et al. *Guide to Cyber Threat Information Sharing*. Tech. rep. National Institute of Standards and Technology, 2016. URL: http://dx.doi.org/10.6028/NIST.SP.800-150.

[180] F. Skopik, G. Settanni, and R. Fiedler. "A problem shared is a problem halved: A survey on the dimensions of collective cyber defense through security information sharing". In: *Computers & Security* 60 (July 2016), pp. 154–176. ISSN: 0167-4048. DOI: 10.1016/j.cose.2016.04.003.

[181] B. Callaway and P. H. C. Sant'Anna. "Difference-in-Differences with multiple time periods". In: *Journal of Econometrics* 225.2 (Dec. 2021), pp. 200–230. ISSN: 0304-4076. DOI: 10.1016/j.jeconom.2020.12.001.

[182] L. Sun and S. Abraham. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects". In: *Journal of Econometrics* 225.2 (Dec. 2021), pp. 175–199. ISSN: 0304-4076. DOI: 10.1016/j.jeconom.2020.09.006.

[183] Simon, Herbert. *A Behavioral Model of Rational Choice*. Feb. 1955. URL: https://www.jstor.org/stable/1884852.

[184] J. E. Sims. *Decision Advantage*. Oxford, England, UK: Oxford University Press, Nov. 2022. ISBN: 978-0-19750804-6. URL: https://global.oup.com/academic/product/decision-advantage-9780197508046.

[185] Inc CrowdStrike. *CrowdStrike and Microsoft Collaborate to Harmonize Cyber Threat Attribution*. June 2025. URL: https://ir.crowdstrike.com/news-releases/news-release-details/crowdstrike-and-microsoft-collaborate-harmonize-cyber-threat.

[186] *Cristin Goodwin on Twitter (Internet Archive)*. Jan. 2019. URL: https://web.archive.org/web/20210510172415/https://twitter.com/CristinGoodwin/status/1181722707500703744.

[187] *A Conversation on Cybersecurity with NSA's Rob Joyce*. Apr. 2023. URL: https://www.csis.org/analysis/conversation-cybersecurity-nsas-rob-joyce.

[188] Yi-Min Wang et al. "Strider Typo-Patrol: Discovery and Analysis of Systematic Typo-Squatting". In: *2nd Workshop on Steps to Reducing Unwanted Traffic on the Internet*. 2006, pp. 31–36.

[189] Tobias Holgers, David E. Watson, and Steven D. Gribble. “Cutting Through the Confusion: A Measurement Study of Homograph Attacks”. In: *USENIX Annual Technical Conference*. 2006, pp. 261–266.