



Delft University of Technology

Quantifying cybercriminal bitcoin abuse

Oosthoek, K.

DOI

[10.4233/uuid:62782ffc-5958-4eff-b429-545c34405200b](https://doi.org/10.4233/uuid:62782ffc-5958-4eff-b429-545c34405200b)

Publication date

2023

Document Version

Final published version

Citation (APA)

Oosthoek, K. (2023). *Quantifying cybercriminal bitcoin abuse*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:62782ffc-5958-4eff-b429-545c34405200b>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

QUANTIFYING CYBERCRIMINAL BITCOIN ABUSE

QUANTIFYING CYBERCRIMINAL BITCOIN ABUSE

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op maandag 6 februari 2023 om 15:00 uur

door

Kris OOSTHOEK

Master of Science in Strategic Management,
Erasmus University Rotterdam,
geboren te Dordrecht, Nederland.

Dit proefschrift is goedgekeurd door de promotoren.

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. ir. R. L. Lagendijk	Technische Universiteit Delft, promotor
Prof. dr. G. Smaragdakis	Technische Universiteit Delft, copromotor

Onafhankelijke leden:

Prof. dr. B. van den Berg	Universiteit Leiden
Prof. dr. M. J. G. van Eeten	Technische Universiteit Delft
Dr. ir. J. A. Pouwelse	Technische Universiteit Delft
Prof. dr. K. P. Gummadi	Max Planck Institute for Software Systems
Prof. dr. B. P. F. Jacobs	Radboud Universiteit
Prof. dr. ir. F. A. Kuipers	Technische Universiteit Delft, reservelid



Keywords: Bitcoin, Cybercrime, Cybersecurity

Front & Back: Kris Oosthoek.

Copyright © 2022 by K. Oosthoek

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

All we are not stares back at what we are.

W.H. Auden

CONTENTS

Summary	xi
Samenvatting	xiii
1 Introduction	1
1.1 Bitcoin	2
1.2 Cybercriminal Abuse of Bitcoin	3
1.3 Transaction Tracking in Bitcoin	4
1.4 Problem Statement	6
1.5 Contribution of the Thesis	7
1.5.1 Outline	8
1.5.2 List of Excluded Publications	8
1.5.3 About the Thesis	9
2 Cyber Security Threats to Bitcoin Exchanges	13
2.1 Introduction	14
2.2 Related Work	16
2.3 Cyber Threat Intelligence	16
2.3.1 Vocabulary for Event Recording and Incident Sharing (VERIS)	17
2.4 Methodology	17
2.4.1 Data Collection	17
2.4.2 Classification of Breaches	18
2.5 Analysis Of Bitcoin Exchange Breaches	19
2.5.1 Analysis of Attack Vectors	21
2.5.2 Analysis of impact	23
2.6 Evolution of Platform Security	26
2.6.1 Multiple services on single IP	27
2.6.2 Vulnerable web server versions	27
2.6.3 Exposed management interfaces	27
2.6.4 Slow patching in general	28
2.6.5 HTTP Security Headers	28
2.6.6 Server-Side Software	28
2.7 Tracing Stolen Funds	29
2.7.1 Mt. Gox	29
2.7.2 Bitfinex Breach	30
2.7.3 Yapizon	32
2.8 Limitations	32
2.9 Conclusion	33

3	Investigating the Ransomware Payments Economy	41
3.1	Introduction	42
3.2	The Ransomware Ecosystem	43
3.2.1	Commodity Ransomware	43
3.2.2	Ransomware as a Service (RaaS)	43
3.3	Methodology	45
3.3.1	Addresses involved in Ransom Payments	45
3.3.2	Ransom Payments and Laundering	46
3.3.3	Ransomware Actors	47
3.3.4	Limitations	47
3.4	Ransom Payment Analysis	47
3.4.1	Ransomware Revenue	48
3.4.2	Ransomware Payment Characteristics	49
3.5	Money Laundering Analysis	51
3.5.1	Laundering Strategies	51
3.5.2	Challenges in Fighting Laundering	53
3.6	Conclusion	55
4	Illicit Revenue of Dark Web Shops	61
4.1	Introduction	62
4.2	Background	64
4.2.1	Tor	64
4.2.2	Bitcoin	64
4.2.3	Bitcoin: Regulation and Market Capitalization	64
4.3	Related Work	65
4.4	Methodology	66
4.4.1	Tor Crawler	66
4.4.2	Bitcoin Address Clustering	68
4.5	Cleansing Methodology	69
4.5.1	Tor Domains	69
4.5.2	Bitcoin Addresses	72
4.5.3	Bitcoin Address Clusters	73
4.6	Quantifying Illicit Revenue	74
4.6.1	Seed Address Revenue per Abuse Type	74
4.6.2	Longitudinal Analysis of Seed Address Transactions	75
4.6.3	Revenue per Abuse Type after Address Clustering	76
4.7	Quantifying Shop vs Marketplace Revenue	78
4.7.1	The Hydra Marketplace and its Take-down	78
4.7.2	Dark Web Shops vs. Hydra Transactions Revenue	79
4.7.3	Dark Web Shops vs. Hydra Bitcoin Address and Laundering Overlap	80
4.7.4	Discussion	80
4.8	Conclusion	81
4.9	Appendices	82

5	Discussion	91
5.1	Exchange Hacks.	91
5.1.1	TTPs and Financial Impact.	91
5.1.2	Limitations.	93
5.1.3	Reflection	93
5.1.4	Future Work	93
5.2	Ransomware Payments	94
5.2.1	Payment and Laundering Analysis	94
5.2.2	Limitations.	94
5.2.3	Future Work	95
5.3	Dark Web Shops.	95
5.3.1	Payment Analysis	95
5.3.2	Limitations.	96
5.3.3	Future Work	96
5.4	Conclusion	96
	Acknowledgements	99
	List of Publications	101

SUMMARY

Cybercrime is negatively impacting everybody. In recent years cybercriminal activity has directly affected individuals, companies, governments and critical infrastructure. It has led to significant financial damage, impeded critical infrastructure and harmed human lives. Defending against cybercrime is difficult, as persistent actors perpetually hunt for soft spots in Internet-connected systems, which exist due to either lax vulnerability management or for convenience, complicating adequate detection and mitigation.

Cybercriminal actors are financially motivated and for their doings and dealings they rely on Bitcoin. Alternatives exist, but Bitcoin has proven to be the most liquid digital currency, meaning it is easy to swap and to conceal illicit transactions. The magnitude of many cybercriminal activities is largely unknown. However Bitcoin runs on a blockchain - an open, decentralized ledger, allowing virtually everyone to analyze financial transactions, as opposed to traditional banking. Furthermore, contrary to popular belief Bitcoin is pseudonymous, not anonymous and several techniques exist to identify illicit activity.

In this thesis, we illuminate three cybercriminal ecosystems that did not receive significant prior research attention: Bitcoin exchange heists, ransomware and single-vendor shops in the Dark Web. For each of these, we gather datasets from open sources. We first focus on the technical behavior and financial impact of attacks on Bitcoin exchange platforms. We also highlight the ransomware ecosystem, showing how it moved from small to large-scale attacks with similar financial impact. We further focus on how small shops in the Dark Web generate significant revenue with niche illicit activity. To understand the financial impact within each of these ecosystems, we analyze associated financial transactions. We also apply heuristics to discover additional Bitcoin addresses controlled by the same actor.

We observe that cybercriminal actors successfully extract millions of funds from Bitcoin exchanges through relatively low-level attack vectors. When compared with traditional financial institutions, the lack of sophistication of attacks and the accompanying financial impact is unprecedented. In our analysis of ransomware, we observe attackers have shifted from attacking individual users resulting in relatively small ransom amounts to targeting large organizations with significant financial resources, resulting in multi-million ransom payments. We also find that with this shift, attackers have also improved their operational security in address usage and money laundering. For Dark Web shops, we found that this relatively uncharted territory of the Dark Web as compared to the bigger marketplaces specializes into niches such as sexual abuse material and various forms of financial crime. To allow for future research in this area, we introduce a methodology to estimate illicit revenue based on web scrape results and cluster these on category.

SAMENVATTING

Cybercrime raakt iedereen. De afgelopen jaren heeft cybercriminele activiteit rechtstreekse negatieve gevolgen gehad voor mensen, bedrijven, overheden en kritieke infrastructuur. Het heeft flinke financiële schade teweeggebracht, kritieke infrastructuur gedwarsboemd en mensenlevens geschaad. Verdediging tegen cybercriminaliteit is ingewikkeld, omdat kwaadwillende actoren continu zoeken naar zwakke plekken in met het Internet verbonden systemen. Die kunnen bestaan door gebrekkig nazicht op kwetsbaarheden of uit gebruiksgemak, wat doeltreffende detectie en mitigatie bemoeilijkt.

Cybercriminele actoren zijn financieel gemotiveerd en hun handel en wandel steunt op Bitcoin. Alternatieven bestaan, maar Bitcoin heeft aangetoond de meest liquide digitale munt te zijn, wat betekent dat het eenvoudig is in te wisselen en om onwettige transacties te verhullen. De omvang van deze cybercriminele activiteit is veelal onbekend. Bitcoin draait echter op een blockchain - een open decentraal grootboek, wat vrijwel iedereen in staat stelt financiële transacties te analyseren, in tegenstelling tot traditioneel bankieren. Anders dan men algemeen aanneemt is Bitcoin pseudoniem, niet anoniem en bestaan er verscheidene technieken om illegale activiteit vast te stellen.

In dit proefschrift werpen we een licht op drie cybercriminele ecosystemen waar onvoldoende substantieel onderzoek naar is verricht: hacks van Bitcoin-handelsplatformen, ransomware en eenmanswinkeltjes in het *Dark Web*. Voor elk daarvan verzamelen we datasets uit open bronnen. Als eerste richten we ons op de technische handelswijze en de financiële impact van aanvallen op Bitcoin-handelsplatformen. Ook belichten we het ransomware-ecosysteem en zetten we de verschuiving van kleine naar grootschalige aanvallen met gelijkaardige financiële impact uiteen. Verder bekijken we hoe kleine winkels in het *Dark Web* aanzienlijke omzet behalen met illegale niche-activiteiten. Om de financiële impact van elk van deze ecosystemen te begrijpen analyseren we bijbehorende financiële transacties. We passen ook heuristische methoden toe om aanvullende Bitcoin-adressen in bezit van dezelfde actor bloot te leggen.

We nemen waar dat cybercriminele actoren erin slagen miljoenen contanten buit te maken bij Bitcoin-handelsplatformen met relatief laagdrempelige aanvalsvectoren. Zeker in vergelijking met traditionele financiële instellingen is het gebrek aan kundigheid en de bijbehorende financiële impact ongekend. In de analyse van ransomware nemen we waar dat aanvallers zijn verschoven van aanvallen op particulieren met relatief lage losgeld-eisen naar grote organisaties met aanzienlijke financiële draagkracht, wat resulteert in afkooksommen van meerdere miljoenen. We nemen ook waar dat met deze verschuiving aanvallers hun activiteit beter afschermen door ander adresgebruik en witwasmethoden. Voor *Dark Web*-winkels zien we dat dit onontgonnen gebied van het Dark Web in vergelijking met de grotere marktplaatsen zich vooral specialiseert in niches als seksueel misbruik en financiële criminaliteit. Om toekomstig onderzoek op dit gebied mogelijk te maken introduceren we een methodologie om illegale inkomsten te schatten op basis van *scrape*-resultaten en die te clusteren per categorie.

1

INTRODUCTION

The financial and societal impact of cybercrime is unparalleled. According to its definition cybercrime is any “*crime or illegal activity that is done using the internet*” [11], which merely conveys a means of communication. The impact of cybercrime generally also stretches beyond that of conventional crime. Ransomware encrypts user files, blackmailing victims to pay for the decryption key. In recent years, it has evolved from a threat primarily impacting consumers to a key concern to the continuity of critical infrastructure such as transportation, hospitals, food and energy supply [20]. Adversaries have evolved from individuals deploying self-authored ransomware to full-fledged professional gangs, and ransom demands have grown accordingly. Nominally, ransomware attacks have cost victims an aggregate of at least 124 million US dollar [30]. With a record ransom demand of 240 million US dollar in 2021 [7] and an average of 170 thousand US dollar [16], ransomware is a cybercriminal business model with low opportunity cost and a high profit opportunity. Yet, it is just one swindle in the cybercriminal playbook.

Exchange platforms where cryptocurrency enthusiasts can exchange fiat currency for their digital currency of choice, are frequently targeted by cybercriminals. And with significant success. In the biggest heist to date, half a billion US dollar worth in Bitcoin was taken [10], in other cases millions were stolen to bankroll nuclear weapons development [35]. Another cybercriminal realm is the commerce taking place in the unregulated and un-indexed corners of the Internet; the Dark Web. Much of this is illicit, demonstrated by 315 million US dollar drugs sales annually [36], with single vendors earning over 34 million US dollar [22]. It is readily apparent that cybercriminal actors do not shy away from data breaches, financial fraud, extortion, online harassment, drug and human trafficking. Cybercrime is void of taboos as long as there is a profit opportunity.

For the realization of its profit, cybercrime depends on cryptocurrencies; digital cash based on blockchain technology. According to conservative estimates 0.15% of the aggregate cryptocurrency transaction value in 2021 [5] is labeled illicit, nominally this still accounts for 23.7 billion USD. Bitcoin is the preferred currency in the majority of illicit activity [31, 5], most probably due to its liquidity. And while illegal activity in Bitcoin is often the subject of vivid media coverage, many of the details are still fuzzy.

1.1. BITCOIN

Bitcoin is a digital currency created by the pseudonymous Satoshi Nakamoto in 2008 [29]. Bitcoin is decentralized, meaning no single person or group has central authority and financial transactions are a collective effort. The type of public ledger used by Bitcoin to record transactions has become known as *blockchain*. Bitcoin's blockchain is distributed over a global peer-to-peer network of nodes. Each node stores a copy of the blockchain, keeping a historical record of all transactions and verifying new transactions. New transactions verified as valid are broadcasted to peer nodes, which in turn also perform verification. If a sufficient number of verifications is performed, the transaction is added to a pool with other valid transactions. A subset of nodes, *mining nodes* or simply *miners*, gathers the transactions from this pool to package them into blocks. This is performed based on *proof of work* [14], in which miners compete to create the next block, incentivized by a block reward for the miner that mines the eventual block. The miner who provides the *proof of work*, generating the hash to mine the block, wins the challenge to mine the eventual block. Approximately every 10 minutes a block with new transactions is mined. Bitcoin attains the 10-minute cadence by considering the available computing power in the network and adjusting the *mining difficulty* accordingly. Once a block is confirmed, the contained transactions are immutable, meaning irreversible. When the 10-minute mining cadence persists, and every 4 years the block reward is halved, the total supply of 21 million Bitcoin will be mined in the year 2140.

While some proponents argue that the introduction of Bitcoin in 2008 was revolutionary, its conception was rather evolutionary. Ideologically, Bitcoin is the brainchild of the *cypherpunk* movement. Various events in the 1990s led to the establishment of the cypherpunk movement, advocating the use of strong cryptography to ensure the privacy of personal communication [21]. On a technical level Bitcoin builds on the work on secure exchange of cryptographic keys by Diffie and Hellman in 1976 [12], as well as the invention of public key encryption by Rivest et al. in 1978 [33] and elliptic curve cryptography by Koblitz and Miller in 1985 [24, 27]. The conception of Merkle trees [26], which allow for efficient verification of cryptographic hashes in large systems was important to integrity checking of data sent by peers. The work in 1990 by Haber and Stornetta worked on a tamper-proof time-stamping mechanism [19] was effectively the conceptualization of a *blockchain* and upgraded two years later to include Merkle trees [2].

Building on these technologies, self-identified *cypherpunks* launched digital cash initiatives such as Digicash by David Chaum in 1989 [6], Bit Gold by Nick Szabo in 1998 [34] and Hashcash by Adam Back in 2002 [1]. According to a reference in the Bitcoin whitepaper, Bitcoin mostly builds on *b-money* by Wei Dai [28]. In a post on the cypherpunk mailinglist in 1998 [9], Dai proposed a decentralized cryptocurrency using *proof of work*. In 2004, Hal Finney prototyped the reusable Proof-of-Work system, where the value of real-world resources used to mint a digital token are linked to its value [14]. Further iterating on Proof-of-Work, the introduction of Bitcoin in 2008 the introduction of the concept of a blockchain. After this, many alternative cryptocurrencies and blockchains were introduced. Though alternative blockchains like Ethereum and stablecoins like Tether (USDT) and USD Coin (USDC) have proven their right to exist, Bitcoin is still the *primus inter pares* in terms of trading volume and market capitalization [8].

1.2. CYBERCRIMINAL ABUSE OF BITCOIN

In the years after its introduction, Bitcoin became a social phenomenon and its use expanded into practically every country around the globe. It has attracted the interest of individual and institutional investors, but also cybercriminals. Cybercrime is an umbrella term for all criminal activity performed using a digital device, usually over the Internet [18], such as identity fraud, theft of financial or payment data, trafficking of illicit material, cryptojacking and extortion (including ransomware). Cybercriminals take to Bitcoin due its liquidity and alleged anonymity, in reality rather pseudonymity. Though alternative cryptocurrencies such as Monero might offer better privacy and anonymity, these are considerably less liquid than Bitcoin, meaning it is harder to *hide in the crowd* and to convert to cash without significant expenditure.

The proportion of illicit to overall transaction revenue for all cryptocurrencies has declined from 3.37% in 2019 to 0.62% in 2020 and 0.15% in 2021 [5]. But with a total transaction volume of 15.8 trillion USD for all cryptocurrencies in 2021, the nominal value of illicit activity is still 23.7 billion USD. Illicit use includes the exploitation of technical vulnerabilities in Virtual Asset Service Providers (VASPs), but also the use of cryptocurrencies of a means of payment for illegal activity.

BITCOIN EXCHANGE HACKS

Soon after Bitcoin's inception in 2008, an ecosystem of Virtual Asset Service Providers (VASPs) started to emerge. Most often VASPs provide exchange services, trading fiat currency for bitcoin and vice-versa, as intermediary between sellers and buyers and as custodian, providing cryptocurrency wallets. VASPs operate independently from the Bitcoin blockchain, running on self-hosted or cloud infrastructure, with tailor-made code bases and many software dependancies, generating significant attack surface. Not surprisingly, the custodial wallets with user funds maintained by these platforms, are an attractive target to financially motivated hackers. The inadequate cyber security of these platforms provides a low-effort and low-risk opportunity to gain significant illicit profits.

RANSOMWARE PAYMENTS

Until a few years ago, ransomware authors primarily targeted home-user systems, acquiring moderate ransom amounts. But attackers have professionalized, capitalizing on the lacking cyber security of many enterprises and governments. The exploitation of technical weaknesses in (legacy) Internet-facing systems, followed by exfiltration of sensitive data used to extort victims, has resulted in multi-million USD ransom demands. The ransomware problem is catalyzed by inadequate regulation and law enforcement struggling with the cross-border nature of cryptocurrency payments. In many cases, reputational risk prevents victims from disclosing ransomware infections. The reporting that exists is usually commercial or regional in nature, providing only partial coverage of the problem. As a result, the actual scope and size of the ransomware problem is unclear.

DARK WEB SHOPS

A significant portion of the Internet, such as message boards, online bank systems and all content behind a paywall, is not indexed by search engines. A subset of this unindexed

part of the Internet depends on The Onion Router (Tor) [32]. Only accessible using a custom browser, Tor hosts a diverse set of onions (domains) with services for privacy-aware merchants and purchasers. In principle, these can be legitimate - Facebook and Protonmail are accessible via Tor. However a majority of onions is illicit, simply because upholding a site is a technical burden for engineers and thus not feasible to many regular organizations. An important share of the illicit services in Tor is made up by storefronts with illicit offerings, commonly designated as Dark Web Shops. In contrast to exchange security breaches and ransomware payments, these shops do not directly monetize security breaches. However they can and do indirectly monetize artifacts from breaches by selling proprietary, confidential or otherwise sensitive information obtained in breaches. The selling of abuse material and drugs in the dark web also rests on inadequate regulation and law enforcement of illicit cryptocurrency payments.

1.3. TRANSACTION TRACKING IN BITCOIN

In order to obtain an informed understanding of the quantity and nature of illicit Bitcoin transactions, insight into the identity of senders and recipients is required. The recipient of a Bitcoin transaction is anonymous, unless his/her real-world identity is connected to the address(es) he/she controls. This deanonymization principle used to fingerprint Bitcoin addresses is best demonstrated using the first Bitcoin transaction as an example.

Hal Finney was a computer scientist known for Reusable Proof of Work [14], discussed in the previous section, and early user of Bitcoin. On January 12, 2009, only 9 days after the *genesis block*, the first ever Bitcoin block mined, Finney received the first Bitcoin transaction from Satoshi Nakamoto [3]. The transaction data is pictured in Figure 1.1. According to Finney, he was the first user other than Satoshi to run the Bitcoin software after its announcement on a cryptography mailing list [13]. Finney has always denied he was Satoshi Nakamoto [15]. This event is important because it signified the concept of transaction analysis in blockchains.

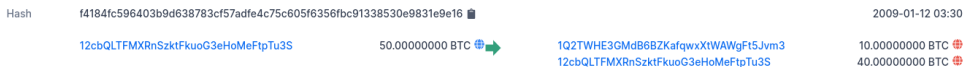


Figure 1.1: The first Bitcoin transaction between to addresses, between Satoshi Nakamoto and Hal Finney, as visible in public Bitcoin explorer *blockchain.com*

Through Finney's self-disclosure as recipient, his identity became affiliated with the Bitcoin address used in the transaction. Effectively he de-anonymized himself, as the owner of the particular Bitcoin address was unidentified before he did so. Because blockchains are immutable, meaning it is not allowed to change data after a block is mined, the evidence is stored forever. After this, plausible deniability of his ownership of the *funds* would only be possible based on evidence recorded in another blockchain transaction. The address itself is linked to him forever. This is the seed of how blockchain analysis works. The privacy risk was already recognized in the original whitepaper [29]:

“As an additional firewall, a new key pair should be used for each transaction to keep them from being linked to a common owner. Some linking is still unavoidable with multi-input transactions, which necessarily reveal that their

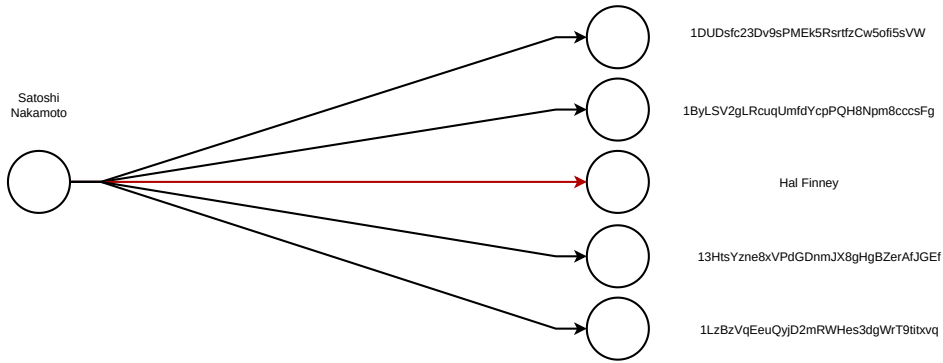


Figure 1.2: Conceptual image of transaction output from Satoshi Nakamoto to Hal Finney and other addresses. All expenses to and from his address are linked to Finney after him claiming ownership of the address.

inputs were owned by the same owner. The risk is that if the owner of a key is revealed, linking could reveal other transactions that belonged to the same owner.”

The pseudonymous Satoshi Nakamoto recognized the privacy risk of reusing an address for more than one transaction. When a user uses a single address for multiple transactions, profiling his/her purchasing behavior is as easy as inspecting all outgoing transactions. Hence Satoshi recommended to create a new address on a per-transaction basis, but even when adhering to this privacy is not ensured. In order to be able to spend funds, the new address still needs funding, which then needs to come from another address or by sending some amount from an exchange to the new address. In both cases, it is pretty easy to reveal the identity.

In addition to address reuse, several alternative methods to backtrack transactions and funds to their owner have been proposed. When generally accepted, such a method is considered a *heuristic*. These heuristics are embedded in the field of blockchain analysis, which has become critical to the success of many law enforcement investigations since then. Blockchain analysis is the forensic activity to trace flows of funds back to a criminal actor. It is usually applied in criminal investigations into ransomware actors, DDoS service operators, phishing crews and other cybercriminal actors. Results of blockchain analysis regularly serve as input to litigation. Blockchain analysis is not to be confused with on-chain analysis, in which transactions are scrutinized to deduce market sentiment to inform cryptocurrency trading strategies.

The heuristic most commonly used for clustering Bitcoin addresses is the *common spend* or *multiple input* heuristic [25]. It is considered safe and widely used in both research and practice, but vulnerable to false positive detection of CoinJoin transactions [23], which thus must be filtered prior to analysis. Another common heuristic is to identify the address which receives change, unspent bitcoin, hence this is called the *change* heuristic [37]. Several individual change heuristics exist, the most popular open source analysis platform including 10 of them [23], but their applicability heavily depends on expert judgment on a per-case basis. Hence change heuristics are less suited for auto-

mated analysis.

The heuristics mentioned hitherto are passive heuristics. A proactive method is *dusting*, in which trace amounts of cryptocurrency (*dust*) are sent to a multitude addresses [17]. The theoretical and practical use cases of dusting are broader than blockchain analysis. The minimum amount of bitcoin per address required for dusting is one *satoshi*, one hundred millionth of a bitcoin, but a single attack usually targets hundreds or thousands of unique addresses. Cybercriminals have used dust to promote scams in the text field of Bitcoin transactions. Theoretically dusting can also be applied to deanonymize users with addresses with large holdings. The sending of funds to unaware and potentially malicious users however raises concerns for its use in academic research.

In addition to heuristics, many commercial blockchain analysis solutions rely on *open source intelligence* to attribute addresses to their real-world owner. Based on both manual and automated scraping of social media and Internet and Dark Web forums, the coverage of labeled addresses is increased. In this thesis, we solely apply the co-spending heuristic due to the lack of dependability of the change address heuristic in bulk analysis and ethical concerns over dusting. In law enforcement investigations, blockchain analysis capabilities are usually supplemented with identifying information enquired from VASPs such as cryptocurrency exchanges.

This thesis investigates three dominant cases of cybercriminal abuse of Bitcoin. We investigate how cybercriminals exploit technical weaknesses in exchange platforms to appropriate millions worth of Bitcoin. We also analyze how cybercriminals have evolved their deployment of ransomware to obtain higher ransom payments. We further estimate proceedings within different cybercriminal revenue categories for actors outside the big Dark Web market places.

1.4. PROBLEM STATEMENT

This thesis addresses Bitcoin in relation to cybercrime by analyzing different types of cybercriminal activity where Bitcoin is the dominant means of payment. It does so based on two perspectives:

- How to perform data collection and the technical innovations required for transaction/blockchain analysis; and based on this:
- The quantification of cybercriminal abuse of Bitcoin.

While previous work on these topic exists, we aim for a fundamental, end-to-end approach by collecting our own datasets and developing analysis methodology. Many previous initiatives are lacking either in transparency on underlying data or the methodology used. In addition to this we also share our datasets to improve community analysis efforts. Related work also tends to look at less representative datasets, or of secondary provenance. Considering gaps in existing research literature focusing on Bitcoin and its cybercriminal use, this thesis will aim to answer to following question:

How can we leverage open data to increase our understanding of cybercriminal usage of Bitcoin and how can we quantify this?

The research question is divided into three sub-questions to narrow down the problem area:

- How do cybercriminal actors abuse security vulnerabilities in Bitcoin exchange platforms and what is the financial impact?
- What is the revenue of ransomware actors utilizing Bitcoin as a means of payment and how has this evolved over time?
- How can we confidently estimate revenue the of Dark Web cyber-criminal actors, and based on that what revenues do we see?

1.5. CONTRIBUTION OF THE THESIS

This thesis contributes to the knowledge-building on cybercriminal abuse of Bitcoin. With technical chapters focusing on novel cybercriminal tactics, techniques and procedures (TTPs) in individual ecosystems, the problem is regarded from various perspectives. The contributions of this thesis are as follows:

- We provide the largest analysis to date of TTPs employed security breaches of Bitcoin exchange platforms in **Chapter 2**. In addition to analyzing their financial impact, we also consider how this compares with security breaches of traditional financial institutions and how specific actors have laundered funds.
- In **Chapter 2** we also consider, based on passive analysis, the attack surface of exchange platforms that were previously targeted and that still exist.
- The analysis in **Chapter 3** is the largest quantification to date of revenues made by ransomware actors. It is also the first analysis that considers differences between authors of commodity ransomware and Ransomware as a Service (RaaS) groups.
- The dataset gathered for the analysis in **Chapter 3** has been made available [4] for future use.
- Our analysis in **Chapter 4** sheds a slight on an underexposed part of the illicit trading in Tor, namely in individually owned storefronts. Based on analysis of many Tor domains, we provide an impression of illicit activity and its financial impact.
- **Chapter 4** also provides a novel, extensive methodology to filter Bitcoin addresses found in open sources, specifically Tor, for illicit activity and to further cleanse on a per-transaction basis.

1.5.1. OUTLINE

The thesis is structured as follows:

CHAPTER 2

CYBER SECURITY THREATS TO BITCOIN EXCHANGES

In this chapter we investigate the most significant security breaches of cryptocurrency exchanges where Bitcoin was stolen. We regard the tools, techniques and procedures (TTPs) used by attackers and compare this with a secondary dataset of TTPs utilized in attacks against traditional financial institutions, such as banks. This chapter has been published as *Cyber Security Threats to Bitcoin Exchanges: Adversary Exploitation and Laundering Techniques* by **Oosthoek, K.** and Doerr, C. in IEEE Transactions on Network and Service Management, vol. 18, no. 2, pp. 1616-1628, June 2021.

CHAPTER 3

A TALE OF TWO MARKETS: INVESTIGATING THE RANSOMWARE PAYMENTS ECONOMY

This chapter highlights our analysis of a large corpus of Bitcoin addresses associated with ransomware actors. Based on this we are able to provide a unique perspective on the evolution of the ransomware ecosystem, from commodity ransomware into RaaS. We also regard how actors launder the illicit proceedings from ransom payments. This chapter will be published as *A Tale of Two Markets: Investigating the Ransomware Payments Economy* by **Oosthoek, K.**, Cable, J. and Smaragdakis, G. in Communications of the ACM (accepted and in print).

CHAPTER 4

QUANTIFYING ILLICIT REVENUE OF DARK WEB SHOPS

The analysis in this chapter focuses on Dark Web Shops, which are smaller storefronts in the Tor network with various illicit offerings, operated by individual merchants. Specifically, we provide an estimate of the revenue of various categories of illicit offerings by these shops. In order to obtain an accurate estimate, we have developed an extensive methodology for the analysis of Bitcoin addresses found in Tor. This chapter is currently in review with the *Journal of Cybersecurity*.

CHAPTER 5

DISCUSSION

Chapter 5 concludes this thesis with a detailed discussion of the research questions. It also discusses opportunities for future work, based on ideas and research gaps identified during the analyses performed for this thesis.

1.5.2. LIST OF EXCLUDED PUBLICATIONS

The list below contains papers that have been published during this Ph.D. but are not included in this work as they are out of the scope of this thesis.

1. **Oosthoek, K.**, and Doerr, C., *SoK: ATT&CK Techniques and Trends in Windows Mal-*

- ware, International Conference on Security and Privacy in Communication Systems (SecureComm 2019).
2. **Oosthoek, K.**, and Doerr, C., *Cyber Threat Intelligence: A Product Without a Process?*, International Journal of Intelligence and CounterIntelligence 34.2 (2021): 300-315.
 3. Griffioen, H., **Oosthoek, K.**, van der Knaap, P., & Doerr, C. (2021, November). *Scan, Test, Execute: Adversarial Tactics in Amplification DDoS Attacks*, in the ACM SIGSAC Conference on Computer and Communications Security (CCS 2021).
 4. **Oosthoek, K.** and Doerr, C., *Inside the Matrix: CTI Frameworks as Partial Abstractions of Complex Threats*, IEEE International Conference on Big Data (Big Data 2021).
 5. **Oosthoek, K.**, *Flash Crash for Cash: Cyber Threats in Decentralized Finance*, 2021, Preprint

1.5.3. ABOUT THE THESIS

This thesis consists of integral copies of three publications over three technical chapters. Each chapter is based on the original peer-reviewed publication with only minor changes. The original full title of the original article is included on the first page of every chapter. Each technical chapter considers a different cybercriminal ecosystem, based on a unique dataset. However as all chapters are integral copies of prior published work, between the chapters certain segments as the introduction and background sections might overlap.

REFERENCES

- [1] Adam Back et al. “Hashcash-a denial of service counter-measure”. In: (2002).
- [2] Dave Bayer, Stuart Haber, and W Scott Stornetta. “Improving the efficiency and reliability of digital time-stamping”. In: *Sequences II*. Springer, 1993, pp. 329–334.
- [3] *Bitcoin Transaction Hash f4184fc596403b9d638783cf57adfe4c75c605f6356fbc91338-530e9831e9e16*. 2009. URL: <https://www.blockchain.com/btc/tx/f4184fc596403b9d638783cf57adfe4c75c605f6356fbc91338530e9831e9e16>.
- [4] Jack Cable. *Ransomwhere: A Crowdsourced Ransomware Payment Dataset*. Version 1.0.0. Zenodo, May 2022. DOI: [10.5281/zenodo.6512123](https://doi.org/10.5281/zenodo.6512123). URL: <https://doi.org/10.5281/zenodo.6512123>.
- [5] Chainalysis. *The 2022 Crypto Crime Report*. <https://go.chainalysis.com/2022-Crypto-Crime-Report.html>. 2022.
- [6] David Chaum. “Security without identification: Transaction systems to make big brother obsolete”. In: *Communications of the ACM* 28.10 (1985), pp. 1030–1044.
- [7] CoinDesk. *Electronics Retailer MediaMarkt Hit by Ransomware Demand for \$50M Bitcoin Payment: Report*. Nov. 2021. URL: <https://www.coindesk.com/business/2021/11/10/electronics-retailer-mediemarkt-hit-by-ransomware-demand-for-50m-bitcoin-payment-report/>.

- [8] CoinGecko. *The Most Comprehensive Cryptocurrency API*. <https://www.coingecko.com/en/api>.
- [9] Wei Dai. "Wei Dai's "b-money" protocol". In: (1998). URL: <http://cypherpunks.venona.com/date/1998/12/msg00194.html>.
- [10] Daily Beast. *Japanese Bitcoin Heist 'an Inside Job,' Not Hackers Alone*. 2017. URL: <https://www.thedailybeast.com/japanese-bitcoin-heist-an-inside-job-not-hackers-alone>.
- [11] Cambridge Dictionary. *Meaning of cybercrime in English*. 2022. URL: <https://dictionary.cambridge.org/dictionary/english/cybercrime>.
- [12] Whitfield Diffie and Martin E Hellman. "Multiuser cryptographic techniques". In: *Proceedings of the June 7-10, 1976, national computer conference and exposition*. 1976, pp. 109–112.
- [13] Hal Finney. *Bitcoin and me (Hal Finney)*. Bitcointalk.org. 2013. URL: <https://bitcointalk.org/index.php?topic=155054.0>.
- [14] Hal Finney. *Reusable Proofs of Work*. rpow.net. 2004. URL: <https://web.archive.org/web/20071222072154/http://rpow.net/>.
- [15] Forbes. *The Little Black Book of Billionaire Secrets Nakamoto's Neighbor: My Hunt For Bitcoin's Creator Led To A Paralyzed Crypto Genius*. 2014. URL: <https://www.forbes.com/sites/andygreenberg/2014/03/25/satoshi-nakamotos-neighbor-the-bitcoin-ghostwriter-who-wasnt/>.
- [16] World Economic Forum. *The average ransomware demand is now \$170K. Here's how we can fight back*. May 2021. URL: <https://www.weforum.org/agenda/2021/05/ransomware-cybersecurity-partnership-cybercrime-fight-back/>.
- [17] Gemini. *What Is a Crypto Dusting Attack?* 2022.
- [18] Sarah Gordon and Richard Ford. "On the definition and classification of cyber-crime". In: *Journal in computer virology* 2.1 (2006), pp. 13–20.
- [19] Stuart Haber and W Scott Stornetta. "How to time-stamp a digital document". In: *Conference on the Theory and Application of Cryptography*. Springer. 1990, pp. 437–455.
- [20] The White House. *FACT SHEET: Ongoing Public U.S. Efforts to Counter Ransomware*. 2021. URL: <https://www.whitehouse.gov/briefing-room/statements-releases/2021/10/13/fact-sheet-ongoing-public-u-s-efforts-to-counter-ransomware/>.
- [21] Eric Hughes. *A Cypherpunk's Manifesto*. URL: <https://www.activism.net/cypherpunk/manifesto.html>.
- [22] Internal Revenue Service (IRS). *United States forfeits millions in cryptocurrency used to launder illicit dark web proceeds*. <https://www.irs.gov/compliance/criminal-investigation/united-states-forfeits-millions-in-cryptocurrency-used-to-launder-illicit-dark-web-proceeds>. 2022.

- [23] Harry Kalodner et al. “BlockSci: Design and applications of a blockchain analysis platform”. In: *USENIX Security Symposium*. 2020.
- [24] Neal Koblitz. “Elliptic curve cryptosystems”. In: *Mathematics of computation* 48.177 (1987), pp. 203–209.
- [25] Sarah Meiklejohn et al. “A fistful of Bitcoins: characterizing payments among men with no names”. In: *ACM IMC*. 2013.
- [26] Ralph C Merkle. “A digital signature based on a conventional encryption function”. In: *Conference on the theory and application of cryptographic techniques*. Springer. 1987, pp. 369–378.
- [27] Victor S Miller. “Use of elliptic curves in cryptography”. In: *Conference on the theory and application of cryptographic techniques*. Springer. 1985, pp. 417–426.
- [28] Satoshi Nakamoto. “Bitcoin whitepaper”. In: (2008).
- [29] Satoshi Nakamoto. “Bitcoin: A Peer-to-Peer Electronic Cash System”. 2008. DOI: [10.1007/s10838-008-9062-0](https://doi.org/10.1007/s10838-008-9062-0). URL: <https://bitcoin.org/bitcoin.pdf>.
- [30] Kris Oosthoek, Jack Cable, and Georgios Smaragdakis. “A Tale of Two Markets: Investigating the Ransomware Payments Economy”. In: *Communications of the ACM*, [to appear], pre-print available at <https://arxiv.org/abs/2205.05028> (2022).
- [31] Kris Oosthoek, Mark van Staalduinen, and Georgios Smaragdakis. “Quantifying Illicit Revenue of Dark Web Shops”. In: *submitted* (2022).
- [32] The Tor Project. *Tor Project | Anonymity Online*. <https://www.torproject.org/>.
- [33] Ronald L Rivest, Adi Shamir, and Leonard Adleman. “A method for obtaining digital signatures and public-key cryptosystems”. In: *Communications of the ACM* 21.2 (1978), pp. 120–126.
- [34] Nick Szabo. *Bit gold*. 2008.
- [35] UN Security Council. *S/2019/171: Report of the Panel of Experts established pursuant to resolution 1874 (2009)*. Tech. rep. 2019.
- [36] United Nations Office on Drugs and Crime (UNODC). *2021 World Drug Report - Global Overview Drug Demand Drug Supply*. https://www.unodc.org/res/wdr2021/field/WDR21_Booklet_2.pdf. 2021.
- [37] Official Bitcoin Wiki. *Blockchain attacks on privacy*. <https://en.bitcoin.it/wiki/Privacy>. 2021.

2

CYBER SECURITY THREATS TO BITCOIN EXCHANGES

Bitcoin is gaining traction as an alternative store of value. Its market capitalization transcends all other cryptocurrencies in the market. But its high monetary value also makes it an attractive target to cyber criminal actors. Hacking campaigns usually target an ecosystem's weakest points. In Bitcoin, the exchange platforms are one of them. Each exchange breach is a threat not only to direct victims, but to the credibility of Bitcoin's entire ecosystem. Based on an extensive analysis of 36 breaches of Bitcoin exchanges, we show the attack patterns used to exploit Bitcoin exchange platforms using an industry standard for reporting intelligence on cyber security breaches. Based on this we are able to provide an overview of the most common attack vectors, showing that all except three hacks were possible due to relatively lax security. We show that while the security regimen of Bitcoin exchanges is subpar compared to other financial service providers, the use of stolen credentials, which does not require any hacking, is decreasing. We also show that the amount of BTC taken during a breach is decreasing, as well as the exchanges that terminate after being breached. Furthermore we show that overall security posture has improved, but still has major flaws. To discover adversarial methods post-breach, we have analyzed two cases of BTC laundering. Through this analysis we provide insight into how exchange platforms with lax cyber security even further increase the intermediary risk introduced by them into the Bitcoin ecosystem.

This chapter has been published as *Cyber Security Threats to Bitcoin Exchanges: Adversary Exploitation and Laundering Techniques* by **Oosthoek, K.** and Doerr, C. in IEEE Transactions on Network and Service Management, vol. 18, no. 2, pp. 1616-1628, June 2021.

2.1. INTRODUCTION

With an average market capitalization of 136 billion USD over the last two years [17], Bitcoin transcends all other currencies in the cryptocurrency market space. Similar to other currencies, security is a critical property in securing its role as a store of value, unit of account and means of exchange. Owners have to be confident that they won't lose their funds or they will withdraw them. While the developers of Bitcoin's reference implementation, Bitcoin Core, acknowledge that certain attack vectors exist [6], their probability is low as long as honest Bitcoin nodes together control more processing power than any group of attacker nodes [59]. Due to its implementation of a stack of cryptographic technologies, Bitcoin is a safe and reliable digital currency in its core. This paper will not review fundamental attacks on Bitcoin's distributed ledger technology, but considers another type of attack that has been proven most lucrative and continues to be.

A high value asset makes a high value attack target. The security of Bitcoin is also dependent on the ecosystem that has emerged around it. This consists of exchange platforms, payment service providers, wallet providers, mining pools and other intermediaries. Each of these is part of a fabric spun around Bitcoin which unlocks its potential to a broader user base, but consequently introduces additional threat vectors.

Cyber criminal actors generally target the weakest points in the ecosystem. In the Bitcoin ecosystem, centralized exchanges make up a large part of these. These act as a broker, allowing users to sell cryptocurrencies for fiat currency (legal tender) or to exchange the latter for cryptocurrency against a commission. Attacks on their platforms are feasible because in contrast to conventional stock exchanges, they do store currencies traded or exchanged by their clients.

This contradicts the original Bitcoin proposition as a decentralized currency, in which ownership depends on knowledge of the public-private key pair. The keys are the money: "not your keys, not your Bitcoin" [2]. However many owners deposit their Bitcoin with the exchange, which acts as a custodian. Although storing Bitcoin with an intermediary is a compelling offer for users such as active traders requiring quick and easy access to their funds, it creates a false sense of security to users less informed about the security aspects of Bitcoin ownership. Control of funds and thus the exercise of ownership is outsourced to or centralized at the exchange. While legal ownership is non-transferable, the public-private key pair that implies ownership of BTC remains with the exchange. According to recent reporting, the biggest exchange holds 966,230 BTC in custody, worth 7.19 billion USD at the moment of writing [15]. Exchanges must implement bank-level security to avoid successful cyber attacks and to safeguard funds, but have failed to do so as proven by the breaches in our analysis.

Bitcoin that are not being actively traded should be stored in cold storage. The hardware wallet is the best-known and most end-user friendly solution for cold storage. While the transfer of user funds to cold storage is a security best practice for exchanges, they regularly have funds available in hot storage in order to provide for quick exchange or withdrawal by legitimate users. With hot wallets being directly connected to the Internet and running as an ongoing process to rapidly meet liquidity requirements, they introduce the risk of exchange platforms losing BTC through exploitation of unknown vulnerabilities in their infrastructure. Cyber security is not top of mind in the development process of many start-up technology companies, which most centralized exchange

platforms are. This has resulted in frequent reports of client funds getting lost due to breaches. According to a March 2019 report from the United Nations Security Council, cryptocurrency exchanges are even targeted by sophisticated nation-state hacking groups in order to fund nuclear weapons programs [75].

With BTC market capitalization growing over time, attention to exchange platform security must grow in importance. Each incident potentially not only has a monetary impact, but potentially affects Bitcoin's credibility as a monetary asset.

This paper is an invited extended version of a paper presented at the 2020 IEEE International Conference on Blockchain and Cryptocurrency [61]. For this paper we have extended our work with an analysis of the security posture of the exchange platforms in our dataset that are still active. In addition to that we also analyze how the Lazarus group and the actor group behind the Bitfinex breach are laundering stolen BTC. Even during our analysis, more than 4 years after the Bitfinex breach, transactions with wallets linked to the hack were still observed.

Our systematic study of Bitcoin exchange breaches provides the following take-aways and contributions:

- We show that most Bitcoin exchanges were breached through relatively straightforward attack vectors.
- We found that while attack vectors overlap with breaches of other financial service providers, the actual exfiltration of funds is unique to Bitcoin exchanges.
- We demonstrate that over recent years the sophistication of the vectors used to breach exchanges has increased.
- We found that while the amount of BTC stolen per breach tends to decrease, the USD yield is higher due to an increased BTC-USD exchange rate.
- We demonstrate that the age of breached exchanges has increased in recent years.
- We show that over recent years more exchanges tend to survive after a breach, but details on the attack vector used are shared decreasingly.
- We demonstrate that while security has improved over time, some platforms still have relative lax web security when held against standards such as OWASP.
- We provide insight into adversary methods to launder stolen BTC through the blockchain and that the conversion of BTC to fiat money has become more complex.

The remainder of this paper is structured as follows: Section 2 provides an overview of related work on Bitcoin exchange security. Section 3 provides an overview of the Cyber Threat Intelligence field and the Vocabulary for Event Recording and Incident Sharing. Section 4 describes the methodology of our analysis. Section 5 presents the results from our analysis. Section 6 provides an overview of the web security posture of exchange platforms. Section 7 describes adversary laundering techniques post-breach. Section 8 outlines the limitations of our research. Section 9 summarizes our findings.

2.2. RELATED WORK

Several authors have focused on theoretical attacks on the Bitcoin network. The extensive research of Conti et al. has delivered a reference article on security and privacy concerns regarding Bitcoin. Their article focuses on various attack types such as double spending, Finney, brute force, Vector 76 and Goldfinger attacks [37]. They also cover the various countermeasures for these attacks. Lim et al. have also focused on security threats to Bitcoin such as DDoS attacks against exchanges, Bitcoin mining malware and extortion [56]. Feder et al. have looked at the impact of DDoS attacks on the now defunct Mt. Gox exchange. They found that on days following DDoS attacks, trading volume significantly decreased, specifically caused by a drop in large volume trades [44].

With regards to the risks introduced by Bitcoin exchanges in particular, Moore et al. have looked at various risk factors that have influenced the closure of Bitcoin exchanges between 2010 and 2015 [58]. They found that nearly half of the exchanges in their dataset have closed due to fraud attempts and security breaches. While they mention Bitcoin exchanges as the scope of their dataset, their analysis does also include services that did not support Bitcoin, e.g. Ripple. Their analysis is particularly useful, as they have analyzed the relationship between the presence of security-related features such as multiple factor authentication, bug bounties and exchange closure. Their dataset is however a bit outdated. They also have more of an economic focus on the exchange ecosystem and focus less on the actual security problems through which breaches have occurred. With regards to the cyber threats to Bitcoin exchanges specifically, several online resources that provide unstructured overviews of breaches exist. [66, 49].

Various authors have focused on the topics Bitcoin exchanges and Bitcoin security independently of each other. However we did not find any peer-reviewed contributions on the cyber security of exchange platforms. As far as we are aware, any significant academic analysis of a corpus of Bitcoin exchange breaches has not been performed, which we deem the main contribution of our research.

2.3. CYBER THREAT INTELLIGENCE

Cyber Threat Intelligence (CTI) is an umbrella term for the analysis of cyber security breaches and their tools, tactics and procedures (TTPs). It aims to provide actionable information to drive cyber security decision-making in order to avoid getting attacked with TTPs that were already disclosed. As there are many cyber threats around which are not all relevant to each organization, CTI aims to provide an understanding of the threats relevant to an organization and its assets. In this paper we focus on cyber threats to Bitcoin exchanges.

The CTI process strives to gain an information advantage on adversarial events to an organization's information systems. Threats are real if they are able to successfully exploit a vulnerability, leading to a normally negative real-world impact. The malicious actor needs to have the capability and opportunity to exploit that vulnerability and the intent to do bad things. Commonly heard attack vectors like ransomware, Denial of Service attacks, SQL injection and phishing can have a different impact for each individual organization as these depend on particularities specific to their technical environment.

Several frameworks to understand cyber threats in context exist, such as STRIDE,

CAPEC, ATT&CK and VERIS. They each have their own distinct use case. STRIDE, a mnemonic for Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege, is useful to understand how threats can impact an information system in several ways [68]. CAPEC is useful for analysis of software exploit methods [71], whereas ATT&CK is proven to be useful to inform security analytics and understand malware trends [62]. For our analysis we have used VERIS, which is primarily useful for post-breach assessments.

2.3.1. VOCABULARY FOR EVENT RECORDING AND INCIDENT SHARING (VERIS)

The Vocabulary for Event Recording and Incident Sharing (VERIS) is a CTI standard open-sourced by Verizon [78]. Of all efforts that exist to systematize the conversation on cyber security and exchange information, VERIS has become the industry standard for strategic CTI. It is targeted towards strategic CTI as it is meant for reporting that informs strategic, longer-term decision making to prioritize security investments based on risk appetite. Four indicators present in every cyber security incident form the basis of how VERIS is structured: the *Action* used to breach the asset, the *Actor* who breached the asset, the compromised *Asset*, the security *Attribute* (confidentiality, integrity or availability) that was affected.

Analysis of TTPs in a set of cyber security breaches can provide an understanding of how attackers target an industry such as Bitcoin exchanges. VERIS provides for structured analysis as it translates the narrative of individual incidents into a structured form. The most well-known example of VERIS in use on a multiple-industry level is the annual Verizon Data Breach Report (DBIR), which has become the industry-standard reference for intelligence on developments in the cyber threat landscape [76]. Breach data from a large body of industry and public sector organizations is used as the source for the report. Publicly disclosed breaches are also recorded in the VERIS Community Database, available on Github [73].

2.4. METHODOLOGY

In this section the methods applied in the collection of our dataset of Bitcoin exchange breaches and their successive classification using VERIS are discussed. In this paper we use *Bitcoin* to refer to the Bitcoin distributed ledger and technology stack. *BTC* is used to designate units of account, e.g. when referring to the amount stolen in a particular breach.

2.4.1. DATA COLLECTION

EXCHANGE BREACHES

We have gathered our dataset in November 2019 using Google Custom Search JSON API queries for *bitcoin exchange breaches* and *bitcoin exchange thefts*. Based on word frequency analysis, we identified 36 incidents of breached exchange platforms, which were cross-checked against media reporting. Our analysis is concentrated on Bitcoin exchanges. Based on our criteria we only included technical *security* breaches of exchanges which focus either on *Bitcoin* trading exclusively or combined with other cryptocurrencies. Exchanges not supporting Bitcoin and exchanges without reported breaches are

not included in our dataset. In some breaches of multi-currency platforms, other currencies than Bitcoin were stolen as well. In those cases we include the amount of BTC stolen according to official reports. Our dataset does not include any decentralized exchanges for peer-to-peer trading, hash-power marketplaces or online wallet services.

FINANCIAL SERVICES BREACHES

In order to compare Bitcoin exchange breaches with breaches of other financial service providers such as banks, we have used the VERIS Community Database (VCDB). This public dataset includes VERIS-formatted, annotated reports of publicly disclosed cyber security breaches in various industries. It is audited by the VERIS Risk team at Verizon and also used as input for their annual report. At the time of writing the database includes 8346 incidents, updated daily. Based on the VERIS taxonomy, it allows to filter for data breaches that occurred with organizations offering financial services. We have used the JSON objects of *validated* incidents, which are manually checked for validity by Verizon [77]. The VCDB captures incidents recorded from 2012 and is thus aligned with the time period covered by our dataset of exchange breaches, allowing for a uniform comparison. We have checked whether incidents from our exchange breach dataset are recorded in VCDB, however no overlap existed.

TRADE VOLUMES

We have gathered data on daily exchange trade volumes from a public API offered by CoinGecko [36]. While CoinGecko is one of the few overview websites that normalizes data in order to account for exchanges reporting fake volume, currently no publicly available dataset exists that fully accounts for exchanges reporting fake volume data. This is a known problem of the exchange ecosystem [15]. For this reason, this data was only used to analyze post-breach impact as reported by exchanges in B5 of section VI.

2.4.2. CLASSIFICATION OF BREACHES

For our analysis we have focused on the Action category of VERIS. The VERIS taxonomy also has an Actor category, but rarely are breaches of Bitcoin exchange platforms attributed to designated actor groups. The Asset category is not used because for each exchange breach, the Server asset would qualify as the platforms in our dataset are online outlets exclusively. In case of another financial service provider like a bank, the breached asset can also be an ATM for example. The Attribute category affected would always be Confidentiality and Integrity, as we have not recorded Denial of Service attacks affecting platform availability.

The analysis of Bitcoin exchange breaches has proven to be onerous as the sharing of information tends to be quite scarce and is getting even more scarce over recent years. For our analysis, security breaches were included in our dataset according to the criteria by Verizon. The entry must be a confirmed security incident, with a loss of confidentiality, integrity, or availability [76]. In the case of our analysis, we also chose to only include officially disclosed breaches, meaning they were announced through official communication channels maintained by the particular exchange. Press releases, but also messages from the official Twitter channel or posts on *bitcointalk.org* from confirmed accounts of exchange staff. We provide references to each source.

Breaches were classified to a threat action category and variety according to the information we had on the *initial* point and means of entry, as this provides for an overview of the attack surfaces of Bitcoin exchanges. As in some cases official sources only reported a successful hacking attempt and lack further detail, we have not identified the sub-variety of Hacking to stay close to the official incident report. Also, in some cases the amount of BTC stolen is not reported. In other cases, a cyber security breach is the official account, but heavy rumors about a cover-up such as an exit scam exist. Because we want our analysis to be a valid but accurate reflection of the current state of the Bitcoin exchange ecosystem, we have included this in *italic* in the Attack Method column of Table 1. We refer to the URLs used for the coding of each breach.

2.5. ANALYSIS OF BITCOIN EXCHANGE BREACHES

In this section we will discuss the observations from our dataset of Bitcoin exchange breaches. We have analyzed 36 incidents of breaches, through which at least 1,156,399 BTC were stolen from their legitimate owners. Table 1 and Figure 1 provide additional insight into the dataset and our analysis.

Figure 1 is a bubble chart representation of the dataset. Exchange breaches are plotted on the X axis by year of compromise. The Y axis indicates the age of the exchange at the time of the incident. Each bubble represents a breach, whereas the line color represents the attack vector used and the bubble diameter the amount of BTC stolen. Figure 1 shows two interesting patterns in particular. The amount of BTC stolen in breaches has decreased in recent years. It also shows that the age at which an exchange gets breached has increased. Furthermore the figure shows that the TTPs deployed in breaches of exchange platforms have developed from trivial exploitation of functionality or vulnerabilities to other hacking vectors.

Table 1 provides an overview of the breaches recorded in our dataset. The *Launch* and *Breach* columns denote when an exchange was first opened and breached subsequently, *BTC* how much BTC were lost, *USD* the loss in USD based on the average exchange rate in the breach month [52], *Action* and *Variety* how the attack vector used is classified within VERIS. In the *Attack* column, we have placed references to the sources used for the VERIS classification of each breach and further analysis in this section. The information in this column is based on reporting on forensic investigation by the breached party. When not available, we have drawn on reports from secondary sources such as media, emphasized in *italic*. *Closed* denotes whether an exchange closed as result of a breach. The asterisk indicates termination after a subsequent breach.

The most occurring varieties in our dataset of Bitcoin exchange breaches are: Unknown (12), Use of stolen credentials (6) and Abuse of functionality (5). In the sections below we will discuss the observations for these breach varieties in more detail, as well as our findings with regards to their impact.

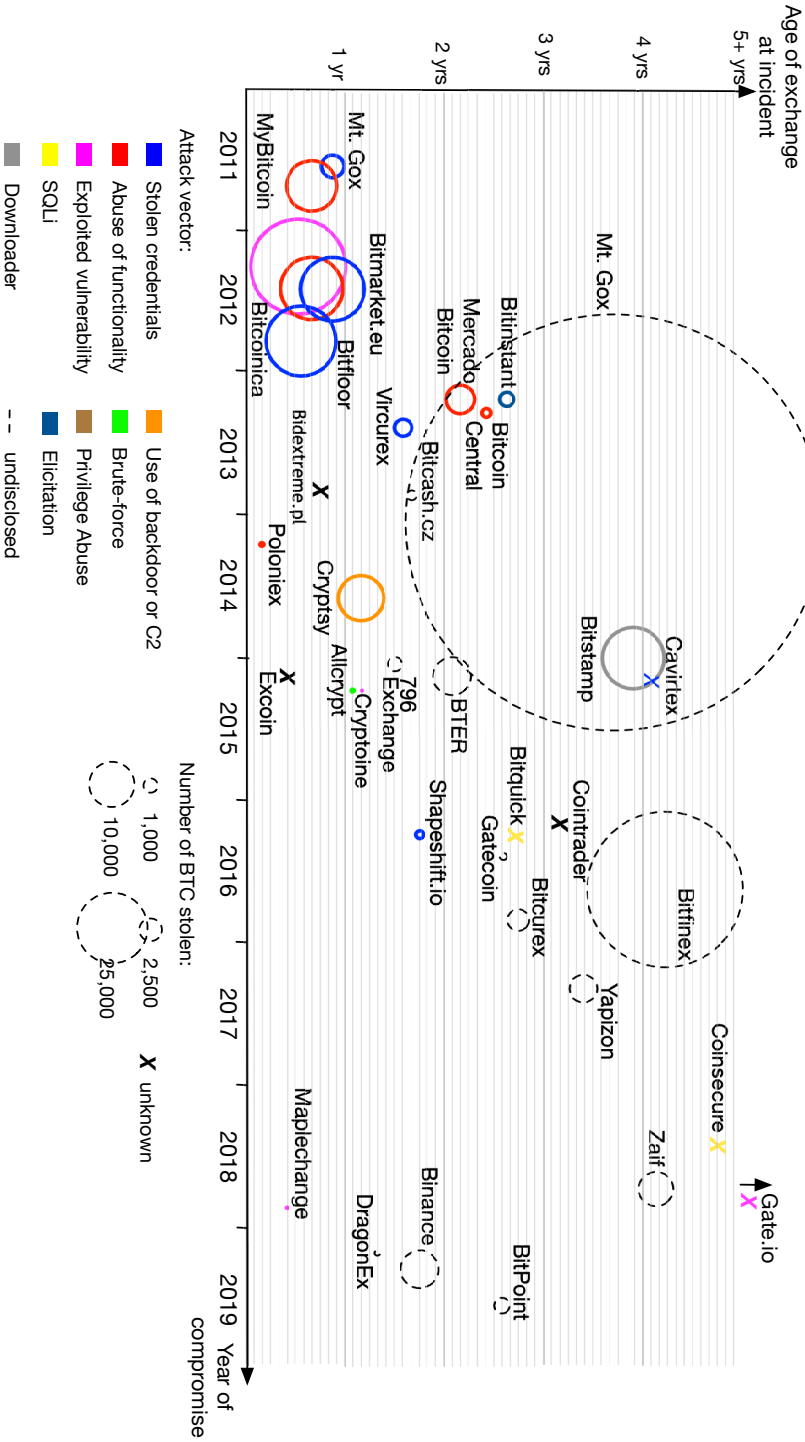


Figure 2.1: Timeline of Bitcoin exchange breaches with respective attack vectors and BTC stolen

2.5.1. ANALYSIS OF ATTACK VECTORS

INCREASE OF UNKNOWN VARIETY INDICATES DECREASE IN DISCLOSURE OF BREACH DETAILS

While we were able to identify the attack vector for most of the breaches in the first half of the time frame covered by our dataset, in recent years the communication of the TTPs used to breach exchanges has gotten more fuzzy. Of the 36 incidents recorded in our dataset, the specific attack vector remains unknown in 15 cases. Of these 15 cases, 9 appear in the last three years. In none of the latter cases, details on the vector through which the exchange was breached was not publicly disclosed.

While over recent years less exchanges terminate after a breach and theoretically better able to provide better incident response information, they often tend not to. While in the early years full details were usually not available due to the exchange going out of business (from 2011 until 2015, only 13.3% of breach methods remained undisclosed), over the last years exchanges survive but do seem to share less details out of concerns for their reputation. It is known that for various reasons, organizations are hesitant to share details on security breaches [65]. Other financial service organizations tend to exert more openness about breaches. One example of this is the hacking operation of Bangladesh Bank, which was disclosed officially by SWIFT [69], as well cyber security vendors [55]. This is partially explained by the fact that traditional financial institutions are subject to strict breach notification regulations such as the Federal Information Security Management Act of 2002 and the European General Data Protection Regulation. However, based on the information recorded in Table 1, a trend of exchanges getting less transparent over the last couple of years can be observed.

DECREASE IN USE OF STOLEN CREDENTIALS VARIETY

The use of stolen credentials (*Stolen Creds* in VERIS terminology) is the method of choice in most of the early attacks in our dataset. In this type of attack the malicious actor breached the exchange platform and consequently exfiltrated funds by using privileged credentials. In many cases, credentials providing elevated (administrator) privileges were obtained through relatively low-level social engineering or unsafely stored. While this type of breach does not involve exploitation of a vulnerability or another form of abuse, it is a cyber security breach because it affects system integrity.

Six exchanges in our dataset were breached through the use of stolen credentials, all of which occurred from 2011 up until 2016. In June 2011 Mt. Gox was breached with a compromised administrator account. More than 24,000 BTC were stolen from Bitfloor after the attacker managed to obtain credentials from their cloud provider to gain access to an unencrypted backup of a wallet used for cold storage. One explanation for the feasibility of this type attack is password reuse, because end users recycle the same password or variants of the same password through multiple online services.

After Unknown hacking, the use of stolen credentials is the biggest hacking vector in our dataset of Bitcoin exchange breaches (17%). The same goes for other financial services recorded in the VCDB, in which for 24 of a total 158 incidents this vector is employed (15%). Based on this data, both verticals are targeted and consequently exploited using the same methodology. However, where in traditional financial services stolen *user credentials* are mostly used to steal funds from individual users, the exchanges in

Table 2.1: Bitcoin Exchange Breaches

Launch	Breach	Exchange	BTC	USD	Action	Variety	Attack method	Closed
2010-07	2011-06	Mt. Gox	2,500	40,250	Hacking	Use of stolen creds	admin account breach [54]	yes*
2010-04	2011-08	MyBitcoin	12,500	102,500	Hacking	Abuse of functionality	programming error [80]	yes
2011-09	2012-03	Bitcoinea	43,554	213,415	Hacking	Exploit vuln	Linode breach [57]	yes*
2011-09	2012-05	Bitcoinea	18,547	96,444	Hacking	Abuse of functionality	Rackspace PW recovery [9]	yes
2011-04	2012-05	BitMarket.eu	19,980	103,896	Hacking	Use of stolen creds	SSH account hacked [10]	yes
2012-02	2012-09	Bitfloor	24,086	298,666	Hacking	Use of stolen creds	unencrypted wallet backup [72]	yes
2011-08	2013-03	BitInstant	999	92,907	Social	Elicitation	DNS hijack, registrar social eng. [14]	yes
2011-01	2013-03	Mercado Bitcoin	4000	372,000	Hacking	Abuse of functionality	coupon functionality hijacked [24]	no
2010-12	2013-04	Bitcoin Central	'few 100'	-	Hacking	Abuse of functionality	account takeover, OVH breach [7]	no
2011-10	2013-05	Vircurex	1,454	187,275	Hacking	Use of stolen creds	PW reset, cloud host social eng. [79]	yes
2011-07	2013-11	Bitcash.cz	485	584,765	Hacking	Unknown	web interface compromised [31]	yes
2013-01	2013-11	Bidextreme.pl	-	-	Hacking	Unknown	[33]	yes
2010-07	2014-02	Mt. Gox	850,000	487,815,000	-	-	insider involvement [41]	yes
2014-01	2014-03	Poloniex	97	43,136	Hacking	Abuse of functionality	race condition [20]	no
2013-03	2014-07	Cryptsy	10,000	5,895,000	Hacking	Use of backdoor or C2	backdoor in dependency [40]	yes
2011-01	2015-01	Bitstamp	18,866	4,122,221	Malware	Downloader	sophisticated malware attack [32]	no
2013-06	2015-01	796 Exchange	1,000	218,500	Hacking	Unknown	compromised "certain weakness" [25]	no
2013-01	2015-02	BTER	7,170	1,821,897	Hacking	Unknown	breach of cold wallet [32]	yes
2011-06	2015-02	Cavirtex	-	-	Hacking	Use of stolen creds	PW hashes, 2FA secrets exposed [23]	no
2014-10	2015-02	Excoin	-	-	Hacking	Unknown	[43]	yes
2014-02	2015-03	AllCrypt	40	9,764	Hacking	Brute force	bruteforced tech staff email [1]	yes
2014-01	2015-03	Cryptoine	6	1,465	Malware	Exploit vuln	race condition in trading engine [81]	yes
2013-01	2016-03	Cointrader	-	-	-	-	[27]	yes
2013-07	2016-04	BitQuick	-	-	Hacking	SQLi	upload feature SQL injection [8]	no
2014-07	2016-04	ShapeShift.io	315	141,278	Hacking	Use of stolen creds	insider involvement [67]	no
2013-07	2016-05	Gatecoin	250	132,225	Hacking	Unknown	multisig cold wallet [47]	yes
2012-01	2016-08	Bitfinex	120,000	68,868,000	Hacking	Unknown	[12]	no
2012-01	2016-11	Bitcurex	2,300	1,707,750	Hacking	Unknown	API signing key exploit [12]	yes
2013-01	2017-04	Yapizon	3,816	5,158,850	Hacking	Unknown	[26]	no
2013-06	2018-04	Coinsecure	438	4,049,354	Misuse	Privilege Abuse	insider, cold storage exposed [70]	yes
2014-06	2018-09	Zaif	5,966	39,585,603	Hacking	Unknown	3 hot wallets hacked [63, 30]	no
2018-05	2018-10	Maplechange	8	50,927	Malware	Exploit vuln	race condition, exit scam [11]	yes
2013-07	2018-11	Gate.io	-	-	Malware	Exploit vuln	supply chain attack [46]	no
2017-11	2019-03	DragonEx	135	553,811	Hacking	Unknown	[34]	no
2017-07	2019-05	Binance	7,000	59,908,100	Hacking	Unknown	API keys and 2FA secrets [5]	no
2016-01	2019-07	BitPoint	1,225	12,350,450	Hacking	Unknown	[28]	no

* the asterisk in the Closed column indicates closure after a subsequent breach

our dataset were breached with *administrative credentials*, which provide instant access to funds of multiple users. The fact that the use of stolen credentials is decreasing over recent years indicates that exchanges have increased their security hygiene.

DECREASE IN ABUSE OF FUNCTIONALITY VARIETY

Abuse of functionality was the attack vector of choice in 5 breaches. Just like *Use of stolen creds*, this method does exploit a platform's access mechanisms. However, rather than the exploitation of a technical vulnerability, the attacker abuses legitimate platform functionality of the exchange platform or its hosting partner. Examples of these are the use of a flawed password recovery or discount modules, which can generally be avoided by thorough unit testing.

This vector was dominant among breaches of early movers in the exchange ecosystem, breached between 2011 and 2014. At the same time this observation is in accordance with a general trend observed in cyber attacks. Over the past years, attackers tend not to deploy malware or custom exploits. If not necessary to accomplish their objectives, they prefer to use functionality native to a target system. This way they are "living off the land" (LOTL). While LOTL attacks are employed both by low-level and sophisticated actors, their feasibility by the misuse of native features usually implies lax monitoring or security audits at the side of the victim. Our dataset records 5 data breaches through abuse of functionality, which is 14% of the total. In VCDB, this is recorded in just 2 incidents, which is only 1.29% of incidents recorded for the financial services industry.

RELATIVELY LIMITED DEPLOYMENT OF ADVANCED METHODS

As discussed in earlier sections, most exchanges in our dataset were breached using relatively straightforward attack vectors. Only three exchanges were compromised through the use of advanced techniques. Cryptsy was breached due to the exploitation of an intentionally placed backdoor in an open-source software dependency. After a malicious actor took over ownership of the development of Lucky7Coin, he was able to place an IRC backdoor into the wallet code base, allowing full and unlimited access to funds stored in the wallet [40]. Gate.io was breached due to a breach at Statcounter, which allowed attackers to place code in the visitor counting script used by Gate.io. Both were targeted attacks, as they were the only instance in which this specific vulnerability in the dependency was exploited. Furthermore the breach of Bitstamp in January 2015 involved multi-staged and targeted malware according to leaked post-mortem reporting. Apart from these cases, the exchanges in our dataset were hacked with relatively straightforward vectors. Especially given the considerable financial impact, this is a characteristic unique to Bitcoin exchanges. If the breach methods are not advanced, it implies the level of technical security is very low. And if security of an exchange platform is low, the company will not be able to keep up against the sophisticated nation state actors by which they are targeted, as mentioned in the introduction.

2.5.2. ANALYSIS OF IMPACT

SAME VECTORS, DIFFERENT OUTCOMES

In the sections above we have found that the vectors used to breach Bitcoin exchanges are similar to those targeting financial institutions. The real world outcomes are however very different. Where attackers targeting Bitcoin exchanges are always motivated to

exfiltrate funds, this is less the case in breaches of traditional financial institutions. According to the DBIR 2019, in 43% of incidents personally identifying information (PII) is exfiltrated and credentials in 38% of breaches [76]. According to the same report, theft of funds mostly happens to physical tampering attacks against automated teller machine (ATMs), which have declined over the last couple of years. Although the breach TTPs overlap, results and implications of breaches hugely differ between these types of organizations.

HOT WALLETS REMAIN THE WEAK SPOT

Exchanges use so-called hot and cold wallets like storage providers differentiate between hot and cold storage. Hot wallets provide liquidity to quickly facilitate transactions that characterize an exchange; depositing and withdrawal of funds to convert fiat to digital assets or exchange between BTC and ERC-20 tokens. The amount of assets stored in hot wallets should be sufficient to provide quick liquidity. The largest share of user funds held in custody should be held in cold wallets, which are stored offline or airgapped, meaning isolated from the regular local network and outside networks and ideally requiring physical access.

Except for 2 cases, in all breaches the funds exfiltrated by the hackers were stored in a hot wallet. Because these wallets are connected to the Internet, private keys can be obtained by breaching the server on which the wallet is stored. Only BTER [29] and Coinsecure [70] reportedly had their cold storage breached. Storing only as much funds as necessary in hot storage is considered good cyber security hygiene, as the offline nature of cold wallets makes them more difficult to breach. This potential financial impact of a breach is significantly decreased if the attacker is only able to compromise hot storage with a constrained amount of funds required to meet liquidity needs.

Our dataset shows that hot storage only provides security when implemented correctly. In the early years, exchanges went insolvent after a breach because they stored practically all funds in hot storage. Recently, regulatory frameworks imposing strict requirements on the custody of exchanges have been introduced in many jurisdictions. The Hong Kong Securities and Futures Commission is one of the first movers among global financial regulators with new regulation. As of November 2019, it requires Hong Kong-based platform operators to store 98% of assets in cold wallets and 2% in hot wallets. It also requires platforms to minimize the number of transactions from cold wallets and to insure funds to cover for a hack [50]. Furthermore, several leaders in the cryptocurrency ecosystem have established the Cryptocurrency Certification Consortium (C4), which has released the CryptoCurrency Security Standard [21]. This standard provides guidance to cryptocurrency companies to implement information security frameworks such as ISO 27001.

The above provides a representation of growing technical security maturity of Bitcoin exchange platforms. It does not require much technical sophistication to put an exchange platform online, as one can build on various widely-available open source components. However keeping such infrastructure secure takes significant resources and only a limited pool of people has the knowledge and experience of exchange security. The sophisticated threat actors targeting exchanges however investigate ample time to find a vulnerability that provides them with a foothold. And they only need one in order to further escalate their access level.

DECREASED EXCHANGE CLOSURE DUE TO BREACHES

Being breached was equal to insolvency and subsequently closing down in most of the first incidents recorded in our dataset of breached Bitcoin exchanges. However over recent years, this is not the case anymore. In most recent cases, exchanges resumed business after a period of ceased trading post-breach. This is only partially good news for owners of Bitcoin. In some cases stolen BTC were reimbursed on a 1:1 basis (Binance, BitPoint), but in other cases the exchange refunded a fixed percentage of BTC (Zaif, DragonEx). More controversial is the issuing of “IOU” (“I Owe You”) tokens by Bitfinex and Yapizon, as these tokens serve as non-negotiable, informal measures of debt and thus are not redeemable for the actual value lost in BTC or USD.

In all cases, it is an improvement that customers are not left absolutely empty-handed after a breach. Our dataset includes exchanges paying out of pocket for this such as Binance’s user asset fund [4], as well as selling the company in order to raise enough money [35]. Moore et al. [58] have argued that high-volume exchanges have a better chance to continue operations after a breach. This is interesting, as the bigger exchanges might have deeper pockets for security spending and thus to fend off cyber attacks. This would drive rational customers to large platforms, which indeed shows both in the trading volume (CoinGecko) and the amount of BTC held in custody [15]. This is however also contrary to the decentralized philosophy described by Satoshi Nakamoto in the Bitcoin whitepaper.

AMOUNT OF BTC SEIZED PER BREACH IS DECREASING, RELATIVE USD YIELD INCREASING

As it can be observed from Figure 1 and Table 1, over the last few years the relative amount of BTC stolen as part of exchange security breaches tends to decrease. Seizures exceeding 10,000 BTC were not uncommon in the early years, however - except for some outliers - this has decreased in 2018 and 2019.

Over recent years the relative USD yield has increased, while the BTC yield has decreased. This is a result of the increased BTC-USD exchange rate. Table 1 also shows that in the early years of our dataset, big amounts of BTC were taken through relatively low-level hacking vectors. Over recent years attack vectors have become more complex and less lucrative when simply considering BTC quantity. This could be the result of improved security best practices by exchange platforms, as well as improved incident response practices. Temporarily terminating withdrawal and trading functionality post-breach has become a common practice, as well as requesting other exchanges to which the hackers diverted BTC to freeze those.

NO IMPACT ON TRADING POST-BREACH

We have also found that in 4 recent breaches, the trading activity on the platform was not impacted. Figure 2 shows the reported trading volume 60 days before and 60 days after breach disclosure.

According to a report by Bitwise for the U.S. Securities and Exchange Commission [15], most Bitcoin exchange platforms operate wash trades or report fake inflated volume in order to increase attractiveness and exposure on market overview websites. This results in many cryptocurrency market overview websites reporting fake volume due to fabricated input data.

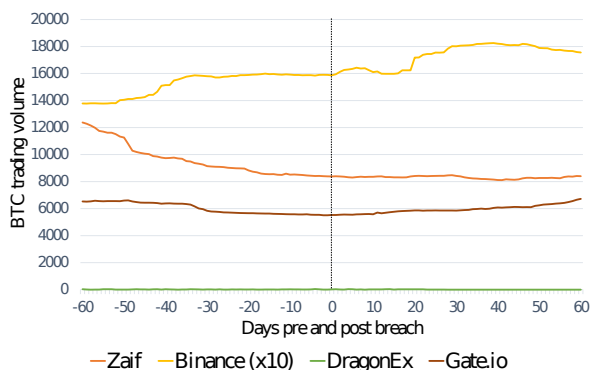


Figure 2.2: Trading volume in BTC 60 days before and after breach.

Historical trading volume data is not available for all exchange platforms, certainly not for those out of business. We have used data available from CoinGecko for exchanges which did not close after a breach as only then data is available. Although we had only BTC trade volume data available from 4 exchange platforms (Zaif, Gate.io, DragonEx, Binance), none of them did show a substantial impact on trading volume post-breach as can be seen in Figure 2.

DragonEx is at the bottom of the diagram, as its daily trade volume is below 100 BTC. The above shows that security breaches do not affect the trade volumes of Bitcoin exchange platforms that continue operations. We have to take into account that some amount of the volume might be fake. However the sustained trading can be explained by the incident response communication by these exchanges. They didn't include much technical detail with regards to breach TTPs and every briefing assured customers that the exchange platform maintained control of the situation. This is coherent with the security industry best practice of frequent post-breach communication in order to avoid customer attrition [22].

2.6. EVOLUTION OF PLATFORM SECURITY

The sheer fact that all organizations in our dataset were vulnerable to attacks in the past, invites the question whether platform security has improved over time. Data on this is scarce as many platforms have suspended their activities. Therefore we have decided to focus on 13 platforms from our dataset which are still operating. We have used scan results from Censys, which continuously scans the entire IPv4 address space since June 2016 for ports used in common services.

Adversaries usually deploy such port scans to acquire information about exposed services and the software versions behind them in a search for potential vulnerabilities for exploitation. To understand what services a website was running at a given point in time, we used a historical passive DNS data source to resolve the domain names of these 13 platforms to the IP addresses used each day between June 2016 and December 2019, and correlated these with the banners grabbed from a particular IP address on that

day. Additionally we have requested the HTTP Response Headers of each platform's *root www*. These are reconnaissance methods which do not negatively impact the platforms in question. To further comply with responsible disclosure of security issues, we do not identify individual platforms.

Based on the aspects observed, we see server security of the platforms investigated has significantly increased from 2016. As highlighted below, the web security of these platforms has matured together with the overall exchange ecosystem, although bad security practices and security lapses are still comparatively widespread.

2.6.1. MULTIPLE SERVICES ON SINGLE IP

In our scans we have found multiple platforms running multiple services on the same IP address as the platform. This is not only bad practice due to availability and reliability, it also increases the attack surface as the potential pivot points for attackers are increased. If the mail relay is used to send spam emails, the website ends up being blacklisted as well. Over the time frame from June 2016 until December 2019 this applied to 4 platforms in total, of which 1 still exposes FTP (which sends user credentials in plain text) and mail services at the moment of writing.

2.6.2. VULNERABLE WEB SERVER VERSIONS

Given the fact that their webpage is the critical frontend to exchange platforms and thus the single point of failure for such platforms, we were surprised to see that so many have been running vulnerable, unpatched versions of web server software for long periods of time.

One of the bigger platforms has been running nginx 1.6.2 from at least June 2016 until February 2019, for which a high severity Denial of Service (DoS) vulnerability was reported already in February 2016 and a privilege escalation vulnerability in November 2016. Another high-volume platform was observed serving its platform's pages via Apache 2.2.22 from at least June 2016 until September 2017, which was then already vulnerable to several medium and high severity threats such as Cross-Site Scripting (XSS), buffer overflow, remote code execution, DoS and authentication bypass. In June 2019, one other platform had Microsoft Internet Information Server 8.5 exposed, for which a medium severity vulnerability was reported in 2014. Given a batch of only 13 platforms, observing 3 of them with major and easily exploited vulnerabilities post breach is astonishing and points to severe shortcomings in security practices such as vulnerability management.

2.6.3. EXPOSED MANAGEMENT INTERFACES

The servers operating the exchange also need to be maintained and updated. These service interfaces would normally not be exposed to the public on the same IP as the website itself, but rather be shielded in a separate compartment and accessible only to select sources. Already above we noted the existence of multiple services running on the platforms. From at least June 2016 until December 2019, one exchange had a Pure-FTPd service and Dovecot email server exposed. In September 2017, the same platform even exposed a vulnerable OpenSSH service version. In addition to configuration interfaces, another platform exposed a database server, MariaDB 5.5.5, a MySQL fork for 7 months

Table 2.2: Detected Improperly Configured Headers

HTTP Security Header	Protects Against	Count
X-Content-Type-Options	Phishing, Cross-Site Scripting	6 of 13
X-Frame-Options	Clickjacking	7 of 13
X-XSS-Protection	Cross-Site Scripting	6 of 13
Strict-Transport-Security	Man-in-the-Middle	6 of 13

over 2018 and 2019. In addition to normal architectural practices where databases would normally be placed deeper in the network and not accessible from a public network, this database was at that time already vulnerable to several high severity remote exploit vectors. Pure-FTPD was also exposed in 2016 and 2017 by another platform.

In principle, the attack surface could be minimized even when exposing such services to the outside, if mitigations such as IP address whitelisting can be applied. However residual attack surface still exists due to spoofing, and whitelisting was obviously not in place (or badly configured) if connections from an Internet scan service were permitted that would make the result available publicly. General security principles such as not to expose such services to untrusted zones like the Internet were thus not followed by the exchanges, even after breaches had happened.

2.6.4. SLOW PATCHING IN GENERAL

Based on our port scan data, it can be concluded that the deployment of security patches by exchange platforms has generally been slow. While over time less services are exposed in general, running vulnerable web services is bad practice for online-only organizations. However room for improvement still exists, as later in this section we will show that as of now one platform is still running a vulnerable scripting engine.

2.6.5. HTTP SECURITY HEADERS

When a web server answers a request from a browser it includes HTTP response headers. A particular category of these are Security Headers, which instruct a browser how to serve that particular page to avoid XSS and clickjacking. Implementing HTTP Security Headers is a straightforward and cost-effective measure, which protects against end users against malicious interventions. We have verified the configuration of the HTTP Security Headers by requesting the webpage via *curl*, which is a non-intrusive method to better understand an organization's recognition of web security. As shown in Table II, of the 13 platforms we investigated, we have found that three different headers were not properly configured for 6 out of the 13, and one type of header incorrectly for 7 out of the 13 platforms.

2.6.6. SERVER-SIDE SOFTWARE

The *X-Powered-By* header is a non-standard HTTP response header, which can be manually adjusted to distract scanners targeting specific vulnerable software versions by increasing the version number to appear as more updated. Decreasing is not common, as it will only increase scrutiny.

For one platform, we found it running PHP/5.3.29, which dates back to 2015 and for which 5 critical vulnerabilities were disclosed in 2019. One of these vulnerabilities can be exploited to cause a buffer overflow, allowing for the server-side execution of attacker-controlled code. Given the history of attacks discussed earlier, this is a serious threat to the integrity of an exchange platform's transactions.

We have notified the platforms in question about open vulnerabilities. Based on our analysis there are however positive things to note. Over the time frame of the scans, we observed more platforms moving behind Cloudflare's reverse proxy service, which protects against several web attacks. As of December 2019, all platforms except 2 have done this. Furthermore, less secondary services are getting exposed, although the practice still exists. While a causal relationship between security breaches which have hit the ecosystem so hard and the improved hygiene can obviously not be established, it seems that some maturing has happened in the overall ecosystem.

None of our findings are significant - any standard penetration test should uncover OWASP Top 10 vulnerabilities. We recommend exchange platforms not only to perform vulnerability scans and penetration tests on a regular basis, but to adhere to implement the OWASP Web Testing Guide in unit tests of their platform code. As simple technical vulnerabilities can threaten business continuity, over recent years many financial organizations have dedicated resources towards collecting threat intelligence on relevant attackers as part of their risk management and overall due diligence [60].

2.7. TRACING STOLEN FUNDS

A question often asked after exchange breaches is how and to where BTC were funneled out. Such intelligence on post-breach TTPs provides an integral view of adversary behavior, however the analysis is delicate due to the privacy aspects inherent to Bitcoin's architecture. Bulk transaction analysis becomes even more complex due to the use of CoinJoin, third-party mixing services and the identification of services where funds terminate into fiat currency, which makes it difficult to establish a ground truth. The analysis is additionally complicated by the fact that details on laundering TTPs only become truly apparent when actors are indicted. In this section, we will analyze the post-breach TTPs to launder stolen funds from the Bitcoin exchanges in our dataset, based on official information shared by victim organizations and prosecutors.

2.7.1. MT. GOX

The second breach of Mt. Gox in 2014 made the platform collapse and 850,000 BTC disappeared. According to US justice department filings from 2017, the Russian-owned exchange platforms BTC-e and Tradehill were used to launder a significant portion of the Mt. Gox funds [53]. According to the Financial Crimes Enforcement Network, BTC-e was used to launder criminal money from miscellaneous origin, such as proceeds from Cryptolocker and Locky ransomware campaigns [45]. BTC-e failed to maintain effective AML measures and did not pursue any form of KYC, essentially favoring money launderers. According to the documents, the actors behind BTC-e allegedly managed to launder 4 billion USD. Among BTC-e's clients was the Russian state hacking group Fancy Bear [42], which is known to have used BTC to sponsor their hacking campaigns [39]. Accord-



Figure 2.3: Force-clustered transactions (green) from wallets (yellow) associated with Bitfinex attack from August 2016 until mid-June 2020.

ing to reporting by the BBC, most of BTC-e's user base vanished to a successor platform called Wex, of which all funds were allegedly seized by the Russian security service FSB [3].

The breach of Mt. Gox and the subsequent laundering has triggered several big investigations, some of which are still active. However the laundering TTP of using questionable exchanges to launder stolen funds, has become obsolete due to the restrictions raised by AML, KYC and FATF regulations.

2.7.2. BITFINEX BREACH

With Bitfinex being a high-volume exchange, it would become the biggest breach after Mt. Gox. In June 2019 two Israeli individuals linked to the Bitfinex hack were arrested [82], but funds are still on the move. Many details on the adversary and its TTPs remain unclear. In August 2020 Bitfinex offered a 5% share of the assets recovered to anyone who puts the company in touch with the attackers and a 25% share to anyone who demonstrates control of the attacker wallets [13]. Just like with Mt. Gox, 4 years after the Bitfinex attack, the incident response is still ongoing.

A list of transactions and wallets associated with the hack was circulated by a Bitfinex director few days after the hack [64]. We have used this to analyze first and second-order movement of funds, in other words movement from wallets included in the list

shared by the victim organization. As adversaries are known to use CoinJoin wallets and third-party mixing services, tracing stolen funds on the blockchain is ambiguous and potentially unreliable. To keep analysis empirical, we focused on movement relating to addresses from the list distributed initially. In general, we have observed laundering the funds has been very laborious to the actors, with relatively limited success.

Figure 3 is a force-directed graph of outbound transactions from Bitfinex to the 410 wallets that have been active until now, based on the list shared by Bitfinex. A yellow dot represents a Bitcoin wallet, a green dot a mainnet transaction. The yellow dot in the center represents Bitfinex, surrounded by the wallets to which funds were exfiltrated as officially reported by Bitfinex. The main takeaway from the graph is how attackers diffuse funds to many wallets through a web of transactions. As the nature and ownership of these wallets can only be speculated, the graph shows how the attackers use the Bitcoin network to obscure movement of yields from the breach. The list shared by Bitfinex contains a total of 2072 transactions associated with the hack, totalling 119,755 BTC. These transactions all took place between 8:54:54 and 12:18:35 on 2 August 2016. In total, we have recorded 1001 transactions associated with the Bitfinex hack, from August 2016 until June 2nd, 2020. Based on our analysis of these transactions, we have observed the following:

Manual obfuscation: In January 2017, the actors collected funds from several small wallets into a single 93 BTC wallet. This was then funneled to a wallet with another 15 BTC of stolen funds and then split into smaller quantities [18]. For the coming months, the funds then got separated into smaller quantities. While from that point onwards it cannot be established if the funds are still in possession of the attackers, apparent manual obfuscation of funds is a widely-documented laundering TTP [51].

Mixing: The attackers can be observed using supposed mixing wallets. One address has 394 incoming transactions from wallets mentioned in the official list [16]. This wallet was created on October 20, 2015 and already handling transactions prior to the Bitfinex hack. While the character of this wallet cannot be established, it has characteristics of a mixing wallet as it exclusively handled transactions worth few satoshis.

Recent activity: In a recent instance of consolidation, between June 1 and June 7, 2020, the attackers moved funds originating from several wallets with smaller holdings, also directly related to the hack [19]. A few days before, on April 28, 2020, the actors emptied a single wallet filled with 168 BTC directly after the attack, which is shown in the cluster just below the center of Figure 3.

The list shared by Bitfinex contains 2072 transactions associated with the hack, in total 119,755 BTC. The transactions took place between 8:54:54 and 12:18:35 on August 2, 2016. From this, 1001 subsequent transactions can be observed taking place from August 2016 until June 7, 2020. At the moment of writing, only 410 wallets of total 2072 have been depleted. The aggregate outbound transactions account for 2663 BTC, which is just 2.2% of the BTC stolen. While it is speculated what is causing this, it is evident that laundering stolen BTC - especially transferring Bitcoin to fiat assets - has become a

Table 2.3: TTPs in Laundering of Stolen BTC

Exchange	Breached	Laundering TTPs
Mt. Gox	2014-02	Fraudulent exchange
Bitfinex	2016-08	Mixing, manual mixing/splitting
Yapizon	2017-04	Bank withdrawals, money mules, gift cards

complex operation for the adversaries behind the hack of Bitfinex.

2.7.3. YAPIZON

The North Korean state-sponsored Lazarus Group is associated with several cryptocurrency-related attacks [48]. The group is suspected of being behind the hack of Yapizon in 2017, recorded in our dataset. Furthermore the group is suspected of stealing altcoin from several other exchanges, as well as malware-based attacks to steal key pairs of unwitting users. In March of 2020, US government authorities indicted two individuals associated with laundering of proceeds from these breaches [74]. From this recent case, it can be observed how adversaries launder stolen BTC in an increasingly regulated ecosystem. Compared to the laundering TTPs discussed earlier, laundering process has become more laborious, as the actors were observed withdrawing 34 million USD of stolen funds from a Chinese bank account that is linked to an exchange account. More interestingly, they were also observed converting 1.4 million USD worth of BTC into iTunes gift cards, which were used to purchase then-laundered BTC.

The practice of laundering through iTunes gift cards is striking, as gift cards are a known money laundering avenue. It is therefore astonishing such volumes have not raised any flags, or that the Chinese authorities did not intervene in the transfer or withdrawal. This shows that persistent actors will seek to maneuver around limitations put up in the ecosystem by bona fide actors.

According the analysis of these 3 cases, it is fair to establish that straightforward conversion of stolen BTC directly into fiat currency is a practice of a time gone by. Most exchanges have flagged wallet addresses associated with security breaches. Laundering criminally obtained BTC is further complicated by KYC and AML regulations. Over 4 years after the fact, few wallets associated with the Bitfinex hack have been emptied, with regulatory scrutiny only increasing.

2.8. LIMITATIONS

We have made significant effort to include all security breaches of Bitcoin exchanges in our dataset. Like any analysis driven by open source data, only publicly disclosed breaches can be included. Breaches may not be reported or remain unknown even to the victim. The composition of the dataset depends on the reporting obligation or generally responsible practice of breached parties. Despite this limitation we believe that our dataset is an accurate representation of Bitcoin exchange breaches over the eight years past. Analysis of an ecosystem as a whole provides a better reflection of reality than analysis of individual breaches. Our analysis serves as an analysis of its current

state, as the trading of Bitcoin is an ecosystem in constant flux.

We have based our analysis on official reports, but in some cases strong rumors of exit scams exist. As indicated in the previous section, we have included these in Table 1 in the interest of completeness. Fear, uncertainty and doubt (FUD) are inherent to the Bitcoin community [38] and hence a tacit limitation to any research of the environment.

Furthermore, our classification of breaches is based on current representation of facts in official sources. The reporting of Bitcoin exchange breaches tends to be very light on technical detail, if any. Almost six years after the fact, is still not publicly known how the Mt. Gox breach, with the highest-ever amount of stolen BTC, could have taken place. With exchange regulations in development in several jurisdictions, as well as court cases on breaches currently ongoing, future work is necessary as more details become available.

2.9. CONCLUSION

With the amount of fiat currency flowing into the Bitcoin market, exchange platforms are an attractive target for cyber criminal actors. We have analyzed 36 instances of cyber security breaches of Bitcoin exchange platforms, cumulatively accounting for at least 1,156,399 BTC stolen from their legitimate owners. Each of these incidents was facilitated by cyber security of the exchange platform and not the negligence of its users, the legitimate Bitcoin owners.

We have found that in recent years exchanges tend to disclose less technical details on the what and why of a breach compared to their earlier victims in our dataset. With regards to the vector used to breach exchanges, both the use of stolen credentials and the abuse of functionality are decreasing. This is good news, as a decrease of easy attack vectors suggest an increase in the levels of technical security of exchange platforms. Other positive developments are decreased exchange closure due to breaches and a decreasing amount of stolen BTC over recent years. This is partially due to other exchange platforms being willing to block funds directed to them by the attackers and returning those. Although the absolute BTC yield per breach has decreased, exchange platforms remain an interesting target due to increased BTC-USD exchange rate. Funds stored in hot wallets remain the primary target for attackers, as only 2 breaches in our dataset involved cold storage. The vectors used to breach Bitcoin exchanges overlap with those used in the broader financial services industry. Actual theft of funds is however rare in traditional financial services, where mostly personal information is targeted.

As of 2019, exchanges are required to comply with Know Your Customer and Anti-Money Laundering regulations in most jurisdictions. Compliance with such legislation is usually accompanied by stricter cyber security. However compared to other organizations in the financial sector, regulatory oversight on exchange platforms falls short in the protection of customer funds. A deposit insurance system, which provides customer protection in case an exchange becomes insolvent, can be a next step for the Bitcoin ecosystem. Exchange platforms can also take the lead, with recent cooperation in freezing and returning funds to breached exchanges serving as an example. The initiation of mutual aid agreements as prevalent in other industries can help formalize such arrangements.

Yet all cyber threats discussed in this paper are a result of centralization of the ecosys-

tem caused by centralized exchanges. Peer-to-peer trading is not vulnerable to these threats as decentralized exchanges do not store user assets. Decentralized trading however shifts security risk and thus responsibility to the user. Service providers keep building propositions on top of the Bitcoin technology stack, each with its own implications on the core attributes of confidentiality, integrity and availability. Whether risk is acceptable remains a responsibility of the user, who votes with his or her (Bitcoin) wallet.

REFERENCES

- [1] AllCrypt. *What happened, and what's going on* - AllCrypt Blog. URL: <https://archive.is/2UY7e> (visited on 11/20/2019).
- [2] Andreas Antonopoulos. *Bitcoin Q&A: How do I secure my bitcoin?* 2017. URL: <https://www.youtube.com/watch?v=vt-zXEsJ6IU>.
- [3] BBC Russia. *Bitcoins in the "FSB fund": how \$450 million disappeared from Wex crypto exchange*. URL: <https://www.bbc.com/russian/features-50420738>.
- [4] Binance. *Binance Is SAFU: 7 Ways We Secure Your Assets* 24/7. 2019. URL: <https://www.binance.com/en/blog/307883269744750592/Binance-Is-SAFU-7-Ways-We-Secure-Your-Assets-247>.
- [5] Binance. *Binance Security Breach Update*. 2019. URL: <https://www.binance.com/en/support/articles/360028031711>.
- [6] Bitcoin Wiki. *Weaknesses*. URL: <https://en.bitcoin.it/wiki/Weaknesses>.
- [7] Bitcoin-Central. *Les Explications De Bitcoin Central*. 2013. URL: <https://bitcoin.fr/les-explications-de-bitcoin-central/>.
- [8] Bitcoin.com. *Names, phone numbers, and emails leaked in BitQuick exchange hack* - Bitcoin News. URL: <https://news.bitcoin.com/names-phone-numbers-emails-leaked-bitquick-exchange-hack/> (visited on 11/20/2019).
- [9] Bitcointalk.org. *[Emergency ANN] Bitcoinica site is taken offline for security investigation*. URL: <https://bitcointalk.org/index.php?topic=81045.0> (visited on 11/20/2019).
- [10] Bitcointalk.org. *BitMarket.Eu has closed down*. URL: <https://bitcointalk.org/index.php?topic=5441.msg1533170.5C#msg1533170> (visited on 11/20/2019).
- [11] Bitcointalk.org. *Minor Crypto Exchange Pulls Off Exit Scam, Steals All User Funds*. 2018. URL: <https://bitcointalk.org/index.php?topic=5059683.0>.
- [12] Bitfinex. *Security breach on Bitfinex*. 2016. URL: <http://archive.is/CnjAn> (visited on 11/20/2019).
- [13] Bitfinex. *Up to US\$400 Million Reward for Return of Stolen 2016 Bitcoin*. URL: <https://www.bitfinex.com/posts/494>.
- [14] BitInstant. *Events of Friday - BitInstant Back Online - Blog - Genesis Block - The BitInstant Blog*. 2013. (Visited on 12/01/2019).
- [15] Bitwise. *Analysis of Real Bitcoin Trade Volume*. Tech. rep. 2019. URL: <https://static.bitwiseinvestments.com/Research/Bitwise-Asset-Management-Analysis-of-Real-Bitcoin-Trade-Volume.pdf>.

- [16] Blockchain.com. “Address”. In: 2019. URL: <https://www.blockchain.com/btc/address/19cj6xavuXErZE9vyob9jsB4AhQRGNZ9z2>.
- [17] Blockchain.com. *Market Capitalization*. 2019. URL: <https://www.blockchain.com/charts/market-cap?timespan=2years>.
- [18] Blockchain.com. “Transaction”. In: 2017. URL: <https://www.blockchain.com/btc/tx/2fb1ad2aceb70d2235e7092559447ec493c21a0bff01f793b8c0161d5f5e92c9>.
- [19] Blockchain.com. “Transaction”. In: 2019. URL: <https://www.blockchain.com/btc/tx/357376f5188054b46cdcf21328d7cbfcef75eface94bcdd8e6e54d4edd7ad8f2>.
- [20] Busoni. *BTC Stolen from Poloniex*. 2014. URL: <https://bitcointalk.org/index.php?topic=499580>.
- [21] C4. *CryptoCurrency Security Standard (CCSS)*. 2019. URL: <https://github.com/CryptoConsortium/CCSS>.
- [22] Tracey Caldwell. “The true cost of being hacked”. In: *Computer Fraud and Security* (2014). ISSN: 13613723. DOI: [10.1016/S1361-3723\(14\)70500-7](https://doi.org/10.1016/S1361-3723(14)70500-7).
- [23] Cavirtex. *Latest News: Update on withdrawals*. 2015. URL: <https://web.archive.org/web/20150220001845/https://www.cavirtex.com/news>.
- [24] Leandro César. *Problema do Mercado Bitcoin*. 2013. URL: <https://bitcointalk.org/index.php?topic=160150.0>.
- [25] Coin Telegraph. *Chinese Exchange Gets 'Goxed' for 1,000 bitcoins*. 2015. URL: <http://cointelegraph.com/news/chinese-exchange-suffers-1000-btc-loss-in-uncertain-service-compromise>.
- [26] Coin Telegraph. *Korean Bitcoin Exchange Yapizon Confirms \$5 mln Hack, All Customers To Pay With Balances*. 2017.
- [27] CoinDesk. *Bitcoin Exchange Cointrader Shuts Down After Alleged Hack*. 2016. URL: <https://www.coindesk.com/bitcoin-exchange-cointrader-shuts-down>.
- [28] CoinDesk. *Bitpoint Exchange Hacked for \$32 Million in Cryptocurrency*. 2019.
- [29] CoinDesk. *BTER Claims \$1.75 Million in Bitcoin Stolen in Cold Wallet Hack*. 2015. URL: <https://www.coindesk.com/bter-bitcoin-stolen-cold-wallet-hack>.
- [30] CoinDesk. *Crypto Exchange Zaif Hacked In \$60 Million Bitcoin Theft*. 2018. URL: <https://www.coindesk.com/crypto-exchange-zaif-hacked-in-60-million-6000-bitcoin-theft>.
- [31] CoinDesk. *Czech bitcoin exchange Bitcash.cz hacked and up to 4,000 user wallets emptied*. 2013. URL: <https://www.coindesk.com/czech-bitcoin-exchange-bitcash-cz-hacked-4000-user-wallets-emptied>.
- [32] CoinDesk. *Details of \$5 Million Bitstamp Hack Revealed*. 2015. URL: <https://www.coindesk.com/unconfirmed-report-5-million-bitstamp-bitcoin-exchange>.

- [33] CoinDesk. *Polish Bitcoin Exchange Bidextreme.pl Hacked, Bitcoin and Litecoin Wallets Emptied*. 2013. URL: <https://www.coindesk.com/hacker-attack-poland-s-bitcoin-exchange>.
- [34] CoinDesk. *Singapore-Based Crypto Exchange DragonEx Has Been Hacked*. 2019. URL: <https://www.coindesk.com/singapore-based-crypto-exchange-dragonex-has-been-hacked>.
- [35] CoinDesk. *Zaif Crypto Exchange Reveals Takeover In New Hack Refund Plan*. 2018. URL: <https://www.coindesk.com/zaif-crypto-exchange-reveals-takeover-in-new-hack-refund-plan>.
- [36] CoinGecko. “Top 100 Coins by Market Capitalization”. In: (2019). URL: <https://www.coingecko.com/en>.
- [37] Mauro Conti et al. “A survey on security and privacy issues of bitcoin”. In: *IEEE Commun. Surv. Tut.* (2018). ISSN: 1553877X. DOI: [10.1109/COMST.2018.2842460](https://doi.org/10.1109/COMST.2018.2842460). arXiv: [1706.00916](https://arxiv.org/abs/1706.00916).
- [38] Barnaby Craggs and Awais Rashid. “Misplacing Trust in Bitcoin Information Sources”. In: (2018).
- [39] Organized Crime and Corruption Reporting Project. *US and Russia Spar Over Accused Crypto-Launderer*. URL: <https://www.occrp.org/en/investigations/us-and-russia-spar-over-accused-crypto-launderer>.
- [40] Cryptsy. *Cryptsy Blog — Announcement*. URL: <http://archive.is/FfECg> (visited on 11/20/2019).
- [41] Daily Beast. *Japanese Bitcoin Heist ‘an Inside Job,’ Not Hackers Alone*. 2017. URL: <https://www.thedailybeast.com/japanese-bitcoin-heist-an-inside-job-not-hackers-alone>.
- [42] Elliptic. *How the DOJ Indictment of Russian Hackers is Supported by Blockchain Analysis*. URL: <https://www.elliptic.co/our-thinking/doj-indictment-russian-hackers-blockchain-analysis>.
- [43] Excoin. *Excoin - Announcement*. URL: <https://web.archive.org/web/20150215200218/https://exco.in/> (visited on 11/20/2019).
- [44] Amir Feder et al. “The impact of DDoS and other security shocks on Bitcoin currency exchanges: Evidence from Mt. Gox”. In: *Journal of Cybersecurity* (2017). ISSN: 20572093. DOI: [10.1093/cybsec/tyx012](https://doi.org/10.1093/cybsec/tyx012).
- [45] FinCEN. *FinCEN Fines BTC-e Virtual Currency Exchange \$110 Million for Facilitating Ransomware, Dark Net Drug Sales*. URL: <https://www.fincen.gov/news/news-releases/fincen-fines-btc-e-virtual-currency-exchange-110-million-facilitating-ransomware>.
- [46] Gate.io. *gate.io will stop using statcounter for traffic stats*. URL: <https://www.gate.io/article/16665> (visited on 11/21/2019).
- [47] Gatecoin. *Gatecoin | Official Statement Regarding Gatecoin Hot Wallet Breach*. URL: <http://archive.is/rcnG4> (visited on 11/20/2019).

- [48] Group-IB. *Group-IB: 14 cyber attacks on crypto exchanges resulted in a loss of \$882 million*. URL: <https://www.group-ib.com/media/gib-crypto-summary/> (visited on 06/09/2020).
- [49] Hackernoon. *A Huge List of Cryptocurrency Thefts*. 2019. URL: <https://hackernoon.com/a-huge-list-of-cryptocurrency-thefts-16d6bf246389>.
- [50] Hong Kong Securities and Futures Commission (SFC). *Position paper on Regulation of virtual asset trading platforms*. Tech. rep.
- [51] Jeff Hu. *Generate and download thousands of Bitcoin wallets in a minute or two*. URL: <https://medium.com/coinmonks/generate-and-download-thousands-of-bitcoin-wallets-in-a-minute-or-two-d42ce73d77d8>.
- [52] Investing.com. *Bitcoin Historical Data*. URL: <https://www.investing.com/crypto/bitcoin/historical-data>.
- [53] US Department of Justice. *Russian National And Bitcoin Exchange Charged In 21-Count Indictment For Operating Alleged International Money Laundering Scheme And Allegedly Laundering Funds From Hack Of Mt. Gox*.
- [54] Mark Karpeles. *Clarification of Mt Gox Compromised Accounts and Major Bitcoin Sell-Off*. 2011. URL: <https://web.archive.org/web/20110919162635> (visited on 11/20/2019).
- [55] Kaspersky Lab. *Lazarus Under The Hood*. Tech. rep. 2018.
- [56] Il Kwon Lim et al. “The analysis and countermeasures on security breach of Bitcoin”. In: *Lecture Notes in Computer Science*. 2014. ISBN: 9783319091464. DOI: [10.1007/978-3-319-09147-1_52](https://doi.org/10.1007/978-3-319-09147-1_52).
- [57] Linode. *Manager Security Incident*. 2012. URL: <http://archive.is/tRQ9> (visited on 11/20/2019).
- [58] Tyler Moore, Nicolas Christin, and Janos Szurdi. “Revisiting the risks of bitcoin currency exchange closure”. In: *ACM Transactions on Internet Technology* (2018). ISSN: 15576051. DOI: [10.1145/3155808](https://doi.org/10.1145/3155808).
- [59] Satoshi Nakamoto. “Bitcoin: A Peer-to-Peer Electronic Cash System”. 2008. DOI: [10.1007/s10838-008-9062-0](https://doi.org/10.1007/s10838-008-9062-0). URL: <https://bitcoin.org/bitcoin.pdf>.
- [60] Kris Oosthoek and Christian Doerr. “Cyber Threat Intelligence: A Product Without a Process?” In: *International Journal of Intelligence and CounterIntelligence* 34.2 (2021), pp. 300–315. DOI: [10.1080/08850607.2020.1780062](https://doi.org/10.1080/08850607.2020.1780062). eprint: <https://doi.org/10.1080/08850607.2020.1780062>. URL: <https://doi.org/10.1080/08850607.2020.1780062>.
- [61] Kris Oosthoek and Christian Doerr. “From Hodl to Heist: Analysis of Cyber Security Threats to Bitcoin Exchanges”. In: *IEEE International Conference on Blockchain and Cryptocurrency 2020*. IEEE. 2020.
- [62] Kris Oosthoek and Christian Doerr. “SoK: ATT&CK Techniques and Trends in Windows Malware”. In: *15th EAI International Conference, SecureComm 2019, Proceedings*. 2019. URL: <https://krisk.io/post/attack/>.

- [63] PR Times. *Report on suspension of deposit and withdrawal of virtual currency and our response*. 2018. URL: <https://prtimes.jp/main/html/rd/p/000000093.00012906.html>.
- [64] Reddit user zanetackett. *Txid and Bitcoin Addresses Connected To The Bitfinex Theft*. 2016. URL: https://np.reddit.com/r/BitcoinMarkets/comments/4wizgv/txid_and_bitcoin_addresses_connected_to_the/.
- [65] Tim Ring. "A breach too far?" In: *Computer Fraud and Security* (2013). ISSN: 13613723. DOI: [10.1016/S1361-3723\(13\)70052-6](https://doi.org/10.1016/S1361-3723(13)70052-6).
- [66] Selfkey. *A Comprehensive List of Cryptocurrency Exchange Hacks*. 2019. URL: <http://selfkey.org/list-of-cryptocurrency-exchange-hacks/>.
- [67] Shapeshift.io. *A Timeline: ShapeShift Hacking Incident - ShapeShift*. URL: <https://info.shapeshift.io/blog/2016/04/19/blog-2016-04-19-timeline-shapeshift-hacking-incident/> (visited on 11/20/2019).
- [68] Adam Shostack. "Experiences threat modeling at Microsoft". In: *CEUR Workshop Proceedings*. 2008.
- [69] SWIFT. *SWIFT customer communication: Customer security issues*. 2016. URL: https://www.swift.com/insights/press-releases/swift-customer-communication%5C_customer-security-issues.
- [70] The Economic Times India. *Bitcoins worth Rs 20 crore stolen from exchange in India's biggest crypto theft*. 2018.
- [71] The MITRE Corporation. *Common Attack Pattern Enumeration and Classification (CAPEC)*. URL: <https://capec.mitre.org/>.
- [72] The Verge. *Bitcoin exchange BitFloor suspends operations after \$250,000 theft - The Verge*. 2012. URL: <https://www.theverge.com/2012/9/5/3293375/bitfloor-bitcoin-exchange-suspended-theft> (visited on 11/20/2019).
- [73] *The VERIS Community Database*. 2019. URL: <https://github.com/vz-risk/Vcdb>.
- [74] U.S. Department Of The Treasury. *Treasury Sanctions Individuals Laundering Cryptocurrency for Lazarus Group*.
- [75] UN Security Council. *S/2019/171: Report of the Panel of Experts established pursuant to resolution 1874 (2009)*. Tech. rep. 2019.
- [76] Verizon. *2019 Data Breach Investigations Report*. 2019.
- [77] Verizon. *VCDB JSON directory README*. URL: <https://github.com/vz-risk/Vcdb/tree/master/data/json>.
- [78] Verizon. *VERIS: The Vocabulary for Event Recording and Incident Sharing*. 2019. URL: <http://veriscommunity.net/>.
- [79] Vircurex. *May 2013 Report*. Tech. rep. 2013. URL: <https://vircurex.com/Reports/2013-05.pdf>.
- [80] Tom Williams. *MyBitcoin Incident Report - August 5th 2011*. 2011. URL: <http://archive.is/LUMzs> (visited on 11/20/2019).

- [81] ZDNet. *Bitcoin exchange Cryptoine hacked*. 2015. URL: <https://www.zdnet.com/article/bitcoin-exchange-cryptoine-hacked/>.
- [82] ZDNet. *Bitfinex hackers arrested after three years*. 2019. URL: <https://www.zdnet.com/article/bitfinex-hackers-arrested-after-three-years/>.

3

INVESTIGATING THE RANSOMWARE PAYMENTS ECONOMY

Ransomware attacks are among the most severe cyber threats. They have made headlines in recent years by threatening the operation of governments, critical infrastructure, and corporations. Collecting and analyzing ransomware data is an important step towards understanding the spread of ransomware and designing effective defense mechanisms. We report on our experience operating Ransomwhere, an open crowdsourced ransomware payment aggregator to collect information from victims of ransomware attacks. With Ransomwhere, we have gathered around 13.5 thousand ransom payments to more than 87 criminal ransomware actors with total payments of more than \$101 million. Leveraging the transparent nature of Bitcoin, the cryptocurrency used for most ransomware payments, we characterize the evolving ransomware criminal structure and ransom laundering strategies. Our analysis shows that there are two parallel ransomware criminal markets: commodity ransomware and Ransomware as a Service (RaaS). We notice striking differences between the two markets in the way that cryptocurrency resources are utilized, revenue per transaction, and ransom laundering efficiency. Although it is relatively easy to identify choke points in commodity ransomware payment activity, it is more difficult to do the same for RaaS.

This chapter will be published as *A Tale of Two Markets: Investigating the Ransomware Payments Economy* by Oosthoek, K., Cable, J. and Smaragdakis, G. in *Communications of the ACM* (accepted and in print).

3.1. INTRODUCTION

Ransomware, a form of malware designed to encrypt a victim's files and make them unusable without payment, has quickly become a threat to the functioning of many institutions and corporations around the globe. In 2021 alone, ransomware caused major hospital disruptions in Ireland [3], empty supermarket shelves in the Netherlands [9], the closing of 800 supermarket stores in Sweden [4], and gasoline shortages in the United States [28]. In a recent report, the European Union Agency for Cybersecurity (ENISA) ranked ransomware as the “prime threat for 2020-2021” [12]. The U.S. government reacted to high profile attacks against U.S. industries by declaring ransomware a national security threat and announcing a “coordinated campaign to counter ransomware” [1]. Other governments, including the United Kingdom [39], Australia [17], Canada [10], and law enforcement agencies, such as the FBI [38] and Europol [14], have also launched similar programs to defend against ransomware and offer help to victims.

To the criminal actors behind these attacks, the resulting disruption is just ‘collateral damage’. A handful of groups and individuals, with names such as NetWalker, Conti, REvil and DarkSide, have received tens of millions of dollars as ransom. But this is just the top of the food chain in an ecosystem with many grey areas, especially when it comes to laundering illicit proceedings. In this article, we will provide a closer look at the ecosystem behind many of the attacks plaguing businesses and societies, known as Ransomware as a Service (RaaS).

Cryptocurrency remains the payment method of choice for criminal ransomware actors. While many cryptocurrencies exist, Bitcoin is preferred due to its network effects, resulting in wide exchange options. Bitcoin's sound monetary features as a medium of exchange, unit of account and store of value make it as attractive to criminals as it is to regular citizens. According to the U.S. Department of Treasury, based on data from the first half of 2021, the “vast majority” of reported ransomware payments were made in Bitcoin [37]. However, significant discrepancies exist between total ransomware revenues reported by industry and government outlets. Law enforcement agencies have started to disrupt ransomware actors by obtaining personal information of threat actors from Bitcoin exchanges. This is realized through anti-money laundering regulations such as Know Your Customer (KYC), which require legal identity verification during registration with a given service. While cryptocurrencies such as Bitcoin enable ransomware, blockchain technology also offers unprecedented opportunities for forensic analysis and intelligence gathering. Using our crowdsourced ransomware payment aggregator, Ransomwhere, we compile a dataset of 7,321 Bitcoin addresses which received ransom payments, based on which we shed light on the structure and state of the ransomware ecosystem.

Our contributions are as follows:

- We collect and analyze the largest public dataset of ransomware activity to date, which includes 13,497 ransom payments to 87 criminal actors over the last five years, worth more than \$101 million.
- We characterize the evolving ransomware ecosystem. Our analysis shows that there are two parallel ransomware markets: commodity and RaaS. After 2019, we

observe the rapid rise of RaaS, which achieves higher revenue per address and transaction, and higher overall revenue.

- We also characterize ransom laundering strategies by commodity ransomware and RaaS actors. Our analysis of more than 13 thousand transfers shows striking differences in laundering time, utilization of exchanges, and other means to cash out ransom payments.
- We discuss difficulties in defending against professionally-operated RaaS and we propose potential manners of tracing back RaaS cryptocurrency activity.
- To enable future research in this area, we make our aggregator, Ransomwhere, and the underlying ransomware payments of our analysis publicly available at [7].

3.2. THE RANSOMWARE ECOSYSTEM

The ransomware ecosystem can be largely divided into two categories: commodity ransomware and ransomware as a service (RaaS).

3.2.1. COMMODITY RANSOMWARE

In the early years of ransomware, the majority of ransomware that spread can be characterized as ‘commodity’ ransomware. Commodity ransomware is distinguished by widespread targeting, fixed ransom demands, and technically-adept operators. It usually targets a single device. Actors behind commodity ransomware are usually technically savvy, as most of the time it is developed and deployed by the same person. Commodity ransomware operators take advantage of preexisting work, often copying and modifying leaked or shared source code, causing the formation of ransomware *families*. Historically, most commodity ransomware campaigns utilized phishing emails as the primary delivery vector and exploited vulnerabilities in common word processing and spreadsheet software, if not directly via malicious executables. The modus operandi was mass exploitation, rather than targeting specific victims or corporations.

Exemplary are the WannaCry and NotPetya ransomware families, which over the course of only two months impacted tens of thousands of organizations in over 150 countries by exploiting a vulnerability allegedly stolen from the NSA [16]. By today’s standards, both families were poorly coded and their payment systems were not ready for business (although allegedly this was on purpose with NotPetya [15]).

Applying the conventional advice of having proper backup and contingency plan was thought to defend against ransomware. The initial philosophy was that a quick ability to restore would make it unnecessary to pay, impairing the financial incentive of ransomware operators. But it turned out that what we now regard as a commodity was just a proving ground for more destructive, widespread forms of ransomware.

3.2.2. RANSOMWARE AS A SERVICE (RAAS)

While the first reports of Ransomware as a Service (RaaS) emerged in 2016, it wasn’t until 2019 that RaaS became widespread, rapidly capturing a large share of the ransomware market. We define RaaS as ransomware created by a core team of developers who license

How ransom payments are executed and laundered

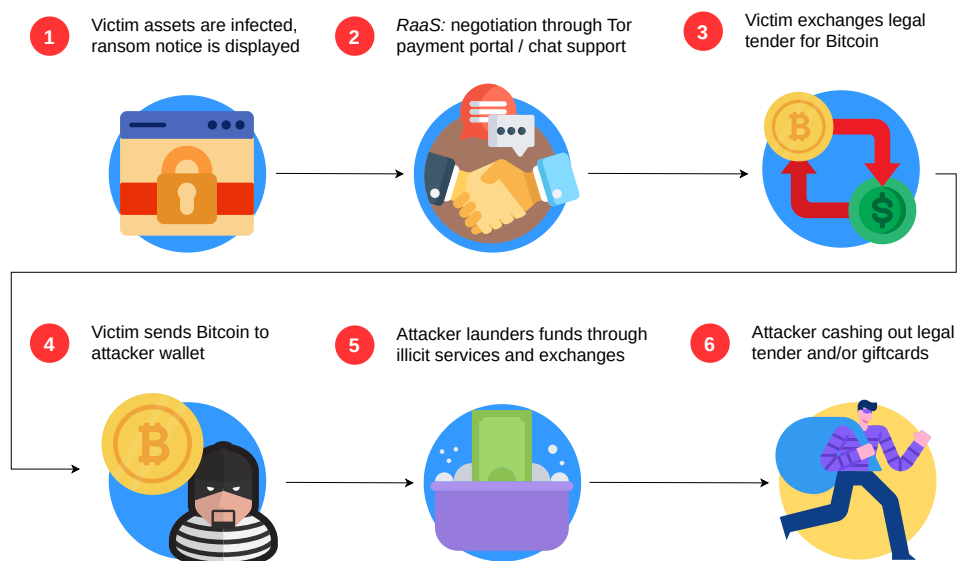


Figure 3.1: General course of a ransom payment and its laundering.

their malware on an affiliate basis. They often provide a payment portal (typically over Tor, an anonymous web protocol), allowing negotiation with victims and dynamic generation of payment addresses (most often Bitcoin). RaaS frequently employs a double extortion scheme, not only encrypting victims' data, but also threatening to leak their data publicly if a ransom is not paid.

The rise of RaaS has enabled existing criminal groups to shift to a lucrative new business model where lower-skilled affiliates can access exploits and techniques previously reserved for highly-skilled criminals. This was exemplified by a leaked playbook from the RaaS group Conti, which enables novice actors to compromise enterprise networks [34]. RaaS affiliates can differ markedly in their approaches. Some scan the entire internet and compromise any victims they can. Once they have identified the victim, they engage in price discrimination based on the victim company's size. Affiliates may even use financial documents obtained in the attack to justify higher prices [27]. Another strategy, known as *big game hunting*, targets big corporations that can afford to pay a high ransom. Darkside is one of most notable RaaS families whose affiliates practice big game hunting, including the notable Colonial Pipeline attack in 2021 [26].

RaaS families often rely on spear phishing over the mass phishing mails utilized by commodity ransomware groups. They also exploit recently disclosed vulnerabilities, taking advantage of vulnerable remote and virtual desktop services [11]. RaaS has lowered the barrier to entry into cyber-criminality, as it has removed the initial expenditure to develop effective ransomware. As a result, attacks can be performed with near zero cost. Combined with high ransom demands, this has led to a low-risk, high reward criminal

Data	Commodity	RaaS	Total
Unique Actors	71	16	87
Bitcoin Addresses	161	7,160	7,321
Received Transactions (Payments)	4,799	8,698	13,497
Transferred Transactions (Laundering)	4,557	8,540	13,097

Table 3.1: Ransomware Dataset Statistics

scheme.

RaaS has effectively *weaponized* the unpatched internet-facing technology of many unwitting organizations. Such organizations have significant financial interest to have systems restored and get back to business after a ransomware attack. Cryptocurrencies enable ransomware actors to directly monetize these vulnerabilities at a scale never seen before. In this paper, we regard the functioning of ransomware actors through what is typically the last mile of the attack.

Figure 3.1 shows the general course of events after a ransomware infection, when the victim decides to pay the attacker (step ❶). In the case of commodity ransomware families, the ransom demand price is fixed and negotiation with the attacker is not necessary. With RaaS, attackers usually run chat-based services to interact with victims and negotiate the final ransom amount (step ❷). After this, a victim will usually exchange fiat legal tender for cryptocurrency such as Bitcoin at an exchange platform (step ❸) and then send it to the attacker's wallet (step ❹). The attacker will then usually launder the obtained Bitcoin through various services (step ❺) in order to obfuscate ownership and reduce the risk of de-anonymization before cashing out (step ❻).

3.3. METHODOLOGY

In this section, we describe how we collected data of ransom payments and ransomware actors in our study.

3.3.1. ADDRESSES INVOLVED IN RANSOM PAYMENTS

We obtain ransomware Bitcoin addresses from our crowdsourced payment aggregator Ransomwhere. The Ransomwhere dataset contains Bitcoin addresses and associated families collected from open-source datasets and publicly-submitted crowdsourced reports. In total, the Ransomwhere dataset contains 7,457 Bitcoin addresses and their corresponding ransomware families.

To seed the dataset, we collected data from several public sources. We imported addresses from Paquet-Clouston et al. [30], who collected 7,222 addresses and labeled families representing approximately \$12.7 million in payments. This dataset provides us with, among other ransomware families, 7,014 addresses belonging to Locky. We further collected 37 addresses and associated families from the AT&T Alien Labs Open Threat Exchange, an open threat intelligence sharing platform [2].

Members of the public may submit reports at our crowdsourced payment aggregator

Table 3.2: Composition of the dataset

Source	Total USD	# BTC Addr.
Ransomwhere reports [7]	\$87M	198
Paquet-Clouston et al. [30]	\$10M	7,222
AlienVault OTX [2]	\$4M	37
Total	\$101M	7,457

Ransomwhere [7]. We received 99 reports containing 198 addresses over a 6-month period from June 2021 to December 2021. While this is a lower number of addresses, they represent the majority of ransomware payment value in our dataset, as seen in Table 3.2. In order to verify reports, the reporter must include the relevant Bitcoin addresses and the associated ransomware family. In addition, they must provide evidence of the ransom demand, such as a screenshot of the ransom payment portal or a ransom message on an infected computer. Some addresses were involved in more than one report. All reports were manually reviewed before being added to the dataset. We did not accept reports that were inaccurate or were not related to ransomware (e.g., addresses involved in extortion scam emails).

All reported ransom addresses were Bitcoin addresses. Due to the transparent nature of Bitcoin it is possible to verify that the collected addresses indeed received payments. Using our own Bitcoin full node, we scraped all transactions for the addresses in our dataset. Overall, 7,323 out of 7,457 Bitcoin addresses were involved in at least one ransom payment. We discarded 134 addresses that did not receive any payment. We have queried Tor using a solution from a peer researcher [33] for all Bitcoin addresses in our dataset to rule out the chance of an address being used for cybercrime purposes other than ransomware. Based on this, we excluded 2 addresses belonging to a cache of Bitcoin seized by the U.S. Department of Justice after the closing of the SilkRoad darkweb market [24], which originally appeared in the Paquet-Clouston et al. dataset. After these steps, the final number of addresses considered for our analysis is 7,321. For a summary of our dataset we refer to Table 3.1. Table 3.2 provides an overview of the sourcing of Bitcoin addresses included in the dataset.

3.3.2. RANSOM PAYMENTS AND LAUNDERING

The transparency of Bitcoin also allows us to collect information about ransom payments, including the amount of Bitcoin received. For each address, we collected the number of incoming (payments) and outgoing (transfers) transactions, their value in Bitcoin, and their timestamp. We calculated the USD value of each transaction using the BTC-USD daily closing rate on the day of the transaction. This serves as an approximate ransom payment and not the exact amount in USD the criminal actors requested or later profited. The total ransom paid to addresses in our dataset is \$101,297,569. The lowest payment received is \$1, and the highest is \$11,042,163. The median payment value is \$1,176.

In collaboration with Crystal Blockchain [6], we tracked the destination of outgoing transactions, i.e., transfers. In order to estimate addresses' potential for illicit use, Crystal Blockchain utilizes clustering heuristics such as one-time change address and common-

input-ownership [41], which allow discovering additional addresses controlled by an actor based on their use in a transaction. When filtering for potential false positives, heuristics and their outcomes are reliable [25]. On top of this, Crystal Blockchain performs manual collection of off-chain data from various cryptocurrency services. In addition to this, Crystal Blockchain scrapes online forums and other Internet services for Bitcoin addresses and their associated real-world entity. Based on this, it is possible to track payments several hops from the original deposit address. To have the most reliable view, in our analysis we have only studied the direct destination of ransom payments (first hop). Based on the characterization of the involved addresses across the path, we are able to study the laundering strategies of ransomware groups as well as the time needed to wash out the money (see Section 3.5).

3.3.3. RANSOMWARE ACTORS

We obtained addresses and labeled families as described in Section 3.3.1. We categorized each ransomware family as used by either commodity ransomware or RaaS actors. Ransomware is generally categorized as RaaS due to the use of an affiliate structure, with the ransomware developer (operator) selling the ransomware to criminal actors either based on a commission for each ransom paid, or a flat monthly fee (*as a service*, like many subscription-based services). As there does not exist any comprehensive public list of RaaS groups, we have labeled a family as RaaS if a reliable industry or law enforcement source claims that a given ransomware is sold *as a service*. A list of commodity and RaaS families in our dataset is presented in Table 3.3.

3.3.4. LIMITATIONS

Our dataset of Bitcoin addresses is the largest public collection of ransomware payment addresses collected to date, based on total USD value. While this allows for a unique view on the ransomware financial ecosystem, it is not exhaustive. An inherent limitation of any research using adversary artifacts is its dependence on the availability of artifacts that bad actors have an interest to hide. Furthermore victims might have an interest to not report addresses, as they prefer keeping attacks undisclosed. We note that certain families, such as NetWalker, may be overrepresented in our dataset due to us having more complete data on these families. Despite this limitation, we believe that our dataset provides a valuable, if incomplete, representation of ransomware payments over many years. This broad view provides a better reflection of the state of affairs than simply focusing on a few families. We hope that this can lay the groundwork for further public data collection in the future, and encourage anyone to submit data at Ransomwhere [7].

3.4. RANSOM PAYMENT ANALYSIS

In this section, we analyze 13,497 payments to the Bitcoin addresses in our dataset (see Table 3.1). A payment is a transaction received by an address in our dataset. Table 3.3 lists the ransomware families used by the actors in our dataset. Our dataset contains Bitcoin addresses associated with 87 commodity ransomware or RaaS actors. For reasons of brevity, families for which our dataset contains just 1 address are excluded from Table 3.3. The 16 actors that are classified as RaaS, highlighted in Table 3.3, account for 7,160

Name	#Addrs.	Name (contd.)	#Addrs.
Locky	7037	DarkSide	3
NetWalker	66	MedusaLocker	3
SamSam	48	NotPetya	3
Ryuk	40	GlobeImposter	3
Conti	27	ThunderCrypt	3
Qlocker	22	Nemucod	3
JigSaw	11	LockBit 2.0	2
CryptConsole	10	Globe v2	2
Egregor	9	EDA2	2
DMALocker v3	9	Flyper	2
Globe v3	7	Black Kingdom	2
REvil	7	CryptoLocker	2
CryptoTorLocker2015	7	AvosLocker	2
HC6/HC7	6	NoobCrypt	2
Globe	5	VenusLocker	2
WannaCry	5	XLocker v5	2
TeslaCrypt	5	Chimera	2
CTB-Locker	5	Badblock	2
Xorist	4	Other Groups/Families*	50

* 50 families with 1 address each. RaaS actors are highlighted.

Table 3.3: Ransomware families in the dataset

out of 7,321 addresses in our dataset. As mentioned previously, for full review our dataset is publicly available [7].

Ransomware victims typically create an account with a reputable exchange platform to buy Bitcoin with fiat currency. Then, victims perform a transaction (payment) to the address provided by the ransomware actor. In our dataset, payment transactions to ransomware addresses tend to originate one to two hops away from reputable exchange platforms, such as Coinbase and Kraken.

3.4.1. RANSOMWARE REVENUE

In Figure 3.2 we list the 15 ransomware families with the highest revenue. The top-grossing families are dominated by RaaS: NetWalker has the highest revenue, \$26.7 million, followed by Conti (\$16.4 million), REvil/Sodinokibi (\$12.1 million), DarkSide (\$9.1 million) and Locky (\$8.1 million). Combined, commodity actors account for a total revenue of \$5.5 million. Although the number of RaaS actors is significantly lower, they together earned \$95.7 million.

Figure 3.3 shows the accumulated revenue of both commodity ransomware and RaaS actors. We see that, from 2015 until 2019, early RaaS actors, primarily Locky, were earning significant but still relatively low revenue. Commodity actors were also active, but with even lower revenue. As seen in Figure 3.3, RaaS revenue reached \$8.2 million in April 2020. This can be primarily attributed to NetWalker, which actively targeted hospitals and healthcare institutions during the first COVID-19 lockdown in that period [20].

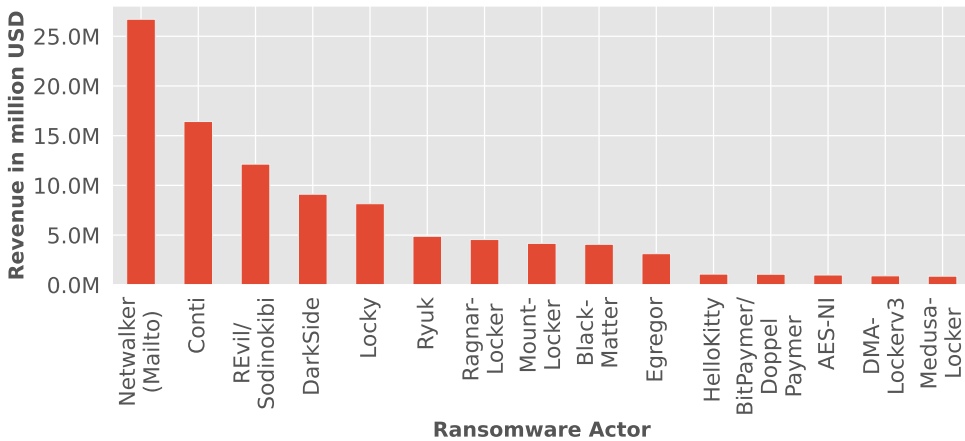


Figure 3.2: Revenue per ransomware actor.

Other revenue peaks caused by RaaS groups are in May and June of 2021, with peaks of \$13.5 million and \$12.8 million respectively. These spikes are caused by large ransom payments by individual victims. One example of this is a payment to REvil/Sodinokibi on June 1st, 2021, accounting for \$11 million. This is a payment by the Brazilian meat processing company JBS, which dominated headlines at the time [19].

Locky has a notorious reputation as one of the biggest ransomware strains in 2016-2017. It is also one of the earliest, if not the first, RaaS families. What stands out apart from its high revenue is its address usage. The actors behind Locky issued new addresses to each victim, a novelty at the time [18]. This is evident in our analysis, with many addresses having only 2 or 3 incoming transactions. According to French court documents, Locky's developer is the same individual who owned BTC-e, a fraudulent exchange [8]. Hence, the actor was able to set up a new address for each payment without raising compliance alarms. Locky is an early, less sophisticated example of a RaaS operation which would serve as an example for many cybercriminals to follow.

3.4.2. RANSOMWARE PAYMENT CHARACTERISTICS

RaaS actors are not only more effective in terms of profits, but also in handling payments. They typically have higher revenue per address, while also generating unique addresses for victims. In Figure 3.4 we show the cumulative distribution of received payments between commodity and RaaS actors. Commodity ransomware actors typically use single wallet addresses to receive hundreds of ransom payments. The highest amount of payments to a single address is 697 to AES-NI, followed by 496 to SynAck and 441 to File-Locker. While these are outliers, Figure 3.4 shows that using a single address to receive upwards of 100 payments is not unusual.

In contrast, RaaS actors almost exclusively use a new wallet address to receive each payment, as observed in Figure 3.4 (right). An outlier is an address associated with Net-Walker which has received 138 payments. This address is likely an intermediate payment address, combining payments from many victims, discovered during McAfee Labs's in-

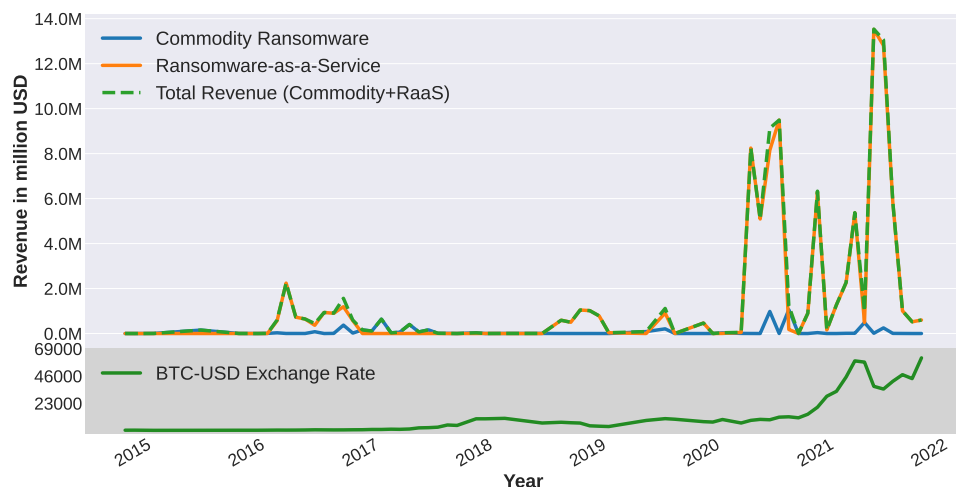


Figure 3.3: USD revenue for commodity and RaaS.

vestigation into NetWalker [35].

The distribution of unique addresses per commodity ransomware and RaaS actor over time is presented in Figure 3.5. In stark contrast to the revenue from ransom activities, presented in Figure 3.3, the number of addresses used in recent years are low, on the order of tens per month. We suspect that RaaS actors prefer to create new addresses for each new ransom payment in order to ensure their pseudo-anonymity, and thus make legal investigations and takedowns more difficult.

Moreover, our analysis shows that RaaS groups apply better operational security practices when using native Bitcoin functionality for wallets (payment addresses). Bitcoin uses Bitcoin Script to handle transactions between addresses. The script type used defines the wallet type. Pay-to-Public-Key-Hash (P2PKH) addresses have the prefix *1*. This is Bitcoin's legacy address format and the most common address format in our dataset with 7,339 addresses. 46 addresses in our dataset are Pay-to-Script-Hash (P2SH) formatted, recognized by the prefix *3*. To spend received payments in Bitcoin, the recipient must specify a redeem script matching the hash. The script can contain functionalities to increase security, such as time-locks or requiring co-signatures. We only observe this for select actors in our dataset: Qlocker, Netwalker, REvil, Ryuk and Phobos. This could mean that these groups have a specific interest in operational security, as transactions usually are not supported by exchange platforms. Another address format is Pay-to-Witness-Public-Key-Hash (P2WPKH), or Segregated Witness (SegWit) protocols, with prefix *bc1q*. In our dataset 72 addresses have this format, belonging to Conti, Netwalker, SunCrypt, DarkSide and HelloKitty. These are all RaaS actors, which could imply deliberate application of SegWit for additional security over traditional address formats.

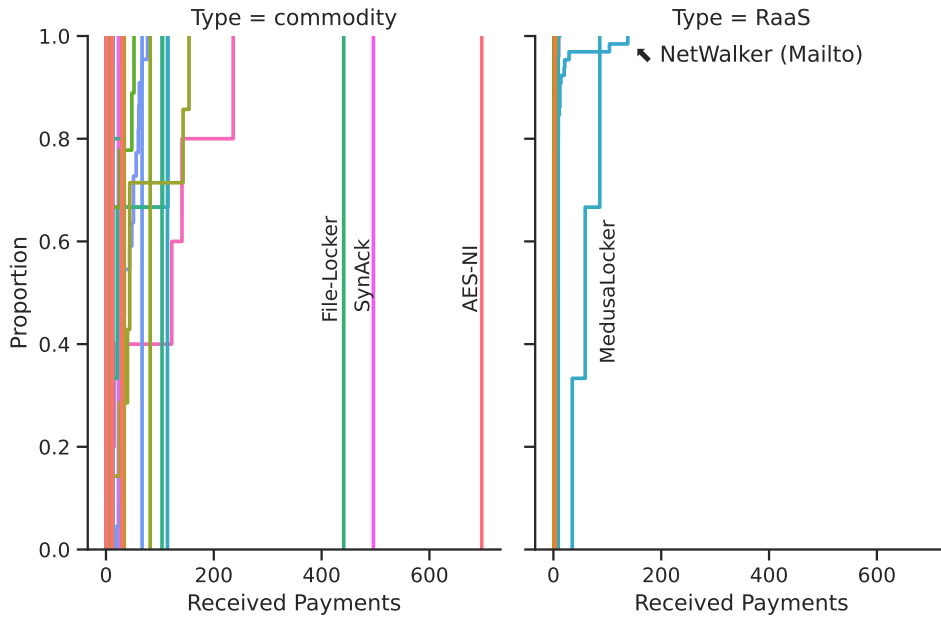


Figure 3.4: ECDF of payments per address for commodity ransomware and RaaS actors.

3.5. MONEY LAUNDERING ANALYSIS

In the previous section, we investigated ransom payments by victims to ransomware actors. In this section, we investigate 13,097 laundering transactions in our dataset (see Table 3.1) to shed light on how these actors liquidate their illicit earnings. For this analysis, we use the methodology introduced in Section 3.3.2.

3.5.1. LAUNDERING STRATEGIES

To avoid exposing their identity, ransomware actors will usually launder their revenue. After routing funds through one or more services to obfuscate the money trail, it is cashed out as legal tender or monetized through the purchase of voucher codes or physical goods. In Figure 3.6 we show the number of transfer transactions per address. The number of transfer (outgoing) transactions provides insights into how actors prefer to initialize their laundering. In short, we see that RaaS actors mostly prefer to empty the deposit address in one transaction, whereas commodity actors prefer multiple smaller transactions – up to hundreds, in some cases more. Hence commodity ransomware actors are less sophisticated. For example, three commodity ransomware actors with the most payments per address (File-Locker, SynAck, AES-NI) also have the most outgoing transactions. While the motivation for this behavior remains unclear, given that law enforcement scrutiny was relatively low, it is likely that the commodity actors took advantage of the ability to cash out more frequently with little risk. This is further supported by their choice of laundering entities.

Almost all ransomware actors in our dataset launder their proceedings entirely. The

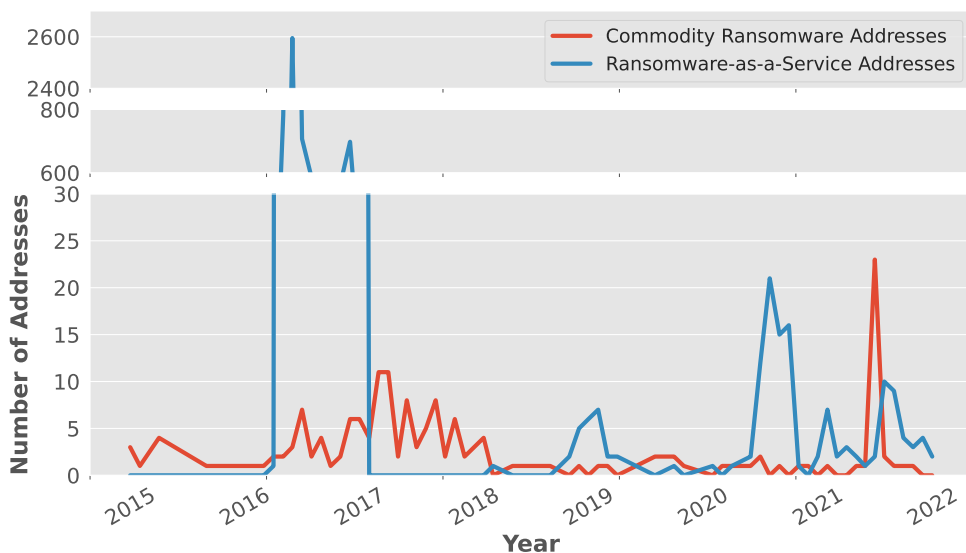


Figure 3.5: Number of unique payment addresses for commodity ransomware and RaaS.

Entity	Description	Evidence
ATM / Payment Provider	Payment gateways for physical/online merchants or ATMs, usually used to launder small amounts	[23]
Dark Market / Illegal Services	Illegal services available on Tor or other Internet services, used to buy illegal server hosting etc.	[13]
Fraudulent Exchange	Exchange platforms officially sanctioned by the US Office of Foreign Assets Control (OFAC)	[8]
Gambling	Online casinos and gambling platforms, used to launder small amounts anonymously	[36]
Low/Moderate ML-Risk Exchange	Exchanges with strict AML/KYC policies might still be used for laundering criminal funds	[32]
Mixers	These services take and 'mix' Bitcoin from various parties to obfuscate ownership	[22]
(Very) High ML-Risk Exchange	Exchanges with lax or no AML/KYC implementations are popular for money laundering	[31]
Wallet Service	Custodial/online wallets, some might have also have privacy features such as mixers.	[36]

Table 3.4: Laundering Entities Overview

speed by which this happens can be inferred from the time between the first incoming payment to and the last outgoing transaction from the deposit address. We define this time duration in which ransomware actors start laundering after having received the payment as *collect-to-laundry time*. Note that this is not the total duration of ransom cash-out, but rather the time spent between receiving the ransom payment and transferring the payment received. Figure 3.7 shows the ECDF of the collect-to-laundry time (in days) for the commodity ransomware and the RaaS actors in our dataset. RaaS actors have a significantly lower collect-to-laundry time compared to commodity actors. Typically, payments to RaaS actors are transferred away from the deposit address in the first minutes to hours after payment. The few outliers in RaaS are caused by NetWalker and individual addresses associated with actors for which we have multiple addresses in our dataset (Ryuk, Conti). As the illicit funds received by RaaS are washed out quickly and, typically, in full, this suggests that it is more difficult to track payments to RaaS, thus lowering the odds of recovery.

Only a small set of families still have significant portions of their proceedings on the original address. This is the case for NetWalker, which has 20.36% still on an address, MedusaLocker (7.98%) and WannaCry (7.92%). In this case, it is likely that the actor has

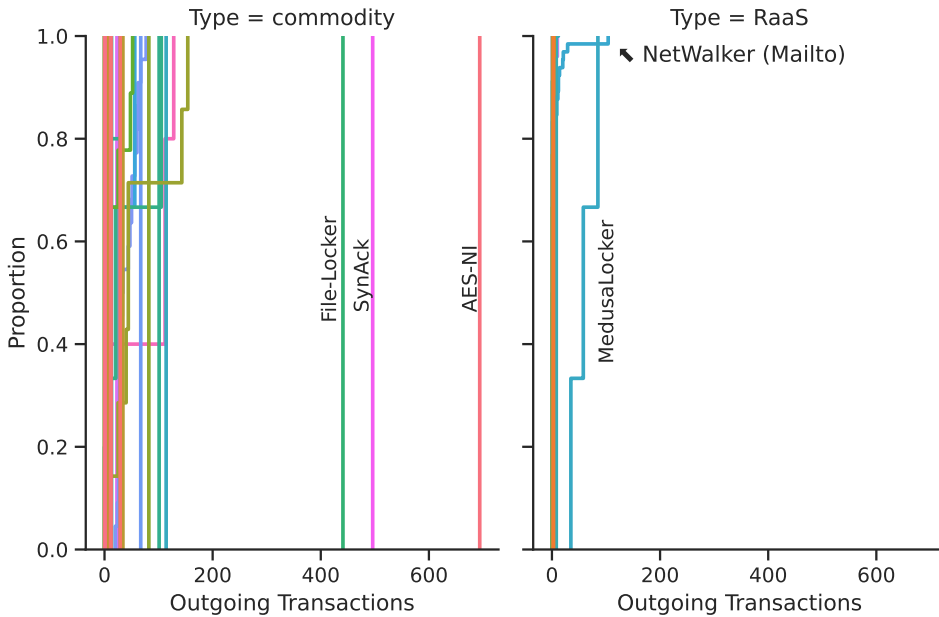


Figure 3.6: Transfer transactions per Address for commodity and RaaS actors.

lost the private key or is incapable to safely launder the ransom, for example due to law enforcement scrutiny. It is known that NetWalker's proceedings have been seized by law enforcement [20], with WannaCry under heavy monitoring and most of the laundering failed [5].

3.5.2. CHALLENGES IN FIGHTING LAUNDERING

Contrary to popular belief, Bitcoin is not anonymous but pseudo-anonymous. Forensic analysis might link a Bitcoin address to a real-world identity, especially when an exchange platform is used to convert between fiat currency and Bitcoin. In most jurisdictions, such platforms are subject to Know Your Customer (KYC) regulations, which require them to verify the identity of every user signing up to their service. During an investigation, when known illicit Bitcoin is routed through an exchange that requires KYC, authorities have a chance to identify the culprit. Law enforcement use blockchain analysis tools in such Anti-Money Laundering (AML) investigations, with technology based on clustering algorithms which can link addresses to a service such as an exchange. As seen in Figure 3.8, we have grouped the data we obtained through Crystal Blockchain in a select set of entities, which are described in Table 3.4.

Laundering can involve routing illicit funds through several hops before cashing out. As it is difficult to know where actual ownership has terminated after several hops, in this analysis we only study the first hop, i.e., the first transfer transaction. This is the service to which actors transfer funds directly after received them from the victim. As this has the closest link to the payment address, this is the first point of investigation for

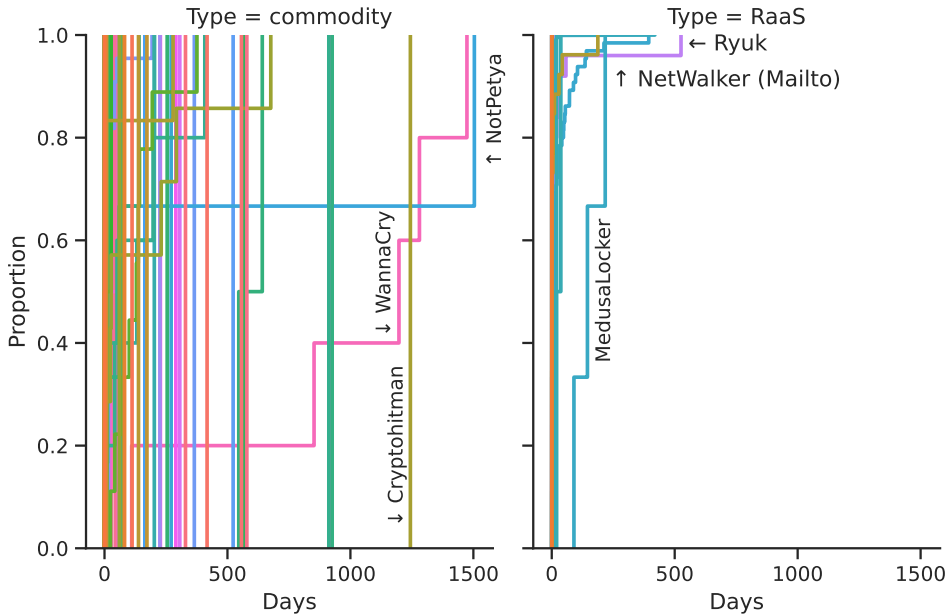


Figure 3.7: Collect-to-Laundry time for commodity ransomware and RaaS actors.

law enforcement. An actor choosing to use a service implies that they trust the service, at least enough not to disclose their identity.

Figure 3.8 shows the proportion of estimated USD value of Bitcoin directly transferred (first hop) to the entities explained in Table 3.4 for commodity and RaaS actors. Due to limitations in reliably establishing (legal) entities behind an address, the direct transactions in our dataset account for a subset of the total revenue generated by the actors in our dataset. Hence we report using percentages, a best practice used with comparable datasets [40].

Our core observation is that commodity actors do not exhibit a specific laundering strategy, while RaaS actors primarily use fraudulent exchanges and mixers. Mixers are services which take in Bitcoin from cybercriminals or privacy-aware users and combine these in many transactions. This hinders the accurate tracking of Bitcoin, as every client gets their initial deposit (minus service fee) back as a mix of other users' Bitcoin. Thus, it is more difficult to trace the laundering activity of RaaS criminal actors.

When considering fraudulent exchanges together with low- and high-risk exchanges, commodity authors tend to prefer exchanges with a low to moderate risk of money-laundering, and thus perhaps cash out to fiat currency or other cryptocurrencies. It is however also known that cybercriminals have wound down the use of fraudulent exchanges [29]. In a sense, commodity actors do not partake in any systematic laundering at all, whereas RaaS actors use fraudulent (non-KYC) exchanges and mixers, a clear laundering strategy. Based on this, we hypothesize that the chances of recovering payments through law enforcement intervention are higher with commodity ransomware

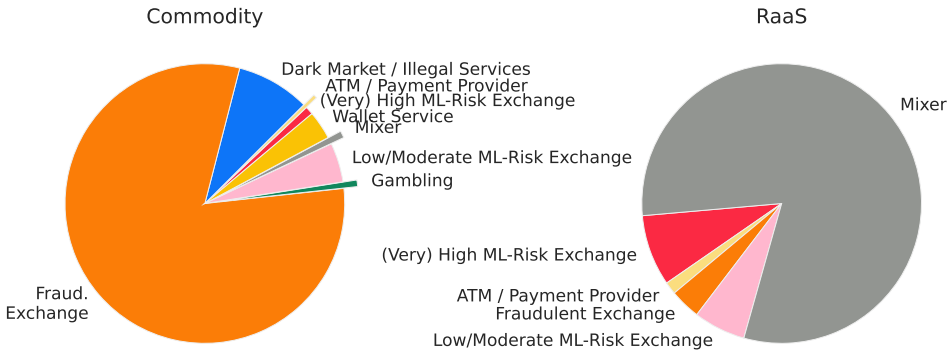


Figure 3.8: Pie chart of one-hop laundering entities.

than with RaaS. The money laundering services they use logically leave more user traces (IP address, login session) than mixer services and fraudulent exchanges with obfuscation of ownership by design.

When an actor's collect-to-laundry time is high, a law enforcement investigation may be able to successfully recover the funds. However, in many such cases there is less incentive to intercept transactions due to the comparatively low ransom amounts. The speed by which RaaS groups transfer funds out suggests criminal sophistication, which is also reflected in their preferred means of laundering. Given this, it is difficult to intercept funds unless law enforcement is already involved at the very moment the payment is made [21].

3.6. CONCLUSION

In this paper, we take a data-driven, “follow the money” approach to characterize the structure and evolution of the ransomware ecosystem. To this end, we report on our experience in operating Ransomwhere, our open crowdsourced ransomware payment aggregator to collect information from victims of ransomware attacks. Our analysis of 13,500 payments unveils that there are two symbiotic, parallel markets: commodity ransomware and (dominant since 2019) Ransomware as a Service (RaaS). The first is operated by individuals or a small group of programmers, the second by professional cybercriminals who offer malware on an affiliate basis to typically less-technical criminal actors. Due to differences in their attack methods, RaaS can demand higher ransom amounts based on the victim at hand. RaaS is also generally more difficult to defend against, with Initial Access Brokers dedicating their time to obtaining access vectors. Their sophisticated pricing models take into account factors such as access level, victims' annual revenue, and impact on critical infrastructure - incentivizing attackers to breach high-value targets.

Our analysis shows that RaaS actors have adopted more sophisticated cryptographic techniques compared to commodity actors in their operation and typically generate one address per victim to hide their identity. This allows RaaS to generate more revenue and

with higher level of protection, attracting more criminal groups to use RaaS to perform high profile attacks in recent years. RaaS actors are also more efficient at laundering ransom payments, as they move to launder funds within hours or days. Lastly, RaaS actors transfer revenue from ransom payments to mixers and other sophisticated laundering entities that increase the difficulty for law enforcement agencies to recover ransom payments.

By providing an extensive overview of ransomware payments and making our data available, we hope to provide insight into a cybercriminal economy that poses a severe threat to many organizations and societies, of which reporting is often fragmented.

REFERENCES

- [1] National Security Agency (NSA). *2021 Cybersecurity Year in Review*. <https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-VIEW/Article/2921744/nsa-releases-2021-cybersecurity-year-in-review/>. 2022.
- [2] AT&T Alien Labs *Open Threat Exchange*. <https://cybersecurity.att.com/open-threat-exchange>. 2021.
- [3] BBC. *HSE cyber-attack: Irish health service still recovering months after hack*. <https://www.bbc.com/news/world-europe-58413448>. 2021.
- [4] BBC. *Swedish Coop supermarkets shut due to US ransomware cyber-attack*. <https://www.bbc.com/news/technology-57707530>. 2021.
- [5] BBC. *Wannacry money laundering attempt thwarted*. 2017. URL: <https://www.bbc.com/news/technology-40826056>.
- [6] Crystal Blockchain. *Crystal Expert*. <https://crystalblockchain.com/crystal-expert/>. 2021.
- [7] Jack Cable. *Ransomwhere: A Crowdsourced Ransomware Payment Dataset*. Version 1.0.0. Zenodo, May 2022. DOI: [10.5281/zenodo.6512123](https://doi.org/10.5281/zenodo.6512123). URL: <https://doi.org/10.5281/zenodo.6512123>.
- [8] Catalin Cimpanu. *BTC-e founder sentenced to five years in prison for laundering ransomware funds*. 2021. URL: <https://www.zdnet.com/article/btc-e-founder-sentenced-to-five-years-in-prison-for-laundering-ransomware-funds/>.
- [9] Bleeping Computer. *Dutch supermarkets run out of cheese after ransomware attack*. <https://www.bleepingcomputer.com/news/security/dutch-supermarkets-run-out-of-cheese-after-ransomware-attack/>. 2021.
- [10] Canadian Centre for Cyber Security. *Ransomware*. <https://cyber.gc.ca/en/ransomware>. 2021.
- [11] U.S. Cybersecurity and Infrastructure Security Agency (CISA). *Alert (AA21-209A): Top Routinely Exploited Vulnerabilities*. 2021. URL: <https://www.cisa.gov/uscert/ncas/alerts/aa21-209a>.

- [12] European Union Agency for Cybersecurity (ENISA). *Threat Landscape report - 2021*. <https://www.enisa.europa.eu/topics/threat-risk-management/threats-and-trends>. 2021.
- [13] Europol. *darkmarket: world's largest illegal dark web marketplace taken down*. 2021. URL: <https://www.europol.europa.eu/media-press/newsroom/news/dark-market-worlds-largest-illegal-dark-web-marketplace-taken-down>.
- [14] Europol. *Ransomware: What you need to know*. <https://www.europol.europa.eu/publications-events/publications/ransomware-what-you-need-to-know>. 2021.
- [15] Andy Greenberg. *The Untold Story of NotPetya, the Most Devastating Cyberattack in History*. 2018. URL: <https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/>.
- [16] The Guardian. *WannaCry, Petya, NotPetya: how ransomware hit the big time in 2017*. 2017. URL: <https://www.theguardian.com/technology/2017/dec/30/wannacry-petya-notpetya-ransomware>.
- [17] Australian Government – Department of Home Affairs. *Australia's Ransomware Action Plan*. <https://www.homeaffairs.gov.au/cyber-security-subsite/files/ransomware-action-plan.pdf>. 2021.
- [18] Danny Yuxing Huang et al. "Tracking ransomware end-to-end". In: *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE. 2018, pp. 618–631.
- [19] The Wall Street Journal. *JBS Paid \$11 Million to Resolve Ransomware Attack*. <https://www.wsj.com/articles/jbs-paid-11-million-to-resolve-ransomware-attack-11623280781>. 2021.
- [20] U.S. Department of Justice. *Department of Justice Launches Global Action Against NetWalker Ransomware*. <https://www.justice.gov/opa/pr/department-justice-launches-global-action-against-netwalker-ransomware>. 2021.
- [21] U.S. Department of Justice. *Department of Justice Seizes \$2.3 Million in Cryptocurrency Paid to the Ransomware Extortionists Darkside*. <https://www.justice.gov/opa/pr/department-justice-seizes-23-million-cryptocurrency-paid-ransomware-extortionists-darkside>. 2021.
- [22] U.S. Department of Justice. *individual arrested and charged with operating notorious darknet cryptocurrency mixer*. 2021. URL: <https://www.justice.gov/opa/pr/individual-arrested-and-charged-operating-notorious-darknet-cryptocurrency-mixer>.
- [23] U.S. Department of Justice. *six charged with crimes related to virtual currency exchange business*. 2021. URL: <https://www.justice.gov/usao-nh/pr/six-charged-crimes-related-virtual-currency-exchange-business>.
- [24] U.S. Department of Justice. *United States Files A Civil Action To Forfeit Cryptocurrency Valued At Over One Billion U.S. Dollars*. <https://www.justice.gov/usao-ndca/pr/united-states-files-civil-action-forfeit-cryptocurrency-valued-over-one-billion-us>. 2020.

- [25] Harry Kalodner et al. “BlockSci: Design and applications of a blockchain analysis platform”. In: *USENIX Security Symposium*. 2020.
- [26] Eric Loui and Josh Reynolds. *CARBON SPIDER Embraces Big Game Hunting, Part 2*. <https://www.crowdstrike.com/blog/carbon-spider-embraces-big-game-hunting-part-2/>. 2021.
- [27] Microsoft. *How cyberattacks are changing according to new Microsoft Digital Defense Report*. 2021. URL: <https://www.microsoft.com/security/blog/2021/10/11/how-cyberattacks-are-changing-according-to-new-microsoft-digital-defense-report/>.
- [28] NPR. *Panic Drives Gas Shortages After Colonial Pipeline Ransomware Attack*. <https://www.npr.org/2021/05/11/996044288/panic-drives-gas-shortages-after-colonial-pipeline-ransomware-attack>. 2021.
- [29] Kris Oosthoek and Christian Doerr. “Cyber Security Threats to Bitcoin Exchanges: Adversary Exploitation and Laundering Techniques”. In: *IEEE Transactions on Network and Service Management* 18.2 (2021), pp. 1616–1628. DOI: [10.1109/TNSM.2020.3046145](https://doi.org/10.1109/TNSM.2020.3046145).
- [30] Masarah Paquet-Clouston, Bernhard Haslhofer, and Benoit Dupont. “Ransomware payments in the Bitcoin ecosystem”. In: *Journal of Cybersecurity* 5.1 (2019).
- [31] Robert Poulsen. *U.S. Accuses Russian of Money Laundering for Ryuk Ransomware Gang*. 2021. URL: <https://www.wsj.com/articles/u-s-accuses-russian-of-money-laundering-for-ryuk-ransomware-gang-11636741333>.
- [32] Reuters. *Crypto Giant Binance Kept Weak Money-Laundering Checks Even As It Promised Tougher Compliance, Documents Show*. 2022. URL: <https://www.reuters.com/investigates/special-report/finance-crypto-currency-binance/>.
- [33] Dark Web Solutions. *Dark Web Monitor*. <https://dws.pm/>. 2022.
- [34] Cisco Talos. *Translated: Talos’ insights from the recently leaked Conti ransomware playbook*. 2021. URL: <https://blog.talosintelligence.com/2021/09/Conti-leak-translation.html>.
- [35] McAfee ATR Operational Intelligence Team. *Take a “NetWalk” on the Wild Side*. 2020. URL: <https://www.mcafee.com/blogs/other-blogs/mcafee-labs/take-a-netwalk-on-the-wild-side/>.
- [36] Financial Times. *the rise of crypto laundries: how criminals cash out of bitcoin*. <https://www.ft.com/content/4169ea4b-d6d7-4a2e-bc91-480550c2f539>. 2022.
- [37] U.S. Department of Treasury Financial Crimes Enforcement Network. *Ransomware Trends in Bank Secrecy Act Data Between January 2021 and June 2021*. https://www.fincen.gov/sites/default/files/2021-10/Financial%20Trend%20Analysis_Ransomware%20508%20FINAL.pdf. 2021.
- [38] U.S. Federal Bureau of Investigation (FBI). *Common Scams and Crimes: Ransomware*. <https://www.fbi.gov/scams-and-safety/common-scams-and-crimes/ransomware>. 2021.

- [39] UK National Cyber Security Centre. *Mitigating malware and ransomware attacks*. <https://www.ncsc.gov.uk/guidance/mitigating-malware-and-ransomware-attacks>. 2021.
- [40] Kai Wang et al. “A Large-scale Empirical Analysis of Ransomware Activities in Bitcoin”. In: *ACM Transactions on the Web (TWEB)* 16.2 (2021), pp. 1–29.
- [41] Official Bitcoin Wiki. *Blockchain attacks on privacy*. <https://en.bitcoin.it/wiki/Privacy>. 2021.

4

ILLICIT REVENUE OF DARK WEB SHOPS

The Dark Web, primarily Tor, has evolved to protect user privacy and freedom of speech through anonymous routing. However, Tor also facilitates cybercriminal actors who utilize it for illicit activities. Quantifying the size and nature of such activity is challenging, as Tor complicates indexing by design. This paper proposes a methodology to estimate both size and nature of illicit commercial activity on the Dark Web. We demonstrate this based on crawling Tor for single-vendor Dark Web Shops, i.e., niche storefronts operated by single cybercriminal actors or small groups. Based on data collected from Tor, we show that just in 2021, Dark Web Shops generated at least 113 million USD in revenue. Sexual abuse is the top illicit revenue category, followed by financial crime at a great distance. We also compare Dark Web Shops' activity with a large Dark Web Marketplace, showing that these are parallel economies. Our methodology contributes towards automated analysis of illicit activity in Tor. Furthermore our analysis sheds light on the evolving Dark Web Shop ecosystem and provides insights into evidence-based policymaking regarding criminal Dark Web activity.

This chapter entitled *Quantifying Dark Web Shops' Illicit Revenue* by **Oosthoek, K.**, Van Staalduinen, M. and Smaragdakis, G. has been accepted for publication in IEEE Access.

4.1. INTRODUCTION

The World Wide Web (shortly Web) has been recognized as one of the greatest achievements of our times. It offers unprecedented opportunities for communication and commerce, and has truly revolutionized our lives. The original design of the Web did not have anonymity as a requirement. Any user browsing the Web leaves digital footprints that can be traced and unveil the user's identity [56, 29]. The public Web is also easy to crawl and index, hence collecting data to profile users. Users and administrators soon realized the privacy risks of the public Web and tried to protect content, user profiles, and communication with passwords and other authentication methods. Together with paywalls restricting access and thus indexing, this created the *Deep Web*, a part of the Web not indexed by search engines.

Over the years, many solutions were developed to offer anonymity to Web users ranging from end-to-end cryptography using public keys [19, 50], to Transport Layer Security (TLS) [1, 2], and anonymous communication [53]. While the first two communicate point-to-point, the latter is relayed, potentially better protecting user identity. The Onion Router (Tor) [48] is the most successful implementation for anonymous communication. Tor started as a US military project to protect the private communication of US military personnel deployed around the globe. Today, Tor is an independent overlay network of 7,000 nodes (relays) globally [58, 57].

Tor also is the infrastructure that supports the *Dark Web*, i.e., the Deep Web content that exists on overlay networks, called *darknets*, that operate on top of the public Internet. Darknets and Dark Web content can only be accessed with specific software, configurations, or authorization and often use a customized communication protocol. Moreover, Darknets can communicate and conduct business anonymously without revealing user information, e.g., the user's location or Internet Protocol (IP) address. The Dark Web became popular among activists as it protects the freedom of speech under duress and activists in different regions of the world, e.g., protesters in Arabic Spring [48], and whistle-blowers such as WikiLeaks [66].

Unfortunately, the anonymity by design facilitated by the Dark Web also was attractive to cybercriminals and terrorists. By some estimates, the illicit activity on the Dark Web exceeds 2 billion USD [13, 12, 60, 6, 59]. However, such reports do not reveal information about their data sources. Usually, they focus on large *Dark Web Marketplaces* that provide a platform for the anonymous distribution of illegal goods, e.g., guns, drugs, sexual abuse material, and stolen financial data. Many Dark Web Marketplaces have been prosecuted and seized by law enforcement agencies, e.g., DarkMarket [25] and Hydra [63].

In recent years, small shops, called *Dark Web Shops*, single-vendor shops run by individuals or small-scale collectives, have been added to the Dark Web ecosystem. There are many reasons these small individually owned shops became popular: (i) Readily available webshop software has enabled Dark Web retailers to sell illicit goods directly, without paying a commission to Dark Web Marketplaces [45]; (ii) Retailers on the Dark Web increasingly avoid affiliation with notorious Dark Web Marketplaces, which are frequently involved in geo-political power games [62]; (iii) The take-down of Dark Web Marketplaces has affected business continuity and trust of some of the retailers, leading them to initiate self-hosted shops [23].

Previous research has analyzed the Dark Web and tried to quantify revenue from illicit trading on the Dark Web. Most of these studies focused on Dark Web Marketplaces as they have been popular during the last years [8, 65, 33, 55, 45, 14, 40, 21, 5, 13, 12, 60, 59]. Furthermore, focusing on a single Tor domain expedites data collection. In this paper, we focus on the evolving ecosystem of individually owned shops, as a specific subset of the whole Dark Web ecosystem. We attempt to understand its structure, operation, payment revenue, and laundering strategies. We also compare the structure and operation of Dark Web Shops with Dark Web Marketplaces and investigate differences and similarities.

The Dark Web Shops ecosystem is a less well-studied portion of the Dark Web that is also fueled with cryptocurrencies, especially Bitcoin [65, 38, 17, 39]. Our study sheds light on the evolving Dark Web ecosystem and is one of the first large-scale studies to estimate the illicit revenue generated by Dark Web Shops and understand the popularity of abuse types such shops facilitate.

We provide timely and valuable insights, as many Dark Web Shop transactions are suspicious. According to forthcoming market regulation legislation, suspicious cryptocurrency transactions must be reported to the authorities. For example, from 2024, the European Union will enforce the new Markets in Crypto-Assets (MiCA) rules [49]. MiCA requires cryptocurrency exchanges and other service providers to identify issuers of cryptocurrency transactions and owners of self-hosted hardware wallets for cryptocurrency transactions over 1,000 Euros. We hope the insights provided in this study contribute to informed policy-making in this area.

The contributions of this paper can be summarized as follows:

- To collect input data for our methodology, we develop a crawler for illicit Tor onions to collect Bitcoin addresses and characterize associated illicit activities.
- We develop a methodology to perform extensive data cleansing on a dataset of illicit Tor domains to filter out non-illicit and duplicate Tor domains, unrelated and incorrectly formatted Bitcoin addresses.
- Our analysis of the Tor crawler data based on our methodology shows that the revenue of Dark Web Shops was at least 113 million USD in 2021.
- Our analysis shows that the top category of illicit offerings by revenue is sexual abuse, totaling close to 94 million USD revenue; followed at large distance by financial crime, accounting for more than 10 million USD.
- Our investigation shows no overlap between Bitcoin addresses we discovered related to Dark Web Shops and those released after the take-down of the largest Dark Web Marketplace, Hydra (that by some measures had 80% of the Dark Market Revenue share). This suggests that shops and marketplaces are parallel Dark Web ecosystems.
- Our analysis shows that cryptocurrency exchange platforms are used by both owners of Dark Web Shop and Dark Web Marketplaces, which motivates the need for continuous monitoring and regulatory intervention.

4.2. BACKGROUND

4.2.1. TOR

Tor is an abbreviation of The Onion Router [18]. It is the most popular software for dark-nets and is widely used for implementing *onion routing*, i.e., relaying traffic through multiple servers (relays) and adding additional encryption at each hop. The Tor core software and Tor Browser are free and open source. As a network, Tor is maintained by many volunteers running Tor nodes, collectively providing an overlay network intended to facilitate increased user privacy over the regular Internet, effectively hiding user IP addresses. Next to many Tor domains (also called *onions*) serving hypertext similar to the regular Hypertext Transfer Protocol (HTTP), the Tor network is also used to facilitate other Transmission Control Protocol (TCP) based services such as email (OnionMail) and instant messaging (Ricochet Refresh), which uses Tor for its peer-to-peer transactions. Many popular browsers are also able to route traffic over Tor for anonymity. The Tor network further provides bridges to the regular Internet to defeat government censorship in several jurisdictions, e.g., during the Arab Spring in late 2010 [48]. Today, more than 7,000 Tor nodes are online [58, 57].

4.2.2. BITCOIN

Bitcoin is a digital currency based on peer-to-peer technology [41]. As opposed to government-issued (fiat) currencies such as the US dollar, the Euro, and the pound sterling, which central banks control, Bitcoin is not overseen by a central authority. Transactions between users and the issuing of new Bitcoin are performed collectively by a global network of close to 15 thousand Bitcoin nodes [6], making it a decentralized currency. Bitcoin transactions, i.e., the transfer of value from one user to another, are effectively data structures broadcasted to the Bitcoin network, composed of at least one input and output. Inputs are quantities of Bitcoin controlled by the sender, with outputs specifying their destination. Every transaction represents a state transition in the blockchain, which is confirmed through mining, which leads to consensus. After confirmation, transactions are irreversible and are stored in the blockchain and propagated to all nodes in the network.

4.2.3. BITCOIN: REGULATION AND MARKET CAPITALIZATION

While Bitcoin was designed to function anonymously, its current mainstream usage has effectively made it pseudonymous. Based on Know Your Customer (KYC) legislation [47] rolled out in many jurisdictions, people are required to legally identify themselves when signing up with an exchange platform to be able to buy Bitcoin. The disclosure of their names makes it difficult to achieve complete anonymity when a transaction shows up in an investigation. Law enforcement investigators can link several steps back to their origin. Suppose this is an exchange platform that is registered as a benign financial service provider in a jurisdiction. In that case, they can order the exchange to disclose the user's identity behind the specific transaction. This opens up possibilities for forensic investigation through blockchain analysis.

In the current bear market (Fall 2022), Bitcoin's market capitalization is 400 billion USD on average [67], which is significantly less than its record market cap of 1,156 billion

USD in November 2021. The illicit activity in Bitcoin is estimated at 2 billion USD, i.e., less than 1 percent as reported (lower bound estimations) by blockchain analytics firms, e.g., Chainalysis [12], the nominal value is still considerable. Especially when taking into account that criminal activity like money laundering usually increases in times of economic downturn [26] and geopolitical tension [61]. Fortunately, Bitcoin's open ledger is a robust forensic tool, enabling unprecedented opportunities to track funds, especially when compared to tracing cross-border bank transactions.

4.3. RELATED WORK

Previous research studied the Dark Web and tried to quantify revenue from illicit trading on the Dark Web. Most authors have focused on Dark Web Marketplaces as they have been popular during the last years [8, 65, 33, 55, 45, 14, 40, 21, 5, 13, 12, 60, 59]. Relatively few studies focused on other parts of the Dark Web [8, 39, 34].

Christin et al. [14] crawled the Silk Road Marketplace and found it was primarily drug-oriented. Meiklejohn et al. [40] purchased items from various Dark Web Marketplaces to obtain seller Bitcoin addresses as input to clustering heuristics. Hiramoto and Tsuchiya [33] have analyzed Bitcoin transactions of addresses associated with seven Dark Web Marketplaces based on Bitcoin addresses gathered via walletexplorer.com [64]. Their analysis, however, didn't check if the addresses appeared on the actual Dark Web Marketplace. Hence they work with an indirect data source, solely relying on a clustering algorithm. Elbahrawy et al. [21] have focused on customer migration between different Dark Web Marketplaces based on pre-processed vendor data.

Bracci et al. [8] studied the selling of COVID-19 products in 194 different Tor outlets, specifically on selling vaccines. In earlier work, authors performed similar focused research into cybercriminal capabilities [65], stolen identity documents [55], firearms [45], and drugs [5].

Lee et al. [39] analyzed Bitcoin transactions to addresses scraped from Tor. The set of addresses was relatively small, but important insights about the Dark Web between 2013 and 2018 could be extracted. The scraped domains were categorized into several categories. Their analysis showed that over 80% of the Bitcoin addresses in the Dark Web were indeed used with malicious intent. Their study estimates the Dark Web revenue in their dataset to be around 180 million USD for the period between 2013 and 2018. Their seed dataset contained 85 Bitcoin addresses.

Paquet-Clouston et al. [46] used the co-spending heuristic [37] to estimate ransomware payments in Bitcoin. Based on an analysis of Bitcoin addresses from 35 ransomware families, they quantify the minimum worth of the ecosystem at over 12 million USD. However, they included addresses that represented 2 million USD in revenue afterward attributed to the Silk Road black market and thus cannot be fully accounted for as ransomware payments. A recent work by Oosthoek et al. [43] analyzed ransomware payments worth around 101 million USD in recent years, and they showed that there is no overlap between the Bitcoin addresses used for ransomware and those used in reported Bitcoin addresses from studies in the Dark Web.

Chainalysis [13, 12] publishes annual reports with estimations about the total revenue of illicit activity on the Dark Web and per category. The estimate for 2021 was 2.1 billion USD. Although the analysis provides valuable policy-making insights, their

methodology is proprietary. Moreover, they focus on Dark Web Marketplaces exclusively. United Nations [60, 59, 35] and Interpol [35] also publish reports for the revenue in the Dark Market, again focusing on notorious Dark Web Marketplaces and illicit activity such as drugs, trafficking, and guns. The sources of the data are also proprietary.

To our knowledge, our analysis is the first to provide a thorough methodology for the analysis of crawled Tor data. In our demonstration of its application, we shed a unique light on the evolving ecosystem of Dark Web Shops, based on a dataset with much higher coverage than previous studies.

4.4. METHODOLOGY

This section describes the Tor crawler we developed and implemented to collect content from onions. It also describes the Bitcoin address clustering methodology we used in our analysis. The cleansing methodology explicitly developed for this analysis is discussed separately in Section 4.5.

4.4.1. TOR CRAWLER

While search engines index the content of the regular Internet, such indexing is not possible on the Dark Web. Access to Dark Web data requires using specialized software, such as the Tor browser and the Tor relay client. Indexing of Dark Web content is further complicated by the fact that Dark Web domains are usually short-lived [52].

The Tor crawler that we utilize for our collection and analysis of Dark Web data was launched in 2013 as part of a research project [54] to increase the coverage of Dark Web that can be indexed beyond a small number of seed Tor domains that can be found on the publicly accessible Web (clearnet). Today, the data collected by the crawler is available as a commercial product, called Dark Web Monitor (DWM), mainly to law enforcement agencies worldwide by CFLW Cyber Strategies. The crawler has provided insights that law enforcement agencies and prosecutors have utilized in recent years.

The crawler maintains a list of onions and adds new domains when they are discovered in the crawling process. Every onion is crawled at least every 18 hours. This ensures that even short-lived domains are crawled and indexed. For each onion indexed, the crawler follows all address paths from pages available within the domain (page tree). If a previously unseen domain is discovered, the crawler will automatically crawl that URL to add it to the archive and schedule for automatic crawling of the new URL. One of the main challenges is to have a complete overview of onions, as this is not facilitated and, on a technical level, not supported by the Tor network itself. This ‘snowballing’ approach of scanning all pages for new URL entries recursively leads to new entries which each crawl.

When a Tor domain is offline, either because it is not active anymore or due to temporal unavailability, e.g., outage or routing issues, the Tor domain is revisited with an inter-visit interval of 1 hour. In the case that the Tor domain continues to be unavailable after three attempts, the crawling schedule for this domain follows an exponential back-off, i.e., the Tor domain is visited after 18, 36, 72 hours up to a maximum revisit regularity of 10 days.

For the content analysis, the crawler uses regular expressions. It automatically ex-

Abuse Type	Description
Cybercrime	DDoS, bulletproof hosting, exploit development
Drugs / Narcotics	Cannabis, synthetic drugs, pharma
Extremism	Extreme right, radical islam, anarchists, neo-nazi
Financial Crime	Carding, hacked accounts, stolen giftcards
Goods and Services	Marketplaces, counterfeit, gambling, firearms
No Abuse	Whistleblower, communities, how-to's, non-profit orgs.
Sexual Abuse	Child sexual abuse, animal sexual abuse, sextortion
Violent Crime	Assasination, hate crime

Table 4.1: Categorization of abuse types used to classify Tor domains based on content for the Tor domains indexed by our crawler.

tracts cryptocurrency addresses, PGP keys, and email addresses that can be used for attribution. The raw data is archived in cloud storage buckets. Since its launch, the data accounts for 25 Terabytes (until end of first semester 2022), 15 Terabytes collected during 2021 alone.

Our analysis shows that multiple cryptocurrencies are used for illicit activity, namely, Bitcoin, Bitcoin Cash, Litecoin, Monero, Ethereum, and Binance Coin. However, our analysis confirms previous results [39] that, by far, the most popular cryptocurrency is Bitcoin. Indeed, around 99% of all addresses discovered in our dataset is Bitcoin.

For page content classification, we use human crowdsourcing to categorize the content. Each newly crawled domain is inspected by a team of analysts that, based on the page content, assigns a label indicating the primary type of abuse observed on that particular domain. A domain may be assigned to more than one human analyst to improve the accuracy of the labeling. An overview of the so-called “abuse types” used in our study is available in Table 4.1. We notice that our categorization and description do not follow other proposed, but not yet standardized categorizations [36].

For our study, we utilize the latest version of our Tor crawler [54], introduced in 2020. The latest version of the Tor crawler establishes over 100 parallel connections and makes it possible to scan all known Tor domains within 24 hours. Since its launch, the crawler is estimated to have indexed about a fifth of Tor domains based on statistics published by the Tor project [57], i.e., approximately 1.5 million unique Tor domains. The un-indexed domains are primarily onions serving non-HTTP protocols. Approximately 100 thousand new unique online domains were crawled and indexed in the first semester of 2022. This figure accounts for the many mirrors used by actors to increase the resiliency of their operations. Such duplicates are aggregated within a single domain ID if the HTML source code is identical to another domain.

Our crawler has certain limitations. Onions may be protected by CAPTCHA [8, 16], making crawling and indexing challenging. This is typically true for Dark Web Marketplaces but not for Dark Web Shops. Indeed, popular Dark Web Marketplaces are usually protected with CAPTCHA or user passwords, e.g., Hydra Market, and are not indexed partially or not indexed at all. A recent study [17] shows that coverage of scrapers of Dark Web Marketplaces is usually low, missing on average 46% of the listings. Due to this,

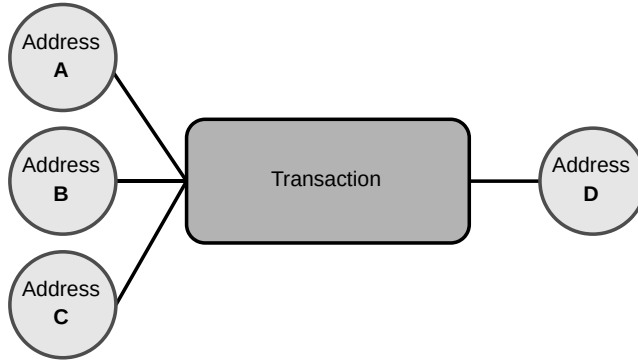


Figure 4.1: Bitcoin address co-spending heuristic: Example of three Bitcoin addresses with common spending in one transaction.

4

the actual revenue of Dark Web Marketplaces is systematically underestimated. On the contrary, due to the implementation of standard off-the-shelf software suits that do not support CAPTCHAs by default, many single vendor shops do not (yet) use CAPTCHAs. The current version of our crawler does not crawl nor index domains protected by a user login. However, it crawls and indexes the front page of the Tor domain. This leads to a partial view, meaning that only non-protected onions are fully indexed. Our analysis provides a lower bound of the estimate of illicit revenue by Dark Web Shops. Moreover, not all the Tor domains are scanned with the same frequency. Thus, it is possible to have a less accurate index for high dynamic content domains when compared with static content domains. This is a limitation of any crawling process and also applies to many crawlers that index the publicly accessible Web.

4.4.2. BITCOIN ADDRESS CLUSTERING

Bitcoin address clustering aims to break pseudoanonymity in blockchain by linking Bitcoin addresses that are controlled by the same entity based on the information available from blockchain transaction analysis. Several heuristics have been proposed to achieve Bitcoin address clustering based on different assumptions of how users transact in a blockchain [40, 42, 31]. To discover whether a Bitcoin address belongs to a cluster of multiple addresses, we use GraphSense [32], which builds on BlockSci [37]. To discover Bitcoin address clusters, called entities, GraphSense exclusively uses the co-spending heuristic, also known as multi-input, which high effectiveness has been shown empirically [40, 3, 31]. The co-spending heuristic recursively queries addresses that were used to combine funds in a transaction. If a transaction has input from multiple addresses, these are all likely controlled by the same actor (individual or group). Figure 4.1 provides a graphical representation of this hypothesis.

While the co-spending heuristic is generally reliable, it might lead to false positives caused by CoinJoin and PayJoin transactions [32]. CoinJoin and PayJoin are privacy-preserving transaction methods that combine payments of multiple parties into one transaction to obfuscate ownership. GraphSense uses the algorithm proposed by Goldfeder

et al. [28] to identify the most common types of CoinJoin transactions and exclude these from input into the clustering heuristic. Another common heuristic, the change address heuristic, isn't implemented in GraphSense as its reliability has been proven inconsistent due to its dependence on critical characteristics in end-user wallet software [37].

4.5. CLEANSING METHODOLOGY

Data crawled from Tor is inherently noisy. Proper filtering will provide a more accurate portrayal of the relevant ecosystem. In this section, we present our methodology to remove corrupted, incorrectly formatted, duplicate, or incomplete data, i.e., to perform extensive data cleansing, resulting in a dataset that can serve as a basis for dependable lower-bound estimates. Our methodology, described in detail below, focuses on cleansing three core aspects of our data set: Tor domains (onions), Bitcoin addresses, and Bitcoin address clusters detected with the co-spending heuristic. For an illustration of the pipeline of our methodology, we refer to Figure 4.2. We hope to contribute to the standardization and replication of analyses like ours by providing a detailed design and evaluation of our methodology.

4.5.1. TOR DOMAINS

A portion of Tor domains is legal, with the facilitation of anonymous, licit services as the sole intention. Our analysis exclusively focuses on *illicit*, i.e., unlawful criminal activity. This means we solely regard pairs of Tor domains and Bitcoin addresses linked to suspicious, or likely illegal, activity, which we confirm through inspection of each pair. This inherently leads to lower-bound results, as the relationship between many domains and addresses needs to be clarified, leading to exclusion from analysis. We only include domain-address pairs which are manually validated as illicit.

The initial stage of our cleansing methodology focuses on filtering out non-illicit or otherwise unwanted domains. Each domain represents a unique address in the `.onion` special-use top-level domain. The key objective of the first cleansing phase is to establish relationships between a Tor page with an illicit offering and a Bitcoin address. These relationships can be one-to-one, meaning an individual domain contains a single valid Bitcoin address or one-to-many, i.e., it contains more than one address.

We focus exclusively on the entire year of 2021, as the latest crawler version was introduced in 2020. From our crawler, we obtained Tor domains, also referred to as *onions*, which appeared online between January 1 and December 31, 2021. The content was collected and indexed for each Tor page crawled. Domain names, Bitcoin addresses, and page titles were parsed from the crawler collection. Other metadata and page sources were stored separately for reference. To each Tor domain, a label was added indicating the abuse type as listed in Table 4.1. These labels are assigned by a team of analysts that manually inspect newly crawled pages. Domains clearly non-illicit, i.e., of civil rights organizations, political parties, or whistle-blower sites, are classified as *No Abuse* in our study. Note that this provides a two-step approach to establish the illicit nature of domains: (i) during the labeling of newly scraped domains and (ii) in our manual analysis of remaining domain-address pairs after completion of all steps of the methodology.

The corpus of collected raw data analyzed for this paper includes 72,595 unique do-

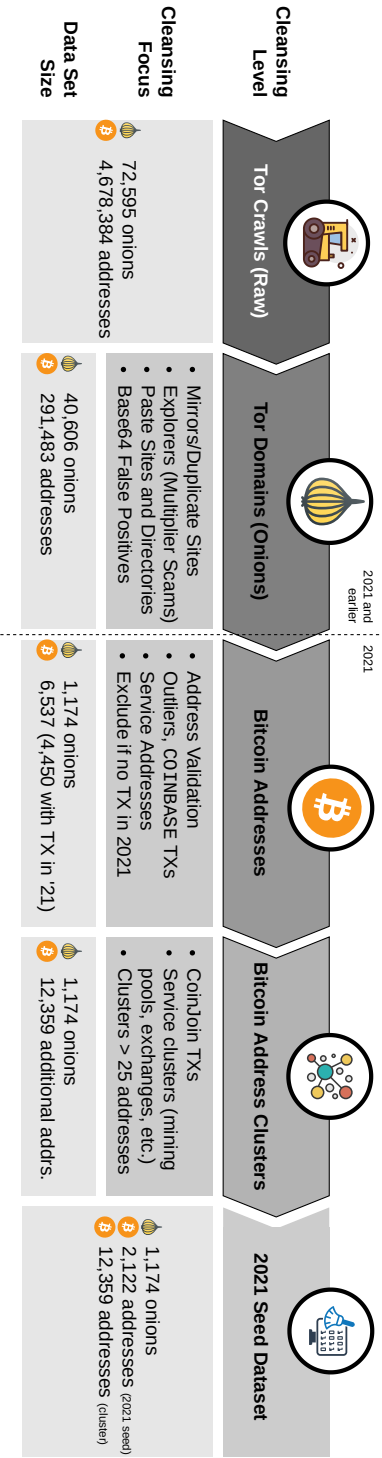


Figure 4.2: Pipeline of our dataset cleansing methodology.

mains which appeared online in the Tor network at some point in 2021. The crawler collected and indexed content from these domains for 710,484 pages (URLs). After analyzing the content, 138,967,218 non-unique cryptocurrency addresses were extracted. A single cryptocurrency address can be detected within multiple Tor domains. This primarily occurs due to mirrored domains and the presence of blockchain explorers, which display recently mined blocks, addresses, and transactions. After our analysis, we identified 4,730,419 unique cryptocurrency addresses, of which the vast majority, i.e., 4,678,384 were Bitcoin addresses. These addresses are unverified, meaning that they are formatted as a Bitcoin address but not yet sanity-checked and confirmed by a Bitcoin node as valid. This happens in a consecutive cleansing phase. With Bitcoin dominating our dataset, with 98.9% addresses being Bitcoin, dominant over other detected cryptocurrencies, we focused on Bitcoin exclusively as the dominant currency in the Dark Web.

MIRRORS

Our raw dataset contains over 70 thousand unique Tor domains. However, owners of Tor sites use multiple redundant domains, and often infrastructure is taken offline and made available again on a new domain. As our crawler saves the page tree with each visit, we were able to filter out full-match duplicates based on the title of the frontpage (`index.html`) and the contents of the page source. Based on this, we identified 51,324 unique onions in our dataset for the year 2021.

NON-ILLICIT/UNWANTED DOMAINS

We excluded Tor domains that did not fit our classification of illicit activity in Table 4.1, focusing on outliers by rank-ordering the domains in our dataset based on the number of Bitcoin addresses per individual domain. This reduces the initial 51,234 domains to 40,606 unique domains, primarily due to the exclusion of three categories:

(i) Explorers: Our crawler output contains domains that automatically post block mining output, similar to blockchain explorers such as *blockchain.com*. These sites are advertised as Bitcoin multipliers, displaying recent transaction data as proof of their supposed capabilities, a tactic also observed by previous studies [22]. While these apparent scams extort money from unaware victims and, thus, are illicit, the Bitcoin addresses advertised are unrelated. Hence we excluded such domains from our analysis. We did this based on a rank order of address quantity per domain and manual inspection.

(ii) Indexes and Directories: We also exclude index sites and Tor directories. These sites, which also exist on the public Web, serve as springboards linking to various Tor hidden services. Some of these host copies of specific pages they are linking to, causing duplicate pages found on different sites. We also removed non-illicit pages that appeared on illicit domains, as these also cause duplicates.

(iii) Paste sites and Forums: The set of Tor domains was manually inspected to remove further sites that weren't clearly illicit. Notable examples of excluded domains are paste sites listing Bitcoin addresses without clear context and forum posts referring to Bitcoin addresses without clear intent. Messages in foreign languages were automatically translated and manually inspected to understand the context, and the corresponding Bitcoin addresses were only preserved when in scope.

FALSE POSITIVES

In this step, we removed false positives caused by domains using inline Base64-encoded images often used to slow down crawlers [14], of which a portion was detected as a Bitcoin address by our crawler. We also checked whether each domain had a label indicating the abuse type attached and additionally checked each abuse type for correctness using a random sample of 100 domains. This first phase of cleansing results in an intermediate dataset of 40,606 unique Tor domains with 291,483 unique Bitcoin addresses.

4.5.2. BITCOIN ADDRESSES

After filtering out non-illicit and unwanted Tor domains, we also need to filter out Bitcoin addresses unrelated to illicit activity. We assume the exact requirement that an address and the majority of its holdings in Bitcoin should be confidently classified as illegal. This isn't straightforward due to Bitcoin's privacy and pseudonymous characteristics. Hence we opted for a lower-bound estimate, excluding all addresses which can be attributed as belonging to a Bitcoin exchange platform. For such addresses, a portion of holdings is likely illicit, but the proportion cannot be reliably established. The domain itself was excluded from further analysis when all Bitcoin addresses detected in a single Tor domain were excluded.

ADDRESS VALIDATION

We checked the remaining 291,483 Bitcoin addresses against a Bitcoin node for validity. 38,212 addresses were reported as invalid, i.e., sanity checks such as for address formatting did not pass the test, or the existence of the address hasn't yet been confirmed in block mining. Out of the 253,271 valid addresses, a remarkable quantity of 246,187 (97.2%) had no transactions, meaning they were never used according to the blockchain data. These addresses cannot represent any illicit activity, so they were also disregarded, resulting in 7,084 valid addresses with one or more transactions.

OUTLIERS

This step excluded outliers based on the number of observations of individual Bitcoin addresses in different Tor domains and the total holdings of these addresses. Based on this, we excluded Bitcoin addresses found in Tor domains with Bitcoin "Rich" Lists, i.e., displaying Bitcoin addresses with the biggest holdings. We also excluded several Bitcoin addresses if they historically only received COINBASE transactions, which indicates they belong to mining pools. COINBASE transactions (not to be confused with the exchange platform of the same name) are newly mined coins issued as a block reward, which cannot be related to illicit activity. This was furthermore validated by excluding mining-related addresses shared by Romiti et al. [51] and GraphSense [30].

SERVICE ADDRESSES

We refer to addresses controlled by centralized exchange platforms such as Coinbase and Kraken as service addresses, as the exchange service owns the private key of the addresses used for deposit and withdrawal. This also includes addresses associated with Bitcoin-accepting payment providers and gambling sites, which store user-owned Bitcoin in custody [44]. Exchange platforms are of great importance to blockchain analysts because they provide an opportunity to identify real-world actors behind Bitcoin

transactions if the exchange adheres to Know Your Customer (KYC) legislation. However, addresses operated by exchanges likely represent the holdings of more than one user. Furthermore, ownership of funds can be transferred without on-chain evidence through paper wallets or shared credentials.

As we cannot reliably classify funds terminating at exchanges as illicit, we have excluded these from our analysis based on two metrics. First identified exchanges using labels from GraphSense [32], walletexplorer.com [64], and BitRank [7] (a commercial service with a free daily allowance). If one or more of these services identified an address controlled by an exchange, it was excluded. Addresses with more than 1,000 incoming transactions were also excluded. In total, 547 addresses were removed, further decreasing our set of addresses to 6,537.

By filtering out exchanges and mining-related addresses, we likely also exclude from our dataset the portion of revenue sent to that address. Filtering out addresses with over 1,000 transactions may also exclude non-exchange addresses. This is a well-considered step in our approach to a conservative but clean estimate. We strive to exclude any funds that cannot reliably be attributed to an illicit offering on Tor.

BITCOIN TRANSACTIONS IN 2021

For our analysis, we focus on the year 2021, which is an entire year with the latest version of the crawler. To get an impression of what 2021 looked like in terms of illicit revenue by Dark Web Shops, we only regarded transactions between January 1 and December 31, 2021. We filtered for addresses ‘active’ in 2021, i.e., with one or more transactions during the above period. This filter reduced the corpus of Bitcoin addresses from 6,537 to 4,450. Tor domains with exclusively Bitcoin addresses that didn’t have any transactions in 2021 were also excluded. As a result of the last filter, the amount of Tor domains included dropped to 1,174.

4.5.3. BITCOIN ADDRESS CLUSTERS

For Bitcoin address clustering, we used GraphSense [32], which builds on BlockSci [37]. GraphSense uses BlockSci’s ability to detect the most common types of CoinJoin and does not detect any when we apply it to our dataset. According to labels from various sources described earlier, using privacy wallets such as Wasabi was also non-existent. Previous reports also mentioned that off-the-shelf Dark Web store frontend software such as Eckmar [20] and TradeMed [4] have become more sophisticated and generate new Bitcoin addresses for each purchase by default. This makes address clustering more challenging.

We excluded probable service clusters if one or more of the following two criteria were met: (i) the cluster contains more than 1,000 addresses and (ii) if one or more of three unique sources (Graphsense [32], walletexplorer.com [64], BitRank [7]) attributes the cluster itself or one or more addresses in a cluster to an exchange platform.

The most significant effect due to this exclusion of service clusters occurred in the Financial Crime category. The identification and subsequent exclusion of clusters of exchanges, Service Clusters, also leads to the exclusion of service addresses in the seed dataset. Because of this, our final number of seed addresses used for analysis is 2,122. This is a significant reduction, the process of which is represented in Section 4.5. The

Dataset Statistics			Filtered Dataset - Transactions and Revenue			
Category	# Tor domains (# pages)	# BTC addresses seed unfiltered / filtered (co-spent unfiltered / filtered)	Transactions		USD Revenue	
			incoming 2021 seed (co-spent)	outgoing 2021 seed (co-spent)	USD received 2021 seed (co-spent)	USD sent 2021 seed (co-spent)
<i>Cybercrime</i>	46 (1,287)	236 / 110 (1,186,952 / 132,092)	577 (6,338)	264 (621)	\$141,329 (\$1,389,204)	\$141,177 (\$1,390,486)
<i>Drugs / Narcotics</i>	17 (392)	104 / 45 (493,877 / 271,362)	1,161 (4,553)	290 (597)	\$330,983 (\$1,594,520)	\$328,764 (\$1,414,285)
<i>Extremism</i>	12 (7,683)	19 / 17 (5880 / 231)	150 (4,001)	33 (246)	\$13,053 (\$577,574)	\$8,461 (\$509,745)
<i>Financial Crime</i>	227 (31,785)	3968 / 397 (35,272,512 / 548,051)	1,948 (45,768)	3,305 (4,337)	\$231,308 (\$10,164,827)	\$213,702 (\$8,062,361)
<i>Goods and Services</i>	41 (2,107)	112 / 67 (41,632 / 18,645)	331 (11,172)	231 (1,072)	\$86,766 (\$1,082,019)	\$85,970 (\$1,028,882)
<i>No Abuse</i>	15 (44)	160 / 79 (501 / 418)	3,303 (152,958)	319 (896)	\$1,795,163 (\$3,845,552)	\$1,791,164 (\$3,845,551)
<i>Sexual Abuse</i>	836 (29,870)	1945 / 1403 (5,532,538 / 1,108,367)	6,636 (61,400)	1,563 (5,151)	\$441,363 (\$94,257,825)	\$390,682 (\$94,241,807)
<i>Violent Crime</i>	3 (41)	29 / 4 (1124 / 7)	13 (15)	3 (4)	\$1,401 (\$10,601)	\$1,109 (\$4,532)
Total	1,197 (73,209)	6537 / 2,122 (42,535,016 / 2,079,173)	14,119 (286,205)	6,008 (12,924)	\$3,041,376 (\$112,922,122)	\$2,961,029 (\$110,497,649)

Table 4.2: Overview of our analysis for Dark Web Shops in 2021. Revenue is in USD, rounded to the nearest whole USD. The initial values correspond to the seed Bitcoin addresses. The values in parentheses correspond to the values after Bitcoin address clustering and data cleansing (“Filtered Dataset”).

illicit revenue represented by this set of addresses is a lower-bound estimate of overall illicit revenue in Tor related to Dark Web Shops. However, due to the various steps taken, we are confident that as opposed to the initial 291 thousand addresses, the 2,122 seed addresses provide a robust representation of payment size, buyer activity, and distribution between different types of illicit activity related to the Dark Web Shops.

4.6. QUANTIFYING ILLICIT REVENUE

Based on the methodology outlined in Section 4.5, in this section, we provide an overview of illicit revenue made by Dark Web Shops in the entire year of 2021. We discuss results from the analysis of incoming and outgoing Bitcoin transactions to the set of Bitcoin seed addresses, as well as based on an expanded set of addresses, using the heuristics discussed in Section 4.4.

4.6.1. SEED ADDRESS REVENUE PER ABUSE TYPE

Table 4.2 provides an overview of the results of our analysis by type of abuse, being the type of illicit activity (in the first column). We refer to Table 4.1 for a description of each abuse type. In the second column, we provide the number of onions and affiliated pages per abuse type. Although the number of domains is in the order of tens, the number of affiliated pages is typically in the order of thousands. Sexual abuse and financial crime are the two categories with the highest number of Tor domains or onions and pages, with around thirty thousand affiliated pages each, i.e., around 82% of the domains are associated with these two categories. Notice also that the No Abuse category is very small.

As discussed in Section 4.5 it only contains a small number of civil rights organizations and whistle-blower sites; onions not evidently non-abusive were not considered in our analysis. As shown in Table 4.2 total, for our analysis, we consider 1,197 Tor domains and 73,209 pages.

The third column presents the number of seed Bitcoin addresses included per abuse type. Unfiltered is the raw crawler result, and the filtered number is after the application of our methodology. For our analysis, we utilize the set of filtered seed addresses after the cleansing data process described in the previous section. In parentheses, we provide the results of Bitcoin address clustering. For completion, we report both the output of the clustering (unfiltered) and the results after cleansing (filtered). In our analysis, we take a conservative approach by only considering the filtered set of Bitcoin addresses and filtered clusters. Again, the popular categories are sexual abuse and financial crime, with more than a million and half a million associated Bitcoin addresses. More than 270 thousand Bitcoin addresses are also associated with the drugs/narcotics category. Overall, in our study, we consider 2,122 seed Bitcoin addresses and, in total, 2,079,173 Bitcoin addresses after address clustering and cleansing.

For the analysis of transactions to and from seed addresses, we focus on the set of transactions without parentheses in columns four and five. Transactions for sexual abuse and financial crime dominate, with about half of the total incoming and outgoing transactions being attributed to these two types of abuse. We also notice that there is a significant imbalance between the number of incoming (14,119) and outgoing transactions (6,008). This is also the case for incoming/outgoing transactions for each and every individual category. This is to be expected as the payments are at a given price of the product, and the outgoing transactions (laundering) are typically aggregated into bulk transactions.

The last two columns of Table 4.2 show the revenue per category for the incoming and the outgoing transactions, respectively. Our estimation of the revenue in USD is based on the daily average Bitcoin-USD exchange rate extracted from CoinGecko's API [15]. All USD values are rounded to the closest USD. We focus again on the values in the parentheses that correspond to the revenues of the transactions of Bitcoin addresses after clustering and cleansing (filtered dataset). For a complete reference, we provide in Table 4.7 (in Appendix A) the results when we consider Bitcoin address clustering without filtering (unfiltered dataset). The total revenue of both the incoming and outgoing transactions exceeds trillions which are totally unrealistic. Even for individual categories, e.g., sexual abuse and financial crime is in the order of hundreds of billions, again not realistic. This further justifies our decision to take a conservative approach and use the filtered data following the cleansing process introduced in 4.5.

4.6.2. LONGITUDINAL ANALYSIS OF SEED ADDRESS TRANSACTIONS

We also have examined the longitudinal revenue of the shops in our dataset per individual abuse category. In Figure 4.3 and 4.4 we plot the revenue per abuse type per month for all the abuse types provided by the Dark Web Shops in our study. Sexual abuse and financial crime again appear as the most high-ranking categories over the entire year, but without significant variation. The contribution of the other categories is relatively stable over time. Regarding overall revenue, although there is more activity during the

Abuse Type	Payments in 2021	Percentile						
		Min	Max	Median	Std Dev	75%	90%	99%
Cybercrime	552	\$0.18	\$50,013.85	\$55.90	2,175.98	\$138.14	\$299.69	\$1,583.66
Drugs / Narcotics	1141	\$0.19	\$6,817.84	\$121.94	534.32	\$285.13	\$617.34	\$2,780.16
Extremism	146	\$0.31	\$2,554.18	\$17.56	260.00	\$62.38	\$105.35	\$1,022.45
Financial Crime	1744	\$0.18	\$9,480.63	\$58.64	418.71	\$102.47	\$178.81	\$1,288.35
Goods and Services	323	\$0.18	\$8,215.41	\$51.28	696.16	\$275.26	\$699.69	\$2,535.59
No Abuse	3303	\$0.16	\$157,043.02	\$3.03	1,637.95	\$4.28	\$11.66	\$640.63
Sexual Abuse	3,946	\$0.17	\$17,957.43	\$31.67	327.51	\$49.23	\$85.56	\$599.64
Violent Crime	13	\$4.92	\$413.59	\$66.67	121.23	\$149.03	\$247.12	\$395.32

Table 4.3: Dark Web payment statistics in 2021 based on seed addresses from our crawler output.

4

first part of the year, an evident seasonal trend is absent. We note that some of the fluctuations may be related to the take-down of shops or the launch of new in some categories that are beyond the scope of this study. One example of such fluctuation is the outlier for Cybercrime in August, which is related to the purchasing of a stolen Bitcoin wallet. We analyzed this and left it in because, based on blockchain transaction data, it seemed authentic.

4.6.3. REVENUE PER ABUSE TYPE AFTER ADDRESS CLUSTERING

The last two columns of Table 4.2 also provide the revenue per abuse type after retrieving additional addresses based on our clustering algorithm, as discussed in Section 4.4. The aggregate estimated incoming revenue is around 113 million USD. The estimated total outgoing revenue is around 110.5 million USD. This shows that although there is an asymmetry in the number of transactions, the incoming/outgoing revenue is rather balanced. Thus, the outgoing transactions are made in bulk, but almost the total incoming revenue is laundered within a year. Notice that some incoming or outgoing transactions may occur in the previous or following year, respectively. Then we focus on the individual categories. Sexual abuse contributes by far the most to the incoming illicit activity revenue of Dark Web Shops. Around 94.2 of 112.9 million incoming revenue is associated with sexual abuse, i.e., more than 83% of the illicit revenue of Dark Web Shops. The second contributor is financial crime, with 10.1 million USD, i.e., around 9% of the illicit revenue. The rest of the contributors in the top 5 list are drugs/narcotics, cybercrime, and goods and services, with approximately 1.6, 1.4, and 1.1 million USD in revenue, respectively.

In Table 4.3, we show the distribution of payments (incoming transactions) to the Bitcoin seed addresses in 2021 per abuse type. We observe that there is a significant difference between the minimum and maximum transaction values. Indeed, the minimum value is typically cents, while the maximum value is multiple thousands of USD. The median values, however, are more representative of the type of business for Dark Web Shops, in the orders of tens of USD. The 75-percentile values are similar to the median values, which is another indicator that the product's price is in the range of 50 to 500 USD. Our observations concur with independent studies for the individual use of drug unit prices and unit prices for other illicit activities [59].

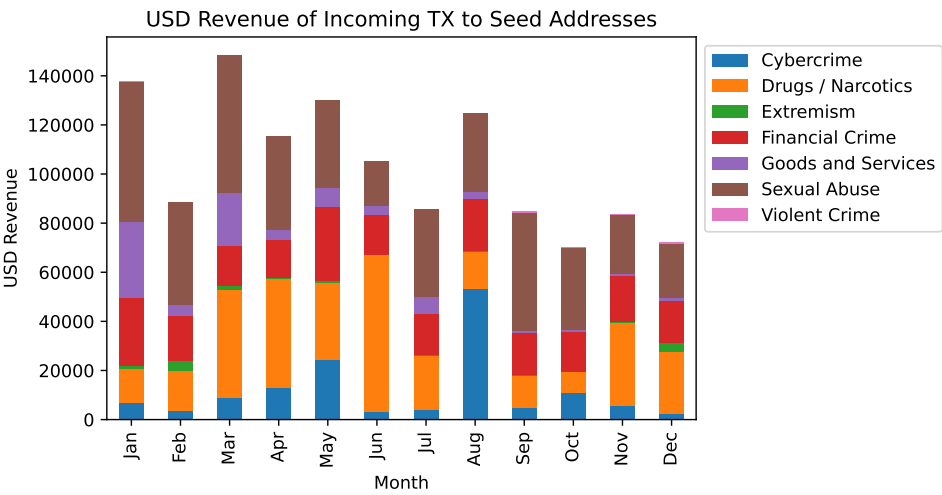


Figure 4.3: USD Revenue of incoming transactions to seed addresses found in the Dark Web in 2021 using our crawler.

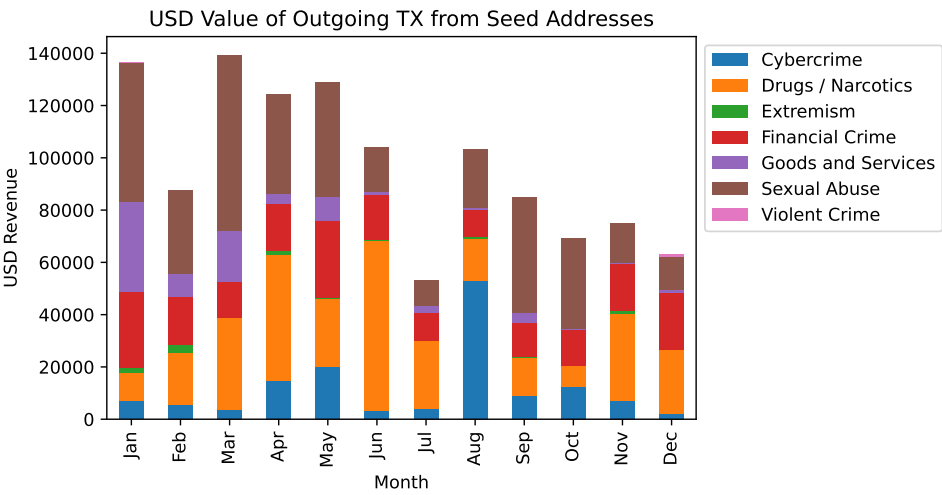


Figure 4.4: USD Value of outgoing transactions from seed addresses found in the Dark Web in 2021 using our crawler.

Lower Bound, Filtered Clustering		
	USD Received	USD Sent
Min	\$1	\$1
Max	\$6,095,231	\$6,095,013
Median	\$9,890	\$3,514
Std Dev	\$164,802	\$79,969
Total	\$485,005,353	\$380,433,794

Table 4.4: Hydra 2021 Revenue - based on 2021 transactions to co-spent clusters of Hydra addresses.

4

4.7. QUANTIFYING SHOP VS MARKETPLACE REVENUE

In this section, we compare the revenue characteristics, operation, and laundering practices of Dark Web Shops with those observed for Dark Web Marketplaces. Recall that Dark Web Shops are run by individual actors and small groups, selling illicit merchandise to customers directly. On the contrary, Dark Web Marketplaces are run by criminal conglomerates, offering themselves and, against a commission, other criminal actors a marketplace to sell, typically, illicit goods.

4.7.1. THE HYDRA MARKETPLACE AND ITS TAKE-DOWN

Hydra was launched in 2015 and has been recognized as one of the largest Dark Web Marketplaces primarily selling drugs in former Soviet bloc countries such as Russia, Ukraine, Belarus, and Kazakhstan. According to an industry report by Chainalysis [13] Hydra was the dominant Dark Web Marketplaces in 2021. This report estimated that the total revenue of Dark Web Marketplace was around 2.1 billion USD, and Hydra's market share was around 80%.

After being the target of law enforcement scrutiny for many years, at least a large part of Hydra infrastructure was taken down in April 2022 by German authorities [10]. The seized server infrastructure reportedly contained more than 17 million user accounts and 19 thousand seller accounts [9]. While many accounts might be superfluous, as Dark Web marketplaces usually do not provide account password reset functionality, these numbers provide an idea of the scale of its customer base. The US Treasury Department publicly released 117 associated Bitcoin addresses associated with Hydra after its take-down by German authorities [62, 63]. The press release by German authorities also claimed Hydra's role as the biggest marketplace [10]. According to data from our crawler, Hydra still partially remains online.

The release of Bitcoin addresses seized by law enforcement allowed us to extract Hydra's transactions in 2021 and use these as input to our clustering algorithm. Based on that, we are thus able to establish a reliable sample of Hydra's revenue in 2021. In filtering, we excluded transactions to the address of Garantex Exchange, also included in the press release [63]. Garantex was an affiliated money laundering service seized simultaneously with Hydra. Inclusion of its Bitcoin address would wrongly multiply reported revenue.

The revenue in USD of incoming transactions to the seed addresses reported by the

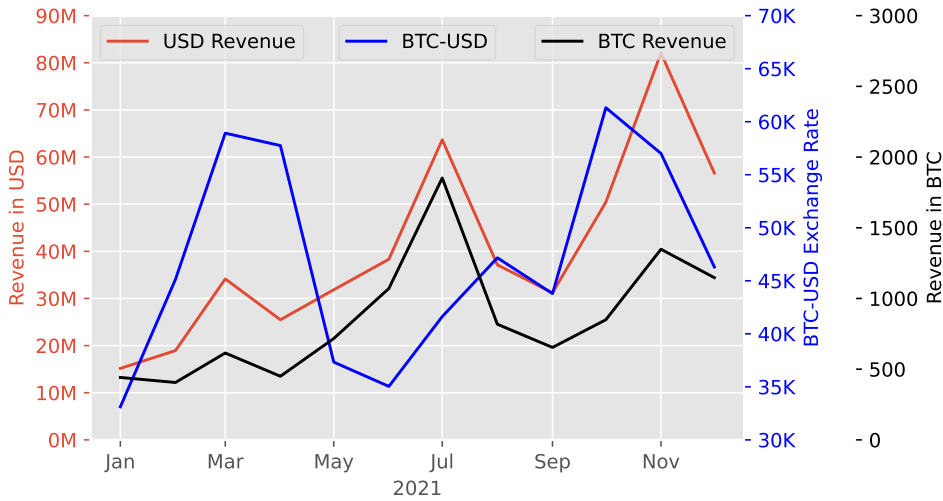


Figure 4.5: Incoming Transaction Revenue to Hydra Address Clusters in 2021.

US Office of Foreign Assets Control (OFAC) was 792.6 million USD, with the earliest incoming transaction on April 25, 2015. Based on the set of seed addresses, 64 Bitcoin address clusters were discovered, of which the largest had over 6,028,684 Bitcoin addresses. 40 Bitcoin addresses reported by OFAC did not belong to a cluster, which means co-spending did not take place.

In Table 4.8 (see Appendix A), we provide the revenue with Bitcoin address clustering without filtering. Again, the number is in the order of multiple billions, and although this is mentioned in some reports [24] as correct, we deem this is caused by address clusters of the Garantex Exchange previously mentioned. The revenue flowing into Garantex can not be fully attributed to Hydra. Without the removal of this cluster, the total incoming payments would have been around 7.6 billion USD.

Some industry reports claim that Hydra was involved in ransomware operations [24]. However, when we compared the Hydra-associated addresses with the publicly available Bitcoin addresses used in ransomware campaigns [43], we did not find any match.

4.7.2. DARK WEB SHOPS VS. HYDRA TRANSACTIONS REVENUE

In Table 4.4, we report Hydra's revenue (in USD) of incoming and outgoing transactions. A first observation is that the median transaction value for Hydra is in the orders of thousands of USD compared to the tens of USD in the Dark Web Shops. The maximum value of Hydra transactions is also multiple orders higher than these of Dark Web Shops, reaching 6 million USD. From these values, we can be confident that the structure and customers of the two markets, namely, the Dark Web Shops and the Dark Web Marketplaces, are quite different. The overall incoming revenue for Hydra during 2021 is around 485 million USD, much higher than our lower-bound revenue estimate of Dark Web Shops of 113 million USD. However the reported Hydra revenue is probably partial,

Entity Label	USD Received	#Transactions
binance.com	15,694,336	29,724
huobi.com	7,324,820	13,548
coincheck.com	966,937	2,664
bitzlato.com	349,483	554

Table 4.5: Outgoing Transactions and USD Value to WalletExplorer Entities (Dark Web Stores).

4

as this is the part of the revenue affected by the take-down. We also notice that there is a substantial imbalance between the incoming and outgoing transaction revenue, most likely due to commissions and other complex transactions that occur in large Dark Web Marketplaces.

In Figure 4.3, we plot the revenue of Hydra per month. Although there is no clear trend, the revenue of Hydra has been increasing over time. This was not the case with the monthly revenue evolution for the Dark Web Shops, see Figures 4.3 and 4.4. The Bitcoin-USD rate seems to have some influence on Hydra's revenue, but there is not always a strong correlation between revenue and the Bitcoin-USD rate. Recall that the crawler did not scrape Hydra as it was protected by CAPTCHA [68]. Thus, we can not analyze the revenue per type of abuse.

4.7.3. DARK WEB SHOPS VS. HYDRA BITCOIN ADDRESS AND LAUNDERING OVERLAP

We also investigate if there is any overlap between the Bitcoin addresses associated with Dark Web Shops that we identified after cluster and cleansing with these identified with the same technique for Hydra. Our analysis shows that there is no overlap, which is another indication that Dark Web Shops and Marketplaces are parallel underground markets. We acknowledge that for our comparison, we take a very conservative approach.

However, when we turn our attention to laundering by Dark Web Shops and Hydra, we notice that they both utilize exchange points. Previous works also confirm that Dark Web Shops utilize sophisticated techniques to laundry money using exchanges and wallets [27]. In Table 4.5, we present the total revenue and number of transactions for one-hop outgoing transactions (laundering) of Dark Web Shops per exchange point in our study. For the analysis of transactions, we used GraphSense [32]. In Table 4.6, we repeat the same for Hydra. We notice that Dark Web Shops and Marketplaces not only utilize exchanges but also share two common ones, namely Huobi and Bitzlato. The two common exchanges have repeatedly reported that they participate in the laundering of illicit activity [11]. We recognize potentially more transactions with exchanges can be uncovered with commercial tools. Our labels were sourced from open sources with outdated, limited datasets.

4.7.4. DISCUSSION

Our analysis shows the necessity of continuously monitoring payments to Dark Web Shops. Our results indicate that based on such monitoring, potentially at least 113 mil-

Entity Label	USD Received	#Transactions
huobi.com	4,248,876	677
bittrex.com	286,292	22
poloniex.com	84,382	8
btc-e.com	57,222	12
localbitcoins.com	15,690	15
bitzlato.com	5,662	7
cryptonator.com	2,897	3
matbea.com	2,724	1

Table 4.6: Outgoing Transactions and USD Value to WalletExplorer Entities (Hydra).

lion USD worth of illicit activity, primarily in sexual abuse and financial crime, can be tackled, which is a significant fraction of the overall estimated Dark Web market, by some measures, 5% to 10% [12] in 2021. Our analysis also shows that Dark Web Shops utilize cryptocurrency exchanges to launder money. Most likely, more advanced laundering mechanisms not (yet) recognized in open-source address labels, such as Bitcoin tumblers, are employed. Our methodology offers a scalable way for cryptocurrency services to monitor illicit activity and exclude them from their operation. It also provides insights about the evolving Dark Web Shops ecosystem to authorities towards evidence-based policymaking.

Identifying legal entities behind a Bitcoin address makes it possible to attribute transactions to human beings. This is accelerated by address clustering technology as well as existing and forthcoming European Union KYC legislation [49]. Based on co-spending, joint ownership of addresses can be established [32, 40, 37]. If the individual or legal entity behind at least one of the addresses in a cluster is known, the ownership of the whole cluster is known. As exchange platforms are bound to the legislation of their particular jurisdiction, most of them nowadays adhere to KYC legislation. Based on this, they require customers signing up for an account to present proof of identity and, in some cases, even share their home addresses. Through this legislation, law enforcement investigators can now request the personal details of someone behind a deposit or withdrawal from an exchange account.

Our ongoing research shows that while the coverage of public labels attributing Bitcoin addresses to their controlling entity is scarce, some coverage in publicly accessible sources does exist. We confirm that we are able to run a similar analysis for some of the large Dark Web Marketplaces. This capability is important for several reasons. These labels not only reveal the exchange platforms that were potentially involved in leading law enforcement to take down Hydra's infrastructure [63] but also show that it is possible to bootstrap our Dark Web crawler to crawl different parts of the Dark Web.

4.8. CONCLUSION

Difficulties with scraping and indexing onions complicate Tor's analysis of illicit offerings. One way to sidestep such challenges in research efforts is to focus on a single

union representing big clusters of illicit activity, namely Dark Web Marketplaces. Many researchers have focused on such marketplaces in the past. Much still needs to be discovered regarding the expanding ecosystem of Dark Web Shops, i.e., single-vendor shops operated by individuals or small groups. For the analysis of this, the difficulties above need to be tackled.

In this paper, we develop and apply a methodology to collect and analyze the content and involved Bitcoin addresses in Dark Web Shop websites. In the process, we rely on experts to annotate the illicit activity associated with each Dark Web Shop page. Part of our methodology is a detailed data cleansing process to reliably estimate a lower bound of the revenue of Dark Web Shops by analyzing their incoming transactions. Our analysis shows that the Dark Web Shop revenue was at least 113 million USD in 2021. The top illicit category facilitated by Dark Web Shops is sexual abuse (with revenue close to 94 million USD, or 83% of the total revenue) and financial crime (with around 9% of the total revenue). Furthermore, our analysis does not show an overlap between Bitcoin addresses associated with Dark Web Shops and those large ones exposed in the (partial) takedown of one of the largest Dark Web Marketplaces, namely, Hydra. This indicates that Shops and Marketplaces are parallel Dark Web economies. However, when we examine the laundering (outgoing) transactions, our analysis shows that both Dark Web Shops and Marketplaces utilize exchanges, in some cases, the same ones (Huobi, Bitzlatto). The insights, tools, and analysis we develop in our work will seed future work in the area and will help computer scientists, economists, and policymakers alike to understand the evolving Dark Web ecosystem.

4.9. APPENDICES

We include Table 4.7 with raw results for full reference. The table complements Table 4.2, which appears in paper Section 4.6 and is also included here. In the inner segment, this table includes seed and cluster revenues before cleansing. Even though problematic domains such as Bitcoin multiplier scams showing unaffiliated Bitcoin addresses are already filtered out, the reported revenues are obviously still heavily influenced by unclean data. With this, we show the importance of thorough cleansing to arrive at a trustworthy estimation of illicit revenue - which is as a result of the filtering a lower bound.

For full reference, Table 4.8 provides Hydra's full 2016-2022 transaction revenue to Bitcoin addresses shared by OFAC [63], instead of the 2021 in Table 4.4 which appears in Section 4.7 of the paper.

Category	# Tor domains (# pages)	# BTC addresses seed unfiltered / filtered (co-spent unfiltered / filtered)	Unfiltered Dataset (excl. explorers)				Filtered Dataset			
			Transactions incoming 2021 seed (co-spent)	Transactions outgoing 2021 seed (co-spent)	USD Revenue USD received 2021 seed (co-spent)	USD Revenue USD sent 2021 seed (co-spent)	Transactions incoming 2021 seed (co-spent)	Transactions outgoing 2021 seed (co-spent)	USD Revenue USD received 2021 seed (co-spent)	USD Revenue USD sent 2021 seed (co-spent)
Cybercrime	46 (1,287)	236 / 110 (1,186,952 / 132,092)	5,721 (1,956,150)	4,271 (26,631)	\$26,670,066 (\$78,176,474)	\$25,854,339 (\$72,241,561)	577 (6,338)	264 (621)	\$141,329 (\$1,389,204)	\$141,177 (\$1,390,486)
Drugs / Narcotics	17 (392)	104 / 45 (493,877 / 271,362)	4,055 (800,979)	1,232 (81,204)	\$470,937 (\$679,114,794)	\$348,272 (\$678,662,966)	1,161 (4,553)	290 (597)	\$330,983 (\$1,594,520)	\$328,764 (\$1,414,285)
Extremism	12 (7,683)	19 / 17 (5880 / 231)	197 (364,433)	153 (5,348)	\$21,129 (\$2,570,640,029)	\$14,152 (\$2,168,525,109)	150 (4,001)	33 (246)	\$13,053 (\$577,574)	\$8,461 (\$509,745)
Financial Crime	227 (31,785)	3968 / 397 (35,272,512 / 548,051)	115,315 (86,634,788)	3,391 (21,916,201)	\$30,106,149,207 (\$446,239,322,493)	\$29,269,595,384 (\$445,208,550,108)	1,948 (45,768)	3,205 (4,537)	\$231,308 (\$10,164,827)	\$213,702 (\$8,062,361)
Goods and Services	41 (2,107)	112 / 67 (41,632 / 18,645)	3,320 (95,602)	2,547 (26,396)	\$421,932 (\$628,863,605)	\$127,401 (\$520,894,173)	331 (11,172)	231 (1,072)	\$86,766 (\$1,092,019)	\$85,970 (\$1,028,882)
No Abuse	15 (44)	160 / 79 (501 / 418)	74,807 (158,413)	926 (1,917)	\$912,433,683 (\$15,290,447,315)	\$832,527,513 (\$14,164,154,157)	3,303 (152,958)	319 (896)	\$1,795,163 (\$3,845,552)	\$1,791,164 (\$3,845,551)
Sexual Abuse	836 (29,870)	1945 / 1403 (5,532,538 / 1,108,367)	88,635 (96,645,013)	51,382 (14,647,957)	\$3,001,567,051 (\$599,737,527,770)	\$2,907,675,309 (\$598,208,550,108)	6,636 (61,400)	1,563 (5,151)	\$441,363 (\$94,257,825)	\$390,682 (\$94,241,807)
Violent Crime	3 (41)	29 / 4 (1124 / 7)	103 (3,990)	43 (1,687)	\$2,501 (\$9,532)	\$1,109 (\$4,607)	13 (15)	3 (4)	\$1,401 (\$10,601)	\$1,109 (\$4,532)
Total	1,197 (73,289)	6537 / 2122 (42,535,016 / 2,079,173)	292,153 (186,859,368)	63,945 (36,707,341)	\$34,047,745,506 (\$1,065,224,102x10 ¹²)	\$33,036,143,479 (\$1,061,021,963x10 ¹²)	14,119 (286,205)	6008 (12,924)	\$3,041,376 (\$112,822,122)	\$2,961,029 (\$110,497,649)

Table 4.7: Overview of the dataset statistics with filtered and unfiltered address clustering.

	Seed Addresses		Upper Bound (all clusters)		Lower Bound (excl. exchanges)	
	USD Received	USD Sent	USD Received	USD Sent	USD Received	USD Sent
Min	\$29	\$0	1	1	1	1
Max	\$153,547,344	\$154,057,456	\$4,671,650,304	\$4,722,737,152	\$155,586,032	\$155,635,424
Median	\$2,573,356	\$2,495,489	\$3,215,183	\$3,301,529	\$3,197,913	\$2,834,768
Std Dev	\$192,324,151	\$19,381,411	\$484,036,683	\$488,907,653	\$22,859,744	\$22,879,493
Total	\$792,325,710	\$792,563,212	\$7,690,539,907	\$7,752,677,732	\$930,687,683	\$930,844,516

Table 4.8: Hydra revenue based on 117 seed addresses from Office of Foreign Assets Control [63].

REFERENCES

- [1] Internet Engineering Task Force (IETF). *The Transport Layer Security (TLS) Protocol Version 1.3*. <https://datatracker.ietf.org/doc/html/rfc8446>.
- [2] Josh Aas et al. “Let’s Encrypt: An Automated Certificate Authority to Encrypt the Entire Web”. In: *ACM CCS*. 2019.
- [3] Elli Androulaki et al. “Evaluating user privacy in Bitcoin”. In: *Financial Cryptography and Data Security*. 2013.
- [4] B0bbyB0livia. *TradeMed*. <https://github.com/B0bbyB0livia/trademed>. 2018.
- [5] Andres Baravalle, Mauro Sanchez Lopez, and Sin Wee Lee. “Mining the Dark Web: Drugs and Fake IDs”. In: *International Conference on Data Mining Workshops*. 2016.
- [6] Bitnodes. *Reachable Bitcoin Nodes*. <https://bitnodes.io/>. 2022.
- [7] BitRank Verified. *BitRank: Crypto Tracking to meet Cryptocurrency Regulations*. <http://www.bitrankverified.com/>. 2022.
- [8] Alberto Bracci et al. “Dark Web Marketplaces and COVID-19: After the Vaccines”. In: *EPJ data science* 10.1 (2021).
- [9] British Broadcasting Corporation (BBC). *Hydra: How German police dismantled Russian darknet site*. <https://www.bbc.com/news/technology-61002904>. 2022.
- [10] Bundeskriminalamt. *Illegaler Darknet-Marktplatz „Hydra Market“ abgeschaltet*. https://www.bka.de/DE/Presse/Listenseite_Pressemitteilungen/2022/Presse2022/220405_PM_IllegalerDarknetMarktplatz.html.
- [11] Chainalysis. *Crypto Money Laundering: How Criminals Cash Out Billions in Bitcoin and Other Cryptocurrencies*. <https://blog.chainalysis.com/reports/crypto-laundering/>.
- [12] Chainalysis. *The 2021 Crypto Crime Report*. <https://go.chainalysis.com/2021-Crypto-Crime-Report.html>. 2021.
- [13] Chainalysis. *The 2022 Crypto Crime Report*. <https://go.chainalysis.com/2022-Crypto-Crime-Report.html>. 2022.
- [14] Nicolas Christin. “Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace”. In: *The Web Conference (WWW)*. 2013.
- [15] CoinGecko. *The Most Comprehensive Cryptocurrency API*. <https://www.coingecko.com/en/api>.
- [16] Bogdan Covrig et al. “Upside Down: Exploring the Ecosystem of Dark Web Data Markets”. In: *IFIP SEC*. 2022.
- [17] Alejandro Cuevas et al. “Measurement by Proxy: On the Accuracy of Online Marketplace Measurements”. In: *USENIX Security*. 2022.
- [18] Naval Research Lab Washington DC. *Tor: The Second-Generation Onion Router*. <https://apps.dtic.mil/sti/citations/ADA465464>.

- [19] Whitfield Diffie and Martin E. Hellman. “New Directions in Cryptography”. In: *IEEE Transactions on Information Theory* 22.6 (1976), pp. 644–654.
- [20] Eckmar Community. *Eckmar (Eckmar’s Marketplace Script)*. <https://github.com/eckmarcommunity/eckmar>. 2022.
- [21] Abeer ElBahrawy et al. “Collective dynamics of dark web marketplaces”. In: *Nature Scientific reports* 10.1 (2020), pp. 1–8.
- [22] Karim Eldefrawy, Ashish Gehani, and Alexandre Matton. “Longitudinal Analysis of Misuse of Bitcoin”. In: *International Conference on Applied Cryptography and Network Security*. 2019.
- [23] Elliptic. *Preventing Financial Crime in Cryptoassets: Typologies Report 2022*. <https://www.elliptic.co/resources/typologies-report-2022>.
- [24] Elliptic. *US Sanctions Garantex Exchange and Hydra Dark Web Marketplace Following Seizure of Hydra by German Authorities*. <https://www.elliptic.co/blog/5-billion-darknet-market-hydra-seized-by-german-authorities>.
- [25] Europol. *DarkMarket: world’s largest illegal dark web marketplace taken down*. <https://www.europol.europa.eu/media-press/newsroom/news/darkmarket-worlds-largest-illegal-dark-web-marketplace-taken-down>. 2022.
- [26] Guilhem Fabre. *Criminal prosperity: Drug trafficking, money laundering and financial crisis after the Cold War*. Routledge, 2013.
- [27] Financial Action Task Force. *Professional Money Laundering*. <https://www.fatf-gafi.org/media/fatf/documents/Professional-Money-Laundering.pdf>. 2018.
- [28] Steven Goldfeder et al. “When the cookie meets the blockchain: Privacy risks of web payments via cryptocurrencies”. In: *PETS*. 2018.
- [29] Matthias Götze et al. “Measuring Web Cookies in Governmental Websites”. In: *ACM Web Science Conference*. 2022.
- [30] GraphSense. *Miner Tagpack*. <https://github.com/graphsense/graphsense-tagpacks/blob/master/packs/miners.yaml>. 2022.
- [31] Martin Harrigan and Christoph Fretter. “The Unreasonable Effectiveness of Address Clustering”. In: *IEEE Conferences on Ubiquitous Intelligence & Computing, Advanced and Trusted Computing, Scalable Computing and Communications, Cloud and Big Data Computing, Internet of People, and Smart World Congress*. 2016.
- [32] Bernhard Haslhofer et al. *GraphSense: A General-Purpose Cryptoasset Analytics Platform*. <https://arxiv.org/abs/2102.13613>. 2021.
- [33] Naoki Hiramoto and Yoichi Tsuchiya. “Measuring dark web marketplaces via Bitcoin transactions: From birth to independence”. In: *Forensic Science International: Digital Investigation* 35 (2020).
- [34] Danny Yuxing Huang et al. “Tracking ransomware end-to-end”. In: *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE. 2018, pp. 618–631.

- [35] Interpol. *Combating Cyber-enabled Financial Crimes in the era of Virtual Asset and Darknet Service Providers*. https://cflw.com/download/20200701_Assessment_Report_Cyber_Enabled_Financial_Crime.pdf. 2020.
- [36] Interpol. *INTERPOL Darknet and Cryptocurrencies Working Group - Abuse Taxonomy*. <https://interpol-innovation-centre.github.io/DW-VA-Taxonomy/taxonomies/abuses>. 2020.
- [37] Harry Kalodner et al. “BlockSci: Design and applications of a blockchain analysis platform”. In: *USENIX Security Symposium*. 2020.
- [38] Jochem van de Laarschot and Rolf van Wegberg. “Risky Business? Investigating the Security Practices of Vendors on an Online Anonymous Market using Ground-Truth Data”. In: *USENIX Security Symposium*. 2021.
- [39] Seunghyeon Lee et al. “Cybercriminal Minds: An investigative study of cryptocurrency abuses in the Dark Web”. In: *NDSS*. 2019.
- [40] Sarah Meiklejohn et al. “A fistful of Bitcoins: characterizing payments among men with no names”. In: *ACM IMC*. 2013.
- [41] Satoshi Nakamoto. “Bitcoin whitepaper”. In: (2008).
- [42] Jonas David Nick. “Data-Driven De-Anonymization in Bitcoin”. MA thesis. ETH-Zürich, 2015.
- [43] Kris Oosthoek, Jack Cable, and Georgios Smaragdakis. “A Tale of Two Markets: Investigating the Ransomware Payments Economy”. In: *Communications of the ACM*, [to appear], pre-print available at <https://arxiv.org/abs/2205.05028> (2022).
- [44] Kris Oosthoek and Christian Doerr. “Cyber Security Threats to Bitcoin Exchanges: Adversary Exploitation and Laundering Techniques”. In: *IEEE Transactions on Network and Service Management* 18.2 (2021), pp. 1616–1628. DOI: [10.1109/TNSM.2020.3046145](https://doi.org/10.1109/TNSM.2020.3046145).
- [45] Giacomo Persi Paoli et al. *Behind the curtain: The illicit trade of firearms, explosives and ammunition on the dark web*. Rand Corporation, https://www.rand.org/pubs/research_reports/RR2091.html. 2017.
- [46] Masarah Paquet-Clouston, Bernhard Haslhofer, and Benoit Dupont. “Ransomware payments in the Bitcoin ecosystem”. In: *Journal of Cybersecurity* 5.1 (2019).
- [47] Fedor Poskriakov, Maria Chiriaeva, and Christophe Cavin. “Cryptocurrency compliance and risks: A European KYC/AML perspective”. In: *Blockchain & Cryptocurrency Regulation 2020* (2020).
- [48] The Tor Project. *Tor Project | Anonymity Online*. <https://www.torproject.org/>.
- [49] European Parliament (press releases). *Crypto assets: deal on new rules to stop illicit flows in the EU*. <https://www.europarl.europa.eu/news/en/press-room/20220627IPR33919/crypto-assets-deal-on-new-rules-to-stop-illicit-flows-in-the-eu>. 2022.

- [50] Ronald L Rivest, Adi Shamir, and Leonard Adleman. "A method for obtaining digital signatures and public-key cryptosystems". In: *Communications of the ACM* 21.2 (1978), pp. 120–126.
- [51] Matteo Romiti et al. "A Deep Dive into Bitcoin Mining Pools: An Empirical Analysis of Mining Shares". In: *WEIS*. 2019.
- [52] Amirali Sanatinia et al. "A Privacy-Preserving Longevity Study of Tor's Hidden Services". In: *arXiv preprint arXiv:1909.03576* (2019).
- [53] Marc Shapiro. "Structure and encapsulation in distributed systems: The proxy principle". In: *IEEE ICDCS*. 1986, pp. 198–204.
- [54] Martijn Spitters, Stefan Verbruggen, and Mark Van Staalduinen. "Towards a comprehensive insight into the thematic organization of the tor hidden services". In: *2014 IEEE Joint Intelligence and Security Informatics Conference*. IEEE. 2014, pp. 220–223.
- [55] Chad M. S. Steel. "Stolen Identity Valuation and Market Evolution on the Dark Web". In: *International Journal of Cyber Criminology* 13.1 (2019), pp. 70–83.
- [56] Steven Englehardt and Arvind Narayanan. "Online Tracking: A 1-million-site Measurement and Analysis". In: *ACM CCS*. 2016.
- [57] Tor Project. *Tor Metrics - Onion Services*. <https://metrics.torproject.org/hidserv-dir-v3-onions-seen.html>. 2022.
- [58] Tor Project. *Tor Metrics - Servers*. <https://metrics.torproject.org/networksize.html>. 2022.
- [59] United Nations Office on Drugs and Crime (UNODC). *2020 World Drug Report - In Focus: Trafficking Over The Darknet*. https://www.unodc.org/documents/Focus/WDR20_Booklet_4_Darknet_web.pdf. 2020.
- [60] United Nations Office on Drugs and Crime (UNODC). *2022 World Drug Report - Global Overview Drug Demand Drug Supply*. <https://www.unodc.org/unodc/en/data-and-analysis/world-drug-report-2022.html>. 2022.
- [61] Office of Foreign Assets Control US Department of the Treasury. *FAQ - 1021. Do the prohibitions of Executive Order (E.O.) 14024 and other Russia-related sanctions extend to virtual currency?* <https://home.treasury.gov/policy-issues/financial-sanctions/faqs/1021?s=09>. 2022.
- [62] US Department of the Treasury, Office of Foreign Assets Control. *Russia-related Designation; Cyber-related Designation*. <https://home.treasury.gov/policy-issues/financial-sanctions/recent-actions/20220405>.
- [63] US Department of the Treasury, Office of Foreign Assets Control. *Treasury Sanctions Russia-Based Hydra, World's Largest Darknet Market, and Ransomware-Enabling Virtual Currency Exchange Garantex*. <https://home.treasury.gov/news/press-releases/jy0701>.
- [64] WalletExplorer.com. *WalletExplorer.com: Smart Bitcoin Block Explorer*. <http://www.walletexplorer.com/>. 2022.

- [65] Rolf van Wegberg et al. “Plug and Prey? Measuring the Commoditization of Cyber-crime via Online Anonymous Markets”. In: *USENIX Security Symposium*. 2018.
- [66] WikiLeaks. *WikiLeaks:Tor*. <https://www.wikileaks.org/wiki/WikiLeaks:Tor?>.
- [67] YCharts. *Bitcoin Market Cap*. https://ycharts.com/indicators/bitcoin_market_cap.
- [68] Guixin Ye et al. “Yet Another Text Captcha Solver: A Generative Adversarial Network Based Approach”. In: *ACM CCS*. 2018.

5

DISCUSSION

The contributions in this thesis improve the current state-of-the-art in research into cybercriminal abuse of cryptocurrencies, specifically Bitcoin. We discuss revenue and laundering methodologies employed by several, disparate actor groups within the cybercriminal ecosystem. We have proposed methodologies to analyze crowd-sourced ransomware payment addresses and to cleanse raw data from the Dark Web, a prolific but noisy cybercriminal ecosystem. This way, we add to the understanding of cybercriminal actors and their practices, contributing to future analysis of cybercriminal activity.

This chapter discusses the findings and limitations of the research considered in this thesis. The challenges introduced in Section 1.2 are addressed based on the research questions introduced in Section 1.4.

5.1. EXCHANGE HACKS

This section explains our contributions to the first sub-question:

RQ: *How do cybercriminal actors abuse security vulnerabilities in Bitcoin exchange platforms and what is the financial impact?*

To answer this question, and consequently better understand the security threats to centralized platforms for Bitcoin exchange, this thesis provides an overview of the TTPs employed in cyber security breaches of Bitcoin exchanges, together with the Bitcoin and USD value of stolen funds. These breaches are compared with cyber security incidents at traditional, established financial service providers to understand the challenges in context.

5.1.1. TTPs AND FINANCIAL IMPACT

Bitcoin exchanges are usually targeted by attackers as they keep Bitcoin owned by their users in custody. A successful attack thus consists of several steps following initial exploitation of a particular vulnerability in an Internet-connected system (*attack vector*)

and, after one or more steps of lateral movement, the eventual exfiltration of the funds (*impact*). Our analysis addresses both the attack vector and its impact, with Figure 2.1 serving as a graphical representation of the attack vector in proportion to the financial impact for each of the incidents in our dataset.

ATTACK VECTORS

The use of stolen credentials to elevate access privileges is the most prolific attack vector observed in our dataset, behind attacks with an unknown vector. For the more recent attacks, we observed an increase of the *Unknown* vector, effectively indicating that exchange operators are more reluctant to publicly share technical details on the techniques employed by the attacker. A relatively limited number of observations of advanced attacked vectors however implies that the technical security level of the breached platforms in our dataset was low. Just one instance of a supply-chain attack was observed [2].

IMPACT

While the attack vectors used in the breaches analyzed are similar to those employed against traditional institutions, their financial impact is strikingly different. According to our analysis of breaches of traditional financial institutions, based on a set of validated and confirmed incident reports from multiple industries [5], such organizations rarely lose funds in cyber attacks - with the exception of a few physical attacks against ATMs.

Furthermore, except for 2 breaches, all attacks against Bitcoin exchanges specifically targeted hot wallets, with funds ready available for exchange. In the case of exchanges this is a business necessity, as users want to obtain purchased funds directly. But when the cybersecurity of a platform is breached, this means these funds are also ready available to the attacker.

PLATFORM SECURITY

For 13 out of 36 previously breached exchange platforms which were still operational at the time of our analysis, we have analyzed passive port scans and historical passive DNS to discover available protocols between 2016 and 2019. We found that 4 exchange platforms had multiple services on a single IP address, 3 were running vulnerable software versions with a high risk of exploitation and several platforms were exposing FTP, email, vulnerable SSH and database management interfaces. On the other hand, all platforms except two were migrated behind Cloudflare, serving as a reverse proxy to mitigate common web attacks.

TRACING STOLEN FUNDS

We tracked funds stolen in the breaches of Mt. Gox and Bitfinex, and considered how these attackers are splitting funds to complicate tracking. The laundering TTPs employed by the attackers were low-level; predominantly manual splitting, few mixers were observed. We also considered the breach of Yapizon, for which two persons, allegedly with ties to the North Korean government, were convicted for money laundering [6]. The laundering of the funds stolen from Yapizon was more complicated yet largely manual, involving money mules and back-and-forth conversion from Bitcoin to gift cards to (laundered) BTC.

5.1.2. LIMITATIONS

We gathered data from open sources, specifically primary-source (official announcements) and secondary-source (media) reports of security breaches of Bitcoin exchanges. Inherent to any analysis based on open source data, it is thus biased towards publicly disclosed breaches. Companies are usually hesitant to acknowledge their security has been breached, if not necessitated by reporting obligations or a negative impact on customer experience. Our analysis focused on Bitcoin exclusively. Breaches were cryptocurrencies other than BTC were stolen, are not included. If BTC was stolen among other currencies, we only included the value of the bitcoin stolen. Hence, the absolute amount of crypto-assets stolen in exchange breaches exceeds the total in Bitcoin as reported by us.

Furthermore, as this analysis was performed almost three years ago and many VASPs have increased their cyber security maturity, our analysis shows that the security of exchange platforms has increased since then. According to an aggregator website that was put online recently [4], after the conference paper of this research was released, attackers are moving to other cryptocurrency-related targets such as non-fungible tokens, cross-chain bridges and flash loans, the latter also recognized in an analysis not in scope of this thesis [3].

5.1.3. REFLECTION

This analysis has considered 36 instances of cyber security breaches of Bitcoin exchange platforms. Together these breaches account for at least 1,156,399 BTC stolen from their legitimate owners, usually platform users who stored their Bitcoin in custody.

The analysis of Bitcoin exchange security was uncharted research territory at the time the conference paper was released. Based on an invitation by the conference chair, it was extended for journal publication with an analysis of existing technical vulnerabilities in exchange platforms and a few case studies on money laundering by high-profile actors. Our analysis aids in understanding cyber threats against centralized exchanges. The comparison with attack vectors employed against traditional financial institutions, in which usually personal information is obtained but almost never funds, shows how Bitcoin exchange security is not yet on par with that of traditional financial custodians.

5.1.4. FUTURE WORK

The current unregulated state of the crypto-asset space also means VASPs are currently not subject to information security laws and standards such as the Payment Card Industry Data Security Standard (PCI-DSS) and the Payment Service Directive 2 (PSD-2). Therefore, IT risks might be recognized, but controls are not required. As long as this is the case, analyses like ours, but focusing on other soft spots in the ecosystems will remain worthwhile.

For example in recent years attackers have prolifically exploited flash loans [3], compared with Bitcoin exchanges a relatively novel feature. The heists are not over, the attackers just move along with innovation. This, together with the observation of relative simple attack vectors in our analysis, might indicate that the actors going after these exchanges form a separate category specific and limited skillset, targeting the weakest points in the ecosystem.

5.2. RANSOMWARE PAYMENTS

This section explains our contributions to the second sub-question:

***RQ:** What is the revenue of ransomware actors utilizing Bitcoin as a means of payment and how has this evolved over time?*

To provide a solid answer to this question, a dataset representative of payments to prominent ransomware actors during a specified period is required. Due to various reasons addressed below, the availability of this data is limited. A dataset of 13.5k payments to 87 unique ransomware criminals, gathered through a combination of collection from the public, as well as from existing data sources, was available to us. Based on this, we performed an analysis of ransom payments and their evolution, as well as the laundering of these funds.

5.2.1. PAYMENT AND LAUNDERING ANALYSIS

We have analyzed 13,497 payments to Bitcoin addresses controlled by 87 ransomware actors. Of these actors, 71 are commodity ransomware actors, 16 are RaaS collectives. According to our analysis, RaaS collectives are the highest grossing ransomware actors. While the total number of RaaS actors in our dataset is lower than commodity actors, the revenue of all RaaS collectives combined is 95.7 million USD, where commodity actors account for 5.5 million USD. We attribute this difference to efficient, more specific targeting of victims by RaaS collectives, as well as persistence and manual attacker intervention to maximize impact. We also observed that RaaS actors are more efficient in their usage of payment addresses, using an address usually for just one payment. Commodity actors on the other hand, tend to re-use addresses, through which they essentially leak information. The larger the amount of transactions, the higher the chance of having one transaction in which the actors slips up, identifying himself.

RaaS actors also tend to be more efficient, or sophisticated in selecting methodologies to launder funds. Commodity actors tend to use a variety of services such as exchange platforms with strict KYC, dark web markets, wallet services and fraudulent exchanges. As these are not necessarily privacy-preserving services, the change of identifying the actor is higher. RaaS collectives tend to be aware of this, primarily using fraudulent exchanges and mixers to cover their tracks - implying finer operational security.

5.2.2. LIMITATIONS

While we have collected the largest set of Bitcoin addresses used in ransomware payments to date, it still has an availability bias. The actual decision of a victim to proceed to pay is often surrounded with reputation concerns and shrouded if possible, or disclosed to a limited set of stakeholders. As certain addresses were collected from third-party research into specific groups, some actors might be better represented than others.

While we are aware that our dataset and thus our analysis outcomes are not complete, we strived to make them as representative as possible by combining different sourcing methods. Sourcing from all existing sources we are aware of, and also including novel submissions by the public, we believe our analysis is timely and important. It adds

a new perspective to existing industry and law enforcement reporting, as the academic output on ransomware and especially its evolution is still limited.

5.2.3. FUTURE WORK

We were able to attribute the Bitcoin addresses in our dataset to ransomware actors, using labels supplied by victims and information security researchers. Based on this and some additional verification steps, the addresses in our dataset can confidently be regarded as actual *ransomware addresses*. The more ransomware addresses are labeled, the better they can be flagged by VASPs to avoid future losses. Hence, having open source address labels helps performing analyses like ours. This also works the other way; we obtained some data from an industry source to identify labels for addresses used in laundering because the open-source availability of address labels for exchanges and other VASPs is low. The only substantial source [7] is outdated. Much of this is contained within proprietary offerings, but open source availability of these labels would increase the security of the ecosystem. As many addresses can be scraped from sources like social media, message forums and Tor, future research in the area would contribute to future analysis based on open sources. Also here, effective regulation and subsequent regulatory compliance is critical, as it expedites the identification of real-world actors, usually one of the last miles in law enforcement investigations. It is imperative that analysis capabilities are matched with adequate regulation to empower law enforcement.

5.3. DARK WEB SHOPS

This section explains our contributions to the third sub-question:

RQ: *How can we confidently estimate revenue the of Dark Web cyber-criminal actors, and based on that what revenues do we see?*

The analysis for to answer this sub-question concentrates on dark web shops, being commercial outlets operated by individual vendors. This is the sole proprietorship alternative to the infamous dark web marketplaces which gathered much research attention in recent years. Where marketplaces serve as centralized platforms connecting vendors and buyers of many categories of illicit offerings, single-vendor shops are usually highly specialized storefronts. Using a scraper, we have collected page contents, being the document body, of most dark web shops and extracted Bitcoin addresses from these. Based on an extensive cleansing methodology developed for this analysis, we were able to quantify a lower-bound of revenue made in several categories of illicit services.

5.3.1. PAYMENT ANALYSIS

Based on our analysis we show that at least 113 million USD revenue was generated by dark web shops just in 2021. Sexual abuse and financial crime are the top revenue categories, making up a significant fraction of the overall estimated Dark Web market, by some measures 5% to 10% of the total revenue in 2021 [1].

The illicit revenue generated by dark web shops in 2021 is a fraction of what is generated at the bigger dark web markets, but it needs to be taken into account that the

figures used for comparison are based on different measurements. A dataset of Bitcoin addresses from a marketplace seized by law enforcement will generally provide a more accurate representation of the entity's revenue than scraped data.

5.3.2. LIMITATIONS

We are aware of the limitations of our analysis, many of which are inherent to analysis of cybercriminal artifacts. As a result of thorough cleansing, we are very certain that the remaining payments in our dataset are indeed illicit. However as a result of the strict cleansing, reducing the set of Bitcoin addresses with a factor 1000, many other illicit payments were probably filtered out. This is deliberate; we rather provide a lower bound, but accurate figure than a higher figure which we cannot fully justify. The use of open source address labels leads to limited identification of exchange platforms, as many open source label sets are not updated anymore. We were however able to identify that exchanges are used in the laundering of proceedings from dark web shops, which is an important observation in the combatting of this criminal activity.

5

5.3.3. FUTURE WORK

We have developed a methodology to filter Bitcoin addresses associated with illicit activity and, based on this, calculate illicit revenue in several abuse types. An interesting opportunity for future work would be to automate this pipeline to keep track of the development of illicit activity in Tor and report periodically. Furthermore, machine learning could be applied to keep track of transactions to and from addresses labeled illicit and automatically report transactions to known VASPs, to speed up law enforcement investigations.

5.4. CONCLUSION

The main research question of this thesis, as introduced in Section 1.4 is:

***RQ:** How can we leverage open data to increase our understanding of cybercriminal usage of Bitcoin and how can we quantify this?*

To address this question, idiosyncratic cybercriminal ecosystems in which Bitcoin is abused have to be examined. It also needs to be considered how Bitcoin addresses and transactions can be used to calculate revenue. This thesis has considered how cybercriminal actors use Bitcoin as a means to transact financially with victims, as well as how it is a target in the hacking of exchange platforms. We have provided insight into cybercriminal activities such as security breaches, ransomware campaigns and dark web shops. The payment handling and other financial transactions related to these criminal activities are often black or grey boxes, characterized by limited academic coverage.

For Bitcoin exchange security breaches, we have found the exploit vectors to be relatively low-level but financially lucrative. Compared to traditional financial institutions, the financial gain of attackers is unique to exchange platforms. We furthermore investigated shortcomings in the technical security of these platforms. Mitigation of disclosed vulnerabilities appears to be slow, with vulnerable services remaining exposed

several months after announcement. With this, we raised awareness on lacking security of middlemen in the Bitcoin ecosystem, which received limited prior academic attention. Our analysis of ransomware payments has shown how this cybercriminal ecosystem has evolved from independent malware authors with marginal revenues to a professional criminal ecosystem with persistent syndicates of specialized actors with corresponding fees. We have also regarded Bitcoin-enabled purchases in the dark web. Instead of focusing on marketplaces, which have received significant prior research attention, we focused on single-vendor shops. While their total revenue appears to be much smaller than the usually drug-focused marketplaces, they tend to differentiate into specific niches, with sexual abuse generating the largest portion of the total revenue.

The most important limitation of the analyses discussed in this thesis, is that it is based on the data available to us. In other words, we might only cover a subset of the total revenue made by illicit actors as both attackers and victims might have an interest to keep events undisclosed. We have strived to account for this by collecting data first-hand and aiming to be complete as possible. Furthermore, by considering different cybercriminal sub-ecosystems that are relevant at the moment, we have aimed to obtain an overview of the cybercriminal ecosystem as a whole.

Cybercrime is here to stay and very likely cryptocurrencies as well. This is an opportunity rather than a problem, as blockchains offer unique opportunities to not only *follow the money*, but also to develop and test analytic capabilities to detect novel cybercriminal techniques. Consequently, the above-mentioned limitations of analysis based on public artifacts is also an advantage, as it offers great potential for analysis to both researchers, law enforcement professionals and anyone with aspirations in these areas.

REFERENCES

- [1] Chainalysis. *The 2021 Crypto Crime Report*. <https://go.chainalysis.com/2021-Crypto-Crime-Report.html>. 2021.
- [2] Cryptsy. *Cryptsy Blog — Announcement*. URL: <http://archive.is/FfECg> (visited on 11/20/2019).
- [3] Kris Oosthoek. “Flash crash for cash: Cyber threats in decentralized finance”. In: *arXiv preprint arXiv:2106.10740* (2021).
- [4] SlowMist. *SlowMist Hacked*. <https://hacked.slowmist.io/>. 2022.
- [5] *The VERIS Community Database*. 2019. URL: <https://github.com/vz-risk/VCDB>.
- [6] U.S. Department Of The Treasury. *Treasury Sanctions Individuals Laundering Cryptocurrency for Lazarus Group*.
- [7] WalletExplorer.com. *WalletExplorer.com: Smart Bitcoin Block Explorer*. <http://www.walletexplorer.com/>. 2022.

ACKNOWLEDGEMENTS

Getting towards a PhD is like a strange loop. After four years of moving towards the end, what I have learned most is that I know even less than I thought I knew. Great achievements are never made in isolation - a half-truth in a pandemic. Nevertheless, in this section I like to thank the people that made the PhD enjoyable, making the acknowledgements the most rigorously peer-reviewed section in any dissertation.

I would like to thank my promoters, Inald and George. Inald, you have not been only a reliable promotor, but also a great motivator throughout my PhD trajectory. Your way of questioning things had me not only question topics in my research, but often also the added value of many conventions in daily cybersecurity practice. I am inspired by your combination of great subject matter expertise, critical thinking and warm personality. Without your adequate interventions when needed, I could not have finished this PhD.

George, you have been a great supervisor and mentor, how we have worked together was a blaze. Your approach to setting up a research project and getting to results quickly is something I will forever carry with me. You have been a great coach, both imposing and facilitating a strong focus on getting towards the finish line, while never losing sight of work-life balance. I also have enjoyed the diverse array of Greek cuisine during our brainstorming and the many barista coffees from Labs. You have been an inspiring mentor and I am glad that our collaboration will continue after the PhD. I also like to thank Christian Doerr for helping getting me into the PhD and supervising the first part of my trajectory.

During the PhD, I have worked with great people. While most of this was from home, I am glad to have experienced working in a lab two days a week in the first year.

Harm, we have had a lot of good chats about research, cyber security, life and food - in particular grilled meat. Even with your crazy amount of output during the PhD, you were always there for a chat, genuinely interested on a personal level and always in for a laugh - a rare combination in this field. You have a great professionalism, talent and intellect which will lead you to great things professionally and personally. Vincent and Mark, I am happy to have enjoyed your company as office mates before the endless lockdowns separated us. A special thanks to Sandra for throwing spontaneous coffee breaks and being a great help with little big things like planning meetings and reimbursing server bills. I would also like to thank my co-authors Mark van Staalduinen and Jack Cable for their collaboration. Together with the both of you, I wrote some of the best papers of my PhD, which had a great impact in our field. I am glad we already have identified interesting future opportunities for collaboration.

I would also like to thank all my team mates at Rijkswaterstaat for their support during the PhD. Hans, if you did not kidnap me from another department while I was bound to leave the organization, I probably would never have ended up in where I am today. You have been a great and inspiring colleague throughout, I enjoyed our conversations on crypto, Hilbert spaces and the future of cyber security. Willem and Matthijs, thanks

for being a great department head and team lead respectively, which I presume is not easy. And Rob, working on everything CTI-related was a great pleasure. You were there when I once again took a week off to focus on this research and you kept the show running as if it was nothing. I can not thank you enough for that.

Outside of work, I have great family and friends. Rik, thanks for being a good friend. You are the smartest person I know. Your personableness, musical appreciation and elephant memory with many corny jokes make you truly delightful company. Jeroen, thanks for being my friend for 20 years already. We only had two years of school together and we took different paths from there, but we always stayed in touch. I have great respect for how you stay true to your principles and juggle your job and busy family life.

I would also like to thank my parents Piet and Anje. You raised me to be inquisitive and hard-working. Without your encouragement of further education, which was not all that obvious from the onset, I probably would not have been here. Thanks Arian for being my brother, an inspiring business man, huge wisecracker and my polar opposite. I am glad you have Sylvia to keep you in check.

Last but the most, I would like to thank Fereshtah for taking this journey with me. You were there when I started this and I am glad you are still here, together with our little *aapje* whom joined us in the meantime. Thank you for your love, patience, support, especially when times were tough and sleep-deprived. No amount of words can describe my love for you.

This thesis is dedicated to my late grandfather Jan Kreuger. If you would not have been there in the 1990s to introduce me to computers, I probably would never have been where I am now. You learned me system architecture by building i386s and i486s together, made me familiar with the command line and provided me with many games, from Perestroika to Commander Keen. The many times we scoured the *HCC-dagen* for cheap hardware desired for a new project, together with the museum outings we had together with grandmother sparked an eternal intellectual curiosity and appreciation of things outside my bubble that I will carry with me for the rest of my life.

This work was supported in part by the European Research Council (ERC) under Starting Grant ResolutioNet ERC-StG-679158.

LIST OF PUBLICATIONS

8. **Oosthoek, K.**, van Staalduinen, M. and Smaragdakis, G., *Quantifying Dark Web Shops' Illicit Revenue*, accepted for publication in IEEE Access.
7. **Oosthoek, K.**, Cable, J. and Smaragdakis, G. *A Tale of Two Markets: Investigating the Ransomware Payments Economy*, accepted for publication in Communications of the ACM.
6. **Oosthoek, K.** and Doerr, C., *Inside the Matrix: CTI Frameworks as Partial Abstractions of Complex Threats*, IEEE International Conference on Big Data (Big Data 2021).
5. Griffioen, H., **Oosthoek, K.**, van der Knaap, P., & Doerr, C. (2021, November). *Scan, Test, Execute: Adversarial Tactics in Amplification DDoS Attacks*, in the ACM SIGSAC Conference on Computer and Communications Security (CCS 2021).
4. **Oosthoek, K.**, *Flash Crash for Cash: Cyber Threats in Decentralized Finance*, 2021, Preprint
3. **Oosthoek, K.** and Doerr, C., *Cyber Security Threats to Bitcoin Exchanges: Adversary Exploitation and Laundering Techniques*, IEEE Transactions on Network and Service Management, vol. 18, no. 2, pp. 1616-1628, June 2021.
2. **Oosthoek, K.**, and Doerr, C., *Cyber Threat Intelligence: A Product Without a Process?*, International Journal of Intelligence and CounterIntelligence 34.2 (2021): 300-315.
1. **Oosthoek, K.**, and Doerr, C., *SoK: ATT&CK Techniques and Trends in Windows Malware*, International Conference on Security and Privacy in Communication Systems (SecureComm 2019).