Outsourcing Cybercrime

van Wegberg, R.S.

DOI
10.4233/uuid:f02096b5-174c-4888-a0a7-dafd29454450

Publication date
2020

Document Version
Final published version

Citation (APA)
van Wegberg, R. S. (2020). *Outsourcing Cybercrime*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:f02096b5-174c-4888-a0a7-dafd29454450

OUTSOURCING CYBERCRIME

ROLF VAN WEGBERG

OUTSOURCING CYBERCRIME

OUTSOURCING CYBERCRIME

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 dinsdag 27 oktober 2020 om 15:00 uur

door

Rolf VAN WEGBERG

Master of Science in Criminologie, Universiteit Leiden
geboren te Voorburg, Nederland.

Dit proefschrift is goedgekeurd door de promotores

Prof. dr. M.J.G. van Eeten
Prof. dr. ing. A.J. Klievink

Samenstelling promotiecommissie:

Rector Magnificus	voorzitter
Prof. dr. M.J.G. van Eeten	Technische Universiteit Delft, promotor
Prof. dr. ing. A.J. Klievink	Universiteit Leiden, promotor
Onafhankelijke leden:	
Prof. dr. P.H. Hartel	Technische Universiteit Delft
Prof. dr. B. van den Berg	Universiteit Leiden
Prof. dr. M. Levi	Cardiff University
Dr. A. Hutchings	University of Cambridge
Dr. mr. E.H.A. van de Sandt	Dutch National High Tech Crime Unit
Reserveleden:	
Prof. dr. ir. M.F.W.H.A. Jansen	Technische Universiteit Delft

This research has been supported by the MALPAY consortium consisting of the Dutch national police, ING, ABN AMRO, Rabobank, Fox-IT, and TNO.

Keywords: Cybercrime, Online anonymous markets, Outsourcing, Policing

Printed by: Gildeprint

Cover image: Courtesy of Dutch Law Enforcement

ISBN 978-94-6419-036-6

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

Live Slow
Ride Fast

ACKNOWLEDGEMENTS

Everyone who says "It is not about the destination, it is about the journey", never tried finishing a PhD-project. Now that my I reached my destination however, I can look back on the journey. Along the way, I was fortunate to receive the support of and got to know some awesome people. Here, I would like to thank a few of them.

First and foremost, Bram and Michel - my supervisors. Without your belief in abilities and capabilities I did not know to have, I wouldn't have been able to start and successfully finish this journey. Bram, you joined the project a couple months after the start, but got right up to speed. You are a fantastic mentor, and brought that different (scientific) perspective just when I needed it - we are missing you in Delft. Michel, you have given me the opportunity to become the scholar I am now. From the start - when we wrote the proposal for this project - I felt very fortunate that I could work on this topic with you and broaden my horizon. That feeling never faded. You patiently, but repetitively and sometimes annoyingly, pushed me to do new things - especially when I was (a little too) comfortable with where I was at or what I thought I knew. I look forward challenging myself in the near future - knowing that you will be just three doors down.

Next, I would like to thank the MALPAY-consortium - the Dutch National Police, ABN AMRO, ING, Rabobank, Fox-IT, and TNO - for making this PhD-project possible. Especially, I am grateful to Eddy and later Myra - my TNO managers - for supporting my efforts to pursue and get the most out of this project.

In Delft, I found myself in a close-knit group of passionate cyber security researchers. You all have been a great source of inspiration throughout the years. Whether it be through out-of-the-box ideas - half of which I could not comprehend in the beginning - sharing academic life hacks or the lunch and coffee small talks. I feel privileged to be part of such a wonderful team - with Arman, Carlos, Orcun, Qasim and Samaneh who were there when I started and the growing number of fantastic new colleagues joining in recent years. Although I did not yet had the pleasure to work with many of you, the POLG-people (now O&G), also made me feel right at home - oftentimes fellow social scientists who lost their way somewhere. And next to that - people tend to forget - a wonderful secretariat (currently with Joy & Jolanda), who actually run the show.

I am also greatly indebted to my roommates, Joyce and Maria, who are a constant reminder that there is more than ones own research. Maria, you are the only person I know who calls me Rolfie - I hope to fuel your intermittent caffeine addiction for times to come. Joyce, you where there on my first day. I am grateful to have been able to share every step of the way with you being there.

Throughout my PhD-project, I had the privilege of collaborating with a diversity of amazing co-authors. All of you have contributed to the research in this dissertation and made my PhD-project lots of fun. You provided me with ideas, structure, supervision, guidance, but most importantly with your often limited time. Nicolas Christin and Kyle Soska - collaborators from afar - I am thankful to work with you and hope we can continue our greatly valued cooperation in the future. Closer to home, my incredible colleagues Arman, Carlos, Fieke, Samaneh and Ugur in addition to Jan-Jaap, Oskar, and Thijmen.

I felt lucky to have some amazing long-term friends - Marten, Ralf, Sanne & Willem - with whom I could spend evenings (and nights) debating or analyzing political standpoints, watch hilarious commercial TV-shows or random episodes from The West Wing, and spend weekends away. Sometimes all simultaneously.

I am also grateful to my fantastic family I could always fall back on. My brother, whom I visited during weekend trips to Berlin, where we could talk about anything but work and try not to get smoked around German pool tables. And my life coaches - my parents - who have taught me well and always have supported me whatever direction I took.

Last, and most importantly, Nicole. This dissertation should be all about me - so you tell me. Just this once, I do know better.

Rolf van Wegberg
Leiden, August 2020

CONTENTS

1

INTRODUCTION

1.1. BACKGROUND

1.1.1. COMMODITIZATION OF CYBERCRIME

Over the years, different scholars have contributed to the now established thought that cybercrime offenders are not all tech-savvy criminals [35, 98, 107, 160]. The initial mapping of technical capabilities used in cybercrime to an offender's skills, has been contrasted by growing empirical evidence that paints a different picture. Replacing offender skill, would be a vast supply of technical capabilities, that came available in the underground economy [73, 136, 137, 145]. Compromised websites, botnets and bulletproof-hosting are all examples of capabilities supplied through criminal markets in the underground economy that can substitute specific offender skill [110, 123, 152].

Forms of cybercrime motivated by financial gain, make use of a unique configuration of such technical capabilities to be successful. These forms of cybercrime, called profit-driven cybercrime, range from carding to financial malware, and from extortion to cryptojacking [24, 79, 90, 144]. We can expect that, given their reliance on technical capabilities, particularly these forms of cybercrime could benefit from a changing crime paradigm: the commoditization of cybercrime. That is, standardized offerings of technical capabilities supplied through structured markets by specialized vendors that cybercriminals can contract to fulfill tools and techniques used in their business model.

Researchers have observed the increasing commoditization of cybercrime. Here, commoditization is referred to as the transformation of a product into a commodity,

and is regarded as such by consumers or the market [45]. In essence, commoditized products lack typical product differentiation. Ultimately, the only difference between offerings of commoditized products, is the price the vendor sets for the product. As a result, different offerings of the commoditized product become mutually interchangeable - since they are identical. Consequently, this allows for the re-use of the same product over time and lowers the knowledge threshold for acquiring this product on the market [108]. Cybercrime commoditization can be observed in the offering of technical capabilities as commodities by specialized suppliers in the underground economy.

For instance, so-called booters or stressers have transformed an act of illegal Internet behavior, into a commodity [86]. With booters or stressers, we refer to professional providers of Distributed Denial of Service attacks (DDoS) [124]. On a standardized platform, cybercriminals are supplied with the means to perpetrate a DDoS-attack by simply pointing the resources of the supplier to a target - e.g., a website or a server. Prices start at $20 per DDoS-attack and some booters and stressers even provide subscription models with unlimited attacks [1]. The Pay-Per-Install (PPI) market can also be described as commoditized, as one knows what it is you get when contracting a criminal vendor [37, 141]. In this case, a specialized vendor sells the distribution of malicious software you provide, through a network of pre-infected machines - i.e., a botnet. You pay per install. Here, commoditization enables outsourcing of components used in cybercrime - i.e., a botnet or cash-out solution. Thus lowering entry barriers for aspiring criminals, and potentially driving further growth in cybercrime.

The market for commoditized cybercrime components is remarkably similar to markets for legal products. Like legal markets, the procurement of a product or service – i.e., cybercrime components – is to be dealt with in a one-off transaction, where no additional communication between buyer and seller should be required to complete the transaction. The major difference to legal markets, would be that online criminal markets are anonymous in nature. The anonymity of the vendor, buyer and the market, requires a market structure wherein anonymity turns into an asset instead of a risk. Review systems, similar to these on eBay and Amazon, have been implemented to allow vendors to build reputation and simultaneously mitigate the risk of scams [47].

The underground economy plays an ever more important role in acquiring and aligning a configuration of technical capabilities. Thereby, allegedly transforming the necessity of expertise on specific capabilities – e.g., a cash-out solution – of a cybercrime scheme into a 'make-or-buy' decision. Arguably this should allow actors with less expertise to

[1] See https://www.europol.europa.eu/newsroom/news/worlds-biggest-marketplace-selling-internet-paralysing-ddos-attacks-taken-down

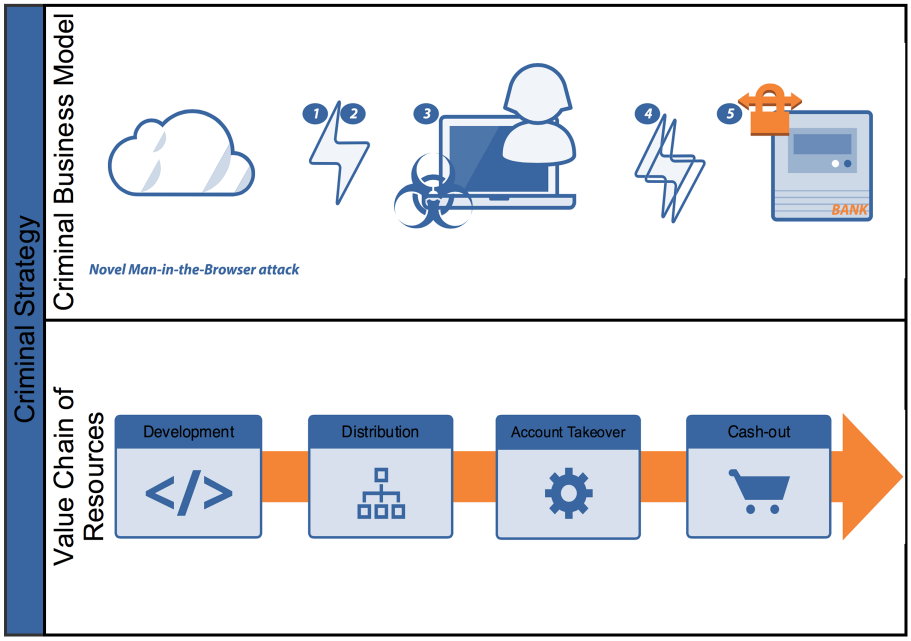

Figure 1.1: Conceptual outline

operate a profitable cybercrime scheme. A cybercriminal business model might rely in part or in full on standardized toolkits and resources available on underground markets, thereby lowering the entry barriers for cybercriminal start-ups. To investigate the impact of commoditization in cybercrime, we first turn to how we can analyze business models in profit-driven cybercrime.

1.1.2. CYBERCRIME VALUE CHAINS

As many cybercriminal entrepreneurs lack the skills to provision certain parts of their business model, this incentivizes them to outsource these parts to specialized criminal vendors. Using underground markets, entrepreneurs have found a new business to business channel to contract vendors and acquire cybercrime components – i.e., technical capabilities – for a range of cybercriminal business models. To investigate the potential for outsourcing in cybercrime schemes, we have to differentiate three levels of analysis. To illustrate this differentiation, we use a financial malware scheme, deploying a Man-in-the-Browser attack (MitB), as an example.

Figure 1.1 gives a visual representation and conceptual outline of three levels of

analysis. First, we identify the criminal business model as an accumulation of all the different criminal activities of a financial malware attack. The specific set of actionable elements makes it economically worthwhile to execute these types of attacks this way. In other words, the line-up of criminal activities in a shape, form, scale and function that potentially generates a profit and outweighs the risks and costs. For instance, operating a financial malware scheme using a MitB-tactic, can line-up the following elements: (1) drive-by-download of (2) financial malware, the clients logs in to online banking environment and the malware activates (3), the account is taken over and funds are sent to money mules (4) and (5) the money mules fulfill the cash-out by withdrawing the funds at an ATM.

Second, the technical capabilities needed to operate such a business model represent the value chain of resources. Herein not the criminal activities themselves, but the components that constitute the technical enablers for all these criminal activities, are described. In this case, we can identify a) development of the malware, b) distribution via compromised websites, c) automated account take-over, and d) cash-out via money mules.

Third, we identify the criminal strategy as a level of analysis. This criminal strategy entails all the strategic and economically motivated decisions within the total criminal scheme. These choices range from the decision to outsource parts of the value chain, scaling up certain activities in the business model to increase profit and to changing the modus operandi – for instance changing the infection vector, adapting target selection or cash-out methods. Unraveling these criminal strategies helps to understand if and how actors chose between setting up the entire scheme themselves or 'outsourcing' parts of the scheme – like leasing a botnet, using malware-as-a-service, pay-per-install or money mule recruitment services.

1.1.3. OUTSOURCING

Meeting places where supply could meet outsourcing demand, are for example carding forums, social media platforms as Facebook and more recently, encrypted messaging services like Telegram. These meeting places serve as a starting point to seek and interact with criminal vendors to do business with. Yet, transactions are dealt with through private channels – e.g., Internet Relay Chat (IRC). That changed with the rise of online anonymous markets: one-stop-shops like Silk Road, where initially and predominantly drugs were traded. In contrast to other online meeting places, all aspects of doing business – from searching to contracting – are handled on and by the market.

Online anonymous markets provide criminal actors with a standardized and anony-

mous marketplace where next to drugs, an increasing supply and demand in cybercrime tools and techniques meet [13, 39, 139]. There, technical skill could indeed be transformed into 'knowing what to buy'. That way, resources needed in a business model are not self-organized, but fulfilled through a transaction on the criminal market.

Acquiring a cybercrime component on an online anonymous market, follows on the decision to outsource parts of a criminal business model. We know however that not all cybercrime components are as easily to acquire in the underground economy as others [20, 64, 71, 136, 145]. As we can derive from previous insights, especially the last but critical step in a successful cybercrime scheme – cashing out – is a tough nut to crack [64, 101]. So, can one really build an entire cybercrime scheme from outsourced commodities, like different scholars hypothesized would be the case [136, 145]. Taking this puzzle as its main focus, *this dissertation investigates the commoditization of cybercrime and its impact on outsourcing technical capabilities through online anonymous markets, enabling profit-driven cybercrime.*

1.2. RESEARCH GAPS

Research on profit-driven cybercrime has focused on unravelling its modus operandi [145]. Much of the work focused on one specific technical capability, i.e., cybercrime component, or one value chain within the cybercrime ecosystem. Grier et al. [71] inspected the (business)model of exploit-as-a-service, where criminals rent out their infrastructures in order to infect systems, e.g. drive-by-downloads, whereas Levchenko et al. [99] have uncovered the spam value chain in an analysis of the full set of resources employed to monetize spam email. Other scholars have analyzed these business models in a similar fashion – e.g., Stone-Gross et al. [142] on fake anti-virus software, Caballero et al. [37] on the pay-per-install market, and Rossow et al. [133] on malware downloaders. This type of work therefore investigates a separate piece of the total cybercrime ecosystem. Hence, we lack the understanding if and how capabilities are unique or common across profit-driven cybercrimes.

To create a comprehensive understanding of how business models in profit-driven cybercrime are impacted by the commoditization of cybercrime, we need to shift focus from a single to successful configurations of components. That is, investigate if and how outsourced components can be used to fulfill which capabilities needed in profit-driven cybercrime. This is where we use an economic lens, and apply the value chain perspective to create an overview of criminal activities, resources and strategies in profit-driven cybercrime.

In turn, this conceptual view helps capture the interactions between outsourcing and

interventions of law enforcement agencies. Knowing how outsourcing fulfils parts of the value chain, can help law enforcement exploit 'chokepoints' – i.e., use the weakest link in the value chain where criminals appear to be vulnerable. These and other policing tactics, create the possibility of studying how different resources involved in these attacks are being combined and interventions in one part of the value chain – e.g., interventions aimed at cash-out strategies – affect the other parts. Understanding these interactions would help creating better evidence-based law enforcement strategies, ideally making certain cybercriminal business models less profitable or even economically unattractive to begin with.

Recent work has partly touched upon how specific outsourcing strategies influence cybercriminal business models – e.g., how criminals organized their so-called cash-out strategy [150] making sure a successful cybercrime scheme ends with a criminal profit. Next, measurement studies on underground markets have shown how specialized vendors enable cybercriminal entrepreneurs in the outsourcing of specific components used in profit-driven cybercrime [75, 82]. These insights already greatly contribute to a better understanding of outsourcing strategies. Still, much of the potential for outsourcing remains unknown. Specifically, we lack the knowledge how outsourcing is facilitated by online anonymous markets and to which extend and how outsourcing can be disrupted.

In sum, we observe three main gaps in current research: a) we lack the understanding if cybercrime components are unique or common across value chains in profit-driven cybercrimes, b) we do not know the extent to which and how cybercrime components in these value chains are amendable to outsourcing and c) we have limited knowledge on how interventions disrupt outsourcing. We will employ five research activities to fill these gaps.

First, we create an overview of value chains in profit-driven cybercrime. Second, we capture how these value chains relate to each other and overlap – i.e., contain different or identical components. Put differently, which cybercrime components are used in only one value chain, and which in more than one. Third, we empirically investigate how the demand for cybercrime components matches the supply. That is, measure and characterize the supply in cybercrime components on online anonymous markets and map this to the demand stemming from profit-driven cybercrime value chains. Fourth, we study which cybercrime components see successful transactions and what the predictors for product performance are. Fifth and last, we take these insights and unravel how we can design interventions to disrupt outsourcing.

To execute these activities, we make use of existing technical insights on the profit-driven cybercrime ecosystem and apply an economic lens. When combining these in-

sights, we can study the overarching value chains in profit-driven cybercrime and investigate how commoditization of cybercrime components leads to the potential to outsource parts of these value chains.

1.3. RESEARCH AIMS & QUESTIONS

This dissertation studies the phenomenon of profit-driven cybercrime through value chains and business models. It investigates how outsourcing is enabled by online anonymous markets and how outsourcing can be disrupted.

Given the above, the following main research question is identified:

How do online anonymous markets facilitate the outsourcing of cybercrime components in profit-driven cybercrime value chains?

The focus of this dissertation is to understand how commoditization of cybercrime on online anonymous markets enables outsourcing in profit-driven cybercrime value chains. We start by discerning business models and value chains, leveraging a mixed-methods approach and a combination of scientific disciplines. The scientific contribution mainly lies in capturing outsourcing strategies in profit-driven cybercrime, specifically by investigating the role of online anonymous markets.

To study the phenomenon of profit-driven cybercrime we require – because of its intrinsic nature – different angles of scientific outlook. As cybercrime itself can be seen as a technical operator, the field of computer science and more specifically the field of information security forms a logical outlook. Looking at the human operator of cybercrime, in this case the criminal or criminal group, a criminological perspective is beneficial in studying the outsourcing strategies in cybercrime [85]. And considering that profit-driven cybercrime truly is a business, looking at the business models involved and the relation to underground markets, we require an economic lens.

Hence, this research is built on an interdisciplinary approach: a combination of computer science tools to perform measurements in the cybercrime ecosystem with criminological and economic theories of information security. Using the approach of economically analyzing cybercrime, supported by large-scale quantitative datasets generated within computer science security research, is a way to add interdisciplinary and empirical insights to this growing field [85]. For this purpose, a combination of both quantitative and qualitative research techniques will be used. As this is a paper-based dissertation, the research methodology will be elaborated upon in each chapter individually.

1

1.4. DISSERTATION OUTLINE

This section outlines the structure of this dissertation. First, this section will present the five studies in this dissertation and which of the formulated research question they try to answer. Second, this section presents an overview of the peer-reviewed papers that are associated with each study and corresponding chapter.

STUDY 1 – VALUE CHAINS IN PROFIT-DRIVEN CYBERCRIME

Fraud with online payment services is an ongoing problem with significant financial-economic and societal impact. One of the main modus operandi uses financial malware, that compromises end-user devices and takes over online banking sessions. Using transaction cost economics, this study illustrates the business model behind financial malware and presents three value chains therein. For this purpose, we use a conceptual synthesis of the state of the art of literature on financial malware, underground markets and cybercrime economics as well as today's banking practice.

The focus of this study is on the following research question:

RQ1: *What are the business models and value chains of profit-driven cybercrime?*

STUDY 2 – COMMODIZATION OF CYBERCRIME COMPONENTS

While there is evidence in the literature of specific examples of cybercrime commoditization, the overall phenomenon is much less understood. Which parts of cybercrime value chains are successfully commoditized, and which are not? What kind of revenue do criminal business-to-business services generate and how fast are they growing? We use longitudinal data from eight online anonymous marketplaces over six years, from the original Silk Road to AlphaBay, and track the evolution of commoditization on these markets. We develop a conceptual model of the value chain components for dominant criminal business models. We then identify the market supply for these components over time.

This study provides an answer to the following research question:

RQ2: *To which extent do commoditized cybercrime components meet the demand for outsourcing on online anonymous markets?*

STUDY 3 – OUTSOURCING THE CASH-OUT OF CYBERCRIME PROCEEDS USING BITCOIN MIXERS

Digital payment methods are increasingly used by criminals to launder money obtained through cybercrime. As many forms of cybercrime are motivated by profit, criminals require a solid cash-out strategy to ensure that crime proceeds are without an incriminating money trail. These cash-out strategies are increasingly facilitated by cryptocurrencies, mainly bitcoin. We examine how cybercrime proceeds can be laundered using bitcoin money laundering services, i.e. mixers, that are offered in the underground economy. Focusing on service-percentages and reputation-mechanisms in these underground services, this study presents the results of a cash-out experiment in which five mixing- and five exchange services are included. We discuss what these findings mean to law enforcement, and how bitcoin laundering can be disrupted.

This study answers the following research question:

RQ3: *How do bitcoin mixers enable the cash-out of cybercrime proceeds?*

STUDY 4 – PREDICTING THE PERFORMANCE OF CYBERCRIME PRODUCTS

Many cybercriminal business models rely on the outsourcing of specific technical capabilities of the underlying value chains of resources. Online anonymous markets, from Silk Road to AlphaBay, have been used to search for these products and contract with their criminal vendors. While one listing of a product generates high sales numbers, another identical listing fails to sell. In this study, we investigate which factors determine the performance of cybercrime products. To answer this question, we analyze scraped data on the business-to-business cybercrime segments of the AlphaBay market (2015-2017). We construct variables to capture price and product differentiators, like refund policies and customer support. We capture the influence of vendor characteristics by identifying five distinct vendor profiles based on latent profile analysis of six properties. We leverage these product and vendor characteristics to empirically predict the number of sales of cybercrime solutions, whilst controlling for the lifespan and the type of solution.

This study provides an answer to the following research question:

RQ4: *What are predictors for successful transactions in cybercrime components?*

1

STUDY 5 – INVESTIGATING THE IMPACT OF ONLINE ANONYMOUS MARKET INTERVENTIONS

Online anonymous markets are facilitators in a wide range of illegal activities. On a structured platform, criminals innovated a growing global market, where both physical goods, predominantly varieties of drugs, and digital goods, like specialized cybercrime toolkits, are traded. Likewise, law enforcement agencies have innovated their operations on and against these markets. In recent years, several police-led interventions have resulted in take-downs and take-overs.

First, we place interventions in a historical perspective, by reconstructing evolving law enforcement intervention strategies. We find that disruption of criminal activities, instead of attribution, has become the focal point in these interventions. Second, we assess the effects of Operation Bayonet, an international policing campaign led by the Federal Bureau of Investigation (FBI) and the Dutch National High Tech Crime Unit (NHTCU) targeting two prominent online anonymous markets. We leverage measurements of the user-base of then market leader, and safe haven: Dream Market. We investigate the effects of the operation on all newly registered vendors on Dream Market during and shortly after Operation Bayonet by mapping their individual and historic characteristics to discern migration patterns and changes in vendor behavior.

This study answers the following research question:

RQ5: *How do interventions aimed at online anonymous markets impact the potential for outsourcing?*

Table 1.1 shows an overview of the different chapters in this dissertation and the peer-reviewed, empirical study that is covered.

Next, we present perspectives for policing. In Chapter 7, we synthesize the state-of-the-art in online anonymous market intervention studies and present best practices to measure the impact of interventions on online anonymous markets. We map these measurements to known aims and tactics of past interventions and present suggestions for novel measurements of future interventions. This way we can design new interventions based on proven historic impact – working towards evidence-based interventions. This dissertation is completed with Chapter 8, which summarizes the main findings, reflects on the results and presents future research directions.

Table 1.1: Dissertation outline

Chapter	Publication(s)
Ch.2	Van Wegberg, R.S., Klievink, B., & Van Eeten, M. (2017). Discerning novel value chains in financial malware. *European Journal on Criminal Policy and Research, 23*(4)
Ch.3	Van Wegberg, R.S., Tajalizadehkhoob, S., Soska, K., Akyazi, U., Ganan, C. H., Klievink, B., Christin, N., & Van Eeten, M. (2018). Plug and Prey? Measuring the commoditization of cybercrime via online anonymous markets. In *Proceedings of the USENIX Security Symposium (USENIX Security 18)*
Ch.4	Van Wegberg, R.S., Oerlemans, J.J., & van Deventer, O. (2018). Bitcoin money laundering: mixed results? An explorative study on money laundering of cybercrime proceeds using bitcoin. *Journal of Financial Crime, 25*(2)
Ch.5	Van Wegberg, R.S., Miedema, F., Akyazi, U., Noroozian, A., Klievink, B., & Van Eeten, M. (2020). Go see a specialist? Predicting cybercrime sales on online anonymous markets from vendor and product characteristics. In *Proceedings of The Web Conference (WWW '20)*
Ch.6	Van Wegberg, R.S., & Verburgh, T. (2018). Lost in the Dream? Measuring the effects of Operation Bayonet on vendors migrating to Dream Market. In *Proceedings of the Evolution of the Darknet Workshop* *Extended with parts of:* Hartel, P., & Van Wegberg, R.S. (2019). Crime and Online Anonymous Markets. In *International and Transnational Crime and Justice*. Natarajan, M. (ed.) Oerlemans, J.J., & Van Wegberg, R.S. (2019). Opsporing en bestrijding van online drugsmarkten. *Strafblad, 17*(5) Verburgh, T., Smits, E., & Van Wegberg, R.S. (2018). Uit de schaduw: Perspectieven voor wetenschappelijk onderzoek naar dark markets. *Justitiële Verkenningen, 44*(5)

2

VALUE CHAINS

Fraud with online payment services is an ongoing problem with significant financial-economic and societal impact. One of the main modus operandi is financial malware that compromises consumer and corporate devices, thereby potentially undermining the security of critical financial systems. Setting up a successful financial malware scheme, requires the aligning of a lot of moving parts. Analysing how cybercrime groups acquire, combine and align these parts into value chains can greatly benefit from existing insights in the economics of online crime. Using transaction cost economics, this chapter illustrates the business model behind financial malware and presents three novel value chains therein. For this purpose, we use a conceptual synthesis of the state of the art of literature on financial malware, underground markets and (cyber)crime economics as well as industry reports.

2.1. INTRODUCTION

Fraud with online payment services has consistently been one of the most damaging forms of cybercrime [67, 144]. The European Central Bank [56] has published fraud statistics for the Single European Payment Area, which puts the total fraud in 2014 at €1,44 billion. Around 66% of the total is "card-not-present" (CNP) fraud, which includes online payments. The overall trend is however undisputed: online payment fraud imposes substantial cost on the economy and has become the dominant form of fraud with payment services [18, 116, 149]. Next to phishing, malicious software - i.e., malware targeting financial service providers worldwide - is an ongoing and continuous threat to

these financial service providers, causing millions in damages in both industrialized and non-industrialized countries [17].

The research covering financial malware has primarily been technical of nature and much of this work focused on only specific parts of the total, overarching malware ecosystem. For example, Grier et al. [71] inspected the business model of exploit-as-a-service [1], where criminals rent out their infrastructures in order to infect systems, e.g. drive-by-downloads [2]. Setting up a successful financial malware scheme however, requires the aligning of a lot of moving parts. Not only having the overview of which parts are needed, but also the expertise to actually set up and operate the total scheme, requires a non-trivial level of skill. Hence that the underground economy, now seen as a sort of criminal Craigslist where these 'parts' – like botnets etc. – are sold or rented out, plays an ever more important role in acquiring and aligning all moving parts. This underground economy transforms the necessity of having expertise on specific parts of a financial malware scheme into 'knowing what to buy', arguably allowing actors with less expertise to operate such a scheme. However, we don't know how these actors chose between setting up the entire scheme themselves or 'outsourcing' parts of the scheme. For instance, leasing a botnet, using crimeware-as-a-service [3], pay-per-install [4] or money mule recruitment services [5]. And if they do outsource, how does this affect both business models – that of the organizer of the total scheme and that of the seller of 'parts'? Which incentives influence this 'outsourcing'?

To address these questions, the existing insights on parts of the financial malware ecosystem will need to be combined with insights from research on the 'economics' of crime. By conceptually synthesizing the literature on financial malware, we will try to shed light on criminal strategies in financial malware schemes. Next to the established economic outlook on crime, transaction cost economics can be of beneficiary value to understand outsourcing incentives within these criminal strategies. Specifically, when looking at economic incentives, the underlying patterns and motivations behind the current modus operandi - i.e., criminal business model - can be unravelled. When combining these insights with the knowledge on the various components of the financial malware ecosystem, the so-called 'value chain' of financial malware can be uncovered.

[1] Exploit-as-a-Service is a service that automates the exploiting of a victim's internet browser [71]

[2] Any download that takes place without the user's authorization or prior knowledge; often initiated already active malicious software [71]

[3] Crimeware-as-a-service (CaaS) is a business model used in the underground market where illegal services are provided to help underground buyers conduct cyber crimes (such as attacks, infections, and money laundering) in an automated manner [137]

[4] Pay-per-install services play a key role in the modern malware marketplace by providing a means for outsourcing the global dissemination of their malware [37]

[5] An example of CaaS, wherein money mules are offered as a commodity [137]

Until now, there is little to no systematic and comparative empirical research that sheds light on the overarching value chains around financial malware. In this paper we therefore aim to unravel value chains in financial malware. These value chains can help extricate the interactions between the strategies of attackers on the one side and the properties and policies of the financial service providers on the other. It would enable us to study how the different resources involved in these attacks are being combined and how interventions in one part of the value chain - e.g., interventions aimed at cash-out strategies - affect the other parts - e.g., targeted payment services. Understanding these interactions would help creating better countermeasures and new security services, ideally making certain fraud models less profitable or even loss making to begin with.

The goal of this paper is to conceptually synthesize the literature on financial malware, underground markets and (cyber)crime economics as well as industry reports, to make a first attempt to discern archetypical and novel value chains in financial malware. In the next part of the paper, sections 2.2 and 2.3 give an overview of the field of economics of crime, respectively on the economics of cybercrime and transaction cost economics. Next, in section 2.4 we give a state-of-art of the literature in financial malware and identify parts of the whole malware ecosystem, which have been studied. In section 2.5 we use both these overviews – (cyber)crime economics and financial malware – to discern three novel value chains in financial malware based on the existing literature, as well as industry reports, followed by our conclusions in section 2.6.

2.2. THEORETICAL BACKGROUND

2.2.1. ECONOMICS & CRIME ANALYSIS

Studying crime in an economic fashion is not new. Famous is the work of Becker [22], wherein he lays the foundation of the economics of crime and punishment. Using a rational choice perspective, he presented the idea that crime and punishment can - to a certain extent - be analyzed on the basis of individual costs and benefits. Knowing these costs and benefits, allows for criminal justice policies to become increasingly effective by raising cost – such as the penal risk – or lowering the benefits - think of bank vaults with time locks in order to lower the immediate reward of robbing a bank.

The work of Becker inspired others to look for an economic approach to study organized crime [53, 69, 70, 93, 104]. Thereof, the work of Levitt became widely popular when he combined earlier work in the best selling book Freakonomics and its successor Superfreakonomics [102, 103]. Literature on the economics of organized crime, let alone financial cybercrime, is quite scarce when comparing this to the growing amount of

economic studies on individual crime and criminal law. Nonetheless, with the attention shifting towards cybercrime more and more, the field of economics of cybercrime – as introduced above – has seen growing amounts of studies with an economic approach to cybercrime from 2006 onwards [8, 17–19, 21, 88, 91, 97, 105, 114, 116, 128, 137, 159].

More in particular both Moore et al. [116] and Thomas et al. [145] made critical, breakthrough attempts to grasp the market structure of online crime - i.e. underground markets. Next, Afroz et al. [8] comparatively studied these underground markets, five to be precise, more in-depth for one of the first times. Furthermore, Kraemer-Mbula et al. [91] have shown the ongoing globalization based on a growing digital ecosystem, in cybercrime and underground markets using credit card fraud and identity theft as exemplary cases. Moreover, Sood and Enbody [137] introduced the model of crimeware-as-a-service, describing and analyzing multiple forms of criminal services purchasable on underground markets. These underground markets thus have a vast supply of specific parts of the malware ecosystem [138]. Matched with a continuous demand for these parts to set up a financial malware scheme, this creates an extraordinary criminal market structure. But how does a criminal actors choose between buying all the parts, buying some parts or even no parts of their financial malware scheme? And why does a criminal actor choose not to buy, but to actually sell parts of a financial malware scheme to others, perhaps even potential competitors?

Just like a regular business, the criminal business that aims for the most profit is one that strives towards the most effective business model, with low operational costs and an optimized net gain. In such an effective model, decisions have to be made on whether to organize specific tasks within the criminal organization itself, or to 'outsource' these to others. The choice of outsourcing can be seen as an economic motivated deliberation on for instance the frequency of this outsourced task and the specificity of this task [53]. In other words, how frequent are the outsourced tasks needed, how specific can the task be described and is this sufficient information to deliver this task as a service to the client in question? For example, a botnet needed to spread malware can be argued to be both specific and frequently used, whereas spear phishing a bank employee to infect computers with Remote Access Tooling (RAT) [6] to hack into – until then – unknown internal bank systems, is lacking both this frequency and sufficient specificity. In consequence the latter is less likely to be outsourced, as the costs do not outweigh the potential benefits. These decisions based on the intrinsic transaction costs, form the basis of the consonant field of economics [163, 164]. Such perspective is essential as the total malware ecosystem in

[6] Remote Access Tooling is software that allows a remote "operator" to control a system, e.g. a computer, as if they have physical access to that system. In that way the operator can have unlimited access to the computer without being in physical contact with that system.

terms of value chains consists of numerous (outsourced) parts, where incentivized decisions form an important part of this generic build-up of parts in an individual financial malware scheme.

2.2.2. TRANSACTION COST ECONOMICS IN OFFLINE CRIME

Originally aimed at contract law, so called transaction cost economics sets out economic principles on and identifies incentives for companies (sub)contracting each other for goods and services [163–165]. This in contrast to keeping all activities in-house, so called vertical integration. The term 'vertical integration' refers to a company which mainly relies on its internal workforce, in contrast to the company who mainly relies on contracted third parties for goods and services needed in the business [162]. In his work on transaction cost economics Williamson [163, 164] describes these different organizational structures on the basis of transaction costs that accompany this differentiation in structures, resulting in a series of institutional implications, such as:

"As uncertainty increases (...) transactions will either be standardized, and shifted to the market, or organized internally." [164, p. 259]

"As generic demand grows and the number of supply sources increases (...) vertical integration may give way to obligational market contracting, which in turn may give way to markets." [164, p. 260]

These propositions imply that when goods or services involved in a transaction can be described as frequent, standardized and do not require highly specialized know-how or skill, these transactions will take place in the market and will not be vertically integrated.

As described above, most of the literature on the economics of organized crime has been focusing on its market structure. In older, but still relevant work by Abadinksy [6] and Reuter [129], next to more recent work of Garoupa [69, 70] and Turvani [146], the importance of transaction costs with regard to the illegal activities of a criminal organization have become mainstream in the economics of organized crime. More specifically, Turvani [146] points out that as most of the activities of a criminal organization are generally illegal, the regular structure of a market economy cannot see to a trustworthy system of transaction monitoring. On underground markets, reviews – like trust in other shadow economies - are therefore a direct substitute for the absence of a transaction monitoring system [75]. However, a viable business relationship is still hard to establish, factoring in absence of such a solid transaction monitoring system. This is for example the reason

why large drug deals often result in rip-offs, because both the drugs and the payment have to be at the same time and place to allow for an immediate exchange of goods.

In a more prominent paper, Dick [53] developed a comprehensive analytical framework in which he shows that transaction costs and not a form of monopoly power, as argued before, primarily determine the (illegal) activities of crime in an organized structure. The paper predicts that when there is a production cost advantage in a specific illegal activity, organized crime regarding that illegal activity will be more successful [53]. When looking at the question Dick asks himself – when does organized crime pay? – he starts with the perspective Williamson laid down. He formulates the hypothesis based on the perspective "that organized crime's activities will be guided primarily by the relative costs of completing illegal transactions within the market versus a downstream firm" [53, p. 28]. With Williamson as a starting point he focuses on a) is the activity suitable for 'large scale production'? b) how specific can the accompanied transaction be described? and c) what is the frequency wherein this transaction would take place. Next, he adds a crucial fourth factor: uncertainty. Compared to legal markets, their illegal counterparts do not have a reliable system of enforcement of transactions and lack the accurate estimation of reputation on such a market [53]. In turn this creates an incentive to not only assume the production cost advantage of outsourcing let's say money-mule recruitment, but also incorporating the risk of uncertainty inherent to the specific transaction. In the case of the money-mule recruitment, this would be the more general notion of the transaction itself – do I get scammed? – and the more specific notion of the risk having undercover police informants pose as mules or the scenario wherein the mules have already flagged bank-accounts and are therefore not useful.

2.2.3. ECONOMICS OF FINANCIAL MALWARE

To help discern value chains in financial malware, the transaction cost economic approach is undeniably very useful. We have briefly touched upon how financial malware schemes exist of different elements, and that many of these parts are purchasable on underground markets [137]. Using the transaction cost economic perspective we illustrated how different incentives have an influence on the choice between 'doing-it-yourself' or 'outsourcing', not only in legitimate but also in the illegitimate business. Whereas organized crime has been the main subject of these illustrations, cybercrime - e.g., financial malware - arguably lends itself even more for this perspective. The underground market is blooming, easily accessible, but above all, nearly anonymous. Which poses the obvious risks of scams, but also allows for a relatively low-risk entry to the market. And with the addition of reliable reputation mechanisms, making headway for traditional criminal rep-

utation behavior. Even potentially diminishing the available options of disrupting such a 'dark network'. Before we can, however, look at financial malware from a transaction cost economic perspective, we have to look in some more detail to our approach of using the state-of-art of existing research on parts of the total malware ecosystem to discern novel value chains in financial malware.

2.3. APPROACH

The following sections of this paper represent the necessary steps towards the actual discernment of the novel value chains in financial malware we present in section 2.5. To provide insight in the used methodology, we describe our approach in the remaining part of this section. First, we clustered and conceptually synthesized literature on financial malware in specific parts of the total financial malware ecosystem. Herein we followed the clustering by Sood et al. [137]. The literature we included in this clustering is published between 2000-2015, is available on Web of Science and has financial or banking malware as keywords. Next, we included literature with keywords related to the concepts per clusters, such as 'infections' or 'botnet' albeit related to the general keyword of financial/banking malware. Thereafter we analyzed the overview of literature, identifying gaps and the extent to which a total view of financial malware ecosystem based on the existing literature can be given. This literature overview thereby served the research goals of discerning the value chains in current-day financial malware schemes. Next, we used the research into financial malware in relation to underground markets to investigate the different underground market alternatives - i.e., outsourcing supply - per cluster. This way, we shed light on the contrast between self-organizing - i.e., vertically integrating - and using underground commodities - i.e., outsourcing. To look at the different current-day practices in financial malware schemes, we used prominent security blogs and reports by security firms. A differentiation in financial malware schemes can be constructed based on the distinguished current-day practices. This differentiation then formed the basis of extricating the novel value chains of these financial malware schemes, wherein we described the specific parts that make up every value chain. Hereafter, we apply the framework proposed by Dick [53] to analyze the different elements of every value chain from a transaction cost economic perspective.

Finally, we therewith can identify both the incentives for vertically integrating and outsourcing per value chain. This results in an answer to the question which elements of a financial malware scheme are most likely to be either vertically integrated or shifted to the underground market. Last, we lever these answers to conclude on potential chokepoints in financial malware schemes, intervention strategies and future research efforts.

2.4. RESEARCH ON FINANCIAL MALWARE

2.4.1. STATE-OF-THE-ART

As stated earlier in this paper, the total puzzle of the malware ecosystem has been recently researched by its separate pieces. Looking not only at separate pieces, but at the entire puzzle, will allow us to assess the different elements of the total malware ecosystem in an integral manner. This integral view will enable us to discern – based on the economics of cybercrime discussed in sections 2.2-2.3 – novel value chains in financial malware. Before we can actually connect the pieces to construct such value chains, we have to put the current state of the art in research on these pieces in the right conceptual perspective. Namely the perspective, where the piece is located within the puzzle or in this case within the overarching financial malware ecosystem. By clustering the different pieces of research a) the total malware ecosystem will become apparent, b) research gaps can be identified and c) value chains in financial malware can be distinguished.

From the mid 2000's onwards mostly computer scientists, but to some extent also social scientists have researched elements of the financial malware ecosystem - which we present in Table 2.1. First, there are studies on the source code and crimeware toolkits [7]. Second, researchers also looked at how malware infections occur and in more detail who is most likely to be infected and how specific online behavior influences these chances. Third, the infrastructure needed for the operation of financial malware is extensively studied, in particular banking botnets and its command and control (C&C) servers [8]. Fourth, the target selecting mechanism that is being operated in the financial malware scheme, e.g. which bank to 'hit' and which not, is being researched. Fifth, the cash-out strategies [9] in financial malware are studied, wherein money mules form the most frequent object of study. Last, the underground markets in relation to financial malware are being separately researched, covering a wide array of studies into underground services.

These specific parts - e.g., crimeware or infrastructure - have been identified before by Sood, Bansal & Enbody [137] and presented as clusters in their work aiming at 'dissecting the state of the underground enterprise'. If we follow their lines of analysis, and stick with the clusters we have described above, we can synthesize the state of the art of research into parts of the financial malware scheme. The studies in each cluster make up a range

[7] In this case, studies aimed at the understanding of the automation of malware source code and toolkits - like the Zeus toolkit, which became the prime monetization model of the infamous financial malware after the source-code became public.

[8] Studies into the automation of the infrastructure supporting cybercrime, such as servers commanding and controlling computers in a botnet used as such an infrastructure.

[9] The term cash-out refers to activities enabling actors to access, remove, and drain funds from bank accounts on and off-line [75]. Here, 'alternative currencies' - like gift cards - are frequently used to transfer funds outside of the traditional financial system.

of divergent concepts as research objects. These concepts are presented in the far right column of Table 2.1.

Part of a financial malware scheme	Literature	Studied concepts
A. Crimeware (Source code & Set-up)	Alazab et al. [10, 11]; Ben-Itzhak [23]; Binsalleeh et al. [25]; Boutin [30]; Criscione et al. [44]; Garcıa-Cervigon and Llinas [68]; Riccardi et al. [131]; Sood and Enbody [137])	Malware source code typologies; Crimeware; Cybercrime toolkits; Web injects
B. Infections (Victimization)	Bossler and Holt [29]; Holt and Bossler [74]	Victimization risk; online routine activities
C. Infrastructure	Ganan et al. [67]; Neugschwandtner et al. [122]; Oro et al. [125]; Park et al. [127]; Riccardi et al. [130]; Watkins et al. [161]	Botnet (detection) Command & Control servers (lifespan)
D. Target selection	Florencio and Herley [65]; Ronchi et al. [132]; Tajalizadehkhoob et al. [144]	Threat model; attack selection; attack vectors
E. Cash-out	Aston et al. [20]; Florencio and Herley [64]	Money mules; Cash-out strategies
X. Underground Markets	Caballero et al. [37]; Christin [39]; Grier et al. [71]; Holz et al. [77]; Miller [111]; Motoyama et al. [120]; Rossow et al. [133]; Sood et al. [136]; Stevens [141]; Zhuge et al. [167]	Cybercrime or financial malware-as-a-service

Table 2.1: Parts of a financial malware scheme and their literature and studied concepts

In this overview, we see the clusters of the state of the art research on financial malware mapped on the before mentioned parts of the financial malware ecosystem. It is noticeable that a lot of research efforts were taken on the malware source code and the specific set-ups, i.e. crimeware toolkits. Next to the source codes and set-up part, the last couple of years has seen an increase in research interest in the infrastructure used in financial malware schemes, aimed both at the botnet itself as at the C&C's. On the other hand, we can observe that both the study of malware infections and the cash-out strategy have little research attention.

2.4.2. MAKE OR BUY?

More specifically, if we look at the identified parts of the financial malware ecosystem in Table 2.1, the literature on underground markets shows plenty of outsourcing opportu-

nities. These parts-for-sale form the underground market counterpart to the option of self-organizing - or in the light of the previous sections, vertically integrating - which as illustrated before requires a significantly higher skill-set. Next to its function as a platform providing alternatives for vertically integrating, the underground market can be seen a facilitator in the search for specialized vendors [140].

Part of a financial malware value chain	**Underground alternative to vertically integrating**
Crimeware (Source code & Set-up)	Exploit-as-a-service, Crimeware-as-a-service, Source code for sale/free, Exploit kits
Infections	Pay-per-install; Drive-by-downloads
Infrastructure	Botnet-lease; C&C-rent
Target selection	Payload, Web Inject/Config-files for sale
Cash-out	Money mule recruitment services; Bitcoin Exchanges; Gift cards; Prepaid Credit Cards

Table 2.2: Underground alternatives for parts of the financial malware scheme

Table 2.2 shows that for every part of the financial malware scheme an underground alternative is available, based on the literature clustering on underground markets in relation to financial malware. In a typical financial malware scheme the choice exist of for instance, using in-house malware developers or an existing crimeware toolkit bought via an underground market. The same choice exists in every other cluster, ranging from choosing between setting up your own botnet and spreading malware or renting out an infrastructure and use a pay-per-install service to recruiting your own money mules or using an underground cash out service. But do all these specific underground alternatives get used in the same composition every time round? Or form the same scheme in every instance? And which of the parts tend to be most likely serviced by an underground service provider?

2.4.3. ARCHETYPICAL VALUE CHAIN

A first value chain in financial malware we can discern is the chain associated with the established and well-researched man-in-the-browser attack. An average citizen, using online banking like many others, first comes into contact with this financial malware scheme when ordinarily browsing the internet or checking up on email. In hindsight, we know that the criminal then already set up the first two parts of the scheme, consisting of (1) the source code and/or crimeware kit of the specific banking malware or trojan and (2) the infrastructure supportive to the specific malware. These both leverage vulnerabilities in for example internet browsers like Internet Explorer or malicious websites, to (3) infect

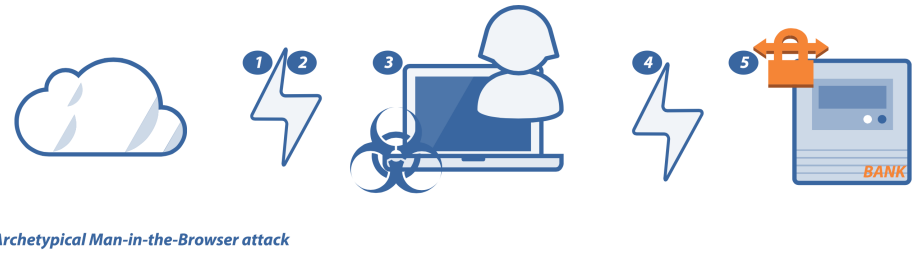

Figure 2.1: Archetypical Man-in-the-Browser attack

these potential victims with the financial malware in question. However, this malware only becomes operational under two conditions: one, the bank the infected client is using, has to be specifically targeted by the cybercriminals and two, the infected client must use the internet browser the malware exploits a vulnerability in. When the infected client then uses his or her browser for online banking with the specifically targeted bank, the cybercriminals use their man-in-the-browser attack to automatically take-over (4) the active banking session to change amounts and bank routing numbers to wire funds to bank accounts under their (in)direct control. Last, the funds stolen will be (5) cashed-out by primarily money mules using ATM withdrawals or the purchasing of high-end or luxury consumer goods. Figure 2.1 shows this man-in-the-browser attack in some more detail.

2.4.4. ONGOING DEVELOPMENTS IN FINANCIAL MALWARE SCHEMES

When we look at publications by known security firms and respected security blogs, we can see that a differentiation in attacks can be observed. First, we still see a continuing momentum of man-in-the-browser (MitB) attacks with evolving modus operandi and ever more sophisticated set-ups. [10] Next, there is a shift observable to increasingly manual and thereby more dynamic, instead of automated, web injects to execute these attacks in the web browser. Furthermore, we see a similar shift to the mobile browser and/or platform as attack vector. [11] These attacks are both scale-able as to some level standardized, allowing

[10] https://blog.kaspersky.com/the-big-four-banking-trojans/
http://krebsonsecurity.com/2015/02/fbi-3m-bounty-for-zeus-trojan-author/
http://www.kaspersky.com/about/news/virus/2014/Kaspersky-Lab-Discovers-Chthonic/
http://blog.trendmicro.com/trendlabs-security-intelligence/new-banking-malware/

[11] http://blog.trendmicro.com/trendlabs-security-intelligence/german-users-hit-by-dirty-mobile-banking-malware-posing-as-paypal-app/
http://www.americanbanker.com/issues/179_114/first-major-mobile-banking-security-threat-hits-the-us-1068100-1.html
https://securelist.com/blog/research/57301/the-android-trojan-svpeng-now-capable-of-mobile-phishing/
https://securityintelligence.com/svpeng-mobile-malware-expanding-to-new-territories/

a higher frequency of attacks. Yet, are not suitable for a more targeted approach. Second, we can distinguish a fairly new trend, wherein criminals use Remote Access Trojans (RAT) to target small and medium sized businesses (SMEs) [12].

In this manner they infect - via spear phishing - business computers, to observe the internal banking or accounting systems. When the criminals have complete insight into the company's financial systems, they hit. For instance, manipulating salary batches that the HR department generates using their financial systems. Then the salary batch is being executed by the bank, like they normally do. The only difference being that not the employees, but the criminals get their monthly pay. As the rewards are high and criminals are moving on to other companies, the eventual detection of the fraud is to be seen as relatively insignificant. Third and last, we note the similar use of RAT, however not aimed at businesses to get to their bank accounts, but aimed at the banks themselves. Therein the same modus operandi is used. Namely, infecting – in this case – bank employees' computers with RAT via spear phishing in order to gain insight of and control over crucial internal banking systems. Once the compromised systems are that familiar to the criminals, they hit. Most famously, the case of Carbanak or Anunak illustrates this scheme as highly targeted and professionally executed with estimations of up to hundreds of millions of dollars in loot. [13]

Even though the last two trends both use more generic malware type RAT, albeit fine-tuned to their specific use, in contrast to the more specialized financial malware used in the first trend, all three show the overall variation of financial malware schemes. Knowing this variation in schemes, takes us back to the original question we asked ourselves. How do criminal actors choose between organizing the tasks in a financial malware scheme themselves – thus vertically integrating the entire operation – or outsourcing (parts of) their total scheme? Only when we look at attackers, victims, targets in a holistic way, we can observe economic mechanisms per type of financial malware scheme. This requires an integral approach trough value chains, based on the previously explained economic perspectives on organized (cyber)crime.

[12] https://www.europol.europa.eu/content/users-remote-access-trojans-arrested-eu-cybercrime-operation
http://krebsonsecurity.com/2014/05/blackshades-trojan-users-had-it-coming/
http://www.symantec.com/connect/blogs/blackshades-coordinated-takedown-leads-multiple-arrests
https://www.europol.europa.eu/content/major-cybercrime-ring-dismantled-joint-investigation-team
http://securityintelligence.com/cybercrime-ecosystem-everything-is-for-sale/

[13] http://www.kaspersky.com/about/news/virus/2015/Carbanak-cybergang-steals-1-bn-USD-from-100-financial-institutions-worldwide
https://www.fox-it.com/en/press-releases/anunak-aka-carbanak-update/
http://krebsonsecurity.com/2015/02/the-great-bank-heist-or-death-by-1000-cuts/

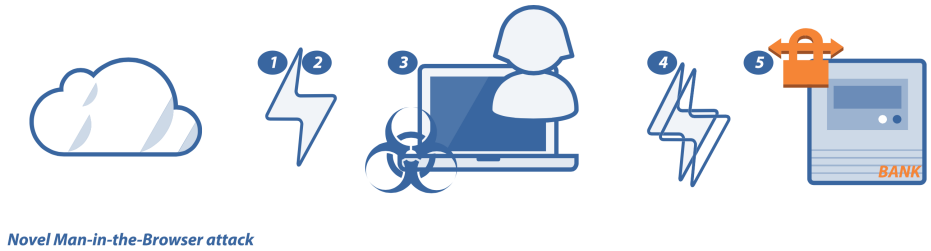

Figure 2.2: Novel Man-in-the-Browser attack

2.4.5. NEW FINANCIAL MALWARE VALUE CHAINS

With the overview of both the state-of-the-art of research in financial malware as well as the three presented differentiations in today's financial malware practices, we can leverage these insights to discern the novel value chains behind those practices. Like previous studies that examined the value chain behind spam, we present the three value chains in today's financial malware practice in the same step-by-step manner [99, 145]. Next, using the transaction cost economic model we presented before in the context of (financial) cybercrime, we can unravel the intrinsic incentives of both outsourcing as vertically integrating, per value chain. In this instance we look at the elements of the value chain and apply the framework of Dick [53]. Finally, we can hypothesize how the underground market will be involved as the 'market-of-choice' when not vertically integrating and thus using market resources to operate an individual financial malware scheme.

NOVEL VALUE CHAIN 1: UNTARGETED CONSUMER-ORIENTED MITB-ATTACK

The first novel value chain in financial malware we can discern, is the chain associated with the already well-known man-in-the-browser attack. Under reference to the described developments in this type of attack, we see a slightly different chain compared to the archetypical one. This novel chain uses near similar steps as its established counterpart (see Figure 2.2).

However, the operated crime ware kit (1) in this case allows the attacker to use dynamic web-inject instead of fully automated versions. The infections (2) are identical to other type of man-in-the-browser attacks. Moreover, the infrastructure (3) has to be set-up for these dynamic web-injects, having human operated scripts to change the web-inject from attack to attack. Again, the malware only becomes operational under two conditions: one, the bank the infected client is using, has to be specifically targeted by the cybercriminals and two, the infected client must use the internet browser the malware

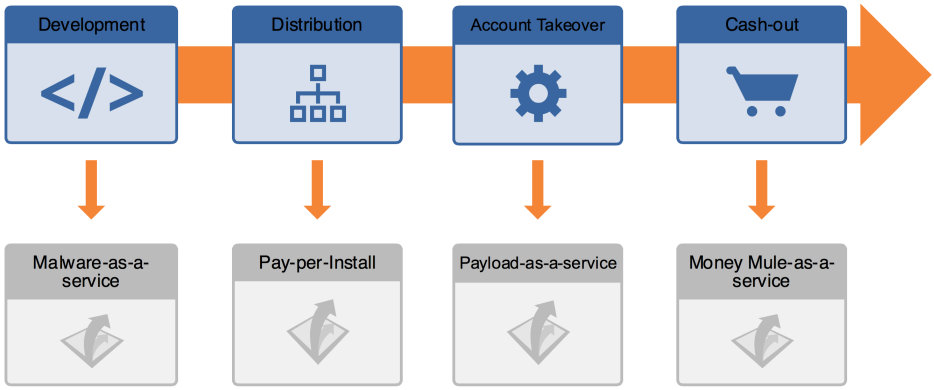

Figure 2.3: Dynamic man-in-the-browser attacks in a value chain perspective

exploits a vulnerability in. When both these conditions are met, the attackers infiltrate the active banking session with a dynamic web-inject (4), varying from pop-up windows for additional login, to creating extra fields in a form. This creates the necessity of an actual human operator to execute these dynamic types of attack. Again, the objective is to manipulate the banking session in such a manner that money is transferred to bank accounts controlled by the cybercriminals, without raising suspicion in the active session. Like the more static man-in-the browser attacks, the funds stolen will be (5) cashed-out by primarily money mules using ATM withdrawals or the purchasing of high-end or luxuries consumer goods. Figure 2.3 shows these more dynamic man-in-the-browser attacks and its resources in a value chain perspective. Here, we observe that all resources can either be vertically integrated or outsourced.

NOVEL VALUE CHAIN 2: SEMI-TARGETED SME-ORIENTED RAT-ATTACK

The second discernable value chain is that of a financial malware scheme using RAT to target SMEs (Figure 2.4).

Unlike the first novel chain, wherein the chances of getting infected are fairly random, here the first contact the potential victims have with the financial malware scheme is nearly always a semi-targeted (1) spear phishing attempt. With this method the criminals single out employees at an exploitable position in the targeted companies, such as the financial administration. Once the often-infected attachment to the spear phishing email has been opened, (2) the RAT sourcecode and/or crimeware kit already in place then has control over the (3) infected client. The criminals, with the RAT having unrestricted access to the infected client, can observe the internal (financial) systems of the targeted SME and

Figure 2.4: Remote Access Tooling targeting SMEs

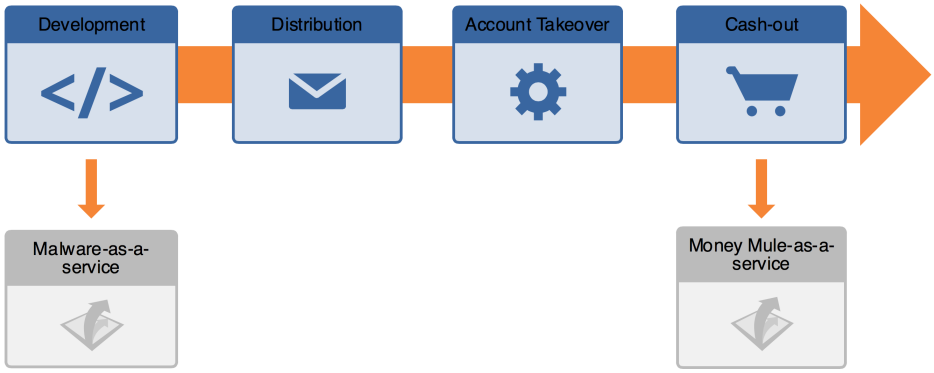

Figure 2.5: Resource value chain of semi-targeted SME-oriented RAT-attacks

spend some time getting familiar with the day-to-day financial practices of the company. By the time they have a full and profound understanding of the systems and know its potential exploitability, they target a (4) specific process in the system. For example, they manipulate salary batches so they get paid instead of the company's employees. Like the other value chains, the last step in the scheme is the (5) cash-out strategy involving money mules to get the stolen money to the criminals.

Figure 2.5 shows the resource value chain of semi-targeted SME-oriented RAT-attacks. Here, we observe that in contrast to the previous value chain, not all resources can be outsourced. Development and cash-out can still either be outsourced or vertically integrated. Yet, distribution and account take-over are highly dynamic in nature, and therefore do not meet the standardized supply, and have to be self-organized.

NOVEL VALUE CHAIN 3: TARGETED (FINANCIAL) BUSINESS-ORIENTED RAT-ATTACK

The third and last novel value chain that can be discerned, is the chain that like the previous one uses RAT. But instead of SMEs this financial malware scheme involves the targeting of banks directly (Figure 2.6).

Figure 2.6: Remote Access Tooling targeting Banks

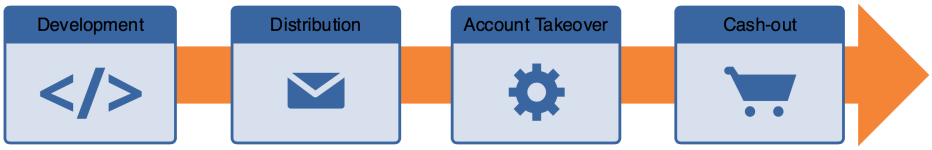

Figure 2.7: Resource value chain of Remote Access Tooling targeting Banks

The first three steps of the chain, are identical to Novel Value Chain 2. With the difference – of course – that the (1) spear phishing emails are in this case send to singled-out bank employees at exploitable position within the bank's internal hierarchy and the (2) RAT thereafter is active on the (3) infected clients within the bank. Again – like the second novel value chain – the criminals first observe the complex internal systems and seek vulnerabilities. However, in this case the cybercriminals operating such as scheme are not in it for the quick buck, but for the long run. Having undetected and unrestricted access to internal bank systems is a potential gold mine. Once they find ways in which they can (4) manipulate the internal banking systems, they shift from observing to acting. The infected clients are used to authorize transactions, create back-to-back loans, hand out mortgages without underlying pledges or trick ATMs in thinking they have a withdrawal. In contrast to the other two value chains, this chain does not primarily rely on money mules as parts of the cash-out strategy. Criminals who have been operating such a scheme relied on the (5) bank systems themselves as their most prominent source of cash-out. Ranging from the use of ATMs (the famous example of ATMs spitting out money on cue) or setting up accounts whereof the bank thinks they do not exist or simply erasing transactions after they have been executed from the bank systems.

Figure 2.7 shows the resource value chain of targeted (financial) business-oriented RAT-attack. Here, we observe that all resources have to be vertically integrated and cannot be outsourced. All resources are highly dynamic in nature and have to be self-organized.

2.5. INCENTIVES FOR SHIFTING TO THE MARKET

Now that we have a view of the three different novel value chains in current-day financial malware schemes, we can apply the framework proposed by Dick [53] to unravel the underlying incentives per chain. The goal of this application is to analyze how the different components of the value chain are suited to be organized within the criminal organization. Thus vertically integrating, or to be shifted to the market, making use of the vast amount of earlier described 'underground market alternatives'. With the last option comes the inevitable financial transaction to be made between the criminal organization and the underground market salesman. As elaborated upon in section II, the transactions characteristics are – in this case – deal maker or breaker in shifting a specific activity to the market. The analysis of these characteristics form the basis of the framework contemplated by Dick [53]. His framework consists of the elements: a) is the activity suitable for 'large scale production'? b) how specific can the accompanied transaction be described? c) what is the frequency wherein this transaction would take place? and d) what is the uncertainty of the transaction? If we map out the different elements of this framework on the three discerned value chains, we can build the following overview (Table 2.3).

Value Chain	Scale	Specificity	Frequency	Certainty
1. MitB	++	+	+	+/-
2. RAT -> SME	+/-	-	-	+/-
3. RAT -> BANK	–	–	–	–

Table 2.3: Value chains in a TCE perspective

NOVEL VALUE CHAIN 1: UNTARGETED CONSUMER-ORIENTED MITB-ATTACK

Looking at the first novel value chain, the first three elements (sourcecode/crimeware kit, infrastructure, infections) score some positive points on the different components of the framework. Starting with the scale, we have shown that MitB attacks primarily rely on infections in bulk, accompanied by thus a large infrastructure of infected clients. Next, the activities in the first elements can be described very specifically due to the standardized way of operation and the high availability of the most popular banking malware toolkits almost all MitB financial malware schemes use. In turn, the interplay between attackers and defenders, i.e., banks and software developers closing security gaps, creates the necessity of updating the more static parts of the MitB financial malware scheme. Resulting in more frequent transactions on activities such as crimeware toolkits and infections. Moreover, all these activities are common commodities available on the underground market and therefore almost guaranteeing a continuous supply of these

activities. As a consequence of potentially doing business with underground market salesmen, the risk of being scammed – the uncertainty of the transaction – is evident. However, in the case of activities being sold by the dozen, by a wide range of sellers and in most cases with an Amazon-like review system in place, the uncertainty is to a large extent downplayed or sometimes even neutralized. This results, for these three elements – sourcecode/crimeware kit, infrastructure, and infections - in an incentive to shift these specific activities to the (underground) market. However, the development of using more dynamic web-injects instead of automated ones, relying on personal interaction and therewith human operators, has the side-effect of diminishing part of the specificity and thereby scale needed for potential outsourcing. Time will tell whether or not we see an ongoing process of shifting from outsourcing back to vertical integration in these MitB attacks.

Yet the elements of target selection and cash-out are somewhat different in relation to the other elements in the first value chain. For target selection, config files are used that instruct the malware to become active when visiting certain predetermined online banking environments, based on the specific domain name of the bank in question. These config files often come with the crimeware toolkit and are not frequently sold separately. The same goes for money mules in the cash-out strategy, which in turn are not being sold by the bulk as frequently as for instance in the flourishing pay-per-install market. It can be argued that maybe these activities are both too valuable and too scarce and therefore not sold as much as other commodities. Acquiring and aligning these parts for your own financial malware schemes seems to be hard enough, let alone to sell these on the underground market. In this case it is not merely the low incentive to shift these activities to the market, as it is the lack of a stable underground market alternative preventing doing so. As a result both these elements form potential chokepoints in the MitB value chain, thereby creating new possibilities for interventions aimed at these elements.

NOVEL VALUE CHAIN 2: SEMI-TARGETED SME-ORIENTED RAT-ATTACK

Moving on to the second novel value chain, we can observe a nearly mirrored mapping on the different components of the framework. As we have demonstrated before, financial malware schemes using RAT coincide with a more targeted approach. The scaling thus depends on size of the criminal organization operating such a scheme as well as the targeted companies in terms of expected return-on-investment. Whereas the MitB scheme uses scale to make itself profitable, the schemes using RAT focusing at SMEs start out at least, low in scale. To go however for the bigger score per attack, takes time, thus lowering the scale of the scheme but on the other hand increasing the reward per attack to maximize the profitability. Using spear phishing for a RAT infection is to some extent a

standardized routine, but the activities carried out after the infection – the observation and identification of potential cash cows in unknown internal financial systems – is not to be called 'specific'. Therefore, resulting in a lower specificity of the activity to be described on forehand. Logically this is also not a high frequency activity, more a high intensity activity. Albeit that the RAT itself is such a prevailing commodity, that it actually is available for free both on the Dark- as on the clear web. So in conclusion, next to the RAT itself, there is little incentive to shift the other specific activities in a financial malware scheme using RAT targeting SMEs to the (underground) market.

NOVEL VALUE CHAIN 3: TARGETED (FINANCIAL) BUSINESS-ORIENTED RAT-ATTACK

Last, with regard to the third novel value chain, we can see a similarly mirrored mapping on the different components of the framework compared to the first value chain. Like the second novel value chain, this chain encompasses a financial malware scheme with a (highly) targeted approach. Again, the scaling depends on size of the criminal organization operating such a scheme as well as – in this case – the continuous and patient efforts to exploit the targeted bank in the long run. With a scheme targeting banks, the rewards can be dazzling if the scheme operates under the radar of security measures implemented at the targeted bank. This requires being able to alter the modus operandi at least from day-to-day and perhaps even from hour-to-hour. That intrinsically creates such an unspecific and infrequent but highly intensive activity, that even the possibility of actually considering shifting this activity to the underground market is likely to be absent. With shifting this activity to the underground market also comes the revealing of a maybe very lucrative financial malware scheme to potential competitors. All in all, next to maybe the RAT itself, the conclusion is that there is no incentive to shift the other specific activities in a financial malware scheme using RAT targeting banks to the (underground) market.

2.6. CONCLUSION

The still evolving current-day financial malware schemes can be captured by three novel value chains. We constructed these value chains based on the conceptual synthesis of the state-of-the-art of the literature on financial malware and the known differentiations in financial malware schemes. A framework of transaction cost economics was used to illustrate the incentives that influence decisions within such a value chain to either vertically integrate or outsource specific parts. Combined with the notion of an increase in underground market activity, this depicts the impact of underground commodities to setting up a financial malware schemes.

The goal of this paper was to integrate literature on financial malware, underground

markets and (cyber)crime economics as well as industry reports, to discern novel value chains in financial malware. These value chains were constructed in a similar fashion to how other researchers reconstructed the spam value chain. The constructed value chains, aided by the framework of Dick [53], allowed us to analyze the economic principles within the underlying criminal business models. This resulted in the answering of the question which elements of a financial malware scheme are most incentivized to be either vertically integrated or shifted to the underground market. We demonstrated that for financial malware schemes using MitB attack vectors there is a clear incentive to shift (parts) of this scheme to the underground market, in contrast to financial malware schemes that rely on RAT. The development of a more dynamic and human operation, tends to diminish part of these economic incentives to outsource.

Next, we believe that in our approach we have shown that a transaction cost economic approach is greatly beneficiary to the series of existing economic perspectives on cybercrime in general and on financial malware schemes in particular. This approach generates new insights as it comes to understanding cyber criminals operating criminal schemes and doing business with other (cyber)criminals in an underground market. We laid down the deliberative considerations and actions that accompany the shifting of a part of a financial malware scheme to the (underground) market. Thereby proving the potential of underground markets in kick starting the opportunity to operate a financial malware scheme.

Furthermore, by conceptually synthesizing the state-of-the-art of literature in financial malware, we have – next to an overview of the current research efforts - also identified research gaps in financial malware research. Moreover, we have made evident that a value chain approach will be of added value when researching financial malware (schemes) or underlying business models of those who operate it. This creates the opportunity to study the important interactions between the strategies of attackers on the one side and the properties and policies of the financial service providers on the other.

Finally, we came to conclude on the different incentives that are apparent in the different value chains. In turn these incentives can be used to analyze chokepoints in the value chain. More specifically, if scarcity of one activity in particular on the underground market influences those incentives, chokepoints derived from these incentives are vital to future interventions. Based on these chokepoints, not only interventions for financial service or security providers but also for – perhaps even more important - law enforcement purposes can be developed.

3

COMMODITIZATION

Researchers have observed the increasing commoditization of cybercrime. That is, the offering of capabilities, services, and resources as commodities by specialized suppliers in the underground economy. Commoditization enables outsourcing, thus lowering entry barriers for aspiring criminals, and potentially driving further growth in cybercrime. While there is evidence in the literature of specific examples of cybercrime commoditization, the overall phenomenon is much less understood. Which parts of cybercrime value chains are successfully commoditized, and which are not? What kind of revenue do criminal business-to-business (B2B) services generate and how fast are they growing?

We use longitudinal data from eight online anonymous marketplaces over six years, from the original Silk Road to AlphaBay, and track the evolution of commoditization on these markets. We develop a conceptual model of the value chain components for dominant criminal business models. We then identify the market supply for these components over time.

3.1. INTRODUCTION

Many scientific studies and industry reports have observed the emergence of cybercrime-as-a-service models, also referred to as the "commoditization of cybercrime." The idea is that specialized suppliers in the underground economy cater to criminal entrepreneurs in need of certain capabilities, services, and resources [73, 98, 136, 145]. Commoditization allows these entrepreneurs to substitute specialized technical knowledge with "knowing

what to buy" - that is, outsourcing parts of the criminal value chain. The impact of this trend could be dramatic: Commoditization substantially lowers entry barriers for criminals, which is hypothesized to accelerate the growth of cybercrime. Prior work found strong evidence for specific cases of commoditization: booters offering DDoS services [86], suppliers in "pay-per-install" markets distributing malware [37], and exploit kit developers supplying "drive-by" browser compromises [71]. The overall pattern is much less clear, however, as not all cybercrime components are equally amenable to outsourcing [64].

This paper answers two core questions: Which parts of cybercrime value chains are successfully commoditized and which are not? What kind of revenue do these criminal business-to-business services generate and how fast are they growing? Addressing these questions requires that we properly define and scope the concept of commoditization. To do so, we turn to transaction cost economics (TCE). We argue that the characteristics of commodities are highly congruent with the characteristics of online anonymous marketplaces. More precisely, the one-shot, anonymous purchases these markets support require suppliers to offer highly commoditized offerings. Conversely, if cybercrime offerings can be commoditized, online anonymous markets should be a highly attractive place to sell them. Indeed, these platforms can reach a large audience and provide risk management services for criminals, e.g., by protecting their anonymity, and featuring reputation systems to root out fraudulent sales and shield sellers from risky interactions with buyers.

While data from online anonymous marketplaces provides a unique opportunity to track the evolution of commoditization, we are not arguing that these marketplaces provide a complete picture. They do not have a monopoly, of course. In fact, certain types of commoditized offerings are not suited for trading on these marketplaces, e.g., affiliate programs, subscription-based offerings, or services requiring a rich search interface may be better served by alternative distribution channels [79, 167]. Yet, on balance, the congruence of commoditized forms of cybercrime and online anonymous markets means that the evolution of commoditization should be clearly observable on those markets.

We analyze longitudinal data on the offerings and transactions from eight online anonymous marketplaces, collected between 2011 and 2017. We first present a conceptual model of the value chain components in dominant criminal business models, and develop a classifier to map cybercrime-related listings across all markets to these components. This allows us to track trends in vendors, offerings and transaction volumes. We then discuss the type of offerings to assess to what extent each component can be outsourced - i.e., to what extent it is successfully commoditized. We make the following contributions:

- We present the first comprehensive empirical study of the commoditization of cybercrime on online anonymous markets. We analyze 44,000 listings and over 564,000 transactions across eight marketplaces. We draw on data from prior work [139] and newly collected data on AlphaBay.
- We find commoditized business-to-business offerings for most value chain components, though many of them are niche products with only modest transaction volumes. Cash-out services contain the most listings and generate the largest revenue. We estimate the lower bound of overall B2B revenue to be around $2 million in 2016 and over $8 million for the whole period.
- We also uncover a surprising amount of revenue in retail cybercrime – that is, business-to-consumer sales rather than business-to-business, similar to the patterns observed for drug sales. The lower-bound estimate for 2016 is over $1 million and nearly $7 million for the whole period.
- We demonstrate that commoditization is a more spotty phenomenon than previously assumed. The lack of strong growth in transactions suggests that bottlenecks remain in outsourcing critical parts of criminal value chains.

The rest of this paper is structured as follows. Section 3.2 defines transaction cost economics, and discusses how the concept applies to cybercrime commoditization. Section 3.3 describes the demand of cybercrime outsourcing. Section 3.4 presents our measurement methodology. Section 3.5 lays down our classification analysis, and section 3.6 identifies the best-selling clusters of cybercrime components. Section 3.7 discusses our findings, and Section 3.8 connects our work to earlier contributions. Section 3.9 concludes.

3.2. COMMODITIZATION AND ANONYMOUS MARKETPLACES

With outsourcing, entrepreneurs can decide to either "make" or "buy" each component of the value chain. Transaction cost economics (TCE) is a mature economic theory that seeks to explain under what conditions economic activity is organized in markets (buy) and when it is vertically integrated (make) – i.e., the entrepreneur develops the component himself or brings someone with that capability into the enterprise. Here, we apply TCE to the context of cybercrime to predict if and when outsourcing takes place.

Williamson [166] distinguishes several asset characteristics that determine if and how outsourcing will occur, as shown in Figure 3.1. *A*, *B*, and *C* are various forms of outsourcing and *D* is vertical integration. Factors such as asset specificity, frequency and

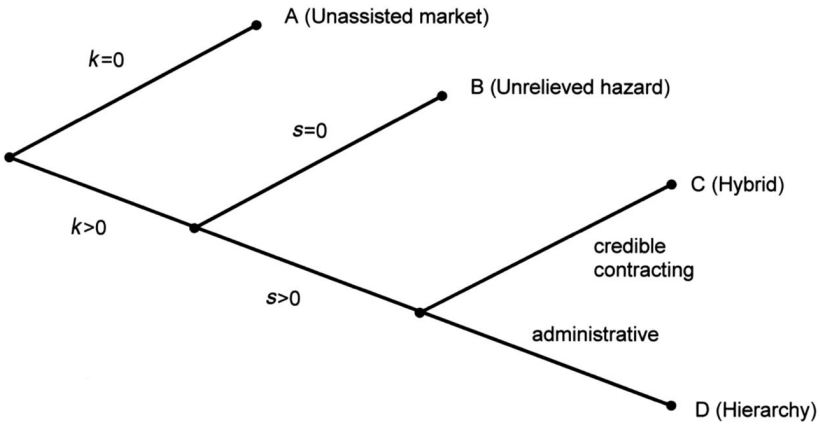

Figure 3.1: **Contracting scheme in the TCE framework.**

uncertainty separate the underlying transactions [164]. k is a measure of asset specificity, referring to the degree to which a product or service is specific to e.g., a vendor, location, control over resources, etc. A key characteristic of commodities is that they are "fungible", meaning that different offerings of it are mutually interchangeable ($k = 0$) – i.e., a booter is a booter [86, 89] – and subject to vendor competition [51]. In commodity markets, buyers can easily turn to other suppliers, and suppliers can sell to other buyers, reducing possible hazards. The more specific an asset is ($k > 0$), the more investments are specialized to a particular transaction.

The second factor, s, refers to contractual safeguards. Transactions where investments are exposed to unrelieved contractual hazards ($s = 0$) will not be traded publicly (i.e., anonymous online marketplaces such as Silk Road or AlphaBay are a poor fit), but on smaller, "invite-only" markets where trust relations are forged among specialized insiders, anonymity is not absolute, and escrow services are less prominent [120]. When $s > 0$, contracts with transaction-specific safeguards are in place.

Commodities are sold via unassisted markets (A). These markets incentivize sellers to reduce asset specificity as much as possible, hence commoditizing the offering. The efficiency gains also work in the other direction: those who offer goods or services that can be commoditized would use these markets to sell them and benefit from the wide reach and high frequency of transactions, without being exposed to risky direct interaction and coordination with buyers.

In terms of TCE, online anonymous marketplaces are unassisted markets – i.e., they are the place to go for commoditized cybercrime. Anonymous markets reduce uncertainty

risks through escrow mechanisms, review systems and strict rules enforced by a market administrator [39, 139]. For transactions where $k = 0$, "no specific assets are involved and the parties are essentially faceless" [165, p. 20], which is precisely the case for anonymous markets. Complex components such as highly customized malware are more likely to be self-supplied or delivered under special contracts, while frequently used, standardized components, like DDoS-services, would be supplied more efficiently by the unassisted market. TCE tells us that the organization of criminal activities will be guided primarily by the relative costs of completing illegal transactions within the market [53, p. 28].

Similar to the prominent drugs-trade on anonymous online markets, we expect two type of commodities on these markets: business-to-business (B2B), e.g., wholesale quantities of credit card details, and business-to-consumer (B2C), e.g., a handful of Netflix accounts. We are primarily interested in B2B, as that is the form of commoditization that is the most worrying and speculated to cause a massive growth in cybercrime, though we will also report the main findings for B2C. To assess the degree to which B2B services are commoditized, the next section develops a framework to identify the different value chains where there is demand for commoditized cybercrime.

3.3. DEMAND FOR CYBERCRIME OUTSOURCING

To empirically assess the commoditization of cybercrime, we first need to establish what capabilities, services and resources criminal entrepreneurs actually need. This provides us with a framework against which to evaluate where commodities are available to meet this demand and where they are not – as measured through listings on anonymous marketplaces. Of course, entrepreneurs might demand an endless variety of goods and services. For this reason, we use as our starting point the dominant criminal business models that were identified in prior work. We look at the value chain underlying each business model and synthesize them in a common set of components that entrepreneurs might want to outsource. Our point of departure is Thomas et al.[145]'s inventory of criminal business models. We update and extend this set with models discussed in related research. Table 3.1 shows this updated overview.

First, we look into the value chain behind spamvertising, which is driven by three resources: a) advertisement distribution b) hosting and click support and c) realization and cash-out [99, 145].

Second, extortion schemes, for instance ransomware or fake anti-virus [43] have a value chain that consists of four distinctive resources: a) development of malware b) distribution, by either exploits or (spear)phishing e-mails, c) take-over and "customer service" and d) cash-out [87, 145].

Table 3.1: **Overview of present-day cybercriminal business models**

Business model	Example Modus Operandi	Source
Spamvertised products	Selling knock-off products	Levchenko et al. [99], Thomas et al. [145]
Extortion	Ransomware	Kharraz et al. [87], Thomas et al. [145]
Clickfraud	Hijacked traffic	Kshetri et al. [92], Thomas et al. [145]
Social engineering scams	Customer support scams	Miramirkhani et al. [112], Christin et al. [41], Thomas et al. [145]
Fraud	Financial malware	Thomas et al. [145], Van Wegberg et al. [154]
Mining	Cryptocurrency mining	Huang et al. [80], Thomas et al. [145]
Carding	Credit card reselling	Holt [73], Thomas et al. [145]
Accounts	Reselling credentials	Holt [73], Thomas et al. [145]

Third, click fraud is supported by four similar, general resources: a) development of a website, malware or a JavaScript, b) distribution through botnets, c) take-over by either malware or JavaScript and d) cash-out [92, 145].

Fourth, the criminal business model in social engineering scams, such as tech support scams [112], or one-click fraud [41] leans on: a) (optional) development of malware or a malicious app, b) distribution by phishing e-mail or website, or through social engineering, c) take-over and setting-up "customer service", and d) cash-out [112, 145]. The boundary between extortion and social engineering scams is fuzzy. Both could well be categorized in the same family. For now, we take the view that extortion (e.g., ransomware) requires development of malware, where social engineering scams do not necessarily rely on anything being installed on the victim's machine (e.g., one-click frauds [41]).

Fifth, cybercriminal fraud schemes, e.g. those enabled by financial malware, build on four general, main resources: a) development and b) distribution of malware or a malicious app, c) take-over, for instance by using web-injects or a RAT,[1] and d) cash-out [145, 154].

Sixth, cryptocurrency mining relies on near-similar resources as click fraud: a) the development of malware or JavaScript, b) distribution of malware by botnets or the

[1] Remote Access Tool, i.e., malware that allows a miscreant to remotely access a victim's machine.

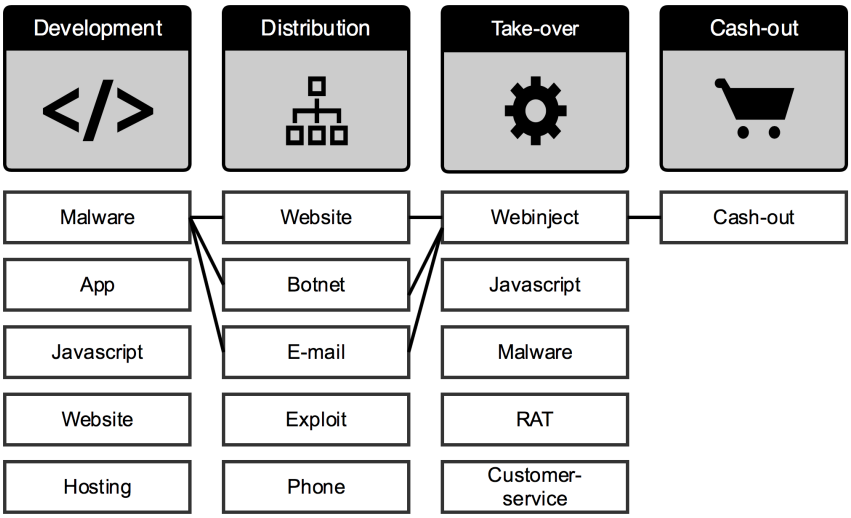

Figure 3.2: **Conceptual model of value chains, showing a representation of the financial malware value chain**

injection of a JavaScript in a compromised websites, c) the take-over, i.e. mining, and d) cash-out [80, 145].

Seventh, the criminal business model that profits from selling stolen credit card details makes use of: a) development of a phishing website, malware or a malicious apps, b) distribution, c) take-over, i.e. the logging of information, and d) reselling and cashing-out [73, 145].

Last, the resale of non-financial accounts leans on the exact same resources as carding [73, 145].

Looking at these value chains, we can see that some components are common among them. All models relate to at least four main resources: development, distribution, take-over and cash-out. We merge these into a single component that belongs to two or more value chains. We can synthesize all value chains in a overall set of 13 components. Some components, e.g., malware, can be used for more than one main resource. Figure 3.2 summarizes our conceptual model and the overall demand for B2B services in cybercrime.

3.4. MEASUREMENT METHODOLOGY

Our measurement methodology consists of 1) collecting and parsing data on listings, prices and buyer feedback from eight prominent online anonymous markets, 2) implementing and applying a classifier to the listings to map them to cybercrime components from our conceptual model of value chains (Figure 3.2) as well as to additional categories

of B2C cybercrime, and 3) using Latent Dirichlet Allocation (LDA, [26]) to identify the best-selling clusters of listings and compare their offerings to the capabilities, resources and services needed for each component of the conceptual model.

3.4.1. DATA COLLECTION

We first leveraged the parsed and analyzed dataset of Soska and Christin [139] to obtain information about item listings and reviews on several prominent online anonymous marketplaces. For each of the over 230,000 item listings, the data include (but are not limited to) titles, descriptions, advertised prices, item-vendor mapping, category classification, shipping restrictions and various timestamps. Additionally, each item listing contains feedback that has been proven to be a reasonable proxy for sales [39, 139]. Each piece of feedback contains a message, a numerical score, and a timestamp.

We then extended this data with an additional 16 complete snapshots of AlphaBay that we collected from May 30, 2016 to May 26, 2017, just two months before its closure in July 2017 [63]. Table 3.2 summarizes the dataset. We merged the new AlphaBay scrapes with the existing dataset by first parsing out the same supported fields and then running a compatible analysis using the categorical classifier from Soska and Christin [139].[2] AlphaBay is important since, according to the FBI [63], by the time of its closure, it had featured over 100,000 listings for stolen and fraudulent documents, counterfeits, and malware in particular. The US Department of Justice (DoJ) also claims that AlphaBay was the largest single online anonymous marketplace ever taken down [3].

Table 3.2: **Markets crawled**

Market	First seen	Last seen	# Snapshots
Agora	2013-12-24	2015-02-11	161
Alphabay	2014-12-31	2017-05-26	33
Black Market Reloaded	2012-11-21	2013-12-04	25
Evolution	2014-01-13	2015-02-18	43
Hydra	2014-04-14	2014-10-26	29
Pandora	2013-11-02	2014-10-13	140
Silk Road 1	2011-06-21	2013-08-19	133
Silk Road 2	2013-11-27	2014-10-29	195

As an important data processing note, some vendors set "holding prices" to their listings when the product or service they are selling is out of stock. Instead of removing the listing, these vendors increase the price (astronomically) to prevent buyers trying

[2] Soska and Christin's dataset included 17 snapshots of AlphaBay, dating back to December 2014, that they did not use in their published analysis [139].

to buy their product. Soska and Christin [139] developed a heuristic that corrects these holding prices, which we applied in the pre-processing of the parsed and labeled dataset. This limits the potential for errors stemming from falsely assuming a certain holding price was associated with a buy.

3.4.2. CLASSIFYING CYBERCRIME LISTINGS

Most listings on these marketplaces are related to drugs and other non-cybercrime activities [39, 139]. Our aim is to classify each item listing into one of the 10 categories of cybercrime components from the conceptual framework (Figure 3.2). Unfortunately, the labels provided by Soska and Christin are not expressive enough to capture these nuanced categories, so we begin by using their labels as a pre-filter and retain only item listings that were identified as being either "Digital goods" or "Miscellaneous" (19% of all listings).

Next, we implemented a Linear Support Vector Machine (SVM) classifier. Manual inspection confirmed our suspicion that the markets also contain retail (B2C) cybercrime offerings, next to wholesale cybercrime offerings. For this reason, we added six product categories to distinguish supply in that part of the market: accounts, custom requests, fake documents, guides and tutorials, pirated goods, and vouchers. A final category, namely, "other", captures the listings that did not fit anywhere else (e.g., scanned legal documents). The classifier is initially trained and evaluated on a sample of listings ($n = 1,500$) from all the markets, where ground truth is created via manual labeling.

Table 3.3: **"Digital Goods" & "Miscellaneous" Listings**

Market	# Listings	# Vendors	Total revenue
Agora	3,240	526	$ 1,818,991
Alphabay	21,350	3,055	$ 13,471,406
Black Market Reloaded	2,069	386	$ 685,108
Evolution	9.551	1.002	$ 6,125,136
Hydra	377	28	$ 242,230
Pandora	1,204	169	$ 394,306
Silk Road 1	4,053	645	$ 2,239,436
Silk Road 2	2,734	441	$ 4,455,339

3.4.3. GROUND TRUTH

For labeling the ground truth, we randomly selected 1,500 items from all listings classified as either "Digital Goods" or "Miscellaneous" ($n = 44,060$), or approximately 3.5% of the data. Only around 30% of the listings in the random sample belonged to one of the ten

B2B cybercrime components. Around 45% belonged to one of the B2C categories and the remaining 25% were labeled as "other." Those were comprised of drug listings that were misclassified as "miscellaneous," as well as luxury items and other physical goods. We also found some incomprehensible listings, which might be test entries by vendors. Labeling the ground truth yielded four more observations. First, we identified listings that contain more than one cybercrime component, e.g., offering both a piece of malware and (access to) a botnet. Second, we identified *package listings*, such as complete cryptocurrency mining schemes. Third, we observed that some vendors add unrelated keywords to their listings, presumably in a marketing effort similar to search engine optimization. Fourth and last, we observed *custom listings*, i.e., listings that are specifically created to be sold only once to one specific buyer. Custom listings contain bespoke products or services ranging from custom quantities to a completely custom-made product such as pre-booked plane tickets.

After labeling our random sample of listings, we can assess whether each category meets our criteria for accurately classifying listings to categories of cybercrime components. To avoid overfitting to a specific component, we ensure the training set for our classifier holds at least 20 listings per category of cybercrime components. Because of the highly skewed distribution of listings in our random sample, we were forced to increase our ground truth by manually adding listings to the following categories: app, botnet, e-mail, exploit, hosting, malware, phone, RAT and website. To that end, we operated a manual search in the filtered portion of data using up to three keywords on those cybercrime components. We manually verified whether the listings with the keyword in the title or description advertise the actual product or was a false positive – e.g., a vendor using the word "malware" in a listing of lottery tickets.

3.4.4. TRAINING AND EVALUATION

Before training the classifier, we excluded three categories of cybercrime components from the classification: JavaScript malware, webinjects, and customer support. For these, we found no listings in our random sample.

The classification phase itself consists of three steps: (i) data cleaning, (ii) tokenizing, (iii) training and evaluation of the ground-truth samples which are the concatenation of the title and description of the item listings. In data cleaning, we removed all English stop words, punctuations, numbers, URLs and accents of all unicode characters. We then lemmatized the words in order to group together the inflected forms of a word so they can be analyzed as a single item, identified by the word's lemma, or dictionary form before being trained and tested. We tokenized each item (assuming all items are in English)

and computed a *tf-idf* (term frequency inverse document frequency) value for each of the resulting 9,629 unique tokens or words. To calculate the *tf-idf*, we used a *max-df* (maximum document frequency) equal to 0.7 – this discards words appearing in more than 70% of the listings. In the classification phase we then used these values as an input for an L2-Penalized SVM under L2-Loss. We implemented this classifier using Python and *scikit-learn*.

The reported imbalance in the distribution of listings among categories causes an imbalance in the labeled categories of our ground truth. On the one hand, we have nearly 25% of listings labeled as "other" and around 45% labeled as one of B2C products or services. On the other hand, we have a large portion of the rest of our ground truth listings (30%) that are labeled as "cash-out" listings (25%). We mitigate the negative impact of this imbalance on our classification results by re-sampling our ground truth listings by the SMOTE (Synthetic Minority Over-sampling Technique) method, thereby increasing the cardinality of each category to match the size of the largest labels; this is a standard technique towards improving algorithmic fairness. Due to the implicit optimization of our classifier, this over-sampling method allows the model to carve broader decision regions, leading to greater coverage of the minority class [38].

Because of the nature of listings that cover multiple categories, e.g. bundled goods, we anticipate some classification errors. It is however important to distinguish between errors where the item listing is classified as "other" (false negative) from acceptable approximations, e.g., a listing that includes access to a botnet bundled with malware and is classified as a botnet. The first example denotes a classification error, while the second is a listing that truly is a combination of multiple cybercrime components. Our main goal is therefore to prevent cybercrime component listings, like malware, from ending up in "other" and vice versa.

We evaluate the performance of our classifier in Figure 3.3. In this normalized confusion matrix, each row represents the instances in an actual category while each column represents the instances in a predicted category. All correct predictions are in the diagonal of the table (numbers denote recall). The average precision is 0.78 and the average recall is 0.76, denoting some confusion between cybercrime components categories. However, the classifier meets our goal of avoiding confusion between cybercrime components and "other" listings.

3.4.5. POST-PROCESSING

The heuristic for dealing with holding prices [139] used in pre-processing does not correct situations where all instances of a listing among our snapshots were either only seen with

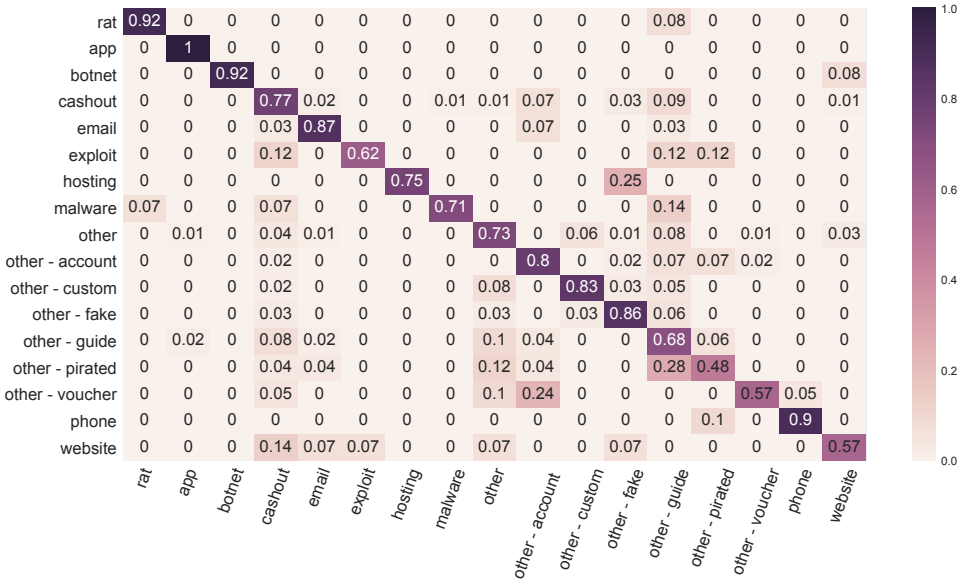

Figure 3.3: **Classifier normalized confusion matrix**

a holding price, or in some cases do not exceed a set maximum of $10,000. To get an idea of how frequently this happens, we looked into items priced above $5,000. We manually identified 12 listings which received a total of 118 pieces of feedback at holding prices. In one case we found the correct price from a customer commenting "good product for $10". The remaining 11 listings seemed clear instances of holding prices, and were removed, as we had no information about the true sales price.

After examining holding prices, we found some instances of misclassified drug listings in categories of cybercrime components (false positives). To correct this, we first removed 12 Xanax listings that we encountered when inspecting the holding prices. To find additional misclassified drug listings, we leveraged the distinctive features of drug listings, namely the unique terminology used to list the quantity of drugs offered, e.g., "grams," "mg," "ug," "lbs," "ml," "pills," etc. Following this process, we automatically identified and removed 82 misclassified drug listings.

3.5. RESULTS

In this section we present the results of the classified listings. At first glance, we can observe the differences in number of listings between the categories. Just over 30% of the listings are in the B2B categories of our conceptual model, listed in the top half of

Table 3.4. The lower half of the table covers B2C cybercrime (around 36% of listings), custom orders (14%) and others (20%).

Table 3.4: **Listings per category.** The top half represents B2B listings, the bottom half, B2C.

Category	# Listings	# Vendors	Total revenue
App	144	75	$ 12,815
Botnet	125	79	$ 46,904
Cash-out	12,125	2,076	$ 7,864,318
E-mail	550	216	$ 97,280
Exploit	115	75	$ 17,603
Hosting	20	15	$ 1,182
Malware	310	162	$ 57,598
Phone	261	148	$ 74,587
RAT	105	65	$ 16,070
Website	664	293	$ 286,405
Accounts	3,759	577	$ 598,491
Fake	3,386	815	$ 2,877,184
Guide	5,049	1,020	$ 2,620,635
Pirated	1,420	338	$ 129,961
Voucher	1,293	386	$ 753,116
Custom	6,310	1,887	$ 5,793,064
Other	8,424	2,652	$ 7,749,788
Total	44,060	5,552	$ 28,997,006

We primarily focus on the B2B categories, though we do report on the B2C categories later in the section. Before we turn to B2B offerings, we take a closer look at the large category of custom listings. These listings are a bit counter-intuitive to the market structure as they concern one-time, buyer-specific products or services. For instance, stolen credit card details from Norway, a modified type of keylogger, or compromised hosts from the Netherlands. Although some of these listings are in fact B2B cybercrime services, they are not fully commoditized, as the listing reflects a one-time sale and a non-standardized product or item.

There are large differences across the categories of B2B offerings. Cash-out stands out: In terms of the number of listings, active vendors, and in total revenue, this category is by far the largest. It also stands out in other ways. Table 3.5 reports the median and mean number of listings for each vendor per category, which reflects the degree in which different products need to be differentiated. We see most products offered do not need differentiation. More specific requests might be handled with custom listings, but are not enough to merit a more permanent listing. Cash-out offerings, on the other hand,

Table 3.5: **Vendors, revenue per category**

Category	Listings per vendor		Revenue per listing
	Median	Mean	Median
App	1	1.97	$24.33
Botnet	1	1.61	$34.44
Cash-out	2	5.88	$60.00
E-mail	1	2.58	$22.85
Exploit	1	1.56	$15.57
Hosting	1	1.33	$31.60
Malware	1	1.95	$22.90
Phone	1	1.80	$30.00
RAT	1	1.66	$20.00
Website	1	2.28	$29.80

Table 3.6: **Price and lifespan per category**

Category	Price per listing				Lifespan in months
	Mediam	Mean	Min–Max	SD	Median
App	$5.70	$18.79	$0–$64	$40.89	0.91
Botnet	$14.73	$106.89	$0–$2,475	$341.13	0.60
Cash-out	$14.85	$72.42	$0–$9,756	$280.20	0.72
E-mail	$7.34	$42.14	$0–$1,606	$139.17	0.52
Exploit	$5.26	$28.64	$1–$500	$80.09	0.36
Hosting	$16.40	$25.14	$3–$99	$25.47	0.32
Malware	$5.45	$37.96	$0–$1,984	$133.68	0.98
Phone	$9.90	$45.13	$0–$3,200	$221.99	0.79
RAT	$5.41	$38.35	$0–$919	$126.78	1.44
Website	$8.72	$51.58	$0–$1,695	$146.42	0.83

contain many more relevant distinctions. A vendor can split up its stock of stolen credit card details into smaller sets of details, for instance differentiated to type of credit card.

The second column in Table 3.5 shows median revenues per listing. Cash-out listings have the highest median revenue. RATs and exploits exhibit, counterintuitively, a similar median revenue. This is a consequence of the generally low-value exploit listed in anonymous marketplaces, e.g., run-of-the-mill Office exploit macros. Rare, high-value exploits, such as iOS or Chrome exploits, would be sold through specialized white or black markets or through private transactions [7]. Other categories have a median between $15 and $34 revenue per listing. As the median revenue is a simple summary of the

underlying distribution, we also show the price range in Table 3.6 – in terms of median, mean, min-max and standard deviation (SD) – for listings in the B2B categories. We see, again, that the cash-out category contains the most expensive set of offerings with very diverse pricing. This diversity in price can also be observed in other categories – in fact, the overall shape of the price distribution function remains relatively unchanged across categories. Moreover, the lifespan of a listing from Table 3.6 also tells us something about the standardization of the product. A listing that receives instances of feedback over multiple months denotes that the associated product remains valuable and has not become outdated or unrecognizable. Like an ecstasy tablet, a RAT will hold its value over time in terms of being a functional solution. In contrast, stolen credentials "go bad" after some time. The first buyer who uses these credentials will in all likelihood set off red flags at the credit card company for irregular spending, making a subsequent purchase of the same credentials worthless. Curiously, the median lifespan of cash-out listings is above average, which could be due to vendors updating the specific product listed, or persistently selling unusable credit card details, or to a slower-than-expected detection of suspicious transactions by credit card companies.

Looking into median lifespan of listings reveals little differences as all but three categories have a median listing lifespan of close to one month. Both exploit and hosting listings have a low median lifespan of around 0.3 months – approximately 10 days. At the other end of the spectrum, we see that RAT listings have a median lifespan of 1.44 months – approximately 40 days. So, a RAT listing has a significant longer lifespan than an exploit listing. The distribution of cybercrime listing lifespan is heavy-tailed and on average, a cybercrime component is offered for 2.7 months. In short, vendors have one or two listings, except for cash-out listings, where that number is higher. Turnover is between $15 and $60 dollars per listing and lifespan is typically less than a month.

3.5.1. LISTINGS AND REVENUE OVER TIME

The claim that cybercrime is commoditizing also implies a growth in transactions and revenue. Figure 3.4(a) shows, per month, the unique number of listings and number of feedback. Figure 3.4(b) shows the corresponding projected revenue. The number of feedback is a proxy for the minimum number of sales, as a buyer can only leave feedback when she buys a product. Feedback does not however yield a one-to-one mapping to sales as customers may leave a single piece of feedback after purchasing a high quantity of an item. Anonymous marketplaces depend on effective reputation mechanisms to mitigate uncertainty in transactions.

Figure 3.4 shows a growth in listings, amount of feedback and revenue for cybercrime

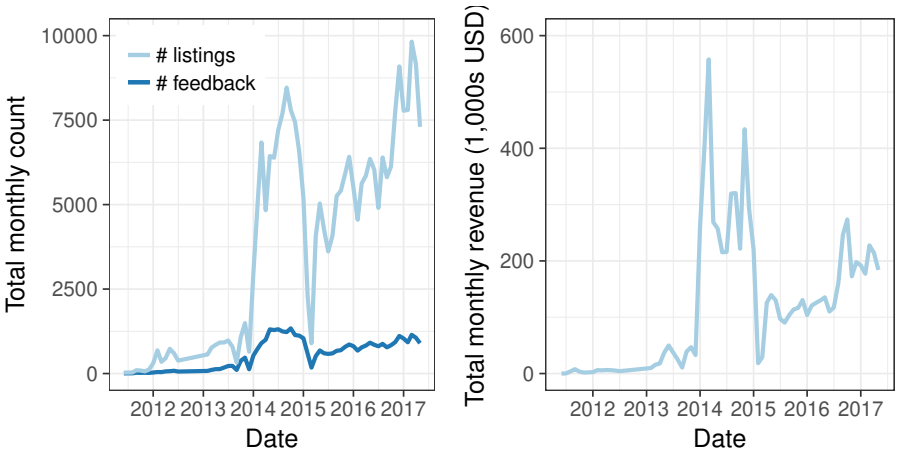

(a) Listings and feedback amount

(b) Revenue

Figure 3.4: **Number of unique listings, feedback and revenue in categories of B2B cybercrime components per month**

components between 2012 and 2017. The drop at the end of 2013 and the beginning of 2014 is partly due to the take-down of Silk Road 1 and Black Market Reloaded. The steep increase thereafter is distributed over four new markets (Agora, Evolution, Hydra and Silk Road 2), but shows that the aggregate pattern is clearly one of rapid growth. The next drop, around the end of 2014, is caused by a combination of the law enforcement operation against Silk Road 2, the exit scam of Evolution and the sudden disappearance of Agora. Right after this volatility, AlphaBay emerged, and subsequently became the largest to date. Their operation halted suddenly in July of 2017, when the FBI together with the Dutch Police shut down AlphaBay. Still, the overall pattern clearly is one of growth. The trade in cybercrime commodities seems resistant to the turbulence across marketplaces.

Figure 3.5 shows that the upward trend in feedback instances is not only caused by an increase in listings, but also to the increase of amount of feedback per listing. In 2011, a listing on average received around five pieces of feedback per month. Over time, this ascended to around eight pieces of feedback per listing in 2017, with intermediate spikes to over ten pieces of feedback in 2012. Those spikes coincide with the period of time in which Silk Road 1 became known by the general public due to extensive coverage by news and media over the course of 2011 [1]. Conversely, the trough at the end of 2013 is primarily due to the Silk Road 1 takedown and the chaotic few weeks that ensued [139]. Overall, we see that the average amount of feedback per listing stabilizes halfway through 2012 and from that moment onwards seems to follow a slow rise.

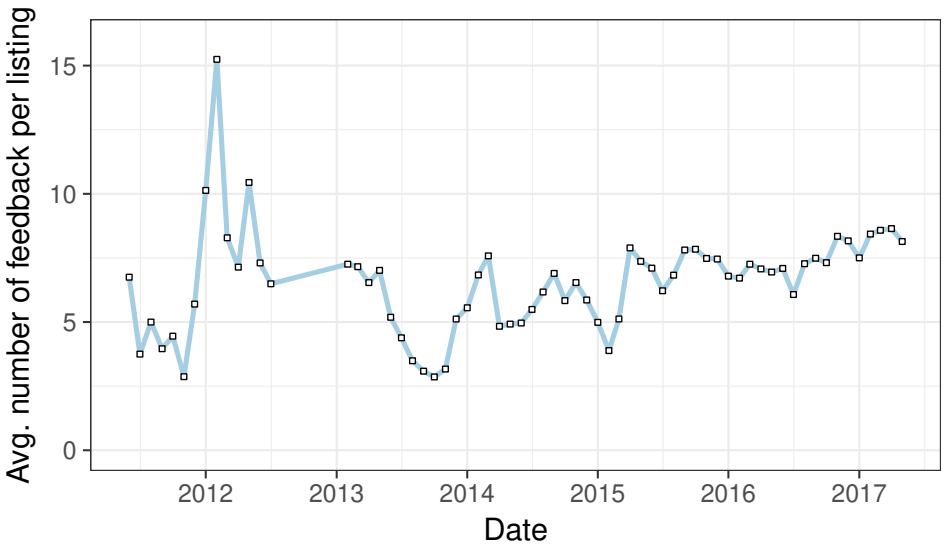

Figure 3.5: **Feedback per listing in categories of B2B cybercrime components per month**

Essential to the understanding of the ecosystem is identifying which categories can be attributed to most of the growth in sales and revenue. For each item listing, revenue is calculated by multiplying each feedback specific to a listing with the dollar-price of that listing at the moment the feedback was generated. The revenue from these listing is then aggregated per month and per category. Figure 3.6 shows revenue per category. The spikes and troughs are, again, the result of marketplace turbulence.

The category of cash-out listings is by far the biggest cybercrime component, in terms of listings, revenue and vendors. We take a closer look to see whether this revenue is driven by a small fraction of listings or whether it represents a broader volume of trade. It turns out the a large portion of the increase between 2014 and 2015 is driven by feedbacks on CVV listings. More specifically, one listing offering "US CVVs" received nearly 700 feedbacks in the first quarter of 2014. From the beginning of 2015 onwards we see a steady growth in revenue alongside the growth of AlphaBay market as a whole. In the early days of the ecosystem we see an increase in cash-out revenue which was primarily driven by a listing offering "10,000 USD CASH," which can be seen as typical money laundering – the customer pays in bitcoin and receives cash.

The revenue of cash-out listings is obscuring the other categories. When we omit it in Figure 3.7, we see that the trend of increasing revenue between 2012 and 2017 becomes apparent yet again. In the second half of 2014, listings in e-mail distribution such as spam tutorials, spam runs or large databases of e-mail addresses generate very high revenue

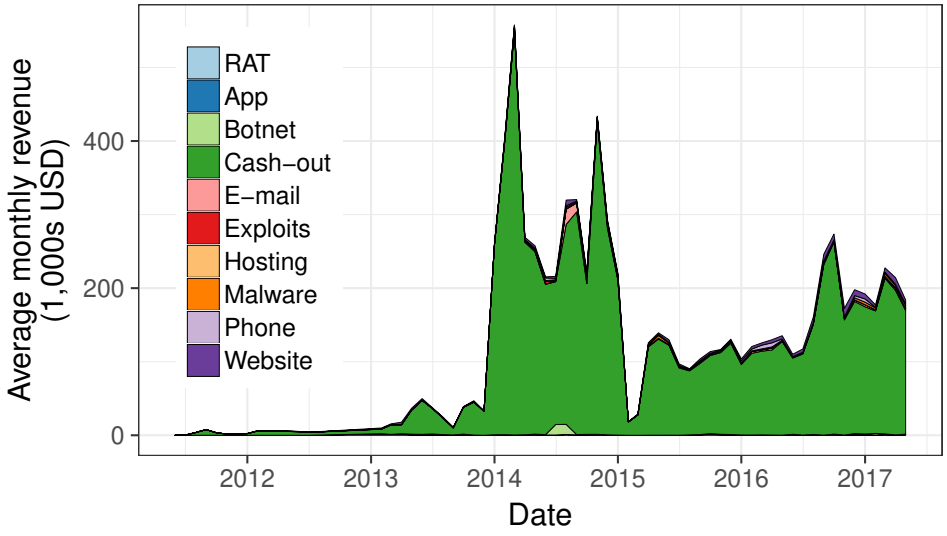

Figure 3.6: **Total revenue per category per month**

numbers. Similarly, we see a spike in botnet-related sales driven by a mysterious listing titled "source," receiving 10 – rather negative – feedbacks in the summer of 2014. The average of $5000 per month in 2013 grows to $15,000 per month in late 2017. Compared to the average monthly revenue of the entire market ecosystem however – nearing $600,000 per month in late 2014, mostly generated by drugs [139] – this is just a fraction.

3.5.2. VENDORS OVER TIME

Another element in the assessment of commoditization is the level of vendor competition. Figure 3.8 shows the number of vendors per category over time. A vendor is defined to be active if she has at least one active item listing and may be instantaneously active in multiple categories.

As with the revenue per listing, the number of unique vendors per category is generally increasing over time, however the increase in vendors is steeper than the increase in listings. Figure 3.9 clearly shows that the increase in vendors from 2014 onwards is due the Evolution and AlphaBay marketplaces. Soska and Christin showed that in the contemporary ecosystem (i.e., after the Silk Road take-down), it is common for each vendor to maintain more than one alias on different marketplaces which may be partially responsible for this observation.

Listings and revenue are not distributed normally across vendors. As in many markets, there are big players and small players. Figure 3.10 plots the cumulative percentage of

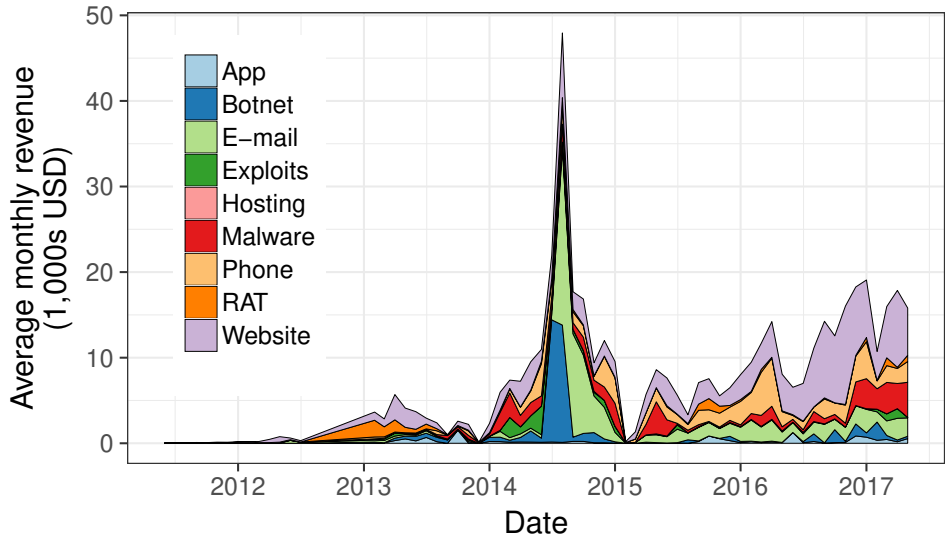

Figure 3.7: **Total revenue per category per month, excluding cash-out category**

listings and revenue of cybercrime components over vendors. A small portion of vendors are responsible for a large fraction of the listings. To be more precise, around 30% of vendors are responsible for 80% of all listings. More interestingly, just under 10% of vendors are responsible for generating 80% of the total revenue. That means that around 174 vendors have sold nearly $7 million worth of cybercrime components. This translates into an average revenue per vendor of around $40,000, but the distribution is wide and skewed. The 174 vendors range from a minimum revenue of $7,355 to a maximum of $1,148,403.

3.5.3. MARKETPLACES

Different marketplaces might develop different profiles or specialties in terms of what they sell – i.e., they attract a different set of vendors, offerings or buyers. To compare the product portfolio of different markets, Figure 3.11 displays the distribution of offerings across different categories. To deal with the large differences in size of the categories, we first take the logarithm of the number of listings in each category and then calculate the percentage of each category in this log-scaled total count of listings. There are minor variations visible, but the more obvious pattern is the similarity between most markets. All except two markets, namely Hydra and Pandora, contain listings in each of the categories. Hydra and Pandora are relatively small markets, with a shorter life-span and the absence of listings in some categories is probably due their comparatively modest size and short

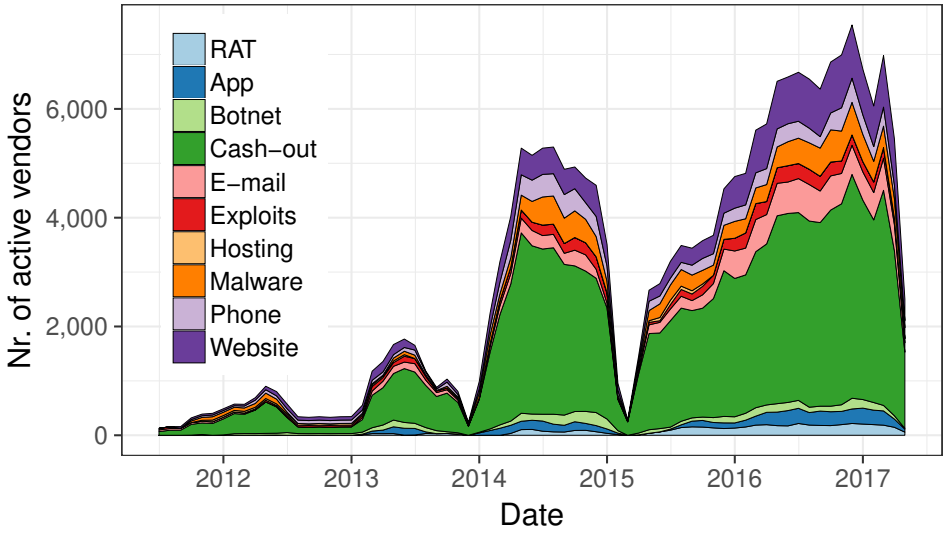

Figure 3.8: **Number of active vendors per month**

existence. In terms of commoditization, all categories of the criminal value chains are consistently offered across markets and time. Moreover, these components see vendor competition.

Another way to evaluate market specialization is by categorical revenue. In Figure 3.12, we show the percentage of revenue – after log transformation – per category of cybercrime component per market. The story is unchanged: there are no major differences between markets. If anything, the picture painted by looking at revenues is even more uniform across marketplaces.

3.5.4. B2C LISTINGS

Finally, we take a look at the listings in retail cybercrime. This covers the categories of "accounts," "fake," "guide," "pirated," and "voucher." We briefly describe the type of listings assigned to those categories. "Accounts" denote listings advertising small batches of accounts from services like Netflix and Spotify. "Fake" contains offerings of fake IDs, counterfeit documents or money. Listings that sell mere instructions or tutorials, are categorized as "guides." The "pirated goods" category encompasses listings that offer pirated movies, software or e-books. Last, the "voucher" category comprise listings that offer discounts at numerous places, ranging from discounted airline tickets to pizza shop

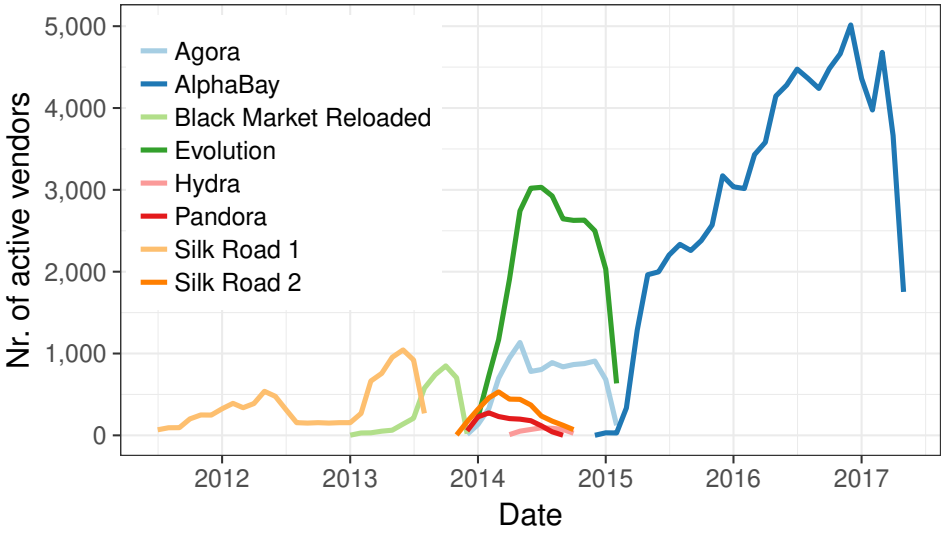

Figure 3.9: **Vendors over time across markets**

gift cards[3]. The retail cybercrime offerings are also forms of commoditization, albeit a slightly atypical one. Indeed, these B2C products are meant to be used or consummed by the buyer, and are not parts for some large value chain with another profit center at the end of it.

The large portion of retail cybercrime is in line with what has been observed on the drugs side of these markets; B2C transactions for consumers of drugs, along with more modest amounts of B2B transactions with larger quantities for lower-level dealers [13].

We do not know however what type of listings within one category are the driving forces for these growing number of listings, feedbacks, vendors and revenue. To understand how commodization of cybercrime components really takes place, we have to look at finer grained information. To do so, we next cluster listings within cybercrime component categories and characterize the supply by analyzing the best-selling clusters within each category.

3.6. CHARACTERIZING SUPPLY

We now want to delve deeper into what is actually being offered in each category and how this supply compares to the overall demand for criminal capability, resources and services

[3] Interestingly, in underground slang, "pizza" may also denote credit card listings—which are sold in "slices." While this vernacular could *a priori* be confusing to an automated classifier, manual inspection suggests misclassification is very rare, as we will discuss later.

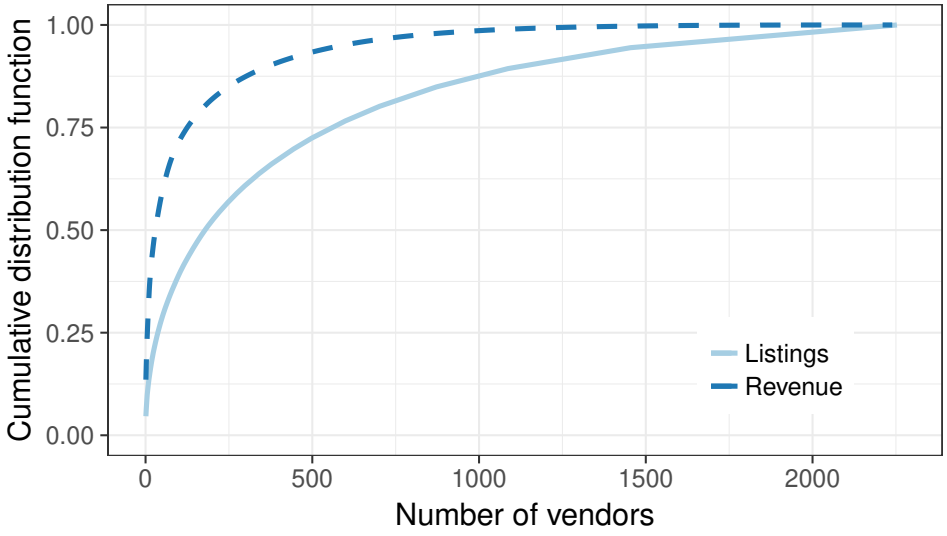

Figure 3.10: **Cumulative distribution function of listings and revenues across vendors**

in that category. We apply unsupervised clustering to the listings in each category and then interpret the three best-selling clusters.

3.6.1. CLUSTERING LISTINGS

The detailed sub-classification is created by identifying clusters within our categories of listings using Topic Modeling. We rely on the Latent Dirichlet Allocation (LDA) [27] clustering algorithm to determine the main topics from a text corpus. We cleaned the data (removing broken fragments and correcting egregious errors) and lemmatized the words before clustering.

Our goal is to extract and analyze the three clusters which represent the "main themes" in each category. A natural choice might be to select the three clusters whose items collectively generate the largest amount of revenue. However we observed that a small fraction of very expensive items tends to obfuscate this analysis, thus we instead opted for identifying these "main theme" clusters based on the number of unique feedbacks. LDA is parameterized by a hyper-parameter that upper bounds the number of clusters to identify. Motivated by the expected heterogeneity of listings in the categories of cybercrime components, combined with the assumed homogeneity in other categories, we set this parameter to 10. As a consequence, it may be the case that LDA will not generate clusters for small categories of listings (when the true number of clusters exceeds 10); and those will instead be projected into larger clusters.

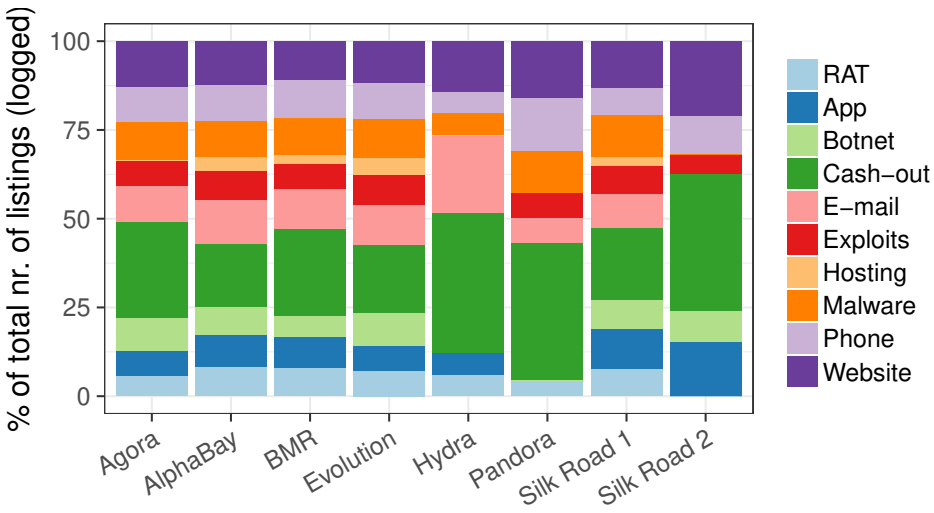

Figure 3.11: **Percentage of total number of listings per market per category (numbers of listings are logged)**

3.6.2. BEST-SELLING CLUSTERS

We identified the three best-selling clusters per category by summing the number of feedbacks of all listings in a specific cluster. We then compute the total revenue generated by the item listings in each cluster. The results are shown in Table 3.7. We excluded three categories from the classification, as explained in Section 3.4.4. For all categories, the three best-selling clusters contain more than 46% of all feedbacks, and in many cases more than 60% of all feedbacks. Looking at revenue, we observe a diffused pattern. The categories "botnet," "website," and "RAT" show lower revenue numbers. Upon manual inspection, we could identify a very small cluster with only a few feedback that was dominated by a few very expensive items.

The second part of this clustering approach aims to understanding which type of products and/or services are transacted in these main clusters. To that end, we can use the output features of our LDA clustering algorithm to label the prominent clusters, sometimes assisted by manual inspection. In the next two sections, we present our findings and elaborate on whether the identified main topic clusters fit the overall demand for criminal capability, resources and services following our conceptual model.

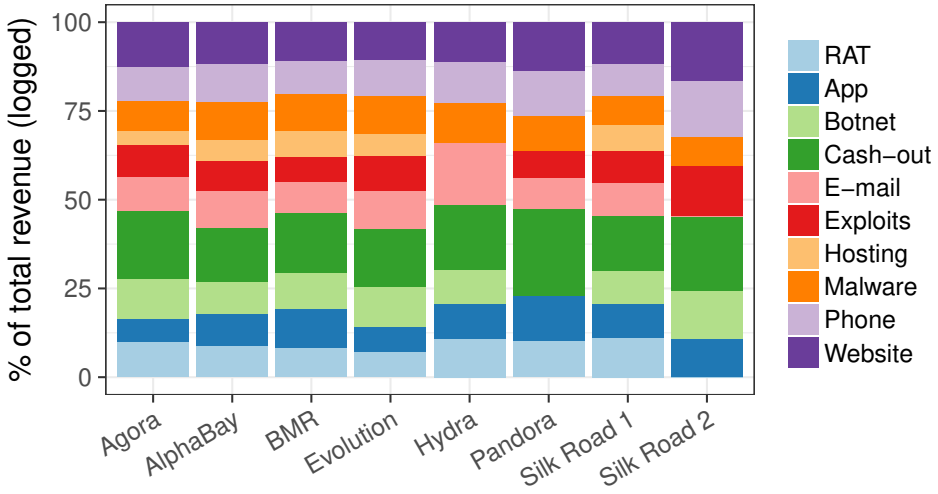

Figure 3.12: **Percentage of total revenue per market per category (revenue is logged)**

Table 3.7: **Best-selling clusters per category**

Category	# Feed-back	Top3 Feedback	Top3 Revenue
App	1,175	784 (67%)	$ 7,083 (55%)
Botnet	968	657 (68%)	$ 8,995 (19%)
Cash-out	236,566	164,124 (69%)	$ 4,991,272 (63%)
E-mail	4,684	2,605 (56%)	$ 64,642 (66%)
Exploit	1,335	936 (70%)	$ 11,514 (65%)
Hosting	120	97 (81%)	$ 829 (70%)
Malware	2,446	1,127 (46%)	$ 30,806 (53%)
Phone	2,731	1,851 (68%)	$ 48,154 (65%)
RAT	768	501 (65%)	$ 4,887 (30%)
Website	8,586	5,044 (59%)	$ 65,111 (23%)
Account	75,469	47,149 (62%)	$ 316,851 (53%)
Fake	34,341	20,568 (60%)	$ 1,386,363 (48%)
Guide	57,361	38,586 (67%)	$ 2,397,006 (91%)
Pirated	11,242	6,093 (54%)	$ 55,864 (43%)
Voucher	22,769	13,643 (60%)	$ 441,572 (59%)

3.6.3. CLUSTERS IN CASH-OUT OFFERINGS

The main clusters of the cash-out category in descending order of size are 1) credit card details, more specifically "computer-generated BINs" [4] - i.e., computer-generated credit

[4] BIN refers to Bank Identification Numbers - the first six numbers on a credit card.

card numbers that pass simple verification, but are not actually issued by banks, 2) so-called "fullz," stolen credit cards, including their full details, such as the CVV number. We can also identify a cluster pertaining to 3) guides on "making money," or money mule recruitment. Next to these three prolific clusters, we explore the seven other clusters in cash-out offerings, ordered by there relative feedback volume. We observe clusters with distinct offerings in 4) carding tutorials, 5) PayPal accounts, 6) Visa and Mastercard card details[5], 7) "bitcoin deals," 8) bank account credentials, 9) Amazon refund guides and 10) Bitcoin exchanges, specialized in cash pay-outs. All in all, we can observe a broad spectrum of cash-out solutions being offered. They range from guides, to actionable solutions, like PayPal or bank account access. Next, we can discern services aimed at cashing out cryptocurrencies, more specifically Bitcoin, through dedicated exchange services. Consistent with what previous studies showed for cybercrime forums [73], carding makes up a big part of cybercrime components transacted on online anonymous markets as well.

3.6.4. CLUSTERS IN OTHER B2B OFFERINGS

In this section we present the best-selling clusters in the other categories of B2B cybercrime components.

App. Prominent clusters of the App category include offers for Android loggers, i.e., malicious keylogger apps, Android bank apps, i.e., malicious banking apps, and Dendroid, a RAT for Android.

Botnet. Prominent clusters of botnet listings feature products and services revolving around Zeus botnets, varying from tutorials, to source-code, to "turn key" setups. We also identified offers on C&C servers and DDoS services.

E-mail. The prominent clusters in the e-mail category contain two types of spam lists, namely basic lists of e-mail addresses, as well as complete databases, including personal details to create personalized (spear) phishing mails. In addition we find a cluster of offerings on spam-related services.

Exploit. Within the exploit category, the two main themes are 1) Microsoft Office exploits, e.g., malicious macros, and 2) browser exploits. We also recorded a non-trivial set of sales for Mac exploits.

Hosting. The prominent "hosting" clusters include hosting through VPS or CPanel-listings. We also find a prominent cluster on hosting of Tor-based websites.

Malware. Within the malware category, ransomware stands out by featuring two promi-

[5] This cluster resembles 1) and 2) but with a focus on Visa and Mastercard brands. It could a priori also include gift cards.

nent clusters. One cluster revolves around the Stampado ransomware, the other on Philadelphia ransomware. We also observed a prominent cluster on miscellaneous (assistive) software tools such as keyloggers or portscanners.

Phone. In the category of phone listings, one prominent cluster comprises listings on bypassing security features on phones. The other two prominent clusters offer respectively hacked Vodafone accounts and lists of usable phone numbers.

RAT. Two out of three prominent clusters in RAT listings contain generic RATs. The third cluster specifically deals with Mac OS RATs.

Website. The website category is composed of three distinct, prominent clusters. One cluster contains website development listings. The second is predominantly VPN-connections and/or SOCKS proxies. The third cluster consists of compromised RDP-servers/hosts listings.

Our analysis suggests that nearly all prolific clusters supply a component that matches B2B demand, but that this supply is incomplete, in that the observed supply fulfills only a niche demand in each category. For instance, we see ransomware dominating the malware category, whereas domain expertise suggests there are, in general, other types of malware in demand. This demand remains mostly unfulfilled in online anonymous marketplaces.

One exception to the aforementioned trend is in the "phone" category, where supply differs from the B2B demand. Research suggests that the actual latent demand is for using phones and social engineering to trick victims into falling for a scam [28]. Yet, the supply is only oriented towards setting-up the necessary phone lines. We observed that guides and tutorials are among the prominent clusters in the botnet and cash-out categories. We however note that selling a guide is not the same as outsourcing a cybercrime component.

In summary, the demand for cybercrime components is frequently met on online anonymous markets in our dataset, but the supply is highly restricted to specific niches and the accompanied revenue is generally modest.

3.6.5. CLUSTERS IN B2C OFFERINGS

In this section we briefly present the prominent clusters in the B2C categories – i.e., retail cybercrime.

Account. In listings that sell accounts, we observed two main clusters that revolve around offerings for single accounts to pornography websites. Next, we see a cluster of listings selling Netflix and Spotify accounts, in quantities between two and ten per listing.

Fake. The three prominent clusters are respectively offering fake passports, fake IDs and counterfeit money.

Guide. The clustering process revealed guides in a) bitcoin ("deals"), b) "making money" or starting a business, and c) "scamming."
Pirated. Miscellaneous pirated software, like the entire Adobe software suite or pirated adult videos, and pirated Microsoft software, e.g. Windows 7, are the prominent clusters in pirated products.
Voucher. In the category of voucher-related listings, we see offers for: a) Tesco vouchers, b) lottery tickets and c) "free" pizzas, of which most are indeed discount vouchers or gift cards for various pizza chains, but a few are in fact credit card offerings, where "slices" refer to groups of accounts.

The nature of products and services in all of the best-selling clusters tells us that we are observing transactions of retail cybercrime. We see that the best-selling clusters within accounts are listings in smaller quantities, ranging from single hacked accounts on a pornography website, to up to ten Netflix or Spotify accounts. It may at first appear to be curious why a single user would want 10 Netflix accounts, but when considering the inherent unreliability (and short lifespan) of stolen accounts, it becomes clear that this demand is plausible for personal use.

3.7. DISCUSSION

In this section, we discuss our approach and results in light of our theoretical assumptions and research design.

3.7.1. VALIDATION

In earlier work, Soska and Christin [139] discuss the validation of measurements on online anonymous markets. They find support for using feedback instances as a proxy for sales by looking at three specific cases where ground truth is available (due to arrests or leaks). However, the online anonymous marketplace ecosystem has grown quite significantly since then - in particular, in 2017, AlphaBay itself grossed, on a daily basis, more than the entire online anonymous marketplace ecosystem did in 2014.

The criminal complaint for forfeiture against the alleged AlphaBay founder and operator [4] estimates that "between May 2015 and February 2017, Bitcoin addresses associated with AlphaBay conducted approximately 4,023,480 transactions, receiving approximately 839,087 Bitcoin and sending approximately 838,976 Bitcoin. This equals approximately US$450 million in deposits to AlphaBay."

The estimates coming from our scrapes yield US $222,932,839 (and 2,223,992 transactions) for the entire time interval (including, this time, all of the goods sold on the

marketplace). We believe the $450 million dollar from the complaint is a slight overestimate, due to currency mixing that could result in double-counting.

On the other hand, our own estimates are on the conservative side. In particular, we have to ignore a small fraction of credit card sales, due to a quirk in the way certain purveyors of credit card numbers do their business: A few stolen credit card number vendors list their items in generic form, with a price of zero, instead leaving the specifics in the shipping costs - presumably to obfuscate their stocks and possibly to reduce the commissions imposed by the marketplace operator. For instance, a listing would be for "credit card dumps," with a price of zero, but with shipping options for various types of cards at various prices. Because we cannot determine which cards are purchased, we simply conservatively ignore such sales.

More importantly, as Soska and Christin point out, it is important to repeatedly scrape online anonymous marketplaces to ensure adequate coverage [139]. This is particularly true when a marketplace is large, as the population of items is more likely to change over small time intervals. Our density of scrapes is lower in mid-2016, meaning that we might have missed a number of transactions occurring then.

All in all, we might be missing a non-negligible number of transactions occurring on AlphaBay; data for the other marketplaces is more complete, as validated in the original paper [139]. We point out, however, that these misses are unlikely to change our analysis beyond underestimating absolute sales volumes: indeed, with the small exception of the vendors using shipping costs for pricing, there are no specific biases in the missing data, so that the items we have in our corpora can be taken as a representative random sample.

3.7.2. LIMITATIONS

We next discuss the limitations of our study in two main areas: first, to what extent our data captures the commoditization of cybercrime and, second, the way we mapped the offerings on these markets onto categories of demand.

Observing cybercrime commoditization starts with knowing where to look. Building on transaction cost economics, we have argued that online anonymous marketplaces are the most logical place to trade cybercrime commodities due to the nature of these transactions. However, what seems logical from a TCE perspective does not necessarily seem logical to the criminal entrepreneur. Trust in a market is to a large extent subjective. This might mean that cybercriminals turn to other platforms with fewer safeguards to trade commoditized cybercrime. Even when criminals do follow TCE, some forms of commoditized cybercrime do not fit well with online anonymous markets: subscription models, affiliate programs, services requiring a rich search interface, or non-English

offerings [79, 167] are all ill-suited to the type of markets we are investigating here. Since we did not study these forms of trade, our picture of commoditization is incomplete. To some extent, the same holds for underground hacker forums, though we would argue that many of the transactions on those forums are not actually commoditized, but forms of contracting (see Section 3.2).

Another limitation relates to how we mapped criminal demand. Successful commoditization is not just a matter of products and services being offered. These offerings also need to meet a demand, as observed in actual sales. To understand the potential demand of cybercriminals, we worked with a scope of known business models. Building on the work of Thomas et al. [145], we have limited ourselves to cybercriminals who aim at making a profit. In other words, there may be cybercrime components that are being offered and that do match cybercriminal demand (e.g., for ideological or tactical purposes, rather than financial pursuits), yet are outside the identified value chains.

3.8. RELATED WORK

Core elements of our paper build on or benefit from recent progress in related research, which we discuss here.

Different researchers have tried to grasp the evolution of criminal activity in the underground economy. Initial work focused on underground forums [77, 120]. After the infamous Silk Road market came into existence, researchers looked closer at online anonymous markets [13, 39] and investigated the evolution of listings and revenue on these markets. Our study is among the first to explicitly leave the predominant drug listings out of scope and focus on a different product type (cybercrime). Most closely connected to our work is the first longitudinal study on the evolution of volumes in products transacted across multiple online anonymous markets by Soska and Christin [139]. Other studies focused on specialized markets or forums, for instance the stolen data and exploit market [15, 73]. They investigated the market for exploits - which turned out to be moderate in size - and the cybercrime-as-a-service market, where growing numbers of new services types were discovered. Furthermore, researchers investigated the increase in online drugs trade, specifically the B2B side of Silk Road 1 drugs offerings, and what factors determine vendor success [13].

In addition to quantitative studies of the evolution of online anonymous markets, our work is related to qualitative studies on buyers and sellers (vendors) on markets and forums. For instance, Van Hout and Bingham [78] looked into the buyers of drugs, and inspected the retail side of the market, as we did. Van Buskirk et al. [148] specifically focused on the motivation of drug buyers in Australia to turn to online anonymous

markets instead of street dealers. They found that a cheaper price and higher quality of the drug are important.

Earlier research into the commoditization of cybercrime found evidence of commoditization of a number of specific products and services. Prominent examples are booters [86], the Pay-Per-Install (PPI) market [37], and exploit kit developers supplying drive-by browser compromise [71]. Thomas et al. [145] provided an overview of the prominent cybercriminal profit centers, based on multiple individual value chains such as spam [99], and clickfraud [92]. We can further identify earlier work on the value chains behind malware [133, 154] and carding [140].

Finally, our work can be tied to studies that aim to understand how and where cybercriminals collaborate. Leukfeldt et al. [98] investigated 40 cybercriminal networks using European and American police cases and interviews, Soudijn and Zegers [140] use data from a seized carding forum to unravel the collaboration between involved actors and Hutchings [81] studied the concept of co-offending in cybercrime and more specifically knowledge transmission amongst cybercriminals and identified distinct typologies of collaboration, ranging from fluid networks to real co-offending. In most cases, they found online meeting places, such as dedicated fora and markets, as the places where to buy tools or to collaborate with co-offenders.

3.9. CONCLUSIONS

We identified key value chain components that criminal entrepreneurs might want to outsource (i.e., purchase on the market) and ordered them in ten categories. In three of them ("javascript," "customer service," and "web inject"), we found no offerings in the large random sample for the ground truth, not even when we searched the whole data with specific keywords. We assume this means there is very little, if any, commoditization of these value-chain components. In the other categories of cybercrime components, we found growing commoditization in terms of listings, vendors and revenue. Cash-out is by far the largest category. Some categories see only modest offerings and transaction volumes. Furthermore, not all offerings reflect the breadth of the demand. In some categories, only niche offerings are available.

In line with what other researchers have observed for the drugs trade on these markets, we see both B2B and B2C transactions in the cybercrime categories. B2B and B2C, a.k.a. retail cybercrime, turns out to be comparable in revenue. Between 2011 and 2017 the revenue of B2C cybercrime was around US $7 million, where B2B cybercrime generated US $8 million in revenue.

In conclusion, we find that, at least on online anonymous marketplaces, commoditi-

zation is a spottier phenomenon than was previously assumed. Within the niches where it flourishes, we do observe growth. That being said, there is no supply for many of the capabilities, systems and resources observed in well-known value chains. There is also no evidence of a rapid growth, and thus of a strong push towards commoditization, contrary to the somewhat alarmist language found in industry reporting and elsewhere.

In terms of generalizability of our findings, we have measured and explained the trends in commoditization of cybercrime on online anonymous markets. Beyond this, our findings only speculatively suggest that the trend toward commoditization might not be as comprehensive as has been claimed elsewhere. Perhaps less commoditized forms of B2B transactions - e.g., collaboration emerging out of forums - are important in the areas absent from the anonymous markets. Also, vertical integration probably remains important for more complex and dynamic forms of cybercrime.

Still, this casts an interesting perspective on the "theory of the commoditization of cybercrime." There is a huge discrepancy between the reported profitability of criminal business models like ransomware (over $1 billion in 2016, according to the FBI [62]) or DDoS-services (one youngster making $385,000 with his booter-service according to local British police [2]) and the marginal markets for cybercrime commodities. The commodities for a ransomware operation seem available in these markets: malware, PPI, cash-out. The huge profits would surely draw in new entrepreneurs to assemble this value chain based on components they can just buy on the anonymous markets. But if that would be the case, should that not cause a more observable rise in the commodities trade on these markets? The lack of strong growth suggests that there are still bottlenecks in outsourcing critical parts of criminal value chains. Entry barriers for would-be criminal entrepreneurs remain. The services that are highly commoditized, like booters, seem to draw in mostly B2C activities – i.e., consumers going after other consumers, as was the dominant finding in a victimization study of commoditized DDoS [124]. A recent takedown of a RAT operation also suggested consumer consumption, rather than B2B transactions [5].

This should not be read to downplay the relevance or danger of commoditization. A better understanding of where commoditization succeeds and fails helps to identify which capabilities, services and resources are still hard to come by, which supports designing better disruption strategies for criminal business models. The absence or scarcity of certain commoditized cybercrime components suggests these are either harder to produce or that they cannot function on their own after a single-shot sale. B2B services that require ongoing coordination among the criminals fall short of full-fledged commoditization. In other words, the scarcity of supply suggests less-scalable and potentially vulnerable

components in criminal value chains. These might be targeted by interventions. Earlier work on interventions that target choke points shows that they can have measurable impact, not via a wholesale shutdown of the business model, but by raising transaction costs [33, 86]. For instance, we found virtually no offerings for customer support services. For a ransomware scheme, the customer service component to guide inexperienced victims through the steps to complete the ransomware payment might be the most vulnerable. Contrast this approach to the series of police actions aimed at the shutdown of whole markets: from our data, these operations seemed to have had only relatively modest effects on the overall trading of commoditized cybercrime. Understanding where commoditization is lagging behind points to alternative disruption strategies.

4

CASH-OUT

Digital payment methods are increasingly used by criminals to launder money obtained through cybercrime. As many forms of cybercrime are motivated by profit, a solid cash-out strategy is required to ensure that crime proceeds end up with the criminals themselves without an incriminating money trail. These cash-out strategies are increasingly facilitated by cryptocurrencies, mainly bitcoin. Bitcoins are already relatively anonymous, but with the rise of specialised bitcoin money laundering services on the Dark Web, laundering money in the form of bitcoins becomes available to a wider audience. We examine how cybercrime proceeds can be laundered using services that are offered on the Dark Web. Focusing on service-percentages and reputation-mechanisms in underground bitcoin laundering services, this chapter presents the results of a cash-out 'experiment' in which five mixing- and five exchange services are included. We discuss what our findings mean to law enforcement, and how bitcoin laundering chains could be disrupted.

4.1. INTRODUCTION

Nowadays, cryptocurrencies - like bitcoin - are commonly used in a variety of cybercrimes. Bitcoins are used in both (1) cyber dependent crime, i.e., crimes that are dependant on computers and the Internet - such as hacking and malware, and (2) cyber enabled crime, i.e., criminal behaviours in which computers and the Internet enables committing crimes - such as drug trade on online forums [35, 109]. In both types of cybercrime, bitcoin can be seen as a facilitator of the digital criminal enterprise. The main instigators of

their popularity among cybercriminals is that they are straightforward to use, relatively anonymous, and their use is unimpeded by borders or legislation [34, 135]. Steadily, bitcoin has proven itself to be a vital part of the criminal enterprises. For instance, ransomware victims are pressed to exchange the ransom from fiat currency to bitcoin and transfer this amount to a specific bitcoin address that is provided by the criminals. On underground markets, large amounts of goods and services - like drugs, weapons and DDoS-attacks - are bought and sold using bitcoin as method of payment. In online underground markets, bitcoins are therefore to be seen as the preferred currency of criminals [113, 120, 137]. And recently, criminals start to embrace bitcoin as a partner in their cash-out strategy and launder money aided by bitcoin [118].

Criminals need a solid cash-out strategy to launder cybercrime proceeds, in this case bitcoin, without getting connected to the associated crime [100]. Regardless of the source, nature, and size of the cybercrime proceeds, the bitcoin ecosystem is utilized as part of the anonymisation or layering process a cash-out strategy entails. Already upon exchanging these proceeds for bitcoin, the money trail becomes obfuscated [59]. On the Dark Web, services are being offered to anonymize bitcoins even further, by mixing them. This way, bitcoin transactions should become untraceable, thereby facilitating the process of money laundering. Two key components in this 'bitcoin laundering' are bitcoin mixers and bitcoin exchanges [32, 42, 115, 118]. Bitcoin mixing services are services that aim to disassociate bitcoins from their often-criminal source. Bitcoin exchange services are services that aim to anonymously convert bitcoins to spendable money. In this paper we investigate both bitcoin mixing and exchange services, to assess the extent to which they facilitate the cash-out and laundering of cybercrime proceeds. Until now, there is little experimental, empirical research into the working of these bitcoins mixers, let alone explorative research into its usability for cash-out strategies. Only Möser, Böhme & Breuker [118] tested five bitcoin mixers to analyze their technical method of operation. But how do you identify and select a reputable service? And what percentage of your crime proceeds do you 'lose' in the laundering process? To answer these and similar questions, we set up an experiment that in addition to testing the mixers itself, illustrates a cash-out strategy using bitcoin mixers. The experiment will allow us to further determine the likeliness of integration of these potential criminal strategies in an actual criminal scheme. Doing so, we were able to analyse these cash-out strategies, as it would provide an opportunity to study how the different elements of such a strategy are connectable and how they operate together as a (successful) cash-out strategy.

The goal of this paper is not only to test the method of operation of bitcoin mixers, but also to make a first attempt to identify the usability of these mixers in a cash-out

strategy. Understanding the strengths and weaknesses in this strategy, will enable creating better evidence-based countermeasures for money laundering, ideally corrupting the underlying business models. The question we attempt to answer is: To what extent are those bitcoin mixing- and exchange services being offered on the DarkWeb reputable and cost-efficient to be used in a money laundering scheme?

The remainder of this paper is divided in five sections. Section 4.2 provides an overview of money laundering, related to bitcoin and underground markets. Section 4.3 considers bitcoins and bitcoin laundering as a part of a cybercrime cash-out strategy. The approach and methodology of the underlying experiment are outlined in section 4.4. The results of the experiment are presented in section 4.5. Finally, we present our conclusion and implications thereof for both law enforcement efforts as for the policies on cryptocurrencies in section 4.6.

4.2. MONEY LAUNDERING & UNDERGROUND MARKETS

Money laundering is not a new criminal phenomenon. It is a constantly changing criminal phenomenon, with updated modus operandi and evolving business models [134]. To the criminal enterprise, a decent cash-out strategy is not easy to achieve. Traditionally, the laundering of crime money is facilitated by (1) money mules, (2) offshore accounts, or (3) luxury products, i.e. art, houses, boats, or a combination of those [20, 64, 100, 101]. Next, alternative payment methods (4), such as WesternUnion or Perfect Money, allegedly have a prominent place in money laundering schemes. Prepaid credit cards, gift vouchers or other easily exchangeable non-traditional value items are also often associated with the laundering of crime money. Today, so called new-payment methods (5) are becoming a more important factor in actual money laundering schemes [57–59, 61]. Within the category of new-payments methods, cryptocurrencies stand out. A shift is apparent, in which criminals more frequently make use of cryptocurrencies in the cash-out of crime proceeds [61]. Bitcoins are also a popular form of payment between criminals. Europol [58, p. 46] even reports that bitcoin is "accounting for over 40% of all identified criminal-to-criminal payments" in cybercrime investigations.

Underground markets can be seen as a facilitator for using cryptocurrencies in current and future money laundering schemes. These underground markets are easily accessible, and are gaining popularity amongst (cyber)criminals to deploy and set-up criminal activities. Anonymous browsing has become available for the general public using the Tor-protocol (The Onion Router). The Tor system routes internet traffic to several Tor nodes, encrypting the network traffic in between [55]. Only the IP-address of the previous Tor node in a chain is visible to a connecting computer. Therefore, Tor makes it possible

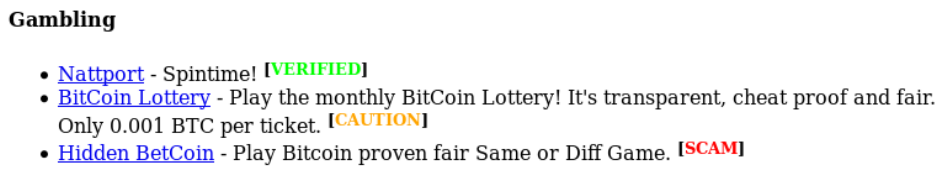

Figure 4.1: Example of gambling site reviews including accompanied labels

to use the Internet without revealing the originating IP address of the computer that is used to access the Internet by a computer user. In this way the so-called DarkWeb came into existence, which attracts growing amounts of criminals. Quantitative research on the DarkWeb indicates that more than fifty percent of all content on the Dark Web is illegal [113, p. 21].

Criminals set-up entire business models and create a whole new, online underground economy [39, 73, 75]. An underground economy has emerged that is based on buying and selling criminal techniques and services on the Internet [77, 120]. Criminals flock to the platform, even those that scam each other. “No honour amongst thieves”, as the saying goes. A review system – comparable with for instance eBay - is put in place to reduce this scamming risk [47, 75]. This review system can be seen as the most prominent element of the reputation-mechanism in underground markets, wherein services are reviewed and flagged as verified and working, a scam or to be cautious. An example of existing, accumulated reviews by customers of a certain type of service, can be found in the figure below (Figure 4.1).

Next, research indicates that these marketplaces are extensively used by criminals to buy and offer illegal goods and services [120, 137, 145]. Van Eeten and Bauer [149] already indicated that a specific underground economy of cybercrime has emerged, where individuals can buy and offer services that play a certain function within a cybercrime scheme. For example, some criminals author the malicious software that other individuals can utilize to infect computer and steal personal or financial information [17, 149]. Many of that software is subsequently offered as a ‘kit’ on the Dark Web that individuals can buy or rent [37, 71, 137]. Other individuals specialize in laundering the proceeds obtained through cybercrime by offering bitcoin mixing and underground exchange services [118]. Europol [57, p. 31] predicts that “individuals with computer skills or other skill that are valuable to criminal organisations are expected to advertise their services for payment in cryptocurrencies”.

The above supports the image that a majority of these marketplaces are involved with criminal activity and that the use of bitcoin - as method of payment and facilitator for

money laundering – is popular among cybercriminals. Our research therefore aims to provide insight in the process of money laundering by the use of the virtual currency and their involvement in cash out strategies used by criminals. Bitcoins are chosen as the virtual currency of interest, because bitcoin is presumably the preferred cryptocurrency among cybercriminals [58, 118]. Yet, little is known about cash-out strategies using bitcoins to launder crime proceeds that are obtained through cybercrime.

4.3. BITCOIN MONEY LAUNDERING

We now shift to the characteristics of the bitcoin ecosystem, to make clear why and how bitcoin facilitates the laundering of cybercrime proceeds. Blockchain is the fundament of bitcoin. The bitcoin blockchain operates as a decentralized bank for the bitcoin cryptocurrency [121]. This means banks are no longer necessary in interpersonal transactions with bitcoins. As a result, bitcoin transactions are made between bitcoin addresses directly [50]. In this sense, the blockchain can be seen as a publicly visible and verifiable ledger. All transactions are logged in the blockchain and can be inspected via public websites, such as blockchain.info and other open source websites. Anyone, anywhere, can see all bitcoin transactions from one bitcoin address to another in real time. The current balance of the amount of bitcoins in a bitcoin address is also visible in the blockchain. Despite its openness, the bitcoin system does provide a high level of anonymity. The reason for this anonymity is that bitcoin addresses are not registered to individuals, in contrast to bank accounts. Comparable with numbered Swiss banking accounts, the bitcoin address itself acts as a unique identifier and the account is only accessible by the owner who has the login details to the bitcoin wallet. Yet, no names are connected to the bitcoin address and wallet. In addition to its high degree of anonymity, bitcoin relies on the instant creating of new bitcoin addresses. This in sharp contrast to bank accounts, which take time to set-up, next to the obligatory registration of personal information to name the account. This level of anonymity explains why bitcoin has become so popular in illegal activities. The total system however – from a criminal perspective - has one 'downside'. Due to the blockchain concept, all historic information on any bitcoin address and transactional information is just one mouse-click away for law enforcement authorities [147]. Figure 4.2 illustrates the workings of the blockchain model for bitcoin transactions.

A bitcoin transaction constitutes of several elements: it has one or more inputs, one or more outputs and holds the cryptographic protection of this information. Each input and output consists of a bitcoin address and a bitcoin amount. As bitcoin transactions are linked to each other, each input is automatically an output of a previous transaction. This way, the value of a bitcoin output can be traced back to previous transactions.

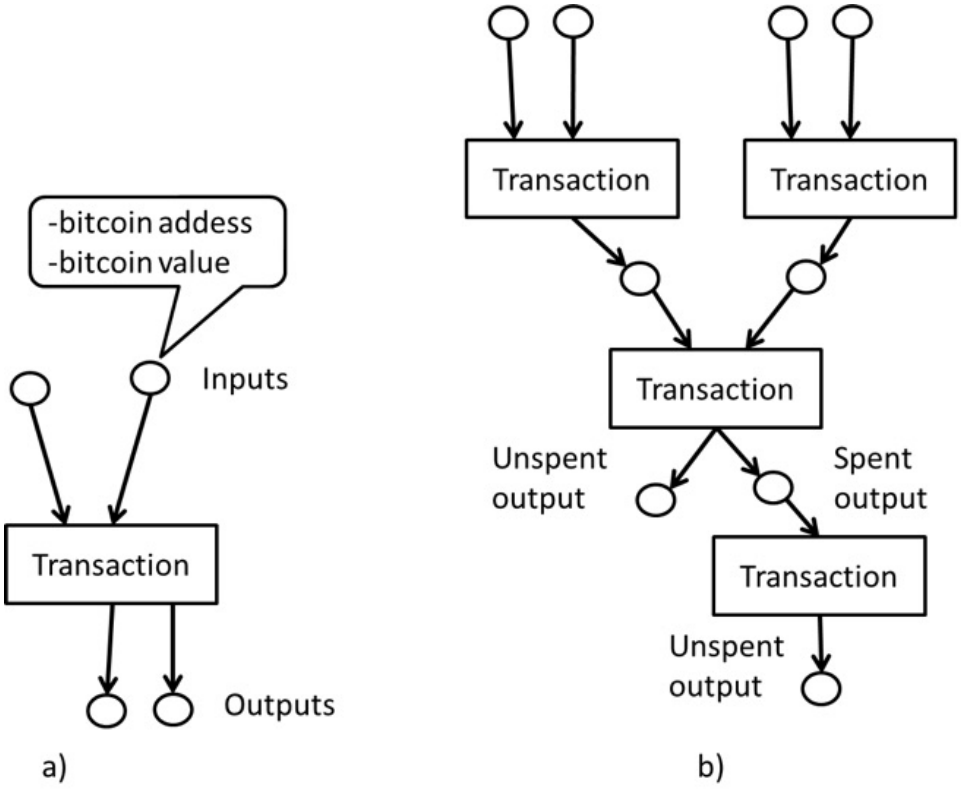

Figure 4.2: A bitcoin transaction has inputs and outputs (a); bitcoin transactions are linked to each other (b)

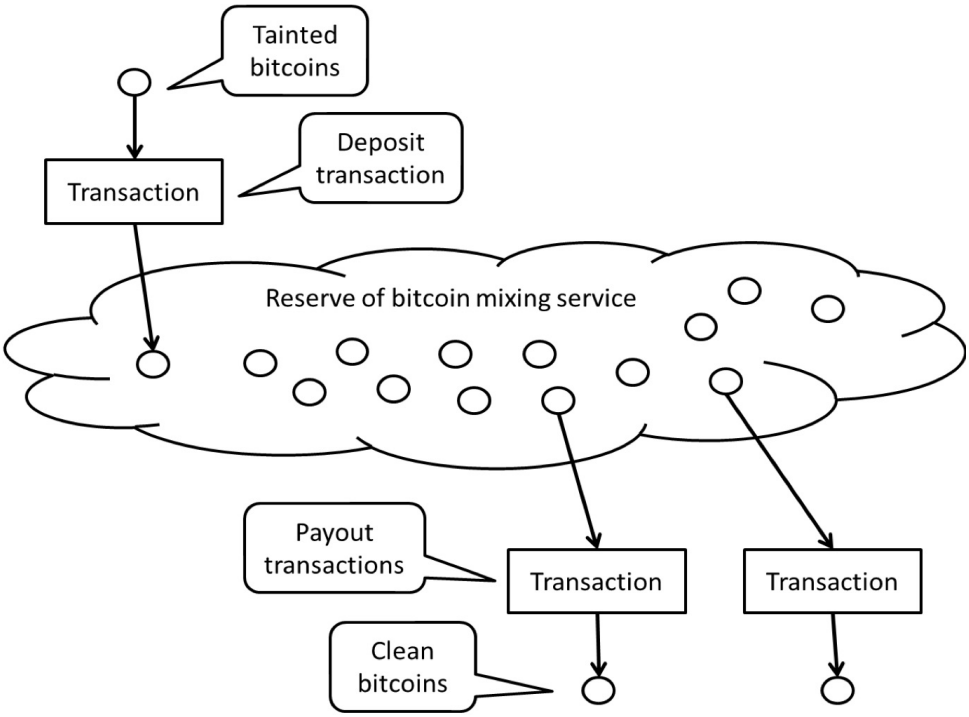

Figure 4.3: A bitcoin mixing service were 'tainted' bitcoins are deposited in the reserve of the mixer, and subsequently 'clean' bitcoins are paid out to the customer

Obviously, this poses a potential risk - from a criminal perspective - as the transaction that is used to cash-out cybercrime proceeds, links back to transaction(s) that are associated with illegal activity. For instance, transactions can potentially be linked to the receipt of ransom, the selling of illegal goods or other cybercrimes. This risk is the instigator of technologies designed to break the transnational link between the bitcoin transaction and illegal activity. For that reason, bitcoin mixing services are created that aim to break the transactional link – or in this case the money trail - of bitcoins [1]. Figure 4.3 illustrates the working of a typical bitcoin mixer.

The typical mode of operation is that bitcoin mixing services provide customers with a newly-generated bitcoin address to make a deposit. The bitcoin mixing service pays out other bitcoins from its reserve to bitcoin addresses provided by the customer, after deducting a mixing fee. In order to provide more anonymity, the pay-outs are spread out over time and some randomness is introduced in the division of amounts and/or the mixing fee.

[1] Bitcoin mixing services are also called -laundering, -tumbling', or –cleaning services

An operational bitcoin mixer makes it virtually impossible to trace back mixed bitcoin to their tainted source. The customer can check the taint of the received bitcoins at the blockchain - e.g., at `blockchain.info`. If the bitcoin mixing is performed correctly, there is no link ("zero percent taint") between the deposited bitcoins and the received bitcoins. Some bitcoin mixing services offer a service to returning customers to ensure that (earlier) deposited tainted bitcoins in their reserve are not accidentally paid out to the same customer in a subsequent use of the mixing service. After each 'mix' the customer is issued a returning customer number. This number can be presented when reusing the mixing service. The mixer then knows which bitcoins in the reserve were earlier deposited and will not pay out these same bitcoins to the client. This service-model is suitable for returning customers, or, said differently: frequent launderers. In this manner a sophisticated strategy can be set-up, using bitcoin as a facilitator in laundering cybercrime proceeds. When individuals successfully use a bitcoin mixer and subsequently an underground bitcoin exchange, only mistakes will leave sparse traces to your true identity or the fruits of your crime.

In summary, we have shown that the cryptocurrency bitcoin is used in many forms of cybercrimes. We have also explained the rise in popularity of bitcoins amongst cyber-criminals. A potentially toxic combination – of anonymous browsing, trading, paying and laundering - becomes apparent. To determine the extent to which bitcoin laundering is feasible and – looking at reputation-mechanisms and service-percentage – actually usable in the criminal enterprise, we set up a bitcoin laundering experiment.

4.4. APPROACH

In this section, we first lay down the approach to our experiment and used to get an overview of both the available mixing and (underground) exchange services. We outline their individual characteristics and the selection criteria for the underground services to be included in the experiment. Thereafter, we elaborate on the set-up of the actual experiment, describing our methodology step-by-step.

4.4.1. SET-UP

To set up the underlying experiment for this paper, we needed to address three conditional methodological questions. First, how do we get an overview of both underground bitcoin mixing and exchange services and its reputation and service-percentage. Second, how do we select bitcoin mixing and exchange services for this experiment? And third and last, how do we determine the effectiveness of these services?

OVERVIEW OF AVAILABLE BITCOIN MIXING AND SHADOW EXCHANGE SERVICES

In order to gain insight into the underground economy of bitcoin laundering services we made use of the TNO Dark Web Monitor [2]. The TNO Dark Web Monitor makes use of a 'crawling' technique to collect and analyse data on the Dark Web. This system therefore provides a solid basis for both exploratory and longitudinal research, as the data is collected over a longer period of time and is independent of the hidden services that are still online. Using this technique, we have discovered over 25.000 hidden services, i.e. Dark Web-sites. By use of the TNO Dark Web Monitor we obtained a solid overview of the total supply of Dark Web services offering bitcoin mixing and exchanges from bitcoin to other non-virtual currency via a diversity of anonymous output platforms [3]. This overview enabled us to see a number of notable differences in the offered services.

First, the analysis revealed that bitcoin mixers differentiate in (1) service percentage, (2) registration and authentication process, (3) reviews and (4) time delay. The service percentage relates to the percentage the service takes for mixing bitcoins. The registration and authentication process relates to how you register for the process and whether there is any form of authentication involved. The reviews relate to how the service is reviewed by other users and the time delay regards the question how much time it takes to receive mixed bitcoins.

Second, Bitcoin exchange services also differ in (1) service percentage, (2) registration and authentication process, (3) reviews and (4) time delay. Importantly, exchange services also show a variation in output platform. In other words, they offer, for instance, PayPal or PerfectMoney, to allow a client to anonymously receive the exchanged bitcoins. Keep in mind that these payment services themselves often do not offer this type of service as they are regulated entities who uphold KYC and AML policies. Using a bitcoin mixer and subsequently an underground bitcoin exchange service, make up together a cash-out strategy for cybercriminals (see Figure 4.4). The aim of the cash-out strategy is to provide a spendable proceeds of the crime that cannot be traced back to its origin.

SELECTION OF SERVICES

Following the first question, we addressed the second question: how do we select the services to be included in our experiment. Starting from the paper of Möser (2013) and desk research into known mixers and exchange services, we added the currently available mixers and underground exchanges and generated a list of still active services. Due to

[2] TNO DarkWeb MoniTOR is an interactive tool designed for indexing and visualizing crawled DarkWeb data. See `https://dwm.pm`.

[3] Here, we differentiate between mainstream, regulated bitcoin exchanges - that have implemented KYC and AML policies - and shadow exchanges. The latter - like bulletproof hosting - purposefully failing to adhere to these regulations [16]

Figure 4.4: Mixing services, combined with underground bitcoin exchanges, from the essential parts of the cash-out strategy

budget restraints, we settled on selecting five mixing services and five exchange services. [4]. The five services were selected on the basis of the three criteria: (1) function and fee, (2) positive/negative reviews and (3) (laundering/mixing) reputation in general. [5] Note that we indeed purposively included negatively reviewed services, e.g. potential scams, into our sample. [6] The selected mixing services can be categorized as follows (see Table 4.1).

Mixer	**Review**	**Advertised fee**
Mixer 1	-	0.5 – 3.5 %
Mixer 2	CAUTION	2 %
Mixer 3	SCAM	0.1 %
Mixer 4	VERIFIED	1 – 3 %
Mixer 5	VERIFIED	1 – 2.5 %

Table 4.1: Overview of included mixing services in the experiment

By use of this sample, we believe we have selected these services that best reflect the total population of services on the Dark Web. The sample covers a variation in reviews, namely from scams to verified services. In addition, this sample provides a differentiation in the charged fee, namely from nearly no fee up to more 'expensive' services. For the underground exchange services however, no such well-documented and reviewed overview of services was present. Therefore, we selected the underground exchanges that we came across in our search for bitcoin mixing services and underlying desk research. Now that we have answered the first of two conditional questions, we can

[4] Allowing us to use 0.5 or 1 BTC per service

[5] In this paper we have chosen not to disclose the names nor the onion-addresses of both the bitcoin mixers and the exchange services we used. As we are aware of the fact that these services can be (easily) found, adding more context information on usability and working cash-out schemes would have turned this paper into a user manual for bitcoin laundering, which is not our intention. Please do not hesitate however to contact us if you do have a scientific interest in the data used in this paper.

[6] The third criteria of 'reputation in general' was determined by the analysis of open sources about bitcoin laundering services

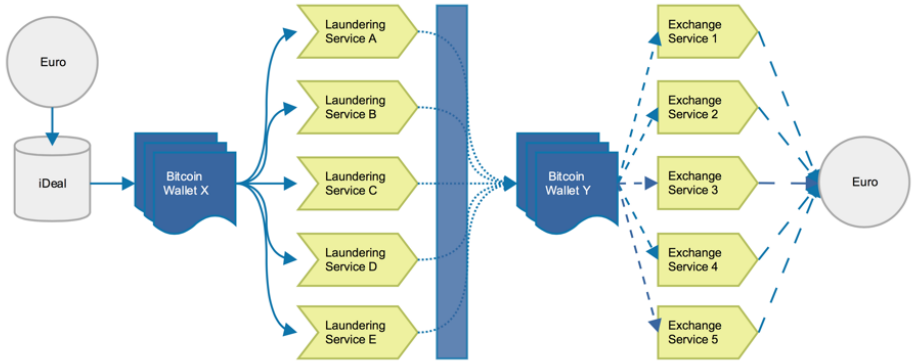

Figure 4.5: Conceptual model of the experiment

create the following model (Figure 4.5) depicting the process of the experiment. [7]

4.4.2. TESTING THE EFFECTIVENESS OF THE SERVICES

Third and finally, we had to decide how to determine the effectiveness of the tested services. For the mixing services we stick with the 'product description' or 'product promise' of many of these services: no taint. In that case, no taint refers to the taint a transaction has between bitcoin addresses in the blockchain. As the blockchain is easily accessible and a taint analysis is just one mouse-click away, we assume that this is a solid identifier for success and is also used or usable by cybercriminals who want to assess the service on its operational excellence. The first indicator for 'success' is the actual transferring of the bitcoins. Being scammed would count as unsuccessful. The underground exchange services are only effective when the bitcoins are (1) exchanged to fiat currency and (2) anonymously transferred to the output platform specified. In that case both the used service as the used output platform (such as WesternUnion and PayPal) needs to be assessed. Again, success counts here as not being scammed.

4.4.3. EXPERIMENT

We operated the experiment on one day, to minimize both the loss of value in bitcoin due to fluctuating currency exchange rates and moreover to test the speed and user-friendliness of the total cash-out strategy. First, we bought the necessary bitcoin via the Dutch payment service iDeal and started the experiment with these bitcoins stored

[7] As we did not have the opportunity of actually starting our money laundering scheme with real crime proceeds – neither in bitcoin nor euros/dollars – we opted to use the research budget to acquire the bitcoins to use in this experiment via a mainstream bitcoin exchange – using the Dutch Payment service iDeal.

in Wallet X on one bitcoin address. Hence, this can be seen as our 'control' address – where all the bitcoins used in the experiment originate from. After this first step, all the following steps are undertaken via the Tor-network to guarantee the most anonymous process. Second, we set-up a Lelantos email-account – which is a form of Tor-mail - for possible future communication with certain services. Third, we created Wallet Y, wherein we generate one or more new, clean bitcoin address per mixing service depending on the variety of services offered by the mixer. Fourth, we used all of the five selected mixing services one by one. Nearly all following the procedure of a) transferring bitcoin to a new – by the mixer provided – bitcoin address from our Wallet X and b) requesting the mixed coins to be transferred to one or multiple bitcoin addresses in our Wallet Y. Fifth, after we confirmed the actual transferring of the (mixed) coins, we analysed the taint of the incoming transaction by cross-checking the traceability to our 'control address'. Sixth, we consecutively used the five exchange services.

In contrary to the mixing services, exchanges services are based on both a service and an output platform. The service refers to the supplier that offers to receive bitcoin and will exchange this to a currency of your choosing. An output platform, such as PayPal or WesternUnion, is used to make sure this exchanged currency ends up in your possession. [8] Following a procedure comparable to the mixing services of a) transferring bitcoin to a new – by the exchange service provided – bitcoin address from our Wallet Y b) selecting the output platform and c) providing details for this platform.

All these steps we could fit in the time-span of one day. The actual cash-out however takes 1 to 3 days, depending upon the output platform used. However, we were able to 'consume' one output platform right away. Therefore we used, without registering an account, the Dutch online food ordering service Thuisbezorgd.nl to order Sushi and paid – of course – with mixed bitcoins. A few days later, we were able to assess the success of the other exchange services and thereby the entire cash-out. We cashed out our mixed bitcoins using the following five service platforms: (1) PayPal, (2) PerfectMoney, (3) WesternUnion and (4) Bitonic [9]

4.5. (MIXED) RESULTS

We present the results of the experiment in two ways. First, we show the results in terms of bitcoin transferring by use of bitcoin mixing services. Second, we show the results of

[8] Note that the platforms used, i.e., PayPal, are not offering these types of services themselves and should be seen as (relative) innocent bystander. However, it is known that WesternUnion is a prominent element in money laundering typologies, as might be expected form its domination of the legitimate financial transfer market globally, see FATF [60].

[9] Bitonic is a Dutch bitcoin exchange. See `http://www.bitonic.nl`

our exchange services. Based on these results, we describe our overarching conclusions based on the experiment and the overall cash-out strategy.

4.5.1. MIXING SERVICES

In this section we show the results of our bitcoins transfers using five bitcoin mixing services. In Table 4.2 we have plotted the results of the experiment for these five specific mixing services.

Mixing service	Review	Advertised fee	BTC IN	BTC OUT	Percentage	Blockchain taint
Mixer 1	-	0.5 - 3.5 %	1.0	0	–	-
Mixer 2	CAUTION	2 %	1.0	0	-	-
Mixer 3	SCAM	0.1 %	0.5	0	-	-
Mixer 4	VERIFIED	1 - 3 %	0.5	0.4931291	-1.37%	0%
Mixer 5	VERIFIED	1 - 2.5 %	0.5	0.497509	-0.5%	0%

Table 4.2: Output and taint of mixing services

Noticeably, we fell for a scam three out of five times. Resulting in an immediate loss of 2.5 BTC. In all of these three cases the bitcoins were successfully transferred from our 'control address' to the bitcoin address the mixing services provided. Yet, we have not received any of the bitcoins we provided to the Mixers 1, 2 and 3.

Next, we used the Numisight [10] blockchain explorer to see what further information we could extract on all the mixing services and what happened to the bitcoins that went 'missing'. We discovered that the deposit address for Mixer 1 had already been used in a transaction three days before our experiment. Reuse of bitcoin addresses should be considered a deadly sin in the world of bitcoin mixing, as it provides correlations that could be used in forensic analysis. [11] Bear in mind that exactly these types of correlations are the main reason of using a bitcoin mixer to begin with. In curious contrast to this sole purpose, this mixer made a vital mistake, thereby annulling our efforts to prevent just these types of correlations. The same deposit address of Mixer 1 was used three more times during the next two weeks. We discovered a similar pattern of bitcoin address reuse by Mixer 2. When looking more closely into the outputs of Mixer 1 and Mixer 2, we discovered that they were combined in a transaction 12 days after our experiment. All of this strongly suggests that Mixer 1 and Mixer 2 are closely collaborating, and that they are very likely to be the same entity. This result, that Mixer 1 and Mixer 2 are closely

[10] See `http://numisight.com`

[11] See the report by UNODC [147] for an overview of forensic techniques used in bitcoin money laundering investigations.

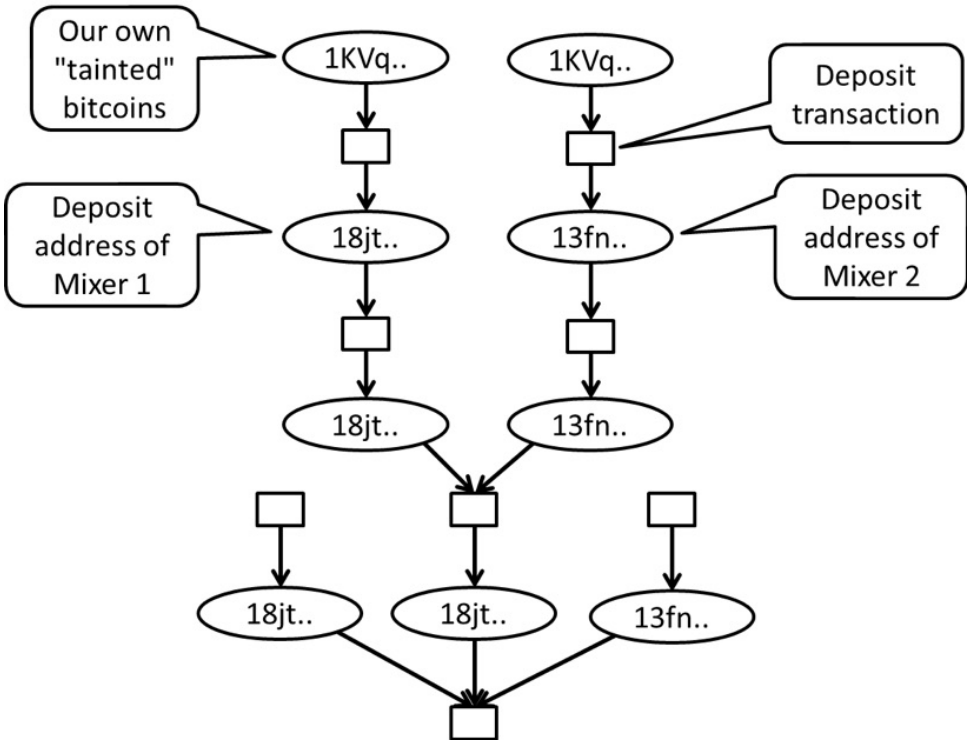

Figure 4.6: Blockchain analysis – using Numisight – strongly suggests that Mixer 1 and Mixer 2 are controlled by the same entity

collaborating and made use of the same bitcoin address at one point, is visualised in Figure 4.6.

To have a complete picture, we performed the same analysis for Mixers 3, 4 and 5. Mixers 3 and 4 each used their deposit address exactly once. Mixer 5 required the payment of an entry fee at the bitcoin address that was also used as deposit address, but there was no other bitcoin address reuse. Another relevant result of the experiment regards to the question how long a deposit remains unspent. This is an indication for how much time worth of reserve a bitcoin mixer has. Or in other words, how large the reserve of a bitcoin mixer is in terms of the amount of bitcoins available for future mixing. When having a large reserve, a bitcoin mixer is not forced to reuse its deposits quickly as output in another 'mix'. For Mixers 1, 2 and 3, the wait time was four days (just after the weekend, as we operated the experiment on a Friday). For Mixer 4, the wait time was 17 hours. For Mixer 5, the wait time was 3 hours.

Furthermore, two out of the five tested services, i.e., all of the successful mixing

attempts, operated conform to their own description. We received the same amount of bitcoins that we transferred to the service minus the percentage for the use of this service. Of both of these services, no taint is observable between the bitcoins we received from the mixers on the addresses we specified and our bitcoin 'control address'. [12] This means that there is – in retrospect – no way of linking these two addresses to each other, i.e., the money trail is obfuscated and in all probability impossible to follow. So despite the fact that we actually transferred bitcoins – via a mixer – from Wallet X to Wallet Y no trail can be found that reconstructs that we did precisely that.

Our results show that some of the examined services provide an excellent, professional and well-reviewed service at competitive cost. Others turned out to be straight scams. As the scams and working services were also reviewed that way, our results suggest that reviews help criminals select reputable and workable service. In reverse, reviews mitigate the risk of getting scammed. Looking at all services, we observe that service-percentages are as advertised. This means, bitcoin mixers have the potential of being a cost-efficient method to dissociate (cyber)crime proceeds from the illegal activity it was obtained with.

4.5.2. EXCHANGE SERVICES

In this section, we show the results of the use of five exchange services we used to cash-out our mixed bitcoins. In Table 4.3 we have plotted the results of the experiment for these five individual services and the platform used by this service.

Exchange service[a]	BTC IN	BTC OUT[b]	Percentage	Anonymous?
1/PayPal	0.233933	$ 52.42	-6%	+
2/PerfectMoney	0.1	$ 0	-6.9%	?
3/WesternUnion	1.220364	$ 197	-10% + $50	+/-
4/Bitonic	0.33740872	$ 71	-0.25%	-

[a] Notwithstanding the fact that we choose not to name the used mixing services by name, a solid overview of the results would be impossible without naming the output platforms as they have very distinct procedures and regulations that make the analysis more in-depth and understandable.
[b] Based on currency rates on 29 to 31 May 2015.

Table 4.3: Overview of output and anonymity of the selected exchange services

At first glance, it is clear that all but one of the exchange services tested was operational and successful in its promise to exchange our mixed coins for currency or food. Only the service-platform combination that makes use of PerfectMoney did not come

[12] Using the blockchain.info web based 'taint analysis' comparing the receiving bitcoin addresses in our Wallet Y with our 'control address' in Wallet X.

through. This could be due to the fact that the service is a scam or that the platform itself – PerfectMoney - has blocked the transaction. [13]

Next, we used PayPal and Western Union as a cash-out strategy. The platform that facilitated this transaction charged a high percentage for the transfer. In exchange, we expected a better anonymization of this part of the cash-out strategy. This turned out to be the case. In order to use PayPal as a cash-out strategy we needed to provide a valid and active PayPal account, which can be bought online on underground markets or generated without leaving anything more than an (TOR) e-mail address [14]. As promised, we received the exchanged funds on the PayPal account we provided. The service-platform combination that used WesternUnion asked us to provide a name and place for the collection of the funds. After we provided this information, we could pick-up the exchanged funds at a WesternUnion office of our choosing without leaving a signature. When comparing both cash-out strategies, it is important to note that the method via PayPal does not include a physical hand-over of the exchanged funds, where this is the case with the WesterUnion method. This is why we rated the PayPal service-platform combination slightly more anonymous than its WesterUnion alternative. From a 'criminal perspective', we did make a "mistake" in our cash-out. Whereas there was no link between our "tainted" bitcoins and "laundered" bitcoins at the day of the experiment, we cleaned up the experiment and retrieved the invested bitcoins four days later. In this clean-up, remaining "tainted" bitcoins and "laundered" bitcoins were sent to our bitcoin address at a regular bitcoin exchange. This 'mistake' linked our identity directly back to the "tainted" bitcoins, negating all our previous efforts to stay anonymous. Of course, this step was intentional – as the experiment was completed and remaining bitcoins were to be collected – it does show how easily it is to make a mistake and render the attempted launderer vulnerable.

4.5.3. OVERARCHING RESULTS

In this section, we discuss the overarching results of our experiment. Our first conclusion is that bitcoin laundering services offered on the dark web are partly scams and partly operational services. Reviews are of the utmost importance to reduce the chance of being scammed. The trustworthiness of this review system is solidified in our experiment,

[13] We described earlier our use of an exchange to buy the bitcoins used in this experiment. This regular exchange we also used to exchange our mixed coins back to euros. This specific service required us fill out a form wherein we had to specify a bank account to which the funds upon exchange should be transferred to. This in turn logically lowered the level of anonymity, as a bank account of course leaves a trace to an actual person or company. But on the other hand this service did not charge us with a high commission for the use of this exchange service.

[14] However, PayPal is getting more pro-active in preventing fraudulent transactions and stopping funds moving through unverified accounts [83]

because the services that were flagged 'orange' or 'red', indicating to be cautious or be warned not to use these services, actually corresponded with our three mixing service scams. [15] In addition, the two operational mixing services were ranked 'green' in reviews. This means that if you read and use reviews of these specific services carefully, the risk of getting scammed is minimized. It would be wise for criminals to use these review systems, since it helps them to keep costs of scams low. The attractiveness of the examined type of case-out strategy is clear, since these services work as advertised and were straightforward to use in an anonymous manner - since no taint was present in the blockchain.

The exchange services show us a different pattern. All exchange services operated as they promised, except one. However, the level of anonymity differs a lot from service to service and from platform to platform. The exchange Bitonic connects the mixed coins directly to a bank account, as this is required to use a regular bitcoin exchange. Of course, it is not obligated and it would not be wise to provide your own bank account for an exchange. A name and bank account number provides an interesting lead for law enforcement authorities. Using output platforms which are often integrated in other cash-out and more physically oriented laundering strategies, like PayPal and WesternUnion, build on an already established reputation in that area. The examined exchange services allow their clients to receive money in any shape and form with minimal registration requirement that may lead back to the identity of individuals.

Finally, it is important to point out that mistakes are made easily in the examined cash-out strategy. For instance, (1) not using the Tor browser, (2) using the same bitcoin address for different purposes or (3) providing an email address are detrimental to an anonymous transaction, are possible mistakes. This is however no different from other cash-out strategies, where mistakes also leave (extra) traces for law enforcement authorities.

In our view, the most significant difference compared to other money laundering strategies is the total percentage that covers the costs of this particular cash-out strategy. In many other cases these percentages add up to nearly 50 percent. In this case, looking at the maximum costs of all services and adjusting for potential value drops regarding exchange rates, we can safely estimate the total percentage of costs will not exceed 15 percent in this specific cash-out strategy. Again, we thereby assess that the strategy itself pre-includes trial-and-error and users reviews to find and use operational mixing and exchange services. Therefore, we do not contemplate the risk of getting scammed in our experiment as one to be very cost-increasing in these strategies as the full scale operation

[15] Figure 1 shows an example of accumulated reviews on a single platform, in this case for gambling services. We have chosen not to show the actual accumulated reviews of the services we used - to demonstrate the importance and accuracy of reviews - as this would neglect all previous efforts not do disclose the actual names of the tested services.

of the strategy will be build on positively-reviewed operational services. Therefore, we conclude that 15 percent is a good estimate of the total cost of this type of cash-out strategy based on the underlying experiment.

Keep in mind however, that intermediary cash out services may block transactions that exceed a certain threshold, for instance 5,000 or 10,000 dollars. Therefore, it is not certain this cash-out method works when larger amounts of money are laundered. Yet, potentially, the low commission of 15 percent for laundering services may make it interesting for cybercriminals to use bitcoins, mixing services and exchanges for money laundering purposes. As a result, this would significantly change the profitability of money laundering models that are now often used by criminals. In conclusion, these operational bitcoin laundering services provide a hassle-free and mostly anonymized exchange of mixed coins in an easy and consumer-friendly manner.

4.6. DISCUSSION

Bitcoin has reportedly established itself as 'a single common currency for cybercriminals within the EU'. The virtual currency potentially provides the means to launder the obtained proceeds without the strict requirements of (international) financial institutions. This paper aimed to examine in which ways cybercrime proceeds that are obtained in the form of bitcoin or converted to bitcoin can be laundered. Following up on previous studies, we focused in particular on money laundering services – both mixing as exchanging - offered for bitcoin on the Dark Web.

In order to examine bitcoin laundering, we used bitcoin mixing services and exchange services and integrated these services in a working cash-out strategy. We examined the usability of these services in a cash-out strategy by analysing the service-percentages and reputation-mechanisms. This aided us in determining the likeliness of integration of these bitcoin laundering services in an actual criminal scheme. We believe that in our experiment we have shown that laundering cybercrime proceeds using bitcoin is a user-friendly and working criminal service-model. However, it is not clear whether the model will work when larger amounts of money are laundered. Yet, for smaller amounts it clearly offers an easy to use and good value-for-money service, as long as criminals keep an eye out for scams.

We conclude that bitcoin money laundering is a practically conceivable concept and has a high degree of likeness to be integrated in current-day and future money laundering schemes. Recent cases and Europol reports support the conclusion that bitcoins are used for money laundering by cybercriminals. Most notably, the ability to lower the cost of laundering, whilst providing more anonymity, make it an interesting money laundering

technique for criminals.

This brings us to the following questions. First, the question arises what does this mean for the profitability of criminal business models. Obviously, we have a limited overview of criminals actually using this very method. However, from this small-scale experiment, we do know that at least in theory the cost of money laundering can be lowered using bitcoin. That could make certain criminal business models potentially more profitable, which is of course attractive for the cybercriminal enterprise.

Second, the question arises how law enforcement should deal with this new money laundering technique. The start- and endpoint of bitcoin money laundering (often) includes the exchange in currencies. Law enforcement authorities may be able to gather evidence at these exchanges. Therein lies the possibility of intervening as current police measures provide sufficient means to do so, but future research is necessary to study this possibility to its full extent. In addition, law enforcement authorities may be able to seize and analyse bitcoin wallets of identified criminals to identify bitcoin addresses in order to trace back bitcoin transfers.

Finally, the question remains on how bitcoin should be treated from a legal perspective. In many countries, bitcoins are not banned, nor regulated. Because governments do not recognize bitcoin as a currency, oftentimes anti-money rules and regulations – such as 'know your customer' (KYC) rules and the reporting of suspicious transactions – do not apply for companies and institutions trading in bitcoin. Ironically, mainstream bitcoin exchanges, are lobbying to become regulated. [16] Recently, the Financial Action Taskforce (FATF) has pushed for the regulation of bitcoin exchange services [61] - which resulted in updated European and American KYC and AML policies, making cryptocurrencies exchanges to be regulated like their traditional counterparts. However, this will likely not stop criminals in using shadow bitcoin exchange services or mainstream exchanges that are located in jurisdictions that have no or less strict regulations. The degree to which banning or regulating will have any effect on the facilitating role bitcoin currently plays in the criminal enterprise, has yet to be determined.

[16] See http://www.cnbc.com/2015/09/18/bitcoin-now-classed-as-a-commodity-in-the-us.html and http://www.cnbc.com/2015/02/25/bitcoin-futures-market-just-changed-the-game-commentary.html

5

OUTSOURCING

Many cybercriminal entrepreneurs lack the skills and techniques to provision certain parts of their business model, leading them to outsource these parts to specialized criminal vendors. Online anonymous markets, from Silk Road to AlphaBay, have been used to search for these products and contract with their criminal vendors. While one listing of a product generates high sales numbers, another identical listing fails to sell. In this chapter, we investigate which factors determine the performance of cybercrime products. Does success depend on the characteristics of the product or of the vendor? Or neither?

To answer this question, we analyze scraped data on the business-to-business cybercrime segments of AlphaBay (2015-2017), consisting of 7,543 listings from 1,339 vendors, sold at least 126,934 times. We construct new variables to capture product differentiators and price. We capture the influence of vendor characteristics by identifying five distinct vendor profiles based on latent profile analysis of six properties. We leverage these product and vendor characteristics to empirically predict the performance of cybercrime products, whilst controlling for the lifespan and type of solution.

5.1. INTRODUCTION

Often, cybercriminals are described as freelancers. Specialized in specific tasks, like malware development or cash-out solutions, they trade self-made products and services to other cybercriminals [15, 37, 154]. These vendors sell their services on forums and

platforms in the underground economy. Online anonymous markets have been found to foster a segment for business-to-business (B2B) cybercrime products or services [152].

Within the B2B cybercrime segments on online anonymous markets, there are significant differences in the types of product offered, sales volume and vendor performance [126, 137, 152]. A small portion of vendors and offerings is responsible for the majority of the revenue. These differences can partially be explained by the heterogeneity in cybercrime solutions. Yet, these dissimilarities remain observable within each product category – e.g., stolen credit card details. Even offerings of a rather specific instance of that product – e.g., credit cards from Canada – show differences in popularity. What drives these differences in sales? Do certain product characteristics determine sales numbers? Or are buyers more focused on vendors and do their characteristics drive the performance of B2B listings?

Law enforcement agencies could greatly benefit from insights into the performance of cybercrime sales, related to both products and vendors characteristics. The understanding of how criminals select reputable and trustworthy partners in crime, sheds light on the economic incentives in criminal B2B trades [47]. This understanding can be used in efforts to disrupt these distribution channels. We build on recent work into interactions and performance on carding forums and extend this interdisciplinary research to study the performance of cybercrime solutions on online anonymous markets [49, 72, 76].

In this paper, we explain the performance of B2B cybercrime listings on AlphaBay (2015-2017) from the associated product and vendor characteristics. Put differently, how do certain products – even in the same category – sell much better than others. We focus on B2B cybercrime sales on online anonymous markets for a number of reasons. First, vendors have incentives to provide their offerings on these online markets, as these platforms provide risk management services for criminals, i.e., reputation systems to protect vendors from treacherous interactions with buyers. Second, the platform lowers entry barriers for cybercriminal entrepeneurs in search for products and service – increasing the potential customer-base of vendors. Third, these markets have the advantage of making relevant aspects of the trade visible. We can observe important interactions in a standardized way. In contrast, a study of underground forums, another location for B2B cybercrime transactions, would only show part of the interaction, as vendor and buyer typically move to private channels to get the deal done.

To study B2B cybercrime sales on online anonymous markets, we adopt an approach that models three constructs – grasping the relative price, product differentiators and distinct vendor profiles – to the sales level of an offering and controls for the lifespan of the offering and the type of product. We make the following contributions:

- We present the first comprehensive study into the performance of B2B cybercrime solutions on online anonymous markets, using measurements from AlphaBay (2015-2017), comprising of 126,934 feedbacks, 7,543 listings and 1,339 vendors related to B2B cybercrime offerings.
- We statistically estimate the influence of product and vendor characteristics and show that these factors can predict up to 47% of the variance in cybercrime sales.
- We develop five vendor profiles that all significantly influence cybercrime sales. Compared to the average 'freelancer', being a 'professional' criminal vendor more than doubles the performance of a cybercrime solution.
- We show that product characteristics correlate significantly with cybercrime sales. Customer support options and refund policies lead to an increase of 43% and 53% in sales, respectively. Branding the product using a vendor's name, nearly doubles the performance of cybercrime solutions.

The rest of this paper is structured as follows. Section 5.2 discusses the structure of and product differentiators in the market for B2B cybercrime solutions. Section 5.3 explains our methodology and presents our approach. Section 5.4 grasps the different product characteristics in newly constructed variables. Section 5.5 lays down our approach to cluster vendors into distinct profiles, and section 5.6 identifies predictors for B2B cybercrime sales. Section 5.7 discusses our findings both in terms of its limitations as well as our public policy take-aways. Section 5.8 connects our work to earlier contributions and section 5.9 concludes.

5.2. ANONYMOUS CYBERCRIME MARKETS

In this section we show how an online anonymous market operates, how sales take place on these markets and how we can observe essential steps in the trading process. That way we can investigate the performance of cybercrime solutions on the market.

5.2.1. B2B CYBERCRIME PRODUCTS

Online anonymous markets – starting out as predominantly drugs oriented markets in 2011 – have become a prominent part of today's cybercrime ecosystem. Their popularity and supply in digital goods, both in quantity as in diversity, has steadily grown over the years [139, 152]. The markets have also matured in business continuity management and in revenue. A single top tier market can turn over more than 200,000 US dollars daily [139]. Apart from drugs, products and services range from physical goods, like

passports, to digital goods, like stolen credit cards or malware packages [139, 152]. Next to retail transactions, aimed at end-users, e.g., drugs in small quantities or a handful of compromised Netflix-accounts, we see a steady portion of the market aimed at wholesale transactions, e.g., drugs sales in bulk or large databases of compromised email-accounts. These two distinctive types of transactions show that criminals also use online anonymous markets as a platform for criminal-to-criminal transactions [14, 152].

Online anonymous markets provide structured data on criminal trading in the underground economy. All listings, from offering stolen creditcards to compromised RDP-hosts, are forced to contain the same information, including a title, description, vendor name and customer feedback on the listing. Earlier work has focused on measuring the volume and nature of trade on these markets in general and in cybercrime solutions in particular. Yet, we do not have insight into why and how criminal B2B customers prefer one specific solution over another. Knowing what sells and which vendor is successful, can help focus police interventions to disrupt cybercrime B2B transactions.

5.2.2. PRODUCT DIFFERENTIATION

In economics, product differentiation is the activity of distinguishing a product or service from its competitors in order to increase its attractiveness. Differentiating characteristics may vary, but generally are: functional features, advertisement, and availability [54]. Here, we apply product differentiation to help us derive product characteristics as potential predictors for the performance of cybercrime solutions. In absence of any market data on the availability of the product, we focus on functional features and marketing-like activities as differentiators.

First, we can identify functional features of B2B cybercrime products. For instance, what terms and conditions are associated with the product? This sounds a bit intriguing, but as a 'consumer-centred' market, online anonymous markets incentivize to be clear about specific terms and conditions of acquiring and/or using the product. Vendors signal the availability of both a refund policy and customer support options and if there are any other terms and conditions associated with acquiring or using the product. We can see this as the functional features of the product [49, 76].

When presenting products to potential buyers, the market shows a grid of titles and pictures – like a supermarket isle. Vendors on online anonymous markets use marketing-like tactics to optimize product performance. For instance, they use capital letters and/or special characters in their title to attract attention. Moreover, some add their vendor name to the title to build on an established brand-name. Next, vendors utilize experiences of buyers as a marketing-tool. Given the consumer-centred aim of markets, a feedback

system is integrated, wherein the products are rated based on buyers' experiences. Accumulated, this gives products ratings that can be used as marketing for a product. We can see all these activities as the marketing of the product [49, 76].

Besides these product differentiators, buyers can distinguish products based on their price. Especially, how cheap or expensive is the product relative to other offerings and is the product worth its price? We can call this the relative price of the product [49, 76]. Next to the price, functional features and marketing, there is one other thing that differs from product to product: who sells it – i.e., the vendor. We know that vendors on online anonymous markets are a rather heterogeneous group based on their different characteristics. Hence, making a meaningful analysis of the sphere of influence of a vendor, requires us to capture this diversity in profiles or subgroups [126]. We will further elaborate on this in Section 5.5.

In short, we set out to explain the performance of a cybercrime solution based on a) the product's functional features, b) the product's marketing, c) the product's relative price and d) the vendor.

5.3. METHODOLOGY

An offering of a cybercrime solution on an online anonymous market can be observed through a product listing, consisting of a title, a description, the price, feedbacks on that product and who sells it. In this section we elaborate on our data and approach to analyze the performance of cybercrime solutions.

5.3.1. DATA

As we aim to understand which factors drive cybercrime sales, we opt to study this on a single market, instead of across multiple markets. If we would study multiple markets over multiple years, our results would be influenced by uncaptured differences among the markets and their evolutionary paths [139]. Thus, we chose to study the performance of cybercrime solutions on one prominent market: AlphaBay. In the underground market ecosystem between 2015 and 2017 AlphaBay was the unchallenged market leader. Until its take-down in 2017 AlphaBay was the most prolific online anonymous market, and – according to the FBI – held 200,000 buyers served by 40,000 vendors [153].

We leverage the parsed and analyzed data set of Van Wegberg et al. [152] spanning eight prominent online anonymous marketplaces, holding cybercrime-related listings (n=44,060) and feedbacks (n=563,223). For each listing, the scraped data includes the title and description of the product, the advertised price, a category classification and the vendor. Additionally, each listing contains feedback that has been proven to be a

reasonable proxy for sales, through internal and external validation [39, 139, 152]. Each feedback contains a comment and a timestamp. The entire data set covers eight markets – ranging from Silk Road 1 to AlphaBay – and spans seven years (2011-2017). AlphaBay is the most recent market in the data set, holding the most listings (n=21,350) and feedbacks (n=288,485) and contains a diversity in cybercrime products.

In their paper, Van Wegberg et al. [152] classified a pre-selection of all listings on the market to ten categories of B2B cybercrime products: malicious apps, botnets, cash-out solutions, compromised email-accounts, exploit kits, hosting services, malware kits, phone banks or details, remote access trojans (RAT) and compromised websites. We leverage their classification and include all AlphaBay listings that have been classified to one of these ten B2B cybercrime product classes (n=7,595). These cybercrime solutions are advertised by 1,346 unique criminal vendors and have received a total of 161,535 feedbacks. This means these solutions have been sold at least that many times, as one can only leave feedback after buying the product or service.

During our manual inspection of the dataset, we found listings that were either classified in the wrong B2B category, or were not a B2B cybercrime product at all. We found four misclassified listings: a listing for renting house cleaning girls (miscellaneous), a listing for 250g ketamine (drugs), a listing for red mastercard ecstasy pills (drugs) and a listing advertising a Beretta and a Glock (guns). These listings and their feedbacks were excluded from the dataset.

Next, we excluded two vendors of credit card data with an amount of feedbacks that is a factor 1000 bigger than the average 16 feedbacks per listing. They received 16,674 and 17,768 feedbacks respectively. There are multiple hypotheses for the size of these numbers. They could have bought from themselves to create many positive reviews or they could have restricted the order amount to 1, forcing buyers to make many purchases to achieve a large order size. Since we can not verify any of these hypotheses, we remove these vendors, their listings and feedbacks from the dataset.

After removing these outliers we have a dataset consisting of 7,543 cybercrime solutions advertised on AlphaBay, sold by 1,339 unique vendors, receiving 126,934 feedbacks.

5.3.2. DESCRIPTIVE STATISTICS

Now, we take a closer look at the performance of cybercrime solutions advertised on AlphaBay. Feedbacks and revenue are not distributed evenly across listings. Figure 5.1 plots the cumulative distribution function of feedbacks and revenues across listings. A small number of the listings is responsible for the majority of the feedbacks and revenue. This is reflected by 20% of the listings receiving 84% of the feedbacks and generating 68%

of the total revenue. These differences between listings can partially be explained by the heterogeneity in cybercrime solutions. We learn from Van Wegberg et al. [152] that large differences between categories of B2B offerings exist. Table 5.1 reports the number of listings, vendors, feedbacks and the total revenue per category of B2B cybercrime solutions on AlphaBay, and Table 5.2 shows the average price, the revenue and lifespan of the listings in that category.

In line with previous research [152], cash-out solutions dominate the cybercrime market in terms of listings, vendors, feedbacks and revenue. Next, we observe a more or less equal distribution of listings and vendors across other categories, with hosting being the smallest and website being the largest category. The number of feedbacks, the revenue and the average price differ from one category to another. App and hosting listings are for example priced relatively low and in turn generate the lowest revenue per listing. The lifespan in days also varies across categories. This means that in some categories listings are removed faster by vendors than in other categories. In total all B2B cybercrime listings generated $3,616,919.45 in revenue on AlphaBay, which is 1.77% of the estimated total revenue of $204,151,800 on the market [40].

Table 5.1: **Listings per category on AlphaBay**

	Total per category			
Category	# Listings	# Vendors	# Feedbacks	Revenue
App	75	48	571	$7,420.53
Botnet	51	37	334	$9,279.99
Cash-out	6,221	1,226	113,897	$3,341,405.44
E-mail	377	151	3,412	$41,191.79
Exploit	54	37	329	$3,922.20
Hosting	7	6	47	$423.56
Malware	149	88	1,140	$34,921.71
Phone	135	80	1,259	$52,457.95
RAT	60	41	425	$7,035.13
Website	414	178	5,520	$118,861.15
Total	7,543	1,339	126,934	$3,616,919.45

Table 5.2: **Listings per category on AlphaBay**

	Average per listing					
Category	Price	*min-max; SD*	Revenue	*min-max; SD*	Lifespan[1]	*min-max; SD*
App	$18.97	0-207; 34.56	$98.94	0-840; 161.67	98.45	1-650; 152.28
Botnet	$116.57	1-1,778; 336.63	$181.96	1-1,778; 395.99	92.20	1-697; 168.43
Cash-out	$71.07	0-6,974; 223.41	$537.12	0-209,124; 3,842.52	85.00	1-798; 140.26
E-mail	$34.44	0-1,100; 108.92	$109.26	0-3109; 319.47	75.48	1-796; 125.24
Exploit	$37.26	0-500; 103.45	$72.63	0-1000; 159.50	117.30	1-708; 182.47
Hosting	$21.06	3-50; 17.64	$60.51	18-136; 45.64	68.14	1-173; 78.57
Malware	$48.51	0-500; 95.21	$234.37	0-5,346; 601.16	91.64	1-762; 149.61
Phone	$67.98	0-3,200; 303.64	$388.57	0-20,910; 1,896.05	92.98	1-745; 146.38
RAT	$42.61	1-648; 122.74	$117.25	0-1,256; 271.12	112.90	1-706; 184.00
Website	$60.49	0-1,695; 158.53	$ 287.10	0-11,088; 919.96	88.75	1-675; 131.70

[1] In days

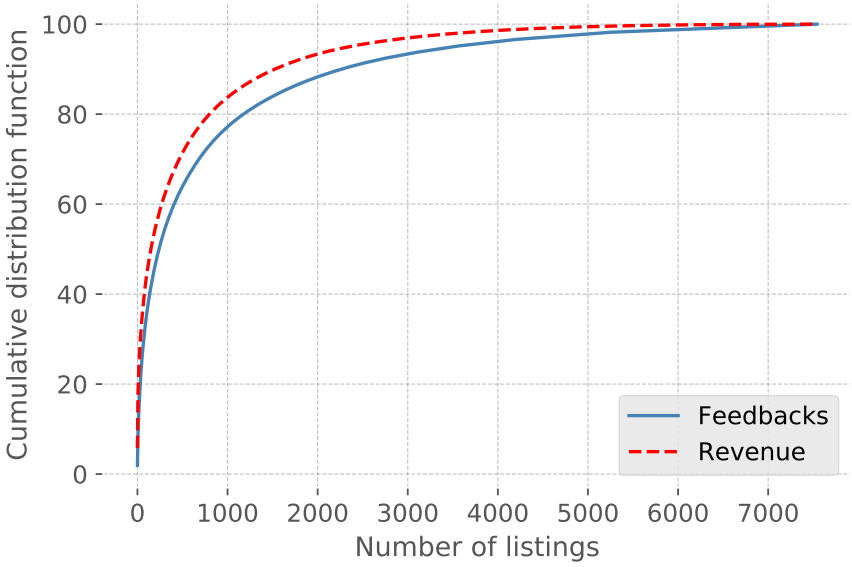

Figure 5.1: **Cumulative distribution function of feedbacks and revenues across listings**

5.3.3. APPROACH

Our methodology to predict the performance of cybercrime solutions on online anonymous markets, consists of four steps:

1. selecting and pre-processing scraped data on B2B cybercrime listings and feedback from AlphaBay (2015-2017)
2. constructing variables that capture product characteristics, i.e., product differentiators and price
3. discerning distinctive vendor profiles by clustering vendor characteristics
4. performing a regression analysis using product characteristics and vendor profiles to predict the performance of cybercrime solutions.

5.4. PRODUCT CHARACTERISTICS

In order to study the performance of cybercrime solutions on AlphaBay in terms of sales, we need to construct variables that grasp the product characteristics introduced in subsection 5.2.2.

Browsing for products and services on the market, AlphaBay presented listings in a grid, showing only a portion of the title and a picture. When a potential buyer clicked on the picture or title, the market would re-direct you to a product page showing all details of the listing – e.g., the price, description, vendor. Therefore, the title is used by vendors for marketing the product. Thus, we will derive the product's marketing from the title. The description of a listing has no character limit and is used by vendors to give a more in-depth product description and to mention how they do business with customers. So we construct variables that grasp the functional features from the description. To do this, we will first have to discern the different ways in which vendors signal or describe certain functional features. In the remainder of this section we will elaborate how we constructed the variables and close with an overview of all variables in Table 5.3.

RELATIVE PRICE

As the data is gathered over a longer period of time, it might hold multiple sales at different prices. Hence, we first need to construct the weighted mean price for a listing. To that end we retrieve for each sale the associated price and sum all these sales prices. Then we divide them by the total amount of sales. This gives us the weighted mean price of a listing. To capture whether a listing is priced 'cheap' or as 'high-end' within its category, we sum all mean prices of the listings in a category and divide them by the amount of

listings in that category. This results in an average listing price per category. We then construct the relative price of a listing by calculating the z-score of the weighted mean price of the listing against the average price of all listings in a category.

CUSTOMER SUPPORT

To find out how vendors signal the availability of customer support, we first manually searched on 'customer support' in the description field of a listing. Examining those listings, we discovered that very few (n=17) listings explicitly mention the term. However, inspecting these listings we discovered 'Jabber', a well-known instant messaging platform, as a way through which the vendor can be contacted for questions. This indicates that vendors might provide their customer support through an external messaging service. Which makes sense, because vendors are active on multiple markets and would want their customers to contact them in one place. Given that vendors mention these platforms and applications as a way for providing support rather than explicitly mentioning customer support, we searched for messaging platforms used. We applied a snowball approach to find which other platforms were mentioned besides Jabber. Starting out with Jabber, this resulted in the list: Jabber (n=537), ICQ (n=361), Skype (n=129), exploit.im (n=106), safe-mail (n=58), jwchat (n=28), Wickr (n=9), protonmail (n=4) and Telegram (n=1). We validated this list by searching for other well-known email services like Gmail, Outlook, Whatsapp and Viber. We found that these are not used as support channels, but are rather part of cybercrime offerings. Finally, using the aforementioned list, we find that 849 listings (≈11%) hold a contact method for providing customer support.

REFUND POLICY

Next to customer support channels, listings often make clear under which conditions one can 'return' the product and get a refund. Manually searching for 'refund' (n=1623) revealed that there are also products such as 'Amazon refunds' and 'refund guides' that contain the word 'refund'. Simply excluding the listings that contain the words 'amazon' or 'guide' would not work, because some 'Amazon refunds' listings ironically also state a refund policy. This calls for a more detailed approach, aimed at reducing false positives – a mention of the word refund while not part of a refund policy – as much as possible. To this end, we separately searched with words or sentences signaling a refund policy ('money back', 'refund if', 'refund after', 'non-refundable' etc.) and with words or sentences signalling a refund related product ('amazon refund', 'double dip', 'refund guide' etc.). We then compare both sets of listings and exclude the listings that are only in the 'refund related product' set and not in the 'refund policy' set. This gave us a set of 1071 listings (≈14%) that explicitly state some kind of refund policy.

TERMS AND CONDITIONS

Besides providing customer support and stating a refund policy, vendors can set other terms and conditions. Some of these terms and conditions are about the anticipated buyer behavior, for example not disputing the sale on the platform or not leaving negative feedback without contacting the customer support first. Another condition is for example that cheating the service will result in being permanently blacklisted. Since these rules differ for each vendor, we will search for signals such as 'condition', 'terms of service', 'terms & conditions', 'rules and terms', 'accept this terms' and 'our rules'. We validated this list in a similar snowball approach as before. We started our search with 'terms and conditions' and manually expanded the list based on the words used by the vendors in the listings, until the addition of more terms did not result in more listings with terms and conditions found. We discovered that 419 listings (≈6%) state that some kind of terms or conditions apply when doing business.

SENTIMENT

On AlphaBay, buyers had the possibility to leave a feedback message after each unique purchase. Many buyers did not use this opportunity and left this field empty, in which case AlphaBay put "No comment" as the feedback message. Of the 126,934 feedbacks, around 45% of the feedbacks has the "No comment" message. The other 55% of the feedback messages contain either a positive experience and recommendation for other buyers, or a negative experience and complaints. To give a score to the negative or positive sentiment of feedback messages, we applied the VADER (Valence Aware Dictionary and sEntiment Reasoner) model for sentiment analysis [84]. VADER uses lexical features and grammatical and syntactical convention rules to express the sentiment with a score ranging from -1 (very negative) to 1 (very positive). This sentiment analysis method has been applied on many different type of texts such as Tweets and performed equal or better to other existing sentiment analysis tools [84]. Since it is domain-agnostic and relies on sentence-level analysis, we do not need to train it using a portion of the feedback data. We apply VADER on all feedback messages and accumulate the sentiment scores of all feedbacks on a listing. The mean feedback sentiment of a listing is 0.14. On a scale from -1 to 1, with each listing having at least one feedback, this means that on average listings receive more positive than negative feedback. Calculating the amount of listings with a sentiment score above 0.0 gives a total of 4,799 listings with an on average positive sentiment (≈64%).

USE OF VENDOR NAME

A vendor name itself can have value for buyers as a well-known and respectable party to do business with. For instance, if a vendor has been active on certain markets under the same name for quite some time. Linking the product that is being offered with the vendor name is therefore used for branding or marketing purposes. An example is the vendor *BHGroup*, who has a listing titled: "BHGroup Fresh Cracked SMTP". Comparing the vendor name with the title text, we discovered that 198 listings (≈3%) contain the name of the vendor offering the cybercrime solution.

RATIO CAPITAL LETTERS

The titles of listings are shown in a grid when a potential buyer searches products on the market. To draw attention, some titles are written with an 'all caps' approach. To quantify the amount of capitals that are used to attract attention, we calculate the ratio of the capital characters to the total amount of characters used in the title. We find that the average ratio of capitals in a title is around 34% and that 7375 listing titles (≈98%) contain at least one capital.

RATIO SPECIAL CHARACTERS

Besides using capital letters, vendors have the option to include special characters in their title, such as a star (⋆), a bow tie (⋈) and many other, different (unicode) special characters. We will use a ratio to express the extent to which a title contains such characters. In order to calculate this ratio, we first remove all words and normal punctuation characters (such as periods, commas, question marks, hyphens, dashes, parentheses, apostrophes, quotation marks etc.) from the titles, in order to calculate the amount of special characters that remain. We discover that the average percentage of special characters is low (≈2.5%) and that 1896 listings (≈25%) contain at least one special character.

CONTROL VARIABLES

To make a meaningful prediction on what drives cybercrime sales, we have constructed several variables that capture product differentiators and the relative price. However, we need to control for factors influencing sales that we expect to have an impact, but do not wish to take into account. As large differences exist in the number of listings within categories of cybercrime solutions, we need to control for the category a listing is in. Otherwise, we end up with a prediction based on the popularity of the product, instead of its differentiators. Next, we need to control for the lifespan of the listing. After all, the longer a listing is on the market, the more time it has to get feedback.

Table 5.3: **Constructed listing variables**

Variable	Mean	Min–Max	SD	Type
Number of feedbacks	16.82	1–2,453	70.09	*Integer*
Relative price	0.00	-1.11–30.90	1.00	*Double*
Customer support	0.11	0–1	-	*Binary*
Refund policy	0.14	0–1	-	*Binary*
Terms & Conditions	0.06	0–1	-	*Binary*
Feedback sentiment	0.14	-0.98–0.98	0.33	*Double*
Ratio of special characters	0.03	0.00–0.89	0.07	*Double*
Ratio of capital letters	0.34	0.00–1.00	0.25	*Double*
Use of vendor name	0.03	0–1	-	*Binary*

5.5. VENDOR PROFILES

We now turn to capturing the influence of the vendor, to assess the impact of the seller on the performance of the product. A vendor name itself has meaning to a buyer as a recognizable force on the market; as it encompasses all the intrinsic characteristics of who he or she is on a market. In turn, these characteristics depict the axes on which vendors differ from one another. Vendor characteristics that have been observed are the amount of listings of a vendor (exposure [126]), its time on the market (experience [49, 76, 126]), the relative pricing of its products (price deviation [76]), having listings in one or multiple categories (diversity [126]), the amount of sales (performance [126]) and the sentiment of the feedback (reputation [49]). From earlier research we know that these vendor characteristics, evaluated separately, influence the performance of products on anonymous cybercrime markets [49, 76].

We do not yet know if there are groups of vendors with distinct configurations of characteristics in the cybercrime segment of AlphaBay. Based on the research of Paquet-Clouston et al. [126] that found three groups exist in vendors selling drugs-related products on AlphaBay, we hypothesize distinct vendor profiles also exist in vendors selling B2B cybercrime solutions. To identify profiles by allowing patterns of characteristics to emerge without assuming ex ante that certain profiles exist, we turn to the person-centered approach of Latent Profile Analysis (LPA).

5.5.1. LATENT PROFILE ANALYSIS

Latent Profile Analysis, a type of Latent Class Analysis (LCA), is a clustering approach that aims to recover hidden groups – called 'latent profiles'– from observed indicators. It is the predominant approach to discern underlying groups in data measuring individuals,

for example criminal actors such as burglars [155], homicide [156] or sex offenders [52]. LPA is a (finite) mixture modelling technique that uncovers continuous or discrete latent variables by estimating the distribution of the latent variable from the data. Because LPA is model-driven, the model is estimated for the population of the study sample, rather than assumed to have some parametric form [157]. With LPA, the indicators can be continuous or mixed-mode and the latent variable is assumed to be discrete, from a multinomial distribution. Since LPA is model-based, information criteria such as the Bayesian information criteria (BIC) and the Coherent Akaike information criterion (CAIC) can be used for model selection.

We construct the aforementioned six vendor characteristics that measure the exposure, experience, performance, reputation, price deviation and diversity of a vendor. The first five characteristics are constructed by aggregating the aforementioned variables for each vendor based on all their listings. We compute the binary characteristic diversity based on whether the vendor has listings that all belong to the same product category (a '0') or to two or more different categories (a '1').

Different Latent Profile models were created based on these six characteristics, using Latent Gold 5.1 software [158], with the goal of analyzing one to five profiles in each. Models with a higher number of profiles could be created, however, these models create profiles with sizes smaller than 5% of the whole vendor population. As we aim to maintain the interpretability and parsimony of the emergent profiles, we limit the amount of profiles to five [52, 66, 155].

Table 5.4 shows the final solutions of models of one to five profiles: the Log-Likelihood (LL), Number of Parameters (Np), Bayesian information criterion (BIC), Corrected Akaike information criterion (CAIC) and entropy values of the model. The ideal model solution has small BIC and CAIC values compared to other models. This means that the model of 5 profiles, as it has the lowest BIC and CAIC, as well as an entropy value equal to the 3 and 4 profile model, is the best fitting model to our data.

Table 5.4: Model output of 1 to 5 profiles

Model	**LL**	**Np**	**BIC**	**CAIC**	**Entropy**
1-Profile	-26353.02	17	52828.43	52845.43	1.00
2-Profiles	-21630.59	33	43498.76	43531.76	0.93
3-Profiles	-20128.14	49	40609.06	40658.06	0.89
4-Profiles	-19274.88	65	39017.73	39082.73	0.89
5-Profiles	-18834.75	*81*	*38252.68*	*38333.68*	*0.89*

To validate that the profiles are clearly differentiated, we conducted one-way ANOVAs using profile membership as the independent variable and the continuous characteristics

as dependent variables, and a Chi-Square test for the nominal characteristic. All profile means are significantly different with a 95% confidence interval on at least four of the six characteristics, except for profiles 1 and 3 – which significantly differ on two characteristics – and profiles 4 and 3 – which significantly differ on three characteristics.

5.5.2. RESULTING PROFILES

To better understand these five profiles, Tables 4-8 show their distinct configuration of characteristics. Per vendor profiles, all six vendor characteristics are reported by the mean score, the delta from the average of all vendors, the median, and the min-max. The revenue is shown separately, as it was not a part of the variables used for LPA, but is useful in comparing and interpreting a profile.

The first profile is the average vendor on the market, with mean values in exposure, diversity and reputation closest to the general average of all vendors (see Table 5.5). We thus name this profile the 'freelancer' profile. It depicts vendors that are neither very successful nor very unsuccessful, but do make some money from offering their cybercrime solutions on AlphaBay. The second profile we encounter is a group of vendors that belong to the established vendors, with a high lifespan (see Table 5.6). As they successfully sell in multiple categories and have many listings, we call this the 'generalist' profile. The third profile is the group of 'specialized' vendors that has a limited exposure, but is still able to generate substantial revenue due to their reputation. They focus on selling expensive products in only one category (see Table 5.7). The products sold by specialists are PayPal accounts and guides, as well as enriched credit card details like BIN's and 'fullz'. The fourth profile holds 'professional' vendors, i.e., established cybercrime facilitators, with both high exposure and experience (see Table 5.8). They sell a diversity of relatively expensive products and services and generate the most revenue of all vendor profiles. The fifth profile is the group of vendors that can be seen as a representation of the 'loafers'. Their exposure and experience are the lowest of all profiles and they generate very little revenue with low-priced listings (see Table 5.9). When examining their listings, Loafers appear to sell mainly 'make money' and cash-out guides.

Table 5.5: Freelancer profile (*n = 305*)

Characteristic	Mean	Δ average	Median	Min	Max
Exposure	2.94	-2.69	3	1	8
Diversity	0.49	0.00	0	0	1
Experience	110.59	-76.61	94	3	336
Performance	17.84	-76.96	11	2	82
Reputation	0.14	+0.03	0.13	-0.43	0.70
Price deviation	-0.24	-0.36	-0.27	-0.49	0.05
Revenue	301.99	-2,399.22	103.97	0.19	3,499.50

Table 5.6: Generalist profile (*n = 339*)

Characteristic	Mean	Δ average	Median	Min	Max
Exposure	9.02	+3.39	8	1	28
Diversity	0.85	+0.44	1	0	1
Experience	382.05	+194.84	372	12	797
Performance	97.33	-2.53	72	2	396
Reputation	0.14	-0.03	0.13	-0.36	0.59
Price deviation	-0.16	-0.24	-0.22	-0.51	0.51
Revenue	2,165.29	+535,92	980	0	28,360.97

Table 5.7: Specialist profile (*n = 205*)

Characteristic	Mean	Δ average	Median	Min	Max
Exposure	2.43	-3.23	2	1	8
Diversity	0.34	-0.15	0	0	1
Experience	110.59	-67.75	58	1	730
Performance	9.87	-84.93	6	1	76
Reputation	0.27	+0.10	0.30	-0.70	0.98
Price deviation	1.63	+1.41	-0.15	-0.36	3.26
Revenue	2,700.87	+0.34	1100	6	131,509.49

Table 5.8: Professional profile (*n = 114*)

Characteristic	Mean	Δ average	Median	Min	Max
Exposure	22.96	+17.33	12.5	1	172
Diversity	0.82	+0.33	1	0	1
Experience	519.63	+332.42	574.50	42	801
Performance	747.66	+652.86	512	18	5,613
Reputation	0.15	-0.02	0.14	-0.14	0.44
Price deviation	0.22	+0.10	-0.15	-0.36	3.26
Revenue	19,336.65	+16,635.44	8,881.39	657.50	281,364

Table 5.9: Loafer profile (*n = 376*)

Characteristic	Mean	Δ average	Median	Min	Max
Exposure	1.26	-4.37	1	1	3
Diversity	0.14	-0.35	0	0	1
Experience	9.83	-177.38	3	1	51
Performance	3.31	-91.49	2	1	28
Reputation	0.16	-0.01	0.17	-0.92	0.95
Price deviation	-0.21	-0.33	-0.27	-0.52	0.35
Revenue	87.03	-2,614.18	22.39	0.00	2,100

5.6. PREDICTING CYBERCRIME SALES

As stated earlier, our objective is to empirically derive how different vendors and product characteristics contribute to B2B cybercrime sales. We employ regression analysis to this end and use our constructed vendor profiles, along with variables that capture product characteristics, as regressors to make predictions about sales.

We now test the extent to which product characteristics and vendor profiles influence the prevalence of sales. We do so by constructing several explanatory regression models on top of the data. Note, that since exact sales figures are not available we use the *number of feedbacks* as a proxy for its sales [39, 139, 152]. This quantity constitutes the dependent variable of our regression models. We model feedbacks via Negative Binomial regression using a logarithmic link function. The specific choice of regression model is due to our dependent variable constituting a count. When modelling count data, linear regression (e.g., Ordinary Least Squares) is less appropriate as it assumes the response variable to be a continuous quantity. Whereas count data are non-negative integer values for which Generalized Linear Models (GLM)s such as Negative Binomial regression are more suited.

Our regression models have the following general structure:

$$ln(d_v) = c_0 + \sum c_i \times v_i + e$$

where d_v is the dependent sales numbers variable and v_i are the product- and vendor-related variables. The extent to which the independent variables influence the dependent variable are captured by regression coefficients c_i. Moreover, c_0 is a constant value setting a baseline for sales and finally e an error term. All variable definitions along with their descriptive statistics were provided earlier in Table 5.1 and Table 5.2.

We have three groups of independent variables. First, we define `Listing lifespan` and product `Category` as control variables. In doing so, we may factor out the effect of having a higher number of sales due to listings having been advertised for longer periods, or belonging to a specific category which may be more popular relative to others. Next, we have our vendor-related variables, and finally the product-related ones as the remaining two groups. Given these groups of variables, Table 5.10 provides an overview of the regression models that we have constructed. It lists the estimated coefficient values, their significance levels, in addition to several other goodness-of-fit quantities of interest per model that we will discuss shortly.

We start by constructing a model which only includes our control variables (model 1) as a baseline to compare against. Next, we construct two more models (models 2-3) by additionally including only those groups of variables pertaining to either vendors or products to independently demonstrate the effects of vendor and product characteristics on sales. Finally model 4 constitutes our complete model of the data, which simultaneously includes control variables, vendor-related variables, as well as the product-related variables.

We first discuss our overall findings based on these 4 models, and subsequently move on to discuss model interpretation and the details of our full model. In terms of our overall findings, we observe that 32% of the variance in sales numbers is purely explainable by our control variables. This may be observed through the pseudo-R^2 value of model 1. That is to say that a non-trivial amount of variance in sales numbers is explainable by either the amount of time a listing has been advertised or the category to which it belongs. In comparison, the pseudo-R^2 value of our full model (model 4) suggests that 47% of variance in sales is explainable by control variables, vendor characteristics and product characteristics together. Since these pseudo-R^2 values are calculated against a baseline model with only a constant baseline coefficient, however, (not shown here), we may compare the pseudo-R^2 values of model 1 and 4 relative to model 1 to characterize how much additional variation the vendors and products explain. The secondary pseudo-R^2

values relative to model 1 that are reported in Table 5.10 suggest that an additional 22% of the variance in sales numbers is purely explainable by vendor or product characteristics. Similar comparisons may be drawn among models 1-2 or models 1-3 to observe the effects of vendors and products independently. Several goodness-of-fit quantities have also been reported for all models, e.g., dispersion, log-likelihood and Akaike Information Criterion (AIC). These indicate that our model 4, our complete model, is a better fit to our data. This may be observed by a dispersion value that is closer to 1, increased log likelihood and a smaller AIC values for model 4 in relation to the others. Next, we move on to discuss model 4, how it may be interpreted and our findings in more detail since it is a better fitting model to our data.

We start by taking a closer look at what effects vendors are suggested to have based on our model. Note that the group of vendor-related variables of our full model, constitute four so-called 'dummy' variables signifying if a particular vendor has a `Generalist`, `Loafer`, `Professional` or a `Specialist` profile. By definition, if the vendor profile is neither of the above, it should be that of a `Freelancer`. Hence, we do not need to include five dummy variables to represent all vendor profiles in our model. As such, our full model captures the effects of vendor profiles on sales, relative to vendors in the `Freelancer` profile which has been left out.

We may examine the effect of vendors, by interpreting the coefficient values associated with each vendor profile. We illustrate by example. For instance, our model suggests that belonging to the `Generalist` vendor profile has a significant positive coefficient value of 0.16 and correlates with a relative increase in sales. As stated earlier, this is an increase relative to vendors in the `Freelancer` profile. More specifically, if all else were held constant, a change of vendor profile from `Freelancer` to `Generalist` is correlated with a $e^{0.16} = 1.17$ multiplicative increase in sales. As such, we expect a `Generalist` vendor's sales to be 17% higher than the `Freelancer`'s sales. `Loafer` sellers on the other hand, exhibit a relatively lowered sales figure if all else were held constant ($e^{-0.53} = 0.59$, i.e., 59% of freelancer sales). Curiously, Model 4 also suggests that 'specializing' does not lead to a significant increase in sales compared to the `Freelancer`. Last, we see that `Professional` vendors perform best, as they appear to have 150% higher sales relative to `Freelancers` ($e^{0.93} = 2.5$). Overall, we observe - apart from the `Specialist` - all vendor profiles to significantly correlate with higher or lower sales figures as we have initially hypothesized.

Next, we examine how product characteristics influence sales. As before, we do so by interpreting the coefficients values of our product-related group of variables. Unlike before, however, these should be interpreted differently since they do not capture effect

sizes relative to a variable that has been left out.

Take the `Customer support` variable for instance which has a significant coefficient value of 0.36. This value suggests that if all else where held constant, products that are sold with customer support sell $e^{0.36} = 1.43$ times more than those that are not. Strictly speaking, this effect should be interpreted as if the customer support variable where to increase by 1 standard deviation from its mean value, while all else where held constant, we should expect to see 1.43 times more sales. Furthermore, we see that the other two functional features, namely `Refund policy` and `Terms & Conditions` show a mixed result. Whereas products that entail refund policy information do see a significant positive correlation (0.43) with product performance, signaling terms & conditions when buying and using the product does not have a significant impact on sales.

As another example, we also see that products that deviate from the mean price of their product category by 1 standard deviation, sell $e^{-0.27} = 0.76$ times less. That is, equivalent products that are listed with higher prices on average sell less. In line with earlier work [47, 117] we find evidence of feedback sentiment influencing product sales. The reported coefficient (0.20) for the `Mean sentiment` variable shows that an increase in sentiment has a significant positive effect on sales. We also find evidence of marketing on products, like the use of either special characters or capitals in their title, influence sales in a positive way. These may be observed via the coefficient values of the `Ratio special characters` and `Ratio capital letters` variables respectively which may be interpreted in a similar fashion to the previously discussed product-related variables. Vendors employing marketing-like techniques, i.e., using their own name in the title of a listing, also appear to positively correlate with higher sales (see the `Use of vendor name` variable).

In summary, we have found evidence in support of both vendor profiles and certain product characteristics positively or negatively influencing sales numbers. That being said, while we have explained a non-trivial amount of the variation in feedback numbers among sold cybercrime products – and by proxy sales prevalence – much still remains to be explained.

5.7. DISCUSSION

In this section, we first discuss inherent challenges within our approach in light of the constructs used in the research design. Second, we will touch upon the public policy take-aways of our findings.

5.7.1. LIMITATIONS

First and foremost, given the fact that we use scraped data from AlphaBay, we have to rely on proxies for a number of variables that are not visible by just observing the market's web interface. Most importantly, we use the number of feedbacks as a proxy for sales, similar to earlier work [39, 139, 152]. Note that not all buyers leave feedback, so the proxy systematically underestimates the sales and thus represents a lower-bound. While this gives us a reliable lower bound proxy, we do not know if this proxy always corresponds similarly to the actual sales volume or whether there it contains bias – i.e., whether for some product type customers are more likely to leave feedback than for other types. The differences in the ratio between feedbacks and sales is, however, only directly observable from the seized backend server of AlphaBay. Future research might shed more light on this and on potential bias. That being said, our findings are in line with other studies that use feedbacks or comments as a proxy for sales and using that same proxy for predicting 'criminal performance' [49, 76].

Second, we should state that our choice to analyze the performance of listings on one market instead of across market, yields valid results for the cybercrime segment of AlphaBay, but leaves the question on generalizability of our findings unanswered. As we argued before, AlphaBay was the most 'complete' market up until now, so any market dynamics identified at AlphaBay's cybercrime segment might well be in play at other markets. Future work could focus on comparing our findings on AlphaBay with the predictors of product performance on other online anonymous markets.

5.7.2. PUBLIC POLICY TAKE-AWAYS

Our findings suggest that simply looking at either successful vendors – in terms of revenue – or popular products – in terms of high feedback numbers – one turns a blind eye towards less obvious 'pathways' into vendor success. As we demonstrated, not all vendors fit the same profile and there are indeed multiple ways to make it big. On average, both a 'specialist' and a 'generalist' turn over near-similar amounts, but between them the amount of listings, feedbacks, price and diversity of the products they sell, differs significantly. Next to interventions on online anonymous markets, like take-over and infiltrations to undermine trust in the market ecosystem or take-downs to simply shutdown certain markets, law enforcement agencies try focusing on big or central players.

Based on our insights, authorities might differentiate interventions in certain market segments, e.g., cybercrime solutions, considering the distinctions in vendor profiles. For instance, an intervention aimed at (professional) facilitators of many aspects of the cybercrime enterprise, might focus on vendors who fit the 'generalist' or 'professional'

typology. Or interventions aimed at specific niche products, might target 'specialists'. In turn, these profiles influence the relative success of a cybercrime solution. Apparently when choosing who to do business with, cybercriminals dislike certain sellers and favor distinctive others. One can imagine the usefulness of these insights when setting-up a sting operation.

Apart from who sells the product, our findings indicate that certain product differentiators significantly influence the performance of cybercrime solutions. Marketing techniques influence the performance, in terms of feedbacks, of offerings. For instance by branding the product using the vendor name in the title of the listing. Next, we have seen evidence that certain functional features influence product performance, specifically customer support and refund policies. Both features hint towards a professional set-up of doing business, which in turn is reflected by higher sales numbers of the product that contain these functional features.

All in all, the aforementioned aspects can give insights into which cybercrime solutions perform better compared to others. This might even give law enforcement agencies the potential to take a more preventive course of action – by looking at popular products and/or vendors early in their life-course. Future work could identify how our model can predict the popularity of certain products spanning a market's complete life-cycle by using early and late stage snapshots and compare the predictions with reality.

5.8. RELATED WORK

Important parts of our paper build on or benefit from recent insights into a number of topics. First, our work can be tied to measurements of the nature, size and volume of trade on online anonymous markets. Second, we can identify similar analysis compared to our vendor profiles in studies into 'criminal performance' in underground markets. Third and last, we benefit from and contribute to the research body on collaboration between cybercriminals. In this section, we discuss related work on these three topics.

MEASUREMENTS ON ONLINE ANONYMOUS MARKETS

The first longitudinal study on the size, nature and volume in sales over time and across multiple online anonymous markets was undertaken by Soska and Christin [139]. Most existing studies include, or even focus on, drugs and physical goods, which represent a large share of the products offered on the markets [13, 14, 151]. In contrast, and most closely connected to our work, Van Wegberg et al. [152] investigated the trade of cybercrime commodities on online anonymous markets, thereby explicitly focusing on a different product type, i.e., cybercrime solutions.

CRIMINAL PERFORMANCE ON UNDERGROUND MARKETS

Next, our work is related to research into the 'criminal performance' of actors and products on underground markets [49, 76, 126]. Both Decary-Hetu & Leppanen [49] and Holt et al. [76] leveraged signaling theory to predict criminal performance on stolen data markets, e.g. carding forums. They show that vendor experience, e.g. lifespan and number of forum posts, and certain product features, like customer support options, predict the performance of carders on forums. Next, Paquet-Clouston et al. [126] investigated 'vendor trajectories' on AlphaBay using group-based trajectory modeling in vendor market share.

CYBERCRIMINAL COLLABORATIONS

Finally, our work can be tied to research efforts aimed to understand the collaboration between cybercriminals. An in-depth analysis of European and American police cases by Leukfeldt et al. [98] yielded relevant insights into the offline contacts of online criminals. This offline angle in collaboration was also investigated by Lusthaus [106, 107] who interviewed over one hundred cybercriminals and unraveled how and where collaborations start. Next, Hutchings [82] studied the sharing of techniques amongst cybercriminals and identified distinct collaboration types, ranging from one-time partners to sustainable partnerships.

5.9. CONCLUSION

In this paper we investigated the performance of products in the business-to-business cybercrime market segments on AlphaBay. As we know that not all products and vendors are equally successful on the market, we aim to predict which characteristics of both the criminal entrepreneur and their product influence the performance of cybercrime solutions. To that end, we constructed new variables to grasp the relative price, functional features and the marketing of the product. Next, we have captured the diversity in vendors on the market in five distinct profiles based on hierarchical clustering of five vendor characteristics: exposure, diversity, experience, performance, reputation and price deviation.

We use our constructed variables and vendor profiles to empirically predict the performance of cybercrime solutions. Since we are not interested in how the type of product or lifespan contributes to the relative success of a solution, we control for the time a listing is on the market and the type of product offered. First, we find that all vendor profiles – either positively or negatively – influence cybercrime sales. Second, in line with what other researchers have observed on carding forums, we identify particular functional features, i.e., refund policy and customer support, to be positively and significantly correlated

with the performance of a cybercrime solution [49, 76]. Third, we show that marketing the product, in terms of using capitals in the title to attract attention when browsing the market, influences the sales numbers of a cybercrime solution in a positive way. Likewise, branding a product, i.e., using the vendor's name in the title, increases the performance of the product.

Furthermore, our findings show that the profile of the criminal entrepreneur is able to predict a relative high degree of variance in the performance of cybercrime solutions, compared to all the product differentiators combined. This suggests that outsourcing is and has remained a 'human process', wherein decisions leading up to acquiring a cybercrime solution literally start and end with who sells it to you. Interestingly, specialized criminal vendors do not significantly out perform 'freelancers'. It seems that rather than specialized vendors of niche-products, buyers on online anonymous markets would rather do business with 'professional' criminal vendors, i.e., experienced facilitators, supplying a wide range of products and services.

In terms of generalizability of our findings, we should point out that our choice to analyze the performance of vendors and cybercrime solutions on AlphaBay, only gives us an accurate picture of the market dynamics on this market. As we argued before, AlphaBay was the most complete market up until now, so any market dynamics identified at AlphaBay might well be in play at other markets. Still, this leaves us unable to extrapolate this picture beyond AlphaBay. Nonetheless, our model explains up to 47% of the variance in feedbacks on listings, whereof 22% stems from our constructs. Future work can therefore try to unravel the factors that influence 'criminal performance' that we do not yet know of.

Yet, we have added light to the black box of dynamics behind the performance of cybercrime products on online anonymous markets. Many studies into the size and nature of trade of drugs and/or digital goods on online anonymous markets observed that not all product nor vendors are equally successful [14, 139, 153]. To the contrary, many products just sell a handful of times, and some vendors make less than a couple of hundred bucks in their entire career on the market. Using the economics lens of product differentiators and taking the profile of the vendor into account, we were able to look at what drives the performance of cybercrime solutions for the first time. It seems that just some differentiators really matter, specifically those that can be seen as signals of a professional operation, e.g., market independent customer support channels and detailed refund policies, and clever marketing, e.g., branding products with a vendor's name.

Likewise, our findings suggest that - apart from product differentiators - being a professional facilitator who sells a variety of relatively expensive cybercrime solutions, is

an important predictor of product performance. However, simply looking at successful vendors by adding up their sales numbers or calculating their revenue still is a rather crude approach to identify big players - as we see reflected by the 'generalist' and 'specialist' profile. We uncovered that cybercriminal entrepreneurs on AlphaBay can be considered a truly heterogeneous group and the 'pathways' into vendor success are rather diverse. Based on these insights, authorities might differentiate interventions in certain market segments, e.g., cybercrime solutions.

Table 5.10: Generalized Linear Regression Model (GLM) for feedback size of the products

		Response Variable: Count of product feedbacks			
		Negative Binomial with Log Link Function			
		(1)	(2)	(3)	(4)
Control Variables	Listing lifespan	0.01**	0.01**	0.01**	0.01**
		(0.0001)	(0.0001)	(0.0001)	(0.0001)
	Category botnet	−0.31	−0.38	−0.37	−0.42
		(0.25)	(0.23)	(0.24)	(0.23)
	Category cash-out	0.62**	0.23	0.46**	0.21
		(0.16)	(0.15)	(0.15)	(0.14)
	Category e-mail	−0.06	−0.44**	0.07	−0.28
		(0.17)	(0.17)	(0.16)	(0.16)
	Category exploits	−0.77**	−0.72**	−0.64**	−0.62**
		(0.25)	(0.23)	(0.24)	(0.23)
	Category hosting	−0.28	−0.19	−0.14	−0.17
		(0.54)	(0.51)	(0.52)	(0.50)
	Category malware	−0.21	−0.30	−0.16	−0.25
		(0.19)	(0.18)	(0.18)	(0.18)
	Category phone	−0.21	−0.45*	−0.25	−0.43*
		(0.20)	(0.19)	(0.19)	(0.18)
	Category RAT	−0.84**	−0.96**	−0.61**	−0.76**
		(0.24)	(0.23)	(0.22)	(0.22)
	Category website	0.07	−0.18	0.09	−0.11
		(0.17)	(0.16)	(0.16)	(0.16)
Vendor Variables	Vendor profile Generalist		0.15**		0.16**
			(0.05)		(0.05)
	Vendor profile Loafer		−0.61**		−0.53**
			(0.07)		(0.07)
	Vendor profile Professional		1.05**		0.93**
			(0.05)		(0.05)
	Vendor profile Specialist		−0.48**		−0.06
			(0.07)		(0.08)
Product Variables	Refund policy			0.65**	0.43**
				(0.04)	(0.04)
	Terms & Conditions			0.02	0.003
				(0.07)	(0.06)
	Price deviation			−0.26**	−0.27**
				(0.02)	(0.02)
	Customer support			0.50**	0.36**
				(0.05)	(0.05)
	Mean sentiment			0.12**	0.20**
				(0.05)	(0.04)
	Ratio special characters			−0.46*	−0.86**
				(0.20)	(0.20)
	Ratio capitals			1.23**	0.85**
				(0.06)	(0.06)
	Use of vendor name			0.76**	0.57**
				(0.09)	(0.09)
	Constant	1.40**	1.37**	0.80**	0.88**
		(0.16)	(0.15)	(0.15)	(0.15)
	Dispersion	6.77	3.6	5.1	3.5
	Pseudo R2	0.32	0.43	0.41	0.47
	Pseudo R2 relative to Model (1)	-	0.16	0.13	0.22
	Observations	7,543	7,543	7,543	7,543
	Log Likelihood	−25,686.18	−25,042.86	−25,181.69	−24,758.57
	Akaike Inf. Crit.	51,394.36	50,115.71	50,401.37	49,563.15

6

INTERVENTIONS

The previous chapters in this dissertation have investigated the role of online anonymous markets in outsourcing crucial resources of profit-driven cybercrime. We have demonstrated the potential for outsourcing via an online anonymous market, yet we know that law enforcement agencies have tried intervening on online anonymous markets to mitigate this outsourcing potential. In this chapter we first explore the enablers of online anonymous markets and thereafter construct a historical perspective of online anonymous market interventions, before focusing on evaluating these interventions. To this end, we take the recent international intervention, dubbed Operation Bayonet, targeting two prominent online anonymous markets - i.e., AlphaBay and Hansa Market - as a case. We investigate the effects of the operation on all newly registered vendors on Dream Market (n=220) during and shortly after Operation Bayonet by mapping their individual and historic characteristics to discern migration patterns and changes in vendor behavior. [1]

[1] Sections 6.1.1 and 6.1.3 are partly based on earlier work published as "Hartel, P., & Van Wegberg, R.S. (2019). 'Crime and Online Anonymous Markets'. In *International and Transnational Crime and Justice*. Natarajan, M. (ed.), 67-72". Section 6.1.2 is based on translated parts of the paper "Verburgh, T., Smits, E., & van Wegberg, R.S. (2018). 'Uit de schaduw: Perspectieven voor wetenschappelijk onderzoek naar dark markets'. *Justitiële Verkenningen*, 44(5) 68-82." Section 6.1.4 is based on translated parts of the paper "Oerlemans, J.J., & Van Wegberg, R.S. (2019). 'Opsporing en bestrijding van online drugsmarkten'. *Strafblad*, 17(5), (pp. 25-31)." Section 6.2 is previously published as "Van Wegberg, R.S, & Verburgh, T. (2018). 'Lost in the Dream? Measuring the effects of Operation Bayonet on vendors migrating to Dream Market'. In *Proceedings of the Evolution of the Darknet Workshop* (pp. 1-5)."

6.1. A CHANGING POLICING PARADIGM

6.1.1. INTRODUCTION

Online anonymous markets have been around since early 2011 and are since then a documented part of today's cybercrime ecosystem. Different researchers have studied how online anonymous markets support the evolution of online criminal activity in the underground economy. The first studies focused on underground forums, e.g. carding forums revolving around the (re)selling of stolen credit card details [120]. After the first market came – i.e., Silk Road – mostly computer scientists started to shift focus to look at these online anonymous markets [39, 139]. Their popularity as markets in illicit goods has steadily grown over the years [139, 152]. With the rise of markets like Silk Road, similar marketplaces came into existence where next to drugs, supply and demand of other products and services could meet: ranging from physical goods, like passports and weapons, to digital goods and services, like carding and cybercrime software [13, 145, 152]. As a result, we can witness an increasing supply of criminal product and services on standardized markets in the underground economy [145, 152].

The general public has become aware of online anonymous markets and their underlying technologies, such as The Onion Router (TOR) and Bitcoin, by their practical application. The cryptocurrency Bitcoin allows 'anonymous' transactions and the TOR-protocol supports anonymous browsing. The perception of anonymity attracted large number of vendors and buyers to online anonymous markets. Over the last few years, these have further matured in business continuity management, in consumer-oriented operations, and in turnover. A single top tier market can turn over around 200,000 US dollars daily [139]. But how do these online anonymous markets challenge traditional law enforcement?

The focus of traditional law enforcement interventions primarily lies in arresting and prosecuting market actors, next to the seizure of profits on anonymous online markets. Ironically, the anonymous or pseudonymous nature of these markets, provided by technologies such as TOR and the use of cryptocurrencies as a means of payment, creates a high level of transparency of the entire ecosystem of markets – i.e., the market landscape, including market turn-over and product portfolio. This gives law enforcement agencies (LEA) the opportunity to prioritize their policing efforts based on existing ecosystem insights. For instance, prioritize markets that facilitate the trade in child sexual abuse material, or go after niche-markets. In turn, this could allow LEA to disrupt the ecosystem and the business models behind the online anonymous markets. Nevertheless, the current criminal justice interventions seems to be relatively ineffective in terms of the overall sales volume in the online anonymous market ecosystem [48, 139].

In this chapter, we look at the enablers of online anonymous markets and analyze how law enforcement agencies have intervened over the past decade - making use of these enablers. We take one recent intervention - Operation Bayonet - as a case, and demonstrate how to evaluate the impact of interventions on online anonymous markets.

6.1.2. ENABLERS FOR ONLINE ANONYMOUS MARKETS

Three underlying security measures enable online anonymous markets: anonymization protocols, cryptocurrencies, and reputation management – i.e., review systems. These security mechanisms are based on two main concepts that facilitate the thriving business model of online anonymous markets: anonymity and trust [12–14]. Anonymity and trust ensure that online anonymous markets became and remain a prominent part of the present-day ecosystem of underground markets. An ecosystem that is known to withstand external events – e.g., law enforcement interventions. In this section we show how anonymity and trust facilitate online anonymous markets in more detail.

ANONYMITY

Since most – if not all – advertised products and services are illegal, anonymity is very important in an online anonymous market. The various techniques that promote this anonymity are discussed here: anonymization protocols, cryptocurrencies and encryption.

First, we discuss the technologies that allow you to surf the web anonymously and to host anonymous sites - also referred to as anonymization protocols or darknets. This technology allows for the creation of a 'meeting place' where all parties are anonymous – i.e., the buyer, seller and the owner of the site. Note that, next to actors on the market, the location of the market itself remains unknown. That is, the anonymization protocol prevents (geo-)attribution of the market's server – i.e., a hidden service. TOR is the best known and most used anonymization protocol. Yet, other protocols are also available such as I2P & Freenet.

Second, the original Silk Road showed a clear development compared to previous online anonymous markets – like Farmer's Markets. Where the previous markets only used TOR, the original Silk Road was the first to combine TOR with the use of the cryptocurrency bitcoin as a means of payment. This allowed users to not only browse, but also contract with vendors through 'anonymous' payments, which meant a huge leap forward in user anonymity [36]. Bitcoin is currently still the most used cryptocurrency on online anonymous markets [119, 139]. However, we do see more privacy-oriented cryptocurrencies, like Monero, supported on some (niche) markets [119, 139].

Third and last, Pretty Good Privacy (PGP) is used on online anonymous markets to further support anonymity [36]. PGP allows users to encrypt messages, so that they can only be read by the sender and receiver. By using PGP on online anonymous markets, criminals try to make sure that the police or administrator cannot access the content of their messages – as they could for instance include delivery address of the physical packages or other personal identifiable information. Because the messages can only be decrypted by one entity, PGP is also used as a kind of proof of identity. By checking whether someone uses the same PGP key at different times, allows you to determine whether you are in contact with the same person. As we know that vendors (pro)actively migrate between markets, PGP is used as a means to transfer reputation alongside with supply from market to market, thereby aiming for business continuity [96].

Combined, anonymization protocols, cryptocurrencies and encryption make these markets anonymous in nature. At the same time, as a by-product of this anonymous nature, these markets have introduced a global and scalable solution for buyers and vendors of illicit goods and services.

TRUST

In turn, anonymity can create distrust. After all, anonymity might fuel a scammer's paradise. That is why market security mechanisms are integrated to ensure trust, regardless of anonymity. Before criminals make a trade, there has to be an adequate level of trust in the market itself, both parts of the transaction – i.e., buyer and vendor – and the payment system. The market's review system is of great importance for the overall confidence in the market and its vendors.

First, comparable to the rating systems of legal sales websites such as Ebay and Amazon, the review system was set up to reduce intrinsic risks of scams [47, 75]. The basis of the system is that trust is created through information sharing. This, of course, does not concern data that jeopardizes anonymity, such as geo-location. It is about public information sharing since the community must have access to the information before trust can be created. A review system is used where both sellers and buyers can be assessed, thus increasing confidence on both sides of the transaction. Buyers leave feedback on vendors and their products or services, allowing vendors to build and grow their reputation and attract returning customers [139, 152].

Second, the review system is backed up by other market control mechanisms. To prevent scams, the market handles payments in a transaction, and keeps the funds in escrow until the buyer indicates the goods or services are received and/or working. If this is not the case, the market will not release the funds and as a consequence, the vendor is not paid. Similarly, buyers can trigger admin interference if they dispute the

transaction – a situation wherein the admin investigates the transaction and can issue a refund. The latter follows on the former, whenever the market still has the funds in escrow and the vendor has repeatedly failed to deliver the goods. As these market controls can also be misused by admins – i.e., exit scams, where market admins disappear with all funds in escrow – markets have implemented so-called multi-sig bitcoin wallets. In contrast to normal wallets, multi-sig wallets require two out of three signatures to initiate a transaction. In this case, the market, buyer and vendor are all signatories, which means that no one can initiate a transaction single-handedly.

As we now have a better understanding of the enablers of online anonymous markets and how they conceptually operate, we can turn to how interventions have impacted the operation of markets and vendors.

6.1.3. POLICING ONLINE ANONYMOUS MARKETS

Being confronted with these new aspects of digital innovations in crime, law enforcement had to come up with at least the same level of innovation in their actions against these online anonymous markets. And they did. Law enforcement agencies intervened on online anonymous markets in two different ways. First, they aimed interventions on complete markets - taking over and/or taking down a market. Second, they focused on prominent market actors - mostly administrators, like Ross Ulbricht, the admin of the original Silk Road. In the remainder of this section, we give an overview of major law enforcement interventions between 2011 and 2019 and summarize all important elements in Table 6.1 based on publicly available information - e.g., law enforcement press releases.

MARKET TAKE-DOWNS

Operation Onymous [2] was a coordinated operation between police forces from 17 countries coordinated by Europol in 2014. It resulted in the take down of around 27 hidden services. The investigation was sparked by a TOR zero-day vulnerability, through which server location of hidden services could be attributed. Amongst others, Operation Onymous shut down the online anonymous markets Cloud 9, Hydra, Pandora & Silk Road 2. Law enforcement agencies arrested 17 actors, of which only one name has been made public: Blake Benthall, known on Silk Road 2 as the main administrator under the pseudonym 'DEFCON'.

Mid-2019, German law enforcement agencies – with the help of the American and

[2] See https://www.europol.europa.eu/activities-services/europol-in-action/operations/operation-onymous and https://www.wired.com/2014/11/operation-onymous-dark-web-arrests

Dutch police agencies – succeeded in the take-down of Wall Street Market [3]. After a two-year investigation, which focused on keeping track of server locations, law enforcement made use of the exit-scam by administrators, to connect the final piece of the puzzle – i.e., tying the market's servers to admin finances and therewith the attribution of actors and infrastructure. The admins where arrested and the market taken down.

MARKET TAKE-OVERS

During Operation Bayonet [4] two leading online anonymous markets took centre stage in a 2017 joint policing effort of the Federal Bureau of Investigation (FBI) and the National High Tech Crime Unit (NHTCU) of the Dutch police. In a coordinated sweep, the FBI succeeded in the take-down of AlphaBay, while the Dutch police took over, briefly controlled, and then shut down Hansa Market. With this take-over and administration of the market, their aim was to undermine trust in the ecosystem. Next, by planning these actions sequentially, the police agencies expected criminals active on AlphaBay, which was shut down first and made look like an exit-scam by the FBI, to make their way to Hansa Market – which at that moment was operated by the Dutch police. This put law enforcement agencies in a perfect position to not only disrupt the ecosystem, but also to collect valuable data on thousands of users. Thereby anticipating users to become worried that their personal information was now in the hands of law enforcement, refraining the future use of these credentials and forcing discontinuity of business.

ACTOR-FOCUSED INTERVENTIONS

Next to interventions aimed at either the ecosystem or specific markets, there have been interventions aimed primarily at the attribution of actors. In traditional policing, actor attribution – or in other words: catching criminals – is the main or sole focus of law enforcement operations. It only seems logical therefore, that the first law enforcement operations against online anonymous markets was predominantly aimed at only one individual: the admin of Silk Road.

After a tip by an informant early 2011, a US inter-agency taskforce – dubbed Marco Polo [5] – was set-up to investigate the original Silk Road, which was then in operation for six months. To get closer to the identity of administrator Dread Pirate Roberts, law enforcement turned to yet another traditional policing tactic: going undercover. After a

[3] See https://www.europol.europa.eu/newsroom/news/double-blow-to-dark-web-marketplaces and https://www.zdnet.com/article/law-enforcement-seizes-dark-web-market-after-moderator-leaks-backend-credentials

[4] See https://www.europol.europa.eu/newsroom/news/massive-blow-to-criminal-dark-web-activities-after-globally-coordinated-operation and https://www.wired.com/story/hansa-dutch-police-sting-operation

[5] See https://www.wired.com/2013/11/silk-road/

year-long investigation the undercover agent was not only able to gain the trust of the Silk Road administrator but also identify Ross Ulbricht as the person behind the moniker Dread Pirate Roberts. As they succeeded in his arrest and subsequent take-over of his admin-account, the taskforce was able to take down the market.

Late-2019 the Italian Guardia di Finanzia used similar tactics to attribute the admins of Berlusconi Market [6]. When they investigated one of the top sellers on the market, they were able to link his online activity to his home address. During his arrest the police seized laptops, on which they found evidence that helped them trace and identify the Berlusconi admins. After their arrest, the police were able to take control of the market and shut down its operation.

[6] See http://www.gdf.gov.it/stampa/ultime-notizie/anno-2019/novembre/arrestati-tre-amministratori-del-black-market-denominato-berlusconi-market-attivo-nel-dark-web and https://securityaffairs.co/wordpress/93603/cyber-crime/berlusconi-market-darkweb.html

Table 6.1: Overview of major online anonymous market interventions

Name	Timing	Lead LEA's	Focus	Type	Vector	Arrests	Seizures
Operation Marco Polo	2011	FBI; DEA; DHS;IRS	Silk Road 1 –admins, moderators and top sellers	Take-down	Infiltration (undercover agent); Pseudo buys; Postal inspection;	~12, among which the two Silk Road administrators Ross Ulbricht & Blake Benthall	$3.6 million in bitcoin
Operation Onymous	2014	Europol; FBI	Hidden services (~27), of which 9 online anonymous markets (including Silk Road 2.0) – admins, moderators and top sellers	Take-down	Server attribution; TOR zero-day vulnerability	~17, mainly top sellers and administrators	$1 million in bitcoin
Operation Bayonet / Gravesac	2017	NHTCU; FBI	AlphaBay; Hansa Market – admins, moderators, top sellers, buyers; Special focus on decreasing trust	Take-over;	Server attribution; Market administration;	~4, the two admins and one top seller on Hansa Market and the admin of AlphaBay	$8 million in bitcoin (AlphaBay); € 2 million in bitcoin (Hansa)
-	2019	Europol; BKA	Wall Street Market; Valhalla – admins and top sellers	Take-down	Server attribution	~5, the three Wall Street admins and two top sellers	'6 digit amounts of Bitcoin and Monero'
-	2019	Europol; Guardia di Finanzia	Berlusconi Market – admins	Take-down	Pseudo buys; Postal inspection;	~3, including the two admins 'Emanuel Macron and Vladimir Putin'	€400.000 in bitcoins

From Table 6.1 we learn that we have witnessed five major law enforcement operations in the past 10 years. We note that US and European law enforcement agencies have been involved in all of them. Oftentimes, the intervention halted the operation of one or two markets. Only Operation Onymous, making use of a TOR-vulnerability, was able to take-down more than two markets simultaneously. Both technical means of server attribution, as traditional policing tactics have been employed to execute interventions. Curiously, although many users are active on the seized markets, in the direct aftermath of an intervention no more than 17 arrests were made. Since 2014, only admins and top sellers were arrested in major interventions. Last, we observe law enforcement agencies successfully seizing criminal profits – mostly the funds in escrow and personal admin funds – in the operation. Now, we move to analyzing which markets where targeted by law enforcement, the tactics used and if and how follow-up interventions where carried out.

TARGET SELECTION

Law enforcement agencies do not seem to randomly select markets to investigate or focus inventions on. Upon inspecting the now defunct website DeepDotWeb [7], which listed markets and ranked them, we uncover that nearly all markets that were subject to take-over and take-down in major operations, where listed as top markets shortly before their seizure [8]. Notable exceptions to this finding would be Valhalla – a regional Scandinavian market – taken down alongside Wall Street Market, and five smaller markets part of Operation Onymous.

TACTICS

Mapping these interventions on the enablers of online anonymous markets - i.e., anonymity and trust - we learn that all listed major interventions started by means of de-anonymiziation efforts. In two out of five interventions - which resulted in the take-down of the original Silk Road and Berlusconi - this de-anonymization was performed using traditional policing tactics, i.e., observations and the use of undercover agents. After the admins of the market, in both cases, where physically identified, linked to their anonymous online alias and arrested, law enforcement was able to use their admin credentials to shut-down the market.

In the other major interventions, de-anonymization was achieved through technical means. In one case, the police caught a lucky break, as admins made several small

[7] See `https://web.archive.org/web/*/www.deepdotweb.com`. Note that DeepDotWeb itself was seized by law enforcement as they were complicit in the online anonymous market ecoystem by posting links to markets in return for money.

[8] DeepDotWeb ranked 3 to 5 markets in their 'Top Market' list.

mistakes - of which the last connected admin payments to specific server infrastructure - enabling the take-down of the market. In Operation Onymous a TOR zero-day vulnerability was used to attribute server locations, thereby de-anonymizing markets and their admins, and subsequently shutting down the operation of nine online anonymous markets.

Last, Operation Bayonet stand out as this - again - started with a de-anonymization effort, but gradually focused on the other enabler of markets: trust. After a tip by a IT security firm on the server location of Hansa Market, law enforcement agencies were able to attribute the server location. In contrast to other interventions, they were just getting started. Only a short while after they attributed the server, they were able copy its content to a server of their own. Coincidentally, the admins were found by tracing payments to the identified servers, arrested and were cooperating fully - so that police could silently take-over their complete operation. From that moment onward, de-anonymization turned into maximizing distrust after the operation. Nearly all market-controlled mechanisms were reverted to play a part in the policing playbook. Multi-sig wallets were disabled, PGP-encryption was turned off, and all physical addresses where packages should be received, were kept and some high-risk packages seized in transit. All without anybody noticing. Nearly a month after the take-over, the operation - and Hansa Market - was shut down. The initial reactions by buyers and vendors on forums like Reddit or Dread was in line with what the police was hoping for: distrust in online anonymous markets.

We conclude that law enforcement agencies primarily rely on de-anonymization efforts to conduct major inventions. Both technical and traditional policing measures are used to de-anonymize. Intriguingly, market anonymity fuelled by the TOR-protocol seems not absolute. In some cases this is due to user mistakes - like we have seen exploited in Operation Bayonet or the take-down of Wall Street Market. In others - like Operation Onymous - a critical vulnerability in the protocol was exploited. Anonymity provided by cryptocurrencies or PGP-encryption, has not yet sparked major interventions. We do observe that payments - like admins paying server infrastructure - can lead to laying the last piece in an attribution puzzle. Likewise, PGP-encryption has not been a sole enabler of interventions. Yet, we do know that law enforcement makes use of the blind trust in this anonymization technology that is often provided by the market. In Operation Bayonet, we have seen that police agencies can pretend to enforce PGP-encrypted communication on the market, but in reality they could read every message written - as they had disabled encryption without anybody noticing.

FOLLOW-UP INTERVENTIONS

Besides the major law enforcement operations, we have summarized in Table 6.1, we have seen small-scale follow-up 'interventions' that actively informed people of the fact that, despite their assumption of being anonymous, were identified by the law enforcement. In three different ways, these users were informed. First, lists of identified usernames and parts of their real name and residence were published on a hidden service - i.e., onion-site managed by the Dutch police. Second, identified users are informed by letter, so-called 'love letters', to announce that the police was onto them. Third, 'knock-and-talks' were carried out, in which the police personally confronted users with their illegal behaviour. In all cases, no arrests nor seizures or prosecutions were made.

6.1.4. EVOLUTION IN ONLINE ANONYMOUS MARKET INTERVENTIONS

From the moment LEA became aware of the existence of online anonymous markets as a prominent meeting place and trading place for criminals, police forces around the world have carried out various operations to close these criminal organizations. In particular traditional policing tactics were used, such as market take-over and the interception of physical packages. This traditional policework is increasingly being combined with technical investigation techniques, such as the deployment of crawlers, scrapers, and the exploitation of server misconfigurations or other vulnerabilities, which are aimed at attribution the server location of markets. The use of a combination of traditional and technical investigation methods makes it possible to take over the infrastructure of these markets, to shut-out the administrators and to make the market either inaccessible or allow the market to be administrated for some period of time - while remaining in possession of the entire market administration on the seized servers. For example, the planning for the Hansa take-over – also known as operation Gravesac – started long before the summer of 2017. The idea for the operation came after a tip about the location of the server reached NHTCU. This tip led to a year-long investigation that eventually ended with the arrest of the administrators in Germany. This allowed the police – using the admin credentials – to migrate the seized servers in Lithuania to the Netherlands and start administrating the market as part of the police operation.

Based on the insights gained from earlier take downs - such as both Silk Road take downs - we know that they often result in the migration of users to other markets, where they continue their illegal activities [48]. Researchers describe this as a "waterbed effect", or crime displacement [94, 95]. In other words, the overall effect on the crime, e.g., the sales volume or number of active vendors on a market, in the entire ecosystem is temporary, and minor [139]. The Dutch police seemed determined to break with the

waterbed effect and change their intervention strategy to mitigate this unwanted side-effect – in this case turning it into an asset. During the first stages of Operation Bayonet, they intentionally used this effect – as they incentived users who were in the dark about of the sudden disappearance of AlphaBay to migrate to Hansa Market. In turn, the userbase of Hansa Market grew and more users than anticipated found their way to the market. Making the blow to the ecosystem potentially even bigger, as users would have felt tricked twice – first by the disappearance of AlphaBay, and second by being draw to Hansa were the Dutch police was 'lying in wait'.

Operation Bayonet, therefore, can be seen as the first major intervention aimed at disruption, rather than only attribution – i.e., arrest, seize, prosecute. While some users were arrested and the market's escrow funds were seized, the operation was primarily aimed at breaking trust in an anonymous setting. If you don't know who the market admins really are, it turns out that the police might be watching your every move, since they took over the market without anyone noticing. As a result, they hope vendors and buyers think twice about trading through anonymous market, disrupting their business.

Disruption as an 'alternative intervention' has been high on the agenda of the Dutch police and the Public Prosecution Service for some time, but since the Dutch Police announced its plans for the years to come, disruption was formalised and named a priority. The Ministry of Justice and Security also includes alternative interventions in 2020 as a performance indicator for the police in its assessment of the fight against digital crime. Since 2017, the Dutch police can be seen as a pioneer in alternative interventions aimed at online anonymous markets. But did this disruption tactic actually have a different effect than previous operations against online anonymous markets?

To see if and how we can evaluate this changing policing paradigm, we take Operation Bayonet as a case. In the remaining part of this chapter, we will investigate how we can leverage measurements in the online anonymous market ecosystem to evaluate this intervention.

6.2. LOST IN THE DREAM?

6.2.1. INTRODUCTION

In a coordinated effort, two leading online anonymous markets - AlphaBay and Hansa Market - were taken down by the Federal Bureau of Investigation (FBI) and the Dutch National High Tech Crime Unit (NHTCU) during Operation Bayonet [9]. The FBI managed to take down AlphaBay, while the NHTCU took over and operated Hansa market for

[9] See https://www.europol.europa.eu/newsroom/news/massive-blow-to-criminal-dark-web-activities-after-globally-coordinated-operation

nearly a month as administrator and thereafter shut down Hansa Market for good. The unexpected and unannounced implosion of AlphaBay left buyers and vendors in uncertainty and despair as the FBI - contrary to previous take-downs - remained completely silent about their involvement. Many AlphaBay users sought refuge to Hansa Market - which at that moment was operated by the NHTCU. Hence, the police agencies were in a perfect position to not only disrupt the ecosystem by creating distrust amongst users on these anonymous markets, but also collect valuable data on thousands of them. As the police agencies changed their intervention strategy distinctively, the question arises: did this intervention in turn result in a change in user strategy? Can we identify changes in behavior of vendors forced to migrate in the aftermath of Operation Bayonet?

In this paper we use measurements of the user-base of Dream Market to investigate the effects of the operation on all newly registered vendors on Dream Market (n=220) during and shortly after Operation Bayonet. To measure changes in user behavior, we identify where these vendors migrated from and observe changes in their 'appearance', i.e. a change in username or PGP-key. First however, we briefly look into Operation Bayonet itself.

In that operation, the FBI took down AlphaBay on July 5th 2017 and the Dutch police forces took down Hansa on July 21st 2017, while informing the world that they had been in full control of the site for 27 days. [10] In a bold move, the NHTCU had also been able to turn off the encryption of personal messages on Hansa, which allowed them to monitor personal information, like street-addresses, passing through the site. On the day of the Hansa take-down, the FBI announced to be responsible for the take-down of AlphaBay a month before. The FBI were able to seize multiple AlphaBay servers and arrest Alexandre Cazes - allegedly one of the administrators of Alphabay, known as Alpha02 - on July 5th 2017 in Thailand. Meanwhile, the planning of the Hansa take-down started long before and originated from a tip about the server's location. This tip led to a yearlong investigation, ultimately ending in the arrest of the administrators in Germany and the police being able to mirror the confiscated servers in Lithuania.

6.2.2. CRIME DISPLACEMENT

Leveraging the insights of previous take-downs - like the Silk Road 1.0 and 2.0 cases - we know that the typical result of a 'simple' take-down is that users migrate to other markets and simply carry on with their illegal business [48]. Researchers from Carnegie Mellon University [139] studied the overall trade-volume on online anonymous markets in the

[10] See https://www.politie.nl/en/news/2017/july/20/underground-hansa-market-taken-over-and-shut-down.html and https://www.justice.gov/opa/pr/alphabay-largest-online-dark-market-shut-down

ecosystem before, during and after both the Silk Road take-downs. Although the trade-volumes changed after both take-downs, they increased instead of decreased, reflecting a ecosystem that not merely recovers from an intervention but continues to grow regardless. Noteworthy are the insights of the sociologist Ladegaard [94] linking the media coverage of both take-downs to this increased sales-volume. Aside from any multiplication effect by extensive media coverage, take-downs often result in a so-called waterbed-effect, or the displacement of crime. In that sense we can draw upon the insights from crime displacement theory in the physical world and assess them in a digital environment [31]. Previously, Decary-Hetu [46] studied the effectiveness of interventions aimed at the warez scene - the hacker community specialized in distributing pirated material, like pirated movies - and concluded that crime displacement was one of the primary results of these interventions.Police agencies nowadays seem determined however, to break with this tradition and changed their intervention strategies to tackle precisely this unwanted side-effect.

Crime displacement is an effect that is frequently encountered in policing online crime [46]. In the case of online anonymous markets, it can be described as buyers and vendors moving on from one marketplace to the next, if one becomes unavailable - due to police interventions or an exit-scam [11]. Ironically, given the anonymous yet transparent nature of these markets, measuring this displacement has become easier as well. Current methodologies investigating the effect of online anonymous market interventions primarily look into the number of listings, number of users and sales-volume in order to determine effect sizes. This paper however, applies a new methodology to measure the impact of the different aspects of Operation Bayonet by not only looking into crime displacement, but also capturing specific vendor migration patterns and changes in vendor behavior during and immediately after the intervention.

6.2.3. MEASUREMENTS ON DREAM MARKET

To observe the initial effects - in terms of crime displacement - of Operation Bayonet, we study the user-base of another market: Dream Market, who became market leader right after Operation Bayonet. This market was established at the end of 2013 and has grown steadily ever since, making it a suitable market for our analysis as we can lever a baseline of pre Operation Bayonet operation.

From 2014 onwards we have scraped the forum of Dream Market, extracting - next to forum posts - its (number of) users. We made daily scrapes between January 2014 and

[11] An exit-scam is the sudden shut down of a market by its administrators, who take off with all funds in escrow. This could be severial millions worth of funds.

September 2017. We use the (number of) users on the Dream Market forum as a proxy for the (number of) users in the Dream Market community. Being active on a forum is not compulsory for trading on a market - so not all vendors are automatically registered on the forum - but is rather incentivized by the nature of online anonymous markets. Building a solid and verifiable reputation as a respectable vendor or honest buyer on a market, goes hand in hand with being active on a forum [47]. Specifically for vendors, reputation is an important part of doing business on an anonymous market. Dealing with all sorts of questions on the forum - ranging from product requests, to mishandled or seized shipments of drugs - be it addressed to individual vendors or not, helps grow ones reputation. That way, vendors apply similar tactics as in the legal economy: companies use approachable 'helpdesks' to increase the brand's reputation. Moreover, users connect their status on the market, to their status on the forum - i.e., a vendor on the market is recognizable as such on the forum. Meaning that we can discriminate between vendors and buyers via their respective status on the forum.

In each snapshot, i.e., scrape, we collect - among other things - the usernames and registration dates of active buyers and vendors from their individual user pages. Note that acquiring accurate registration dates - and not derived from first seen vendor activity in a certain scrape - is only possible through forums. Parsing the information on these pages gives us an aggregate of the total number of users. We can calculate the daily influx of users by taking the registration date as a time stamp and cumulate all new registrations on a certain date. The accuracy of the scraped information does not hinge completely on regular interval scrapes - which can prove difficult some days - because we can collect information, e.g. the registration date of a user, in retrospect from the individual user page. All in all, we are confident that our scraped data gives an accurate picture of the lower bound of users active in the Dream Market community between January 2014 and September 2017.

At the beginning of 2017 we measured that on Dream Market around 10,000 users were active. In terms of daily influx in the year 2017, Dream Market saw about 20 new users registering per day. That changed significantly from July 2017 onward. From that moment on, Dream Market started taking in more than 60 new users per day, with some days where even 180 new users registered (Figure 6.1). As a consequence, Dream Market nearly doubled its user base to almost 20,000 users in only nine months' time (Figure 6.2). Looking at the exact timing of this sudden rise in daily influx in July 2017, we can state that Operation Bayonet - where AlphaBay went down on July 4th 2017 and Hansa Market was shut down on July 20th 2017 - was probably the direct cause. Due to this rise in new users, Dream Market became the leading market right after the operation.

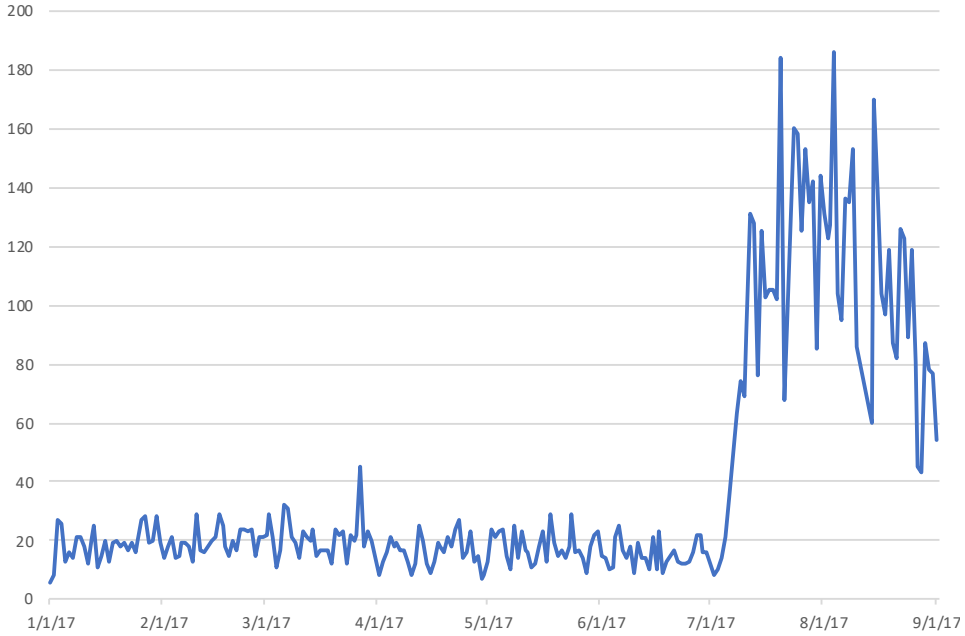

Figure 6.1: **Daily new users on Dream Market in 2017**

However, looking at the user-base of AlphaBay - according to US Attorney General Jeff Sessions over 40,000 vendors were selling to more than 200,000 buyers [12] - and Hansa Market - which had marginal presence in the ecosystem - prior to their take-down, certainly not all users migrated to Dream Market [13]. Yet, the increase of users is consistent with earlier take-down effects. To properly assess the detailed effects of Operation Bayonet on vendors migrating to Dream Market, we specifically look at all newly registered vendors on Dream Market (n=220) between July 1st and September 1st 2017. We investigate their background in terms of earlier presence on online anonymous markets. That way, we can identify specifics in crime displacement, i.e. vendor migration patterns to Dream Market, during and shortly after Operation Bayonet. Given the nature of the Hansa Market take-over and take-down, we expect less linkable vendor migration from Hansa Market to Dream Market.

[12] See https://www.justice.gov/opa/pr/alphabay-largest-online-dark-market-shut-down

[13] There can be multiple explanations for this relatively low figure of newly registered users on Dream Market. We do not know for instance, how many users had already registered at Dream Market prior to the AlphaBay and Hansa take-down. From a business continuity standpoint, it makes sense for vendors to spread risk and be active on multiple markets simultaneously.

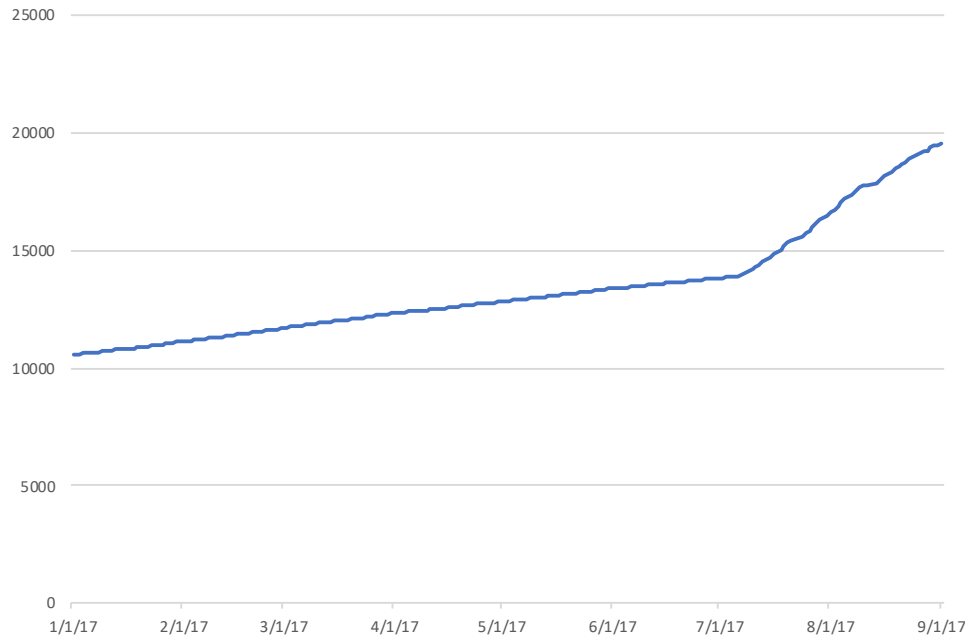

Figure 6.2: **Users on Dream Market in 2017**

6.2.4. MIGRATION PATTERNS

After obtaining the usernames of all vendors (n=220) that registered on Dream Market between July 1st and September 1st 2017, we used online anonymous market search engine *Grams*[14] to map specific (historic) characteristics of these vendors, for instance on which markets they were previously active. The search engine allowed 'informed customers' to track down vendors of products and services to assess their track record using previous sales and accompanied feedback. In turn, this allowed us to investigate the newly registered vendors on Dream Market and analyze their past and present behavior, i.e. their behavior before and after the intervention. *Grams* made it possible to search for vendors using either their username or PGP-key. For each vendor we executed a *Grams*-search with their Dream Market-username. The output of this search always was at least the combination 'username-market' of that user on Dream Market. Hence, we were able to validate our initial assumption that all 220 newly registered vendors were indeed active on the market and were not merely active on the forum. Next, the output of Grams would show us any other 'username-market' combinations that either use the same username

[14] On December 12th 2017 the administrator of *Grams* placed a message on Reddit announcing the discontinuation of Grams later that week. Fortunately, we finished our analysis before *Grams* became unavailable.

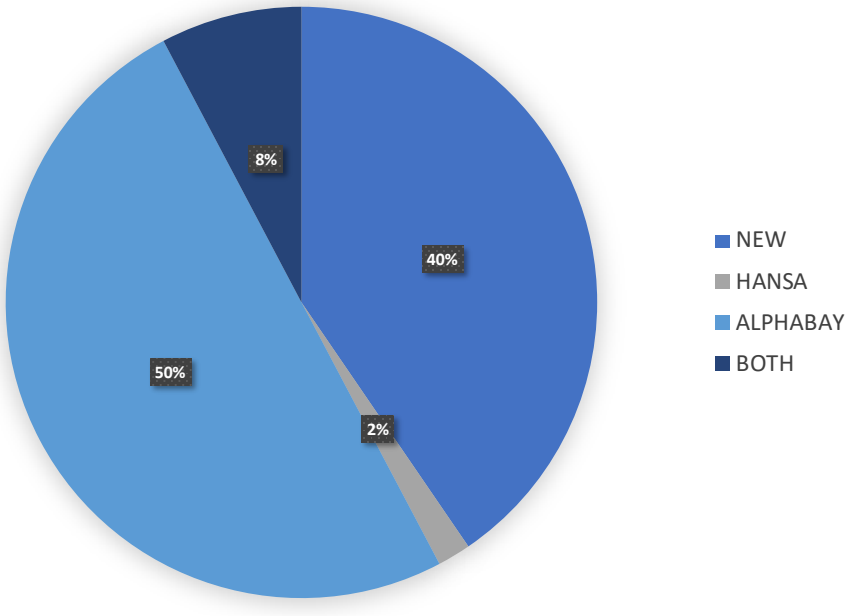

Figure 6.3: **Breakdown of newly registered vendors on Dream Market (*n*=220)**

or are connected through the same unique PGP-key. That way, we determined where the vendor migrated from: AlphaBay, Hansa Market, or that the vendor was active on both markets before migrating to Dream Market.

Figure 6.3 shows the breakdown of newly registered vendors on Dream Market. First, we can observe that many vendors migrating to Dream Market came from AlphaBay (40%) - at that time the largest market in the ecosystem. Curiously, the migration path from Hansa Market to Dream Market is near absent (2%). The latter is particularly interesting given the one major difference between the two take-downs in Operation Bayonet. Where AlphaBay was a take-down like many others, the Hansa take-down followed on nearly a month of complete control. This breakdown shows that there is a striking difference in migration patterns directly after that take-down. Second and rather unexpected, many of the newly registered vendors are completely 'new' and are without any previous reputation or track-record. This can mean two things: 1) real 'new' vendors picked this exact moment to start their online business and chose to do so on Dream Market or 2) vendors that were previously active on AlphaBay, Hansa or other markets took the rather drastic measure to completely start over - throwing away months or even years' worth of reputation and changing their identity by switching username and PGP-key. To investigate the effects of Operation Bayonet further, we look closer at the migrated

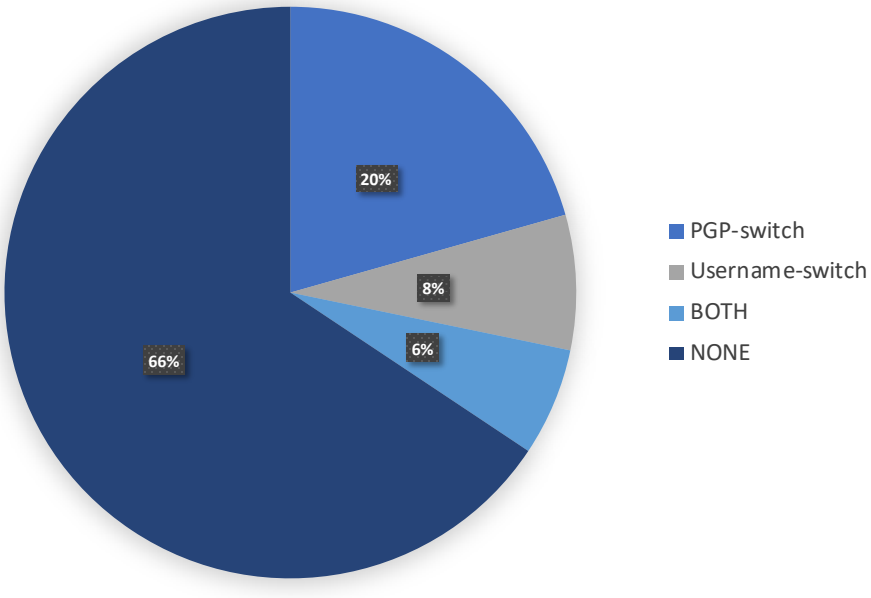

Figure 6.4: **Breakdown of evasive strategies of migrated vendors to Dream Market (*n*=131)**

vendors, so the 131 users that were active on AlphaBay, Hansa or both, as the question arises: did they put any effort into evasive measures after both take-downs?

6.2.5. VENDOR BEHAVIOR

To measure changes in vendor behavior in the group of migrated vendors (*n*=131) we turn to the online anonymous market search engine *Grams* again. Using the search engine, we identified vendors that changed usernames, but stuck to their PGP-key, or vendors that stuck to their username but changed PGP-keys. Because of the fact *Grams* uses both usernames and PGP-key to connect vendors, we leverage this output to see if the Dream Market username is the same as other usernames on other markets but has a different PGP-key. Or that the Dream Market username is different from earlier used usernames, but all have the same PGP-key connected to it. Figure 6.4 shows that two-thirds of the migrated users did not take any noticeable evasive measures. However, we can see that respectively 20% of users changed their PGP-keys, 8% changed their usernames and 6% did both. We were able to identify a handful (the 6%) of newly registered vendors on Dream Market who tried to start over completely - by changing both their username and PGP-key - but failed in some respect. For instance, they used the same e-mail address to register their new PGP-key as they used to register their old ones. Allowing us to

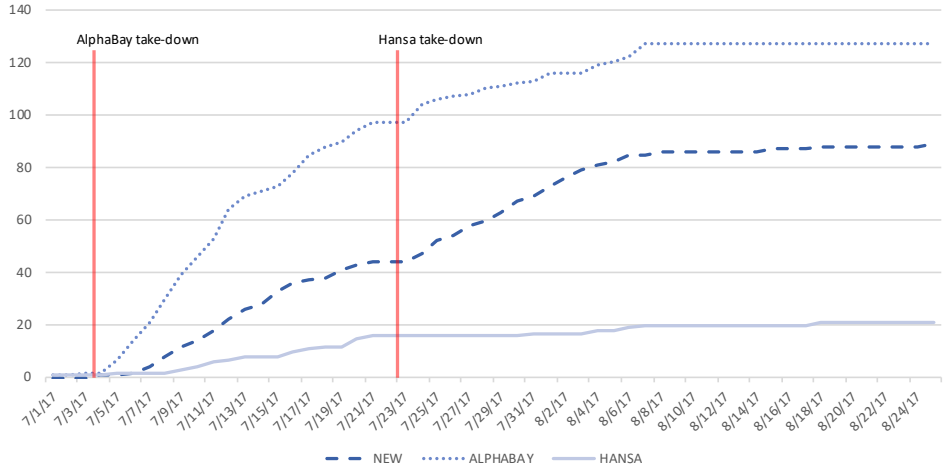

Figure 6.5: **Cumulative number of newly registered vendors on Dream Market per origin on date (*n*=220)**

deduce that these vendors at least tried to start over completely and provided us with the understanding that others might have successfully did so.

Both the number of evasive measures and the share of 'new' vendors are a strong indicator that this intervention has more than meets the eye. Knowing that a username and PGP-key are valuable assets in an anonymized setting - like underground markets - users do not change PGP-keys or usernames unless they really have to [12, 47]. Looking beyond the influx of users to Dream Market, one could see a scenario of a 'panicking' community or at least a community wherein vendors feel forced try to change their identity, be it with a new username, new PGP-key or even start over completely.

6.2.6. LONGITUDINAL ANALYSIS

We can assess this scenario even further by looking at these elements, i.e. the migration pattern and evasion measures, longitudinally. That way, we can see if the behavior of these vendors after the AlphaBay take-down differs from the Hansa take-over - where the police infiltrated, disabled encryption on personal messages and could see everything being said and done for three weeks without arising any suspicion. We expect that a 'simple' take-down would result in similar vendor behavior as reported by earlier studies [39, 48, 139]: migrate and carry-on. As the Hansa Market take-down coincided with the public statement of NHTCU that they operated the market for more than three weeks, and have gathered information about the true identity of thousands of users, we hypothesize that this would result in a different vendor response compared to 'simple' take-downs.

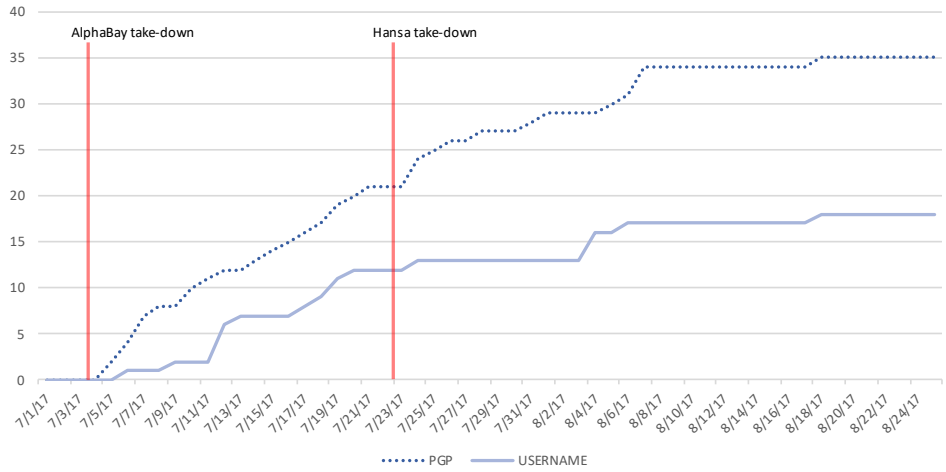

Figure 6.6: **Cumulative number of evasive measures by newly registered vendors on Dream Market on date (*n*=53)**

Figure 6.5 shows the cumulative number of newly registered vendors during our period of analysis in July and August 2017. Noticeably, the influx of AlphaBay migrants starts steep - right after the take-down on July 4th. The number of AlphaBay migrants stays relatively stable from the last week of July onwards. Even more interestingly, the number of 'new' vendors increases whilst the number of Hansa migrants stagnates at precisely the same time, namely around the 22nd of July: right after the Hansa Market take-down. This again builds to the scenario of a community seeking refuge to completely new identities, right after the Hansa take-down.

Looking at the evasive measures (Figure 6.6), this scenario finds more support. Right after the Hansa take-down, the username-switch stagnates. In turn, migrating vendors apparently turn to a more drastic measure: starting over.

6.2.7. DISCUSSION

In this paper we presented a new methodology to discern the different effect-types of online anonymous markets interventions. By taking into account the (longitudinal) changes in behavior on a vendor-level, i.e. specific migration patterns and changing vendor behavior - we were able to see beyond the waterbed-effect for the first time.

We have to stress however, that the methodology and measurements in this paper have some limitations. First, our methodology is partly based on a third-party search engine to make connections between vendors across markets. Leaning on that defunct

service to identify changing vendor behavior, in terms of migration patterns and evasive measures, means that replicating our findings - using that same service - has become rather impossible. Second, our measurements contain data on users of the Dream Market forum - not the market - for a long period of time before, but only depict a relatively short window after, Operation Bayonet. This could mean that we only witness the first and not final effects of this intervention. Third, as we employ a novel methodology to measure changes in vendor behavior, we cannot compare these results one-on-one with previous research efforts into police interventions. Hence, more research efforts therefore should be taken to a) investigate the long-term effects of this intervention in terms of crime displacement and changes in vendor behavior, b) how to identify migration patterns of vendors across markets and c) leveraging the before mentioned research efforts to investigate previous and/or future interventions using this novel methodology to better compare the effects of these respective operations.

Notwithstanding these limitations, if we apply our methodology to measure the effects of Operation Bayonet on migrating vendors to Dream Market, we see the first signs of a game-changing police intervention. Compared to 'simple' take-downs, like the AlphaBay take-down, the Hansa Market take-down stands out - in a positive way the police might add, as users do not just move along after the Hansa Market shutdown. Few simply migrate, some take evasive measure - like changing their username and/or PGP-key - but many start with a clean slate on Dream Market. This may sound as a minor detail, but the opposite is the case. When a vendor starts over he/she loses their track-record, reputation and customer-base. Like a Michelin star restaurant moving cities, whilst changing its name, website and phone-number: nobody will recognize the fancy restaurant from before and the chef will likely be forced to start (re)building a reputation from scratch. We have to see if the effects of this innovative intervention hold in the long run, but for now the initial effects are remarkable in the light of earlier interventions aimed at online anonymous markets.

7

MEASURING INTERVENTIONS

The previous chapters in this dissertation have shown how online anonymous markets play an important role in fulfilling resources in value chains of profit-driven cybercrime. We see that – although supply and demand are not completely matched – most forms of profit-driven cybercrime can make use of the potential for outsourcing that an online anonymous market provides. Leveraging measurements in the ecosystem, we have investigated commoditizaton of cybercrime components, the predictors for product success and how one particular law enforcement intervention impacted market users. We have not however touched upon how measurements can help evaluate policing practice. That is, how do we leverage measurements of online anonymous markets to evaluate interventions? [1]

7.1. INTRODUCTION

Since their inception, online anonymous markets, such as Silk Road, Hansa Market and Alphabay, have been of interest to both the research and law enforcement community. International policing efforts have resulted in the take-down of many prominent markets. Coordinated interventions, as Operation Onymous or Operation Bayonet can be seen as examples of an evolving law enforcement strategy. 'Simple' take-downs are integrated in a sophisticated strategy aimed at breaking trust in the underground market ecosystem, by taking over and infiltrating prominent markets. But how do we know if these novel

[1] Sections 7.1 – 7.4 are based on parts of the paper "Schwedersky, A., Van Wegberg, R.S., & Verburgh, T. (under review). 'Victory versus fiasco: A framework to measure the effects of online anonymous market interventions'."

strategies yielded any impact and can be seen as successful? And, on what indicators do we measure success?

Over the recent years the underground ecosystem seems to be quite robust when confronted with law enforcement interventions or other significant events, i.e., exit-scams or animosity between rivaling markets leading to DDoS-attacks [139]. Studying online anonymous market interventions is a complex endeavor. An online anonymous market's structure, the tactics used during law enforcement interventions and the online anonymous market ecosystem all evolve fast and are hardly ever constant. Furthermore, the methods, background and approaches used by researchers in order to study the impact of interventions, also differ from case to case. For instance, due to data availability or the chosen research methods, different variables or proxies might be chosen to measure the effects of online anonymous market interventions. Yet, studying the impact of law enforcement interventions is essential for both law enforcement agencies and policy makers. Law enforcement agencies will be able to improve the efficiency of their interventions when taking insights on historic interventions into account. Policy makers can take these insights to asses current policing strategies objectively and independently.

Recent work has identified different ways in which to investigate and evaluate the impact of online anonymous market interventions [48, 94, 95, 153]. Due to the variety of methodologies used and differences in the types of research, a common standard to study the effects of interventions remains absent. In this chapter we try to fill that gap. We synthesize the state-of-the-art in online anonymous market intervention studies and present best practices to measure the impact of interventions. We provide indicators and constructs to adequately measure the effects of interventions and allow LEA to design new interventions based on proven historic impact – working towards evidence-based interventions.

The rest of this chapter is structured as follows. First, section 7.2 provides context to measurements in the online anonymous market ecosystem. Thereafter, we synthesize the state-of-the-art in intervention studies and identify common standards to measure interventions on online anonymous market in section 7.3. Last, section 7.4 discusses in-depth how these measurements map to known aims and tactics of past interventions and presents suggestions for novel measurements of future interventions.

7.2. MEASUREMENTS OF ONLINE ANONYMOUS MARKETS

Over the past few years several studies have investigated the impact of online anonymous market interventions [48, 94, 95, 139, 153]. To synthesize the state-of-the-art in online anonymous market intervention studies, we investigate core elements of the aforemen-

tioned research efforts in-depth. First however, we briefly introduce these intervention studies one by one and discuss their research design, measurement and analysis methodology, and results.

Using longitudinal measurements on five online anonymous markets (Agora, Cloud 9, Evolution, Hydra and Silk Road 2.0), Décary-Hétu and Giommoni [48] studied the impact of law enforcement interventions in the time-frame 2014-2015. For their analysis, they used price, listings, new vendors, active vendors, and feedbacks, in the given time-frame. Based on the analysis of the changes in these indicators over time, they conclude that online anonymous market users remain active after law enforcement interventions. They consider the effects of Operation Onymous to be limited in terms of the impact on price, supply and consumption of traded products.

Ladegaard [94] analyzed feedback data of two online anonymous markets – Evolution and Agora – based of daily scrapes over a ten-month period in the time-frame 2014-2015. He investigated if and how market revenue was affected by media coverage on online anonymous market interventions, as well as the highly publicized conviction of Dread Pirate Roberts – the administrator of the original Silk Road. Contrary to the anticipated deterrence effects – i.e., both events will reduce crime and trade – the results show that trade increased, instead of decreased. In conclusion, deterrence is not a likely effect of the online anonymous market interventions in that time-frame.

A second and more recent study by Ladegaard [95] discusses the question how crime displacement can be observed in online crime – in particular in the online anonymous market ecosystem. Leveraging scraped data on Silk Road 2, Evolution and Agora from October 2014 to September 2015 – and combining this with publicly available data such as DeepDotWeb archives - his results indicate that law enforcement interventions impact the weekly revenue of vendors negatively, yet temporarily. After online anonymous market arrests and take-downs, drug trade displaces to other vendors and to other, or novel online anonymous markets. Furthermore, his results indicate that the number of online anonymous markets is quite stable over time. After an intervention, new markets emerge in the same ratio compared to the markets taken down. The relatively stable number of markets is a function of platform economics in the online anonymous market ecosystem – where the critical mass of vendors and buyers makes only a certain number of markets feasible. This underlines the temporary effects of interventions and resilience of the ecosystem, in terms of the number of markets before and after an intervention.

Soska and Christin [139] present a longitudinal analysis on the online anonymous market ecosystem. Meticulously scraping data on 16 different online anonymous markets over a long period in time (2013-2015) the evolution of the online anonymous market

ecosystem is captured in terms of market turnover. Their results depict the fluctuation of the daily transaction volume during 'events' such as the take-down of Silk Road and Operation Onymous, and exit-scams. Yet, their results suggest that - within a few months - trust in the ecosystem seems restored and markets reach an equilibrium. This indicates that the online anonymous market ecosystem is remarkably resilient to take-downs and scams.

Van Wegberg and Verburgh [153] analyzed the user-base of Dream Market to assess the effects of Operation Bayonet. By investigating all newly registered vendors on Dream Market (n=220) shortly before and after the intervention, they identify migration patterns and detect changes in vendor behavior. This set of newly registered vendors is acquired through daily Dream Market forum scrapes between 2014 to 2017 – which also serves as the market's user-base baseline to compare against. Results show that while there are a lot of vendors who simply migrate from AlphaBay or Hansa Market to Dream Market, a large portion thereof does take the trouble to change their PGP-key and/or their username. This makes the effects of the intervention stand out in a positive way, as this means that vendors have opted to change their username, PGP-key or both. That way, vendors have accepted that buyers cannot identify them as the 'reputable vendors' they once were on AlphaBay or Hansa Market. Logically, this would have a negative impact on their trade volume for a period of time after their registration on Dream Market (see Chapter 6).

Based on this brief overview we can already see a lot of differences in the approach towards studying the effects of online anonymous market interventions. In general, we see that measuring the impact of interventions often happens on a high level, such as the impact on total revenue in the ecosystem. Yet, many interventions consist of numerous aspects, which can be measured on as many levels. Differentiating the impact for these aspects is difficult but essential, as it might change the insights into the overall impact. For example, the Hansa take-down consisted of several aspects such as a media strategy, the infiltration and administration of Hansa Market, the take-down of Hansa Market and shutting off parts of the encryption. Each of those aspects might interfere with another. As an example, it has been hypothesized that the enormous media attention after the initial Silk Road 1.0 take-down did not achieve a deterrence effect, but instead had a promotional effect. People who previously were unaware of the existence of online anonymous markets might have been drawn towards them, because of the media attention [94]. In an ideal situation we want to be able to differentiate the impact of each aspect of an intervention, to be able to improve interventions or create new ones. A first step is to create an understanding of the state-of-the-art and turn that into a framework to measure online anonymous market interventions.

7.3. SYNTHESIZING THE STATE-OF-THE-ART

Although measuring the effects of interventions on online anonymous markets is complex, we have identified five studies that have reported on the impact of (an) intervention(s). We synthesize this state-of-the-art based on two aspects: a) the research approach of the study and b) the analysis used. We elaborate on these two aspects for all papers, and find their common ground.

7.3.1. RESEARCH APPROACH

The first and maybe even most important take-away from previous studies, is to use a longitudinal research design. Collecting data over a relatively long period of time seems crucial, especially when aiming to identify and study changes – like sales volume or number of users. Changes in behaviour of the online anonymous market community range from enormous and sudden, like market shutdowns or show-ups, to being unable to see changes with the naked eye, such as PGP key changes [153]. As some of these changes in user behaviour only show in due time, a longitudinal research design is necessary. For example, Décary-Hétu and Giommoni [48] suggest that after an intervention, it is possible that vendors wait till the dust settles and only thereafter start selling again. The data collection periods of the state-of-the-art ranges from nine months [153] to two-and-a-half years [139]. A relatively long period of time can provide more certainty when measuring changes of any kind.

In addition to the data collection period, we identify a pro-active approach to be an important aspect from the state-of-the-art. Online anonymous markets are quite unstable in terms of up-time [139]. Some markets are just hard to reach, others disappear suddenly at an unexpected moment in time. In contrast to the traditional research approach - where research questions are formulated before gathering data – data collection might be the first step. This traditional approach might even impose severe limitations on the potential of the intervention study. In order to overcome this limitation a more pro-active approach is required.

Knowing data collection should be carried out over a longer period of time and handling a pro-active method doing this – we now want to elaborate some lessons the data collection itself. In contrast to more a traditional data collection approach, like using surveys, a more technical approach is necessary. Namely, the method of crawling, scraping and parsing data. The state-of-the-art relies predominantly on these methods. Some researchers build their own tools [94, 139] where others rely on public tools or databases [48]. Combinations are also common – Van Wegberg and Verburgh [153] and Ladegaard [95] combine their own tools with publicly available datasets. The first lesson

is that combining data leads to more data points and validation possibilities. Yet, datasets can become voluminous. For example, the dataset of Soska & Christin [139] consists of almost 2000 online anonymous market scrapes resulting in a dataset of 3,2 TB. Soska and Christin [139] show that it is in fact possible to adequately analyse this amount of data in a structured way. Note that this reflects a high level of understanding in both the online anonymous market ecosystem as well as in data collection and analysis. The second lesson we identify is scraping in a stealthy manner [139]. By that we mean that researchers do not want to alert the online anonymous markets and its actors that data is being collected, as this might lead to being denied access to the markets or influences the data collection in other ways.

Considering the fact that scraping and crawling is the common standard in in the state-of-the-art, parsing this data can be identified as an adjacent common practice. In parsing, raw data is transformed into structured data, e.g., listings, usernames, prices. This leads to the question of how to extract entities – for instance usernames of vendors and buyers. Despite the fact that entity determination remains complex, it is an essential element of studying the effects of online anonymous market interventions. For the determination of entities on online anonymous markets and across the online anonymous market ecosystem, different methods can be used. Most methods however include the use of username and/or PGP-key as key attributes [48, 139, 143].

7.3.2. DATA ANALYSIS

For research in the online anonymous market area, the use of proxy variables is inevitable. For example, when measuring the market revenue. Scientifically, the ideal situation would be to use the price and number of sales per listing. Unfortunately, this data is almost never available in scrapes of the front-end of a market. As a proxy, other metrics can be used. Instead of the number of sales, the number of feedbacks can be used [13, 39, 95, 139, 152]. When using the feedback system as a proxy, the estimate of sales is lower than when the actual number of sales is used since not all transactions result in a feedback. This is why researchers oftentimes refer to feedback as a structural underestimation, or lower-bound of the amount of sales. Furthermore, this proxy does not only affect the interpretation of the results but also the method of data collection, pre-processing and analysis. This is because the available information regarding the feedback differs between online anonymous markets. Some online anonymous markets do not show feedback per sale, but only the feedback total per vendor.

Next, the time-frame of the scrape. As scraped data always shows something in the past, this is an important aspect. For example, when analyzing prices. The value of

cryptocurrencies is subject to change and when analyzing prices over a longer period of time, this might cause problems. Prices should be converted to the same currency at the right point in time. Another aspect that needs to be considered are the so-called holding price, which is related to the supply of a certain product by a vendor. When a product is out of stock and the vendor wishes to halt sales, an extremely high price is applied. That way the vendor does not need to remove the listing and keeps the associated feedbacks – thereby not impacting his/her reputation [139]. Also, other factors associated to the element time, might bias the dataset. For example, when data collection covers a period wherein a central online anonymous market actor is arrested [94, 95]. This could also be the case with other events such as detected law enforcement activity on online anonymous markets.

Another aspect is the number of users. The total number of users can paint a distorted picture since researchers also create numerous accounts on online anonymous markets to either access the content of the market or to be able to scrape markets. Next to researchers, law enforcement and fascinated individuals also register on online anonymous markets and thereby bias the amount of actual active users. Next to this, after a shutdown or intervention, bulk migrations or registrations can be expected and need to be considered to explain sudden influxes of users. Also note that one account is not equal to one person. One person can hold one account or ten accounts, and one account can be managed by a group.

7.4. LESSONS LEARNED

In this section we introduce our synthesis of best practices to measure the effects of interventions on online anonymous markets. The synthesis is based on our analysis of the state-of-the-art in online anonymous market intervention studies, more specifically their research design, data collection, data pre-processing and data analysis. We structure our overview of best practices along the lines of constructs and measurements. We define a measurement as an approach to capture a certain construct. For example, the number of listings is one of the ways to capture the size of the market. Similarly, revenue also captures the market size, but based on a different variable – population versus a monetary value.

Measurements can be used in a micro, meso or macro context. This reflects the different levels of effect that interventions generate. A micro level effect, for example the arrest of one vendor, has an initial impact on just one individual. Meso level effects, for example the take-down of one online anonymous market, are effects at the market level. Effects on the entire online anonymous market ecosystem, we call macro level effects.

For instance, Operation Bayonet can be considered an intervention with macro-level effects since the effects of the intervention can be identified and measured throughout the ecosystem, i.e., on a macro level.

We cover three different constructs that arise from the synthesis of the state-of-the-art in online anonymous market intervention studies. These constructs are the following: (1) market size, (2) product supply and demand, and (3) business continuity. In the remainder of this section, we try to provide a clear description of each construct.

MARKET SIZE

The first construct is market size. This can be the size of an online anonymous market, or the relative size of a market or several markets within the online anonymous market ecosystem - depending on the impact level one is analyzing. Studying the effects of an online anonymous market intervention using the size of a market – or multiple markets – as a measurement methodology is a common practice and therefore is often used as a first vantage point. Market size can be determined by measurements such as the amount of listing, user or (active) vendor volume and the overall revenue of the market. Next, we can assess a market's position in the ecosystem - which can be determined by measurements such as the number of sales, users, or vendors relative to other markets in a certain time-frame.

PRODUCT SUPPLY & DEMAND

An important element of online anonymous markets are the products traded. The supply of products on a market is a parameter to assess the scope of the market. For instance, is the market one that supplies drugs and digital goods, or does is specialize in one of the two? The type of products offered also partly determines the potential interest of law enforcement. Sometimes this might even be the focal point of an intervention. For example, when law enforcement agencies wish to halt sales of a product like fentanyl. Several measurements are important here. We identify the following: listings per product type, average price per product and sales per product. This of course can be used to study the popularity of a certain product, but it can also be used to get a better understanding of the products that are popular compared to others on a certain online anonymous market or even across the online anonymous market ecosystem. The demand side of online anonymous markets is also of interest when studying the effects of interventions. Demand is strongly related to the popularity of products, the market itself and the active vendor volume. By investigating wanted products we can learn more about the demand side. For example: is a certain product really popular but not available on most markets? Or is a product only available on specific markets?

BUSINESS CONTINUITY

This construct comprises the strategies that market owners and its users employ in order to assure their business to run continuously. For example, vendors apply different strategies when migrating to new or other markets. Some vendors use the same username and/or PGP-key, others start over completely – erasing their reputation and starting with a clean slate. Reputation management is therefore the first measurement to consider. Other measurements include consumer confidence and risk mitigation mechanisms. Respectively, reported scams – i.e., disputes – on the market and mechanisms to mitigate these. Risk mitigation strategies can be identified by looking at feedback, escrow and/or multi-sig payment systems. Leveraging this information allows investigating business continuity of either the market administrators or its vendors. Note that, these elements also interact with each other. For example, without a feedback system or the possibility to use escrow, there is probably less consumer confidence – and more reported scams – on the market.

7.5. FROM MEASUREMENTS TO EVALUATING INTERVENTIONS

In this chapter we have explored how the current measurement approaches create insights into the impact of online anonymous markets interventions. The state-of-the-art contains measurements on three elements: a) market size, b) product supply/demand and c) business continuity. In this section we will investigate how these map to known aims and tactics of past interventions (see Chapter 6) and present directions for novel measurements to help evaluate future interventions.

De-anonymization tactics by law enforcement agencies have produced market take-downs and take-overs, as well as the arrest of (top) sellers and administrators. To assess the effects of a market take-down on the ecosystem, measurements of market size, product supply and business continuity by actors do suffice to a large extent. Initial effects of interventions in terms of the overall volume of crime can be adequately tracked by leveraging measurements of market size - e.g., amount of markets, listings, vendors, etc. This way, we can evaluate if an intervention has impacted the overall trade in the ecosystem, or that the intervention had a short impact and the overall trade volume bounced back to pre-intervention transaction levels. Similarly, tracking the evolution of product portfolios - i.e., product supply - provides a clear view on the impact of an intervention on a specific market segment - e.g., B2B cybercrime products. Last, to assess the impact of an intervention on individual users, measurements of business continuity - like tracking migrating patterns across markets, making use of public PGP-keys as a unique identifier (see Chapter 6) - create insights into crime displacement and the level

of desistance. Displacement here refers to users migrating from market to market. In contrast to desistance, where we refer to users desisting from criminal activity after an intervention. Both displacement and desistance form a valuable outlook to evaluate an intervention, beyond the markets and funds seized and the arrested actors.

Next, these measurements can also highlight certain (un)expected effects of law enforcement prioritization of policies and even suggest new avenues for 'alternative' interventions. For instance, the continued and coordinated fight against child sexual abuse material - and the policing measures that are allowed to be used in this fight - has resulted in a widespread ban of this material on online anonymous markets. Likewise, the focus of US law enforcement on countering the trade of the drug fentanyl, coincides with the scarcity of this drug on many if not all markets. This way, measurements of product supply reveal shifts in product portfolio across markets. Tying these shifts to concrete changes in police priorities and practices, can even help shape new interventions options.

Yet, measurements can be deceiving and even actively deceived. First, we note that when simply counting the number of users on a market, one has to understand what this measurement captures. Everyone from security researchers, to fascinated individuals have registered one or multiple accounts on these markets. This results in an overestimation of market size - since simply counting users not only includes those who trade on the market, but also encompasses a significant amount of innocent bystanders. A better approach therefore should be to focus on the number of active users - vendors and buyers with at least one completed transaction. Second, we know that some markets provide vendors with the option to delete feedbacks after 30 days. This way, vendors can seem less active then in reality. Especially since researchers - by lack of back-end data on actual transactions - use feedbacks as a proxy for sales. Note therefore, that the number of feedbacks displayed on the market does not have to match the number of feedbacks a vendor has received. Third, measurements oftentimes concern users - e.g., vendors. We do not know however if one vendor maps to one individual. It can be the case that one criminal entrepreneur holds multiple vendor accounts on or across multiple markets without a way to link these. In reverse, one vendor account can be used by a group of criminal entrepreneurs. In any case, we should make clear that we are measuring users, not individuals.

Tactics aimed at creating distrust in the online anonymous market ecosystem, are somewhat harder to evaluate with current measurements. Ideally, one would like to get an impression about the level of distrust that is generated by intervention. For instance, actors being cautious to start or continue to do business. In other words, how do we measure an actor's confidence in the anonymzation technologies and trust systems a

market makes use of. We can of course measure the amount of reported scams on a market or count the amount of exit-scams. But in essence, this does not tell you if this is business-as-usual or not. Let alone, that we can distinguish the effects of an intervention versus other external events, like DDoS-attacks that lead to markets being unreachable - which in turn might generate distrust in a certain market or the entire ecosystem. Other platforms, like Reddit or Dread, do provide anecdotal evidence on the level of trust by actors in the ecosystem. Yet, these feelings of (dis)trust seem all but robust measurements. As they certainly do not reflect the feelings of (dis)trust throughout the community.

We see that current measurement studies have sufficient potential to help evaluate interventions that rely de-anonymization tactics. In contrast, interventions that aim to maximize distrust, can only be partly evaluate based on current measurement approaches. However, there are ways - that do not rely on anecdotal posts of users on Dread or Reddit - to measure trust in the online anonymous market community: to look beyond them. The fragmentation of criminal activity relying on anonymization technologies and trust mechanisms is an important indicator to evaluate an intervention a certain platform. For instance, we know that some prominent vendors on online anonymous market - reportedly being fed up with moving from market to market - have decided to move their activities to a so-called single-vendor shop. Measuring the amount of these shops provides an indicator to assess the amount of trust professional vendors have towards the standardized markets. Likewise, certain criminal entrepreneurs have started moving their business to encrypted messaging platforms - like Telegram. Similarly, measurements on Telegram channels that contain cybercriminal activity should present an angle to evaluate the impact of interventions elsewhere in the cybercrime ecosystem.

8

CONCLUSION

This dissertation studied the phenomenon of profit-driven cybercrime by reconstructing value chains and cybercriminal business models. We have presented five peer-reviewed studies (see Chapters 2 - 6), that together with Chapter 7, aimed to investigate how outsourcing is enabled by online anonymous markets and how outsourcing can be disrupted. The focus of this dissertation lies in understanding the outsourcing potential in profit-driven cybercrime by examining the role of online anonymous markets in supplying commoditized cybercrime components. We set out to answer the following main research question:

How do online anonymous markets facilitate the outsourcing of cybercrime components in profit-driven cybercrime value chains?

We started by mapping value chains and business models, leveraging financial malware as a case study in Chapter 2.

CHAPTER 2 – VALUE CHAINS

The goal of this chapter is to discern value chains in financial malware. These value chains were constructed in a similar fashion as other researchers reconstructed the spam value chain. The constructed value chains allowed us to analyze the economic principles within cybercriminal business models. This resulted in the understanding which elements of a financial malware scheme are most suitable to be either vertically integrated or

fulfilled through the underground market. We demonstrated that, for financial malware schemes using man-in-the-browser attack vectors, there is a clear incentive to outsource cybercrime components via underground markets.

We have demonstrated that a value chain approach is evidently useful when researching financial malware or other cybercriminal business models. Next, this creates the opportunity to study the important interactions between the strategies of attackers on the one side and the properties of cybercrime components and policing tactics on the other.

8.1. EMPIRICAL FINDINGS

CHAPTER 3 – COMMODITIZATION OF CYBERCRIME COMPONENTS

In this chapter we studied the phenomenon of commoditization of cybercrime. Drawing on Transaction Cost Economics, we argued that anonymous markets are a good way to study this phenomenon. Using six years of longitudinal scraped data from eight online anonymous markets to measure the evolution of commoditized offerings for the dominant criminal value chains, as well as the volume of transactions and revenue of actual sales to criminal entrepreneurs. Criminal suppliers can commoditize their offerings, using these marketplaces since they provide a wide reach and numerous risk management mechanisms.

In conclusion, we find that, at least on online anonymous marketplaces, commoditization is a spottier phenomenon than was previously assumed. Within the niches where it flourishes, we do observe growth. That being said, there is no supply for many of the capabilities, systems and resources observed in well-known value chains. There is also no evidence of a rapid growth, and thus of a strong push towards commoditization, contrary to the somewhat alarmist language found in industry reporting and elsewhere.

CHAPTER 4 – CYBERCRIME CASH-OUT THROUGH BITCOIN MIXING

This chapter aimed to examine ways in which cybercrime proceeds can be laundered. Building on earlier work, we focused particularly on bitcoin money laundering services – i.e., bitcoin mixers – offered in the underground economy.

To examine bitcoin laundering, we used bitcoin mixing services and exchange services and integrated these services in a cash-out experiment. We examined the usability of these services in a cash-out strategy by analyzing the service percentages and reputation mechanisms. This aided us in determining the likeliness of integration of these bitcoin laundering services in an actual criminal scheme. The results of our experiment suggest that laundering cybercrime proceeds using bitcoin is a user-friendly and working criminal service-model. However, it is not clear whether the model will work when larger amounts

of money are laundered. Yet, for smaller amounts it clearly offers an easy to use and good value-for-money service, as long as criminals keep an eye out for scams.

We conclude that bitcoin money laundering is a practically conceivable concept and likely to be integrated in current-day and future money laundering schemes. The ability to lower the cost of laundering, whilst providing more anonymity, make it an interesting money laundering technique for criminals.

CHAPTER 5 – OUTSOURCING VIA ONLINE ANONYMOUS MARKETS

In this chapter we investigated the performance of products in the business-to-business cybercrime market segments on AlphaBay. As we know that not all products and vendors are equally successful on the market, we aim to predict which characteristics of both the criminal entrepreneur and the product they sell, influences the performance of cybercrime solutions. To that end, we constructed new variables to grasp the relative price, functional features and the marketing of the product: so-called product differentiators.

First, we find that all vendor profiles – which capture the characteristics of the vendor – influence cybercrime sales. Second, in line with what other researchers have observed on carding forums, we identify particular functional features, i.e., refund policy and customer support, to be positively and significantly correlated with the performance of a cybercrime solution. Third, we show that marketing the product, in terms of using capitals in the title to attract attention when browsing the market, influences the sales numbers of a cybercrime solution in a positive way. Likewise, branding a product, i.e., using the vendor's name in the title, increases the performance of the product.

Overall, our findings show that the profile of a professional criminal entrepreneur is able to predict a relative high degree of variance in the performance of cybercrime solutions. Similarly, signals of a professional operational – product branding, refund policies and customer support, explain a significant portion of the variance in performance. Yet, the vendor of a product – rather than its nature or features – is the most powerful predictor for its success.

CHAPTER 6 – INTERVENTIONS

This chapter presented a new methodology to discern the different effect-types of online anonymous markets interventions. By considering the longitudinal changes in behavior on a vendor-level, i.e. specific migration patterns and changing vendor behavior, like username and PGP-key changes - we were able to see past the waterbed-effect for the first time.

Our results suggest that Operation Bayonet is to be considered a game-changing police intervention. Compared to 'simple' take-downs, like the AlphaBay take-down, the

Hansa Market take-over stands out – as users do not just move along after the Hansa Market shutdown. Few simply migrate, some take evasive measure - like changing their username and/or PGP-key - but many start with a clean slate on Dream Market. This may sound as a minor detail, but the opposite is the case. When a vendor starts over that likely results in the loss of their historic track-record – i.e., reputation and customer-base. Inevitable, this means that vendors have to build their reputation and customer-base from scratch, as they cannot refer back to their previous activity. This would render the effect of the clean slate useless. One could imagine that a clean slate impacts business in a negative way – at least in the short run. We have to see if the effects of this innovative intervention hold in the long run, but for now the initial effects are remarkable in the light of earlier interventions aimed at online anonymous markets.

8.2. COMMODITIZATION OF CYBERCRIME

All the empirical findings in this dissertation together, cast - next to their intrinsic insights - a perspective of the phenomenon of cybercrime commoditization. We identify three main reflections on the commoditization of cybercrime from this dissertation that fill the aforementioned research gaps (see Section 1.2).

First, we have uncovered that profit-driven cybercrimes rely on multiple, common technical capabilities - i.e., cybercrime components. The value chains of resources that enable these profit-driven cybercrimes therefore overlap to a large extent. This means that certain technical capabilities - like bullet-proof hosting, compromised websites or cash-out solutions - serve as crucial catalysts to multiple profit-driven cybercrimes. In turn, we observe that these capabilities also see more commoditization. One the one hand, aspiring cybercriminals know that these capabilities are essential enablers of their business model - growing the demand-side of the market in these commodities. On the other, it is possible to produce these capabilities at scale. Moreover, they are well-know enough to sell them in a one-off transaction.

Second, we have found that online anonymous markets are able to supply cybercrime components on a standardized, one-shot platform. Next to for instance carding forums and Telegram channels, we have found an additional distribution channel for cybercrime commodities. Yet, compared to forums and encrypted messaging platforms, online anonymous markets have a different impact on the commoditization of cybercrime. Since forums and messaging platforms rely on more interaction between buyer and vendor - and thus technical offender skill - online anonymous markets have lowered the knowledge threshold for acquiring cybercrime commodities substantially. Furthermore, we know that forums and encrypted messaging platforms are places where specialized vendors

trade cybercrime components - take carding forums as an example. In contrast, on online anonymous markets we see a broader product portfolio, where next to a hand full of 'specialists' selling just one capability, mostly 'generalists' supply a diverse range of cybercrime components. It seems online anonymous markets are home to vendors playing into the simultaneous demand for more than one capability - serving as professional facilitators in the profit-driven cybercrime ecosystem.

Third, we find there to be a large incongruity between industry and policing reports on the acclaimed profitability of criminal business models like financial fraud - e.g., phishing - or extortion - e.g., ransomware - and the size of markets for cybercrime commodities. The lucrativity of these business models should attract new entrepreneurs to build their value chain based on off the shelf cybercrime components. Yet, if this would be the case, we should see more growth in the market for cybercrime commodities. The lack of proliferation in size and expansion of product portfolio of market for cybercrime commodities, suggests that there are still bottlenecks in outsourcing crucial parts of criminal value chains. It seems that substantial entry barriers still remain for aspiring criminal entrepreneurs.

In the niches of cybercrime commodities on online anonymous markets we do observe some modest growth. However, it seems that only a few 'professional' vendors profit from this growth. The product portfolio, profit and competition in cybercrime commodities do not reflect the mature and lively market it ought to be. Hence, industry reports that communicate about the commoditization of cybercrime by referring to natural phenomena, like tsunamis or avalanches, to describe the growing size or sales volumes of these commoditized cybercrime components, highly overstate this trend. Surely, only one commoditized cybercrime component - e.g., an exploit - can do grave damage to governments, businesses or individuals a like. So we must also not underestimate the seriousness of the trend in transacted cybercrime components we can observe on online anonymous markets. Yet, we have to remain true to its relative size, which is modest on all accounts - especially when comparing it to the trade in illegal narcotics.

That is why, a better understanding of where commoditization succeeds and fails helps us to identify cybercrime components for which outsourcing is less than trivial. In turn, this enables designing better disruption strategies for criminal business models. The unavailability or scarcity of certain commoditized cybercrime components suggests these are either tougher to produce at scale or are not suitable for distribution in a one-shot platform. Cybercrime components that rely on ongoing interaction between criminals do not reach full-fledged commoditization. The lack of supply of those components suggests potentially exploitable parts of criminal value chains.

In reverse, where full-fledged commoditization is found, a different impact on profit-driven cybercrimes can be expected. The static nature of these commodities, does not always match with the dynamic character of popular attack vectors. In recent years, industry reports have suggested that most successful attacks revolve around dynamic campaigns. Ransomware can be seen as an example, where the attack vector differs between campaigns and is oftentimes unique. This means that actors who rely on commoditized cybercrime components to perform profit-driven cybercrimes, might have less of a dynamic campaign than actors who vertically integrate their business model. These and similar perspectives might be kick starters for interventions as understanding where commoditization impacts dynamic campaigns or is lagging behind, points to alternative disruption strategies.

8.3. IMPLICATIONS FOR GOVERNANCE AND POLICING

Next to empirical insights, this dissertation has touched upon several practical implications. We have briefly discussed these implications in all chapters. In this section we comprise and explore the main public policy take-aways of this dissertation. First, we discuss how this dissertation adds to the common understanding of profit-driven cybercrime. Second, we discuss the practical implications of our insights into the impact of online anonymous market interventions. Third and last, we will elaborate on how future interventions can be evaluated.

UNDERSTANDING PROFIT-DRIVEN CYBERCRIME

Both industry reports and scientific inquiry have resulted in a comprehensive understanding of how the most prominent profit-driven cybercrimes operate – i.e., ransomware, crypto jacking and financial malware. This dissertation adds a conceptual view on profit-driven cybercrime by analyzing value chains of resources and the recurrence of cybercrime components among more than one form of profit-driven cybercrime. This way, we came to understand that many – by means of a description of their modus operandi or attack vector – different forms of profit-driven cybercrime, do in fact use congruent resources. Similar technical capabilities are therefore used in a wide range of profit-driven cybercrimes. Hence, when analyzing profit-driven cybercrime it is vital to include technical capabilities as a level of analysis. Cybercrime policing can thus focus on more than just the phenomenon – that in the eye of an ordinary policemen might not have much in common with another type of cybercrime, but relies in full or in part on identical resources. In essence, this would allow law enforcement agencies to pinpoint where an intervention might be of best use.

In addition, Chapter 2 concludes that different incentives can be identified in profit-driven cybercrime value chains. These incentives can be used to analyze chokepoints in the value chain. More specifically, if the scarcity of one activity in particular on the underground market influences incentives, chokepoints derived from these incentives are vital to future interventions.

ONLINE ANONYMOUS MARKET INTERVENTIONS

Many law enforcement agencies have made online anonymous markets a focal point in their fight against cybercrime. Various interventions have taken place, yet the online anonymous ecosystem – although everchanging and evolving - still survives. These interventions take a lot of effort and therefore costly. Here we elaborate on current and future interventions. We discsuss three 'alternative' intervention strategies and discuss how these tie back to the enablers of online anonymous markets (see 6.1.2).

INFILTRATING MARKETS

First, we can - based on the documented and successful attempts - identify a seemingly effective strategy aimed at individual markets: infiltration (see Chapter 6). This strategy has already been implemented in Operation Bayonet where it served as a testcase for alternative interventions. Law enforcement agencies pivoted the anonymity of the market admins to their advantage, as users did not have any clue they were dealing with a law enforcement controlled market. From Chapter 6 we know that vendors migrating to other markets during and after the operation, seem to be in distress about attribution efforts to be undertaken with personal information at the Hansa back-end servers, like postal or cryptocurrency addresses. This led a large portion of vendors migrating to Dream Market to take serious measures in an anonymous setting – e.g., changing usernames and/or PGP-keys, or starting over completely. Trust in the online anonymous market ecosystem was significantly lowered right after the intervention took place. We can also point to Reddit or Dread threads where users were in despair about the operation and questioned if they could ever trust an online anonymous market again. Yet, we do not know for how long this distrust will last. What we do know however, is that employing a tactic that makes use of and impact both enablers of online anonymous markets – i.e., anonymity and trust – seems to have become the new way forward.

CHOKING SUPPLY IN SPECIFIC MARKET SEGMENTS

From Chapters 2, 3 and 4 we learn that not all resources in a cybercrime value chain are equally likely to be outsourced as others. On the one hand, this relates to the fact that some products are scarce in the underground economy. On the other, some are that

important, that outsourcing creates external dependencies and therefore turns into a business continuity risk. Both form an interesting outlook to alternative intervention strategies.

First, we identify an alternative intervention that focuses on crucial resources in cybercrime value chains – i.e., chokepoints. As we know that for instance cash-out solutions are one of these vital elements in many – if not all – cybercrime value chains, interventions can be aimed at this specific market segment. For instance, one could make use of the trust mechanisms on the market to lemonize the market [9]. That is, advertise bogus cash-out solutions, buy them yourselves and leave positive feedback. This way, we potentially disrupt the total supply of cash-out solution, as one cannot distinguish a mock positive review, from a true positive review. In reverse, one can leave negative feedbacks to other cash-out solutions. Thereby, theoretically disrupting the total supply in a slightly different way, as one cannot distinguish a mock negative review, from a true negative review. In both cases trust in products and vendors in the cash-out market segment will be impacted. Moreover, as we know that vertically integrating cash-out solutions is complex, we can expect that the effects of this intervention reach beyond market supply, trust in advertised cash-out solutions and cash-out sellers. It could lead to a situation wherein not only trade in cash-out solutions is choked, but also self-engineered cash-out solutions will potentially be vulnerable to exploitation in different ways.

Second, instead of targeting a certain market segment, we can identify an alternative intervention that attempts to choke the trade of best-selling products. Chapter 5 identifies important predictors for product performance on the market. Product features that signal a professional operation – like refund polices, customer support and product branding – lead to higher sales numbers. Leveraging these insights, one could focus similar intervention tactics as described above – frustrating the review system – as others, to specific products instead of entire market segments. That way, scarce law enforcement capacity can be used to maximize impact as these specific products see a lot of sales and are potentially used in a variety of cybercriminal business models.

BREAKING BUSINESS CONTINUITY

Next to market or product focused intervention, we can identify targeting specific vendors as an alternative intervention strategy. From Chapter 5 we learn that distinctive vendor profiles emerge in the cybercrime segment of online anonymous markets – each with their own specific configuration of characteristics. One can imagine that for instance 'generalists' or 'professionals' form an intriguing group of facilitators that can supply a range of value chain resources. Focusing an intervention on these vendor profiles can potentially frustrate outsourcing, as scarce value chain resources become unfulfillable by

these facilitators. Interventions can use similar tactics as described above – making use of the review system. A different tactic can also be employed. As these profiles stand out as experienced vendors, with a long multi-market track record and solid reputation, one could piggy-back their reputation.

Since we know that markets come and go, a long-term vendor would – to ensure business continuity – migrate from market to market. To secure reputation, whilst staying anonymous, vendors make use of their username and public PGP-key to be remain identifiable after a migration. Both of these features however, are public information, which means that one could take a specific name and PGP-key and register on competing markets before the true vendor can. That way, individual business continuity is frustrated, as buyers do not know who is who.

EVALUATING ALTERNATIVE INTERVENTIONS

Disruption as an 'alternative intervention' has been high on the agenda of the Dutch police and the Public Prosecution Service for some time, but since the Dutch Police announced its plans for the years to come, disruption was formalised and named a priority. The Ministry of Justice and Security also includes alternative interventions in 2020 as a performance indicator for the police in its assessment of the fight against digital crime.

Although we have seen evidence (see Chapter 6) that 'alternative' interventions yield results, standardized measurements of disruption tactics is lacking. Especially when governments set out to move away from attribution – and with that, catching bad guys – and move towards disruption, the question arises: how do we know if it worked? This is not only important for law enforcement agencies to innovate policing tactics, but is essential to how we organize checks and balances in spending tax payer money. In other words, how do we know we get the best value out of tax payer money in the fight against profit-driven cybercrime?

Up until quite recently, an important performance indicator of policing aimed at attribution, was simply counting the number of arrests and prosecutions. Although this is rather crude approach that does not tell you if central players were amongst the arrestees, this still provides a uniform way to present the impact of policing. Now that governments tend to prioritize policing efforts differently, we need a similar, standardized way to evaluate the current 'alternative' interventions.

8.4. FUTURE WORK

Each chapter in this dissertation has presented directions for future research. In this section we discuss these and other avenues that arise from the work in this dissertation.

EXPLORING THE INSIDER'S PERSPECTIVE: BACK-END MARKET DATA

In Chapters 3 and 5 we have explored the cybercrime segments of online anonymous markets. Based on scraped data, we could perform longitudinal analyses on the commoditization of cybercrime (Chapter 3) and the predictors of product success (Chapter 5). We have argued our focus on online anonymous markets to measure commodization and thus the potential for outsourcing. For instance, all relevant aspects of the trading process can be observed – in contrast to forums, where a large portion of the trading process is foggy at best. That being said, studying online anonymous markets using scraped data is not perfect. Most importantly, we use feedbacks as a proxy for sales. In Chapter 5 we want to predict sales, but in fact we predict feedbacks -as this is our proxy for sales. Assuming that buyers of cybercrime products will, on average, leave as many feedbacks on one specific solution than on another, no bias is introduced. However, we do not know if that is really the case.

To mitigate this pitfall, future research can make use of back-end data – i.e., seized servers of online anonymous markets. That way, researchers can leverage the administration of markets – where next to listings, vendors and feedbacks, much more important details are available.

CAPTURING THE FRAGMENTATION OF OUTSOURCING

Next, our focus on online anonymous markets to measure commoditization, means turning a blind eye to different platforms where less standardized outsourcing takes place. The main driver behind this fragmentation of the outsourcing potential, is the steady rise in messaging platforms that embrace anonymization technologies. Platforms like Telegram offer a wide range of services that – be it in a less standardized way – supply aspiring criminals with tools and techniques, like phishing kits. In addition, we witness niche platforms in the underground economy – similar to online child sexual abuse platforms – were high-end, specific tool and techniques are traded. The forced security features on some of these sites make it next to impossible to perform measurements, only leaving room for a qualitative assessment of the nature of these channels. Others, like single-vendor shops, were reputable vendors that were previously active on online anonymous markets have set up shop, can be investigated longitudinally.

To map the trend of outsourcing accurately, future research efforts should also broaden

their scope and include messaging platforms, single-vendor shops and ideally niche platforms. Only then, we can try to fully understand and comprehend how the market for outsourcing develops and interventions to disrupt outsourcing can be executed.

MAPPING OUTSOURCING TO CYBERCRIME CAMPAIGNS

The last and critical piece of the outsourcing puzzle, is how cybercrime campaigns rely on technical capabilities fulfilled by cybercriminal markets. Now that we have a firm understanding of how outsourcing is supported by the commoditization of cybercrime on online anonymous, the next step is to look into how these commodities are actually used. Insights in how campaigns make use of capabilities, sheds light on which cybercrime components do not only see transactions, but are also put to use.

In turn, we can use these insights to analyze the repetitive use of certain capabilities. When we are able to map these capabilities to certain vendors, it makes intervening 'further up the chain' more worthwhile than a game of 'whack-a-mole' on campaigns that make use of the specific capability. Especially when that certain capability, upon inspection, turns out to have a unique vulnerability – let's say the standard password of the admin panel is 'password' – that can be used in an offensive way. Future work can therefore focus on matching the available commodities to value chains of cybercrime campaigns, thereby tracking the use of technical capabilities over time and across profit-driven cybercrimes.

In sum, future work can build on this dissertation and extend our knowledge on the differentiating of outsourcing platforms and the real-life use of commodities in cybercrime campaigns.

BIBLIOGRAPHY

[1] 2011. From marijuana to LSD, now illegal drugs delivered on your doorstep. *http://www.ibtimes.com/marijuana-lsd-now-illegal-drugs-delivered-your-doorstep-290021* (2011). http://www.ibtimes.com/marijuana-lsd-now-illegal-drugs-delivered-your-doorstep-290021

[2] 2016. Student pleads guilty to mass cyber attack. *https://www.bedsalert.co.uk/da/158731* (2016). https://www.bedsalert.co.uk/da/158731

[3] 2017. AlphaBay, the Largest Online 'Dark Market', Shut Down. *https://www.justice.gov/opa/pr/alphabay-largest-online-dark-market-shut-down* (2017). https://www.justice.gov/opa/pr/alphabay-largest-online-dark-market-shut-down

[4] 2017. United States of America vs. Alexandre Cazes – Verified compaint for forfeiture in rem. United States District Court, Eastern District of California. Case 1:17-at-00557.

[5] 2018. International Crackdown on Anti-spyware Malware. *https://www.europol.europa.eu/newsroom/news/international-crackdown-anti-spyware-malware* (2018). https://www.europol.europa.eu/newsroom/news/international-crackdown-anti-spyware-malware

[6] H Abadinksy. 1987. McDonald's-ization of the Mafia. In *Organized crime in America: concepts and controversies*, Timothy S Bynum (Ed.). 43–54.

[7] Lillian Ablon, Martin C. Libicki, and Andrea a. Golay. 2014. Markets for Cybercrime Tools and Stolen Data. *National Security Research Division* (2014), 1–85. https://doi.org/10.7249/j.ctt6wq7z6

[8] Sadia Afroz, Vaibhav Garg, Damon McCoy, and Rachel Greenstadt. 2013. Honor Among Thieves: A Common's Analysis of Cybercrime Economies. In *eCrime Researchers Summit (eCRS)*. 1–11.

[9] George A. Akerlof. 1970. The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics* 84, 3 (aug 1970), 488. https://doi.org/10.2307/1879431

[10] Ammar Alazab, Jemal Abawajy, Michael Hobbs, Robert Layton, and Ansam Khraisat. 2013. Crime Toolkits: The Productisation of Cybercrime. In *IEEE International Conference on Trust, Security and Privacy in Computing and Communications*.

[11] Mamoun Alazab, Sitalakshmi Venkatraman, Paul Watters, Moutaz Alazab, and Ammar Alazab. 2012. Cybercrime: The Case of Obfuscated Malware. In *ICGS3/e-Democracy 2011, LNICST 99*, R Bashroush (Ed.). 204–2011.

[12] Judith Aldridge and Rebecca Askew. 2017. Delivery dilemmas: How drug cryptomarket users identify and seek to reduce their risk of detection by law enforcement. *International Journal of Drug Policy* 41 (mar 2017), 101–109. https://doi.org/10.1016/j.drugpo.2016.10.010

[13] Judith Aldridge and David Decary-Hetu. 2014. Not an 'Ebay for Drugs': The Cryptomarket 'Silk Road' as a Paradigm Shifting Criminal Innovation. *SSRN Electronic Journal* 564, October (2014). https://doi.org/10.2139/ssrn.2436643

[14] Judith Aldridge and David Décary-Hétu. 2016. Hidden wholesale: The drug diffusing capacity of online drug cryptomarkets. *International Journal of Drug Policy* (2016). https://doi.org/10.1016/j.drugpo.2016.04.020

[15] Luca Allodi. 2017. Economic Factors of Vulnerability Trade and Exploitation: Empirical Evidence from a Prominent Russian Cybercrime Market. In *CCS'17*. https://doi.org/10.1145/3133956.3133960 arXiv:1708.04866

[16] Ross Anderson. 2018. Making Bitcoin Legal (Transcript of Discussion). In *Cambridge International Workshop on Security Protocols*. Springer, 254–265.

[17] Ross Anderson, Chris Barton, Rainer B??hme, Richard Clayton, Michel J G van Eeten, Michael Levi, Tyler Moore, Stefan Savage, Rainer Boehme, Richard Clayton, Michel van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. 2012. Measuring the cost of cybercrime. In *Workshop on the Economics of Information Security*. 265–300. https://doi.org/10.1007/978-3-642-39498-0_12

[18] Ross Anderson, Rainer Böhme, Richard Clayton, and Tyler Moore. 2008. Security economics and the internal market. *Network* (2008), 114.

[19] Ross Anderson and Tyler Moore. 2006. The Economics of Information Security. *Science* 314 (2006), 610–613.

[20] Manny Aston, Stephen McCombie, Ben Reardon, and Paul Watters. 2009. A Preliminary Profiling of Internet Money Mules: An Australian Perspective. In *Symposia and Workshops on Ubiquitous, Autonomic and Trusted Computing*.

[21] Johannes M. Bauer, Michel J G van Eeten, and Michel van Eeten. 2009. Cybersecurity: Stakeholder incentives, externalities, and policy options. *Telecommunications Policy* 33, 10-11 (2009), 706–719. https://doi.org/10.1016/j.telpol.2009.09.001

[22] Gary S Becker. 1974. Crime and Punishment: An Economic Approach. In *Essays in the Economics of Crime and Punishment*, Gary S Becker and William M Landes (Eds.). 1–54.

[23] Yuval Ben-Itzhak. 2009. Organised cybercrime and payment cards. *Card Technology Today* (2009), 10–11.

[24] Hugo L J Bijmans, Tim M Booij, and Christian Doerr. 2019. Inadvertently Making Cyber Criminals Rich: A Comprehensive Study of Cryptojacking Campaigns at Internet Scale. In *Proceedings of the 28th USENIX Security Symposium*.

[25] H Binsalleeh, T Ormerod, A Boukhtouta, P Sinha, A Youssef, M Debbabi, and L Wang. 2010. On the Analysis of the Zeus Botnet Crimeware Toolkit. In *International Conference on Privacy, Security and Trust*.

[26] David M Blei, Blei@cs Berkeley Edu, Andrew Y Ng, Ang@cs Stanford Edu, Michael I Jordan, and Jordan@cs Berkeley Edu. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022. https://doi.org/10.1162/jmlr.2003.3.4-5.993 arXiv:1111.6189v1

[27] David M Blei, Blei@cs Berkeley Edu, Andrew Y Ng, Ang@cs Stanford Edu, Michael I Jordan, and Jordan@cs Berkeley Edu. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022. https://doi.org/10.1162/jmlr.2003.3.4-5.993

[28] Nathaniel Boggs, Wei Wang, Suhas Mathur, Baris Coskun, and Carol Pincock. 2013. Discovery of emergent malicious campaigns in cellular networks. In *Proceedings of the 29th Annual Computer Security Applications Conference on - ACSAC '13*. ACM

Press, New York, New York, USA, 29–38. https://doi.org/10.1145/2523649.2523657

[29] Adam M Bossler and Thomas J Holt. 2009. On-line Activities, Guardianship, and Malware Infection: An Examination of Routine Activities Theory. *International Journal of Cyber Criminology* 3, 1 (2009), 400–420.

[30] Jean-Ian Boutin. 2014. The Evolution of Web Injects. In *Virus Bulletin Conference*.

[31] Kj Kate J. Bowers and Shane D. Sd Johnson. 2003. Measuring the geographical displacement of crime. *Journal of Quantitative Criminology* 19, 3 (2003), 275–302. https://doi.org/10.1023/A:1024909009240

[32] Christian Brenig, Rafael Accorsi, and Gunther Muller. 2015. Economic Analysis of Cryptocurrency Backed. *ECIS 2015 Proceedings* Ecb 2012 (2015), 1–18.

[33] Ryan Brunt, Prakhar Pandey, and Damon McCoy. 2017. Booted: An Analysis of a Payment Intervention on a DDoS-for-Hire Service. In *Workshop on the Economics of Information Security (WEIS)*.

[34] Danton Bryans. 2014. Bitcoin and money laundering: Mining for an effective solution. *Indiana Law Journal* 89 (2014), 441–472.

[35] Kit Burden and Creole Palmer. 2003. Internet crime: Cyber crime - A new breed of criminal? , 222–227 pages. https://doi.org/10.1016/S0267-3649(03)00306-6

[36] Julia Buxton and Tim Bingham. 2015. The Rise and Challenge of Dark Net Drug Markets. *Global Drugs Policy Observatory Policy Brief* January (2015), 24. https://doi.org/2054-1910

[37] Juan Caballero, Chris Grier, Christian Kreibich, and Vern Paxson. 2011. Measuring Pay-per-Install: The Commoditization of Malware Distribution. In *Usenix Security Symposium*.

[38] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. 2002. SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research* 16 (2002), 321–357. https://doi.org/10.1613/jair.953

[39] Nicolas Christin. 2013. Traveling the Silk Road: a measurement analysis of a large anonymous online marketplace. In *Proceedings of the 22nd international conference on World Wide Web*. 213–224. https://doi.org/10.1145/2488388.2488408

[40] Nicolas Christin. 2017. An EU-focused analysis of drug supply on the online anonymous marketplace ecosystem. *European Monitoring Centre for Drugs and Drug Addiction (EMCDDA)* (2017), 1–34. http://www.emcdda.europa.eu/system/files/attachments/6624/EU-focused-analysis-of-drug-supply-on-the-anomymous-online-marketplace.pdf

[41] N. Christin, S. Yanagihara, and K. Kamataki. 2010. Dissecting One Click Frauds. In *Proc. ACM CCS'10*. Chicago, IL, 15–26.

[42] Catherine M. Christopher. 2014. Whack-a-Mole: Why Prosecuting Digital Currency Exchanges Won't Stop Online Laundering. *Lewis & Clark Law Review* 18, 1 (2014), 1–36. http://papers.ssrn.com/sol3/papers.cfm?abstract{_}id=2312787

[43] M. Cova, C. Leita, O. Thonnard, A. Keromytis, and M. Dacier. 2010. An Analysis of Rogue AV Campaigns. In *Proc. RAID 2010*. Ottawa, ON, Canada.

[44] Claudio Criscione, Fabio Bosatelli, Stefano Zanero, and Federico Maggi. 2014. ZARATHUSTRA: Extracting WebInject Signatures from Banking Trojans. In *Annual Conference on Privacy, Security and Trust*.

[45] Thomas H. Davenport. 2005. The coming commoditization of processes. , 100–108 pages. https://doi.org/10.1108/14637151211225207

[46] David Décary-Hétu. 2014. Police Operations 3.0: On the Impact and Policy Implications of Police Operations on the Warez Scene. *Policy & Internet* 6, 3 (2014), 315–340. https://doi.org/10.1002/1944-2866.POI369

[47] David Décary-Hétu and Benoit Dupont. 2013. Reputation in a dark network of online criminals. *Global Crime* 14, 2-3 (2013), 175–196. https://doi.org/10.1080/17440572.2013.801015

[48] D. Décary-Hétu and L. Giommoni. 2017. Do police crackdowns disrupt drug cryptomarkets? A longitudinal analysis of the effects of Operation Onymous. *Crime, Law and Social Change* 67, 1 (feb 2017), 55–75. https://doi.org/10.1007/s10611-016-9644-4

[49] David Décary-Hétu and Anna Leppänen. 2016. Criminals and signals: An assessment of criminal performance in the carding underworld. *Security Journal* (2016). https://doi.org/10.1057/sj.2013.39

[50] Christian Decker and Roger Wattenhofer. 2013. Information propagation in the Bitcoin network. *13th IEEE International Conference on Peer-to-Peer Computing, IEEE P2P 2013 - Proceedings* (2013), 1–10. https://doi.org/10.1109/P2P.2013.6688704

[51] Fadi P Deek and James A M McHugh. 2007. *Open source: Technology and policy*. Cambridge University Press.

[52] Nadine Deslauriers-Varin and Eric Beauregard. 2010. Victims' routine activities and sex offenders' target selection scripts: A latent class analysis. *Sexual Abuse* 22, 3 (2010), 315–342.

[53] Andrew R Dick. 1995. When does organized crime pay? A transaction cost analysis. *International Review of Law and Economics* 15, 1 (1995), 25–45.

[54] Peter R Dickson and James L Ginter. 1987. Market segmentation, product differentiation, and marketing strategy. *Journal of marketing* 51, 2 (1987), 1–10.

[55] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. Tor: The second-generation onion router. *SSYM'04 Proceedings of the 13th conference on USENIX Security Symposium* 13 (2004), 21. https://doi.org/10.1.1.4.6896

[56] ECB. 2015. *Fourt Report on Card Fruad*. Technical Report. European Central Bank. 1–28 pages. https://doi.org/10.2866/22534

[57] Europol. 2015. *Exploring Tomorrow's Organised Crime*. Technical Report. Europol. https://www.europol.europa.eu/sites/default/files/Europol{_}OrgCrimeReport{_}web-final.pdf

[58] Europol. 2015. *The Internet Organised Crime Threat Assessment*. Technical Report. Europol. https://www.europol.europa.eu/sites/default/files/publications/europol{_}iocta{_}web{_}2015.pdf

[59] Europol. 2015. *Why Cash is Still King?* Technical Report. Europol. https://www.europol.europa.eu/sites/default/files/publications/europolcik.pdf

[60] FATF. 2010. *Money Laundering Using New Payment Methods*. Technical Report. http://www.fatf-gafi.org/dataoecd/4/56/46705859.pdf

[61] FATF. 2015. *Virtual Currencies: Key Definitions and Potential AML/CFT Risks*. Technical Report. FATF. http://fatf-gafi.org/media/fatf/documents/repor

ts/Virtual-currency-key-definitions-and-potential-aml-cft-risks.pdf

[62] FBI. 2017. *2016 Internet Crime Report.* Technical Report. FBI. https://pdf.ic3.gov/2016_IC3Report.pdf

[63] FBI. 2017. Darknet Takedown, Authorities Shutter Online Criminal Market AlphaBay. https://doi.org/news/stories/alphabay-takedown

[64] Dinei Florencio and Cormac Herley. 2010. Phishing and money mules. In *International Workshop on Information Forensics and Security (WIFS)*.

[65] Dinei Florencio and Cormac Herley. 2013. Where Do All The Attacks Go? In *Economics of Information Security and Privacy III*, Bruce Schneier (Ed.). 13–33.

[66] Bryanna Hahn Fox and David P. Farrington. 2012. Creating Burglary Profiles Using Latent Class Analysis: A New Approach to Offender Profiling. *Criminal Justice and Behavior* 39, 12 (2012), 1582–1611. https://doi.org/10.1177/0093854812457921 arXiv:https://doi.org/10.1177/0093854812457921

[67] Carlos Ganan, Orcun Cetin, and Michel van Eeten. 2015. An empirical analysis of ZeuS C&C lifetime. In *ACM Symposium on Information, Computer and Communications Security*. 97–108.

[68] Manuel Garcıa-Cervigon and Manel Medina Llinas. 2012. Browser Function Calls Modeling For Banking Malware Detection. In *International Conference on Risks and Security of Internet and Systems (CRiSIS)*.

[69] Nuno Garoupa. 1997. The Economics of Organized Crime and Optimal Law Enforcement. In *Annual Conference of the European Association of Law and Economics*.

[70] Nuno Garoupa. 2007. Optimal law enforcement and criminal organization. *Journal of Economic Behavior & Organization* 63 (2007), 461–474.

[71] Chris Grier, Lucas Ballard, Juan Caballero, Neha Chachra, Christian J Dietrich, Kirill Levchenko, Panayiotis Mavrommatis, Damon McCoy, Antonio Nappa, Andreas Pitsillidis, Niels Provos, M Zubair Rafique, Moheeb Abu Rajab, Christian Rossow, Kurt Thomas, Vern Paxson, Stefan Savage, and Geoffrey M Voelker. 2012. Manufacturing Compromise: The Emergence of Exploit-as-a-Service. In *ACM Conference on Computer Communications Security*.

[72] Andreas Haslebacher, Jeremiah Onaolapo, and Gianluca Stringhini. 2017. All your cards are belong to us: Understanding online carding forums. In *2017 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 41–51.

[73] Thomas J. Holt. 2013. Exploring the social organisation and structure of stolen data markets. *Global Crime* 14, 2-3 (2013), 155–174. https://doi.org/10.1080/17440572.2013.787925

[74] Thomas J Holt and Adam M Bossler. 2013. Examining the Relationship Between Routine Activities and Malware Infection Indicators. *Journal of Contemporary Criminal Justice* 20, 10 (2013), 1–17.

[75] Thomas J. Holt, Olga Smirnova, Yi Ting Chua, and Heith Copes. 2015. Examining the risk reduction strategies of actors in online criminal markets. *Global Crime* 16, 2 (2015), 81–103. https://doi.org/10.1080/17440572.2015.1013211

[76] Thomas J. Holt, Olga Smirnova, and Alice Hutchings. 2016. Examining signals of trust in criminal markets online. *Journal of Cybersecurity* (2016). https://doi.org/10.1093/cybsec/tyw007

[77] Thorsten Holz, Markus Engelberth, and Felix Freiling. 2009. Learning More About the Underground Economy: A Case-Study of Keyloggers and Dropzones. In *Computer Security–ESORICS*.

[78] Marie Claire Van Hout, Tim Bingham, Marie Claire Van Hout, and Tim Bingham. 2013. "Surfing the Silk Road": A study of users' experiences. *International Journal of Drug Policy* 24, 6 (2013), 524–529. https://doi.org/10.1016/j.drugpo.2013.08.011

[79] Danny Yuxing Huang, Maxwell Matthaios Aliapoulios, Vector Guo Li, Luca Invernizzi, Kylie Mcroberts, Elie Bursztein, Jonathan Levin, Kirill Levchenko, Alex C Snoeren, and Damon Mccoy. 2018. Tracking Ransomware End-to-end. In *IEEE Symposium on Security and Privacy (S&P)*.

[80] Danny Yuxing Huang, Hitesh Dharmdasani, Sarah Meiklejohn, Vacha Dave, Chris Grier, Damon McCoy, Stefan Savage, Nicholas Weaver, Alex C. Snoeren, and Kirill Levchenko. 2014. Botcoin: Monetizing Stolen Cycles. In *Proceedings 2014 Network and Distributed System Security Symposium*. https://doi.org/10.14722/ndss.2014.23044

[81] Alice Hutchings. 2014. Crime from the keyboard: Organised cybercrime, co-offending, initiation and knowledge transmission. *Crime, Law and Social Change* 62, 1 (2014), 1–20. https://doi.org/10.1007/s10611-014-9520-z

[82] Alice Hutchings and Thomas J Holt. 2014. A crime script analysis of the online stolen data market. *British Journal of Criminology* 55, 3 (2014), 596–614.

[83] Alice Hutchings and Sergio Pastrana. 2019. Understanding eWhoring. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 201–214.

[84] C.J. Hutto and Eric Gilbert. 2014. VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text. *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, 216–225.

[85] Colin C Ife, Toby Davies, Steven J Murdoch, and Gianluca Stringhini. 2019. Bridging Information Security and Environmental Criminology Research to Better Mitigate Cybercrime. *arXiv preprint arXiv:1910.06380* (2019).

[86] Mohammad Karami, Youngsam Park, and Damon McCoy. 2016. Stress testing the booters: Understanding and undermining the business of DDoS services. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1033–1043.

[87] Amin Kharraz, William Robertson, Davide Balzarotti, Leyla Bilge, and Engin Kirda. 2015. Cutting the Gordian knot: A look under the hood of ransomware attacks. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 9148. 3–24. https://doi.org/10.1007/978-3-319-20550-2_1

[88] Won Kim, Ok-Ran Jeong, Chulyun Kim, and Jungmin So. 2011. The dark side of the Internet: Attacks, costs and responses. *Information Systems* 36 (2011), 675–705.

[89] Peter Kollock and E. Russell Braziel. 2006. How not to build an online market: the sociology of market microstructure. , 283–306 pages. https://doi.org/10.1016/S0882-6145(06)23011-X

[90] Radhesh Krishnan Konoth, Rolf van Wegberg, Veelasha Moonsamy, and Herbert Bos. 2019. Malicious cryptocurrency miners: Status and Outlook. *arXiv preprint arXiv:1901.10794* (2019).

[91] Erika Kraemer-Mbula, Puay Tang, and Howard Rush. 2013. The cybercrime ecosystem: Online innovation in the shadows? *Technological Forecasting & Social Change* 80 (2013), 541–555.

[92] Nir Kshetri. 2010. The Economics of Click Fraud. *IEEE Security & Privacy Magazine* 8, 3 (may 2010), 45–53. https://doi.org/10.1109/MSP.2010.88 arXiv:arXiv:1011.1669v3

[93] Maurice Kugler, Thierry Verdier, and Yves Zenou. 2005. Organized crime, corruption and punishment. *Journal of Public Economics* 89 (2005), 1639–1663.

[94] Isak Ladegaard. 2017. We Know Where You Are, What You Are Doing and We Will Catch You. *The British Journal of Criminology* (2017). https://doi.org/10.1093/bjc/azx021

[95] Isak Ladegaard. 2019. Crime displacement in digital drug markets. *International Journal of Drug Policy* (2019). https://doi.org/10.1016/j.drugpo.2018.09.013

[96] Isak Ladegaard. 2020. Open Secrecy: How Police Crackdowns and Creative Problem-Solving Brought Illegal Markets out of the Shadows. *Social Forces* (03 2020). https://doi.org/10.1093/sf/soz140 arXiv:https://academic.oup.com/sf/advance-article-pdf/doi/10.1093/sf/soz140/32956743/soz140.pdf soz140.

[97] Monica Lagazio, Nazneen Sherif, and Mike Cushman. 2014. A multi-level approach to understanding the impact of cyber crime on the financial sector. *Computers & Security* 45 (2014), 58–74.

[98] Rutger Leukfeldt, Edward Kleemans, and Wouter Stol. 2017. The Use of Online Crime Markets by Cybercriminal Networks: A View From Within. *American Behavioral Scientist* (2017), 000276421773426. https://doi.org/10.1177/0002764217734267

[99] Kirill Levchenko, Andreas Pitsillidis, Neha Chachra, Brandon Enright, Mark Felegyhazi, Tristan Halvorson, Chris Kanich, Christian Kreibich, He Liu, Damon McCoy, Nicholas Weaver, Vern Paxson, Geoffrey M Voelker, and Stefan Savage. 2011. Click trajectories: End-to-end analysis of the spam value chain. In *IEEE Symposium on Security and Privacy*. IEEE. https://doi.org/10.1109

[100] Michael Levi. 2015. Money for Crime and Money from Crime: Financing Crime and Laundering Crime Proceeds. *European Journal on Criminal Policy and Research* 21, 2 (2015), 275–297.

[101] Michael Levi and Peter Reuter. 2006. Money laundering. *Crime and justice* 34, 1 (2006), 289–375. https://doi.org/10.1086/501508

[102] Steven D Levitt and Stephen J Dubner. 2005. *Freakonomics: A Rogue Economist Explores the Hidden Side of Everything*.

[103] Steven D Levitt and Stephen J Dubner. 2009. *SuperFreakonomics: Global Cooling, Patriotic Prostitutes, and Why Suicide Bombers Should Buy Life Insurance*.

[104] Steven D Levitt and Sudhir Alladi Venkatesh. 2000. An Economic Analysis of a Drug-Selling Gang's Finances. *The Quarterly Journal of Economics* 115, 3 (2000), 755–789.

[105] Zhen Li, Qi Liao, and Aaron Striegel. 2009. Botnet Economics: Uncertainty Matters. In *Managing Information Risk and the Economics of Security*, Eric Johnson (Ed.). 245–267.

[106] Jonathan Lusthaus. 2018. *Industry of Anonymity: Inside the Business of Cybercrime*. Harvard University Press.

[107] Jonathan Lusthaus and Federico Varese. 2017. Offline and Local: The Hidden Face of Cybercrime. *Policing: A Journal of Policy and Practice* (2017). https://doi.org/10.1093/police/pax042

[108] Thomas W Malone, Joanne Yates, and Robert I Benjamin. 1987. Electronic Markets and Electronic Hierarchies. *Commun. ACM* 30, 6 (1987), 484–497. https://doi.org/10.1002/9781118290743/wbiedcs158

[109] Samuel McQuade. 2002. *Encyclopedia of Cybercrime*. Vol. 10. 214–216 pages. https://doi.org/10.1353/sym.2002.0018

[110] Per Håkon Meland, Yara Fareed Fahmy Bayoumy, and Guttorm Sindre. 2020. The Ransomware-as-a-Service Economy within the Darknet. *Computers & Security* (2020), 101762.

[111] Charlie Miller. 2007. The Legitimate Vulnerability Market: Inside the Secretive World of 0-day Exploit Sales.

[112] Najmeh Miramirkhani, Oleksii Starovxi, and Nick Nikiforakis. 2017. Dial One for Scam: A Large-Scale Analysis of Technical Support Scams. In *Proceedings of the 24th Network and Distributed System Security Symposium (NDSS)*.

[113] Daniel Moore and Thomas Rid. 2016. Cryptopolitik and the Darknet. *Survival* 58, 1 (jan 2016), 7–38. https://doi.org/10.1080/00396338.2016.1142085

[114] Tyler Moore. 2010. The economics of cybersecurity: Principles and policy options. *INTERNATIONAL JOURNAL OF CRITICAL INFRASTRUCTURE PROTECTION* 3 (2010), 103–117.

[115] Tyler Moore and Nicolas Christin. 2013. Beware the middleman: Empirical analysis of Bitcoin-exchange risk. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 7859 LNCS (2013), 25–33. https://doi.org/10.1007/978-3-642-39884-1_3

[116] Tyler Moore, Richard Clayton, and Ross Anderson. 2009. The Economics of Online Crime. *The Journal of Economic Perspectives* 23, 3 (2009), 3–20.

[117] Carlo Morselli, David Décary-Hétu, Masarah Paquet-Clouston, and Judith Aldridge. 2017. Conflict management in illicit drug cryptomarkets. *International Criminal Justice Review* 27, 4 (2017), 237–254.

[118] Malte Möser, Rainer Böhme, and Dominic Breuker. 2013. An Inquiry into Money Laundering Tools in the Bitcoin Ecosystem. In *Procedings of the 2013 e-Crime Researches Summit*. 1–14. https://doi.org/10.1109/eCRS.2013.6805780

[119] Malte Möser, Kyle Soska, Ethan Heilman, Kevin Lee, Henry Heffan, Shashvat Srivastava, Kyle Hogan, Jason Hennessey, Andrew Miller, Arvind Narayanan, and Nicolas Christin. 2018. An Empirical Analysis of Traceability in the Monero Blockchain. *Proceedings on Privacy Enhancing Technologies* 2018, 3 (jun 2018), 143–163. https://doi.org/10.1515/popets-2018-0025 arXiv:1704.04299

[120] Marti Motoyama, Damon McCoy, Kirill Levchenko, Stefan Savage, and Geoffrey M Voelker. 2011. An Analysis of Underground Forums. In *ICM*.

[121] Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. *Consulted* (2008), 1–9. https://doi.org/10.1007/s10838-008-9062-0 arXiv:43543534534v343453

[122] Matthias Neugschwandtner, Paolo Milani Comparetti, and Christian Platzer. 2011. Detecting Malware's Failover C&C Strategies with SQUEEZE. In *Annual Computer Security Applications Conference*. 21–30.

[123] Arman Noroozian, Jan Koenders, Eelco Van Veldhuizen, Carlos H Ganan, Sumayah Alrwais, Damon McCoy, and Michel Van Eeten. 2019. Platforms in everything: analyzing ground-truth data on the anatomy and economics of bullet-proof hosting. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 1341–1356.

[124] Arman Noroozian, Maciej Korczyński, Carlos Hernandez Gañan, Daisuke Makita, Katsunari Yoshioka, and Michel van Eeten. 2016. Who gets the boot? analyzing victimization by DDoS-as-a-service. In *International Symposium on Research in Attacks, Intrusions, and Defenses*. Springer, 368–389.

[125] David Oro, Jesus Luna, Toni Felguera, Marc Vilanova, and Jetzabel Serna. 2010. Benchmarking IP Blacklists For Financial Botnet Detection. In *International Conference on Information Assurance and Security*. 62–67.

[126] Masarah Paquet-Clouston, David Décary-Hétu, and Carlo Morselli. 2018. Assessing market competition and vendors' size and scope on AlphaBay. *International Journal of Drug Policy* 54 (apr 2018), 87–98. https://doi.org/10.1016/j.drugpo.2018.01.003

[127] ChulWoo Park, HyoSung Park, and KiChang Kim. 2014. Realtime C&C Zeus Packet Detection Based on RC4 Decryption of Packet Length Field. *Advanced Science and Technology Letters* 64 (2014), 55–59.

[128] Justin M Rao and David H Reiley. 2012. The Economics of Spam. *The Journal of Economic Perspectives* 26, 3 (2012), 87–110.

[129] P Reuter. 1983. *Disorganized Crime - The Economics of the Visible Hand*. Ph.D. Dissertation.

[130] Marco Riccardi, David Oro, Jesus Luna, Marco Cremonini, and Marc Vilanova. 2010. A Framework For Financial Botnet Analysis. In *eCrime Researchers Summit (eCrime)*. 1–7.

[131] Marco Riccardi, Roberto Di Pietro, Marta Palanques, and Jorge Aguila Vila. 2012. Titans' revenge: Detecting Zeus via its own flaws. *Computer Networks* 57 (2012), 422–435.

[132] C Ronchi, A Khodjanov, M Mahkamov, and S Zakhidov. 2011. Security, Privacy and Efficiency of Internet Banking Transactions. In *World Congress on Internet Security (WorldCIS)*. 216–222.

[133] Christian Rossow, Christian Dietrich, and Herbert Bos. 2013. Large-Scale Analysis of Malware Downloaders. In *Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 42–61.

[134] Ernesto Savona. 2014. Organised crime numbers. *Global Crime* 15, 1-2 (2014), 1–9. https://doi.org/10.1080/17440572.2014.886512

[135] Sergii Shcherbak. 2013. How Should bitcoin be regulated? *European Journal of Legal Studies* 7, 1 (2013), 45–91. http://cadmus.eui.eu/bitstream/handle/1814/32273/183UK.pdf?sequence=1

[136] Aditya K Sood, Rohit Bansal, and Richard J Enbody. 2013. Cybercrime: Dissecting the State of Underground Enterprise.

[137] Aditya K Sood and Richard J Enbody. 2013. Crimeware-as-a-service—A survey of commoditized crimeware in the underground market. *INTERNATIONAL JOURNAL OF CRITICAL INFRASTRUCTURE PROTECTION* 6 (2013), 28–38.

[138] Aditya K Sood, Richard J Enbody, and Rohit Bansal. 2013. Dissecting SpyEye – Understanding the design of third generation botnets. *Computer Networks* 57 (2013), 436–450.

[139] Kyle Soska and Nicolas Christin. 2015. Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem. *24th USENIX Security Symposium* August (2015), 33–48. https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/soska

[140] Melvin R J Soudijn and Birgit C H T Zegers. 2012. Cybercrime and virtual offender convergence settings. *Trends in Organized Crime* 15, 2-3 (2012), 111–129. https://doi.org/10.1007/s12117-012-9159-z

[141] Kevin Stevens. 2009. The Underground Economy of the Pay-Per-Install (PPI) Business. (2009).

[142] Brett Stone-Gross, Ryan Abman, Richard A Kemmerer, Christopher Kruegel, Douglas G Steigerwald, and Giovanni Vigna. 2013. The Underground Economy of Fake Antivirus Software. In *Economics of Information Security and Privacy III*, Bruce Schneier (Ed.). 55–78.

[143] Xiao Hui Tai, Kyle Soska, and Nicolas Christin. 2019. Adversarial Matching of Dark Net Market Vendor Accounts. https://doi.org/10.1145/3292500.3330763

[144] Samaneh Tajalizadehkhoob, Hadi Asghari, Carlos Ganan, and Michel van Eeten. 2014. Why Them? Extracting Intelligence about Target Selection from Zeus Financial Malware. In *Workshop on the Economics of Information Security (WEIS)*.

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

[146] M Turvani. 1997. *Illegal markets and the new institutional economics*.

[147] UNODC. 2014. *Basic Manual on the Detection And Investigation of the Laundering of Crime Proceeds Using Virtual Currencies*. Technical Report.

[148] Joe Van Buskirk, Amanda Roxburgh, Raimondo Bruno, Sundresan Naicker, Simon Lenton, Rachel Sutherland, Elizabeth Whittaker, Natasha Sindicich, Allison Matthews, Kerryn Butler, and Lucinda Burns. 2016. Characterising dark net marketplace purchasers in a sample of regular psychostimulant users. *International Journal of Drug Policy* 35 (sep 2016), 32–37. https://doi.org/10.1016/j.drugpo.2016.01.010

[149] Michel van Eeten and Johannes Bauer. 2008. *Economics of malware: Security decisions, incentives and externalities*. Technical Report.

[150] Gert Jan Van Hardeveld, Craig Webber, and Kieron O'Hara. 2016. Discovering credit card fraud methods in online tutorials. In *OnSt 2016 - 1st International Workshop on Online Safety, Trust and Fraud Prevention*. https://doi.org/10.1145/2915368.2915369

[151] Marie Claire Van Hout and Tim Bingham. 2013. 'Silk Road', the virtual drug marketplace: A single case study of user experiences. , 385–391 pages. https://doi.org/10.1016/j.drugpo.2013.01.005

[152] Rolf van Wegberg, Samaneh Tajalizadehkhoob, Kyle Soska, Ugur Akyazi, Carlos Hernandez Ganan, Bram Klievink, Nicolas Christin, and Michel van Eeten. 2018. Plug and Prey? Measuring the Commoditization of Cybercrime via Online Anonymous Markets. In *Proceedings of the 27th USENIX Security Symposium*.

[153] Rolf van Wegberg and Thijmen Verburgh. 2018. Lost in the Dream? Measuring the effects of Operation Bayonet on vendors migrating to Dream Market. In *Evolution of the Darknet Workshop at the Web Science Conference (WebSci 18)*. Association for

Computing Machinery (ACM), New York, NY, USA, 1–5. https://www.narcis.nl/publication/RecordID/oai:tudelft.nl:uuid:8c080055-37fb-4f53-a949-099110f91659

[154] R. S. van Wegberg, A. J. Klievink, and M. J. G. van Eeten. 2017. Discerning Novel Value Chains in Financial Malware. *European Journal on Criminal Policy and Research* 23, 4 (dec 2017), 575–594. https://doi.org/10.1007/s10610-017-9336-3

[155] Michael G. Vaughn, Matt DeLisi, Kevin M. Beaver, and Matthew O. Howard. 2008. Toward a Quantitative Typology of Burglars: A Latent Profile Analysis of Career Offenders. *Journal of Forensic Sciences* 53, 6 (2008), 1387–1392. https://doi.org/10.1111/j.1556-4029.2008.00873.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1556-4029.2008.00873.x

[156] Michael G. Vaughn, Matt DeLisi, Kevin M. Beaver, and Matthew O. Howard. 2009. Multiple murder and criminal careers: A latent class analysis of multiple homicide offenders. *Forensic Science International* 183, 1 (2009), 67 – 73. https://doi.org/10.1016/j.forsciint.2008.10.014

[157] J.K. Vermunt and J. Magidson. 2002. Latent class cluster analysis. In *Applied latent class analysis*, J. Hagenaars and A. McCutcheon (Eds.). Cambridge University Press, United Kingdom, 89–106. Pagination: 476.

[158] J.K. Vermunt and J. Magidson. 2016. *Guide for Latent GOLD 5.1: Basic, Advanced, and Syntax*. Technical Report. Statistical Innovations Inc., Belmont, MA.

[159] Simon Walker. 2012. Economics and the cyber challenge. *Information Security Technical Report* 17 (2012), 9–18.

[160] David S. Wall. 1998. Catching Cybercriminals: Policing the Internet. *International Review of Law, Computers & Technology* 12, 2 (1998), 201–218. https://doi.org/10.1080/13600869855397

[161] Lanier Watkins, Christina Kawka, Cherita Corbett, and William H Robinson. 2014. Fighting Banking Botnets By Exploiting Inherent Command and Control Vulnerabilities. In *International Conference on Malicious and Unwanted Software: The Americas (MALWARE)*. 93–100.

[162] Michael D Whinston. 2000. On the Transaction Costs Determinants of Vertical Integration. *Journal of Law, Economics, and Organization* 19, 1 (2000), 1–23.

[163] Oliver E Williamson. 1971. The Vertical Integration of Production: Market Failure Considerations. *The American Economic Review* 61, 2 (1971), 112–123.

[164] Oliver E Williamson. 1979. Transaction-Cost Economics: The Governance of Contractual Relations. *Journal of Law and Economics* 22, 2 (1979), 233–261.

[165] Oliver E Williamson. 2005. Transaction cost economics and business administration. *Scandinavian Journal of Management* 21, 1 (2005), 19–40.

[166] Oliver E Williamson. 2007. Transaction Cost Economics: An Introduction. *Economics Discussion Paper* (2007), 0–33. https://doi.org/10.1017/CBO9781107415324.004

[167] Jianwei Zhuge, Thorsten Holz, Chengyu Song, Jinpeng Guo, Xinhui Han, and Wei Zou. 2009. Studying Malicious Websites and the Underground Economy on the Chinese Web. In *Managing Information Risk and the Economics of Security*, M E Johnson (Ed.). 225–244.

SUMMARY

Many scientific studies and industry reports have observed the emergence of so-called cybercrime-as-a-service. The idea is that specialized suppliers in the underground economy cater to criminal entrepreneurs in need of certain capabilities – substituting specialized technical knowledge with "knowing what to buy". The impact of this trend could be dramatic, as technical skill becomes an insignificant entry barrier for cybercrime.

Forms of cybercrime motivated by financial gain, make use of a unique configuration of technical capabilities to be successful. Profit-driven cybercrimes, as they are called, range from carding to financial malware, and from extortion to cryptojacking. Given their reliance on technical capabilities, particularly these forms of cybercrime benefit from a changing crime paradigm: the commoditization of cybercrime. That is, standardized offerings of technical capabilities supplied through structured markets by specialized vendors that cybercriminals can contract to fulfill tools and techniques used in their business model. Commoditization enables outsourcing of components used in cybercrime - i.e., a botnet or cash-out solution. Thus lowering entry barriers for aspiring criminals, and potentially driving further growth in cybercrime.

As many cybercriminal entrepreneurs lack the skills to provision certain parts of their business model, this incentivizes them to outsource these parts to specialized criminal vendors. With online anonymous markets - like Silk Road or AlphaBay - these entrepreneurs have found a new platform to contract vendors and acquire technical capabilities for a range of cybercriminal business models. A configuration of technical capabilities used in a business model reflects the value chain of resources. Here, not the criminal activities themselves, but the technical enablers for all these criminal activities are depicted.

To create a comprehensive understanding of how business models in profit-driven cybercrime are impacted by the commoditization of cybercrime, we investigate how outsourced components can fulfill technical capabilities needed in profit-driven cybercrime. This is where we use an economic lens to deliver an overview of criminal activities, resources and strategies in profit-driven cybercrime. In turn, knowing how outsourcing fulfils parts of the value chain, can help law enforcement exploit 'chokepoints' – i.e., use the weakest link in the value chain where criminals appear to be vulnerable.

Hence, this dissertation studies the phenomenon of profit-driven cybercrime through value chains and business models. It investigates how outsourcing is enabled by online anonymous markets and how outsourcing can be disrupted. This leads to the main research question:

How do online anonymous markets facilitate the outsourcing of cybercrime components in profit-driven cybercrime value chains?

The following chapters report on the multiple research avenues that collectively answer the main research question above.

Chapter 2 aims to discern value chains in financial malware. These value chains were constructed in a similar fashion as other researchers reconstructed the spam value chain. We find that, for financial malware schemes using man-in-the-browser attack vectors, there is a clear incentive to outsource cybercrime components via underground markets. We have demonstrated that a value chain approach is evidently useful when researching financial malware and other cybercriminal business models. Next, this creates the opportunity to study the important interactions between the strategies of attackers on the one side and the properties of cybercrime components and policing tactics on the other.

Chapter 3 studies the phenomenon of commoditization of cybercrime. Using six years of longitudinal scraped data from eight online anonymous markets we measure the evolution of commoditized offerings for dominant cybercriminal value chains, and report on the volume of transactions and revenue of actual sales to criminal entrepreneurs. Our results show that on online anonymous marketplaces, commoditization is a spottier phenomenon than was previously assumed. We find no supply for many of the capabilities, systems and resources observed in well-known value chains. Likewise, there is no evidence of rapid growth, and thus of a strong push towards commoditization, contrary to the somewhat alarmist language found in industry reporting and elsewhere.

Chapter 4 aims to examine ways in which cybercrime proceeds can be laundered. Building on earlier work, we focus bitcoin mixers in the underground economy. We investigate bitcoin mixing and shadow exchange services in an cash-out experiment. We examine the usability of these services by analyzing service percentages and reputation mechanisms. The results of our experiment suggest that laundering cybercrime proceeds using bitcoin mixers is a user-friendly and operational criminal service-model. The ability to lower the cost of laundering, whilst providing more anonymity, make it a practically conceivable concept and therefore likely to be integrated in criminal business models.

Chapter 5 investigates the performance of products in the business-to-business cybercrime market segments on AlphaBay. As we know that not all products and vendors are equally successful on the market, we aim to predict which characteristics of both the criminal entrepreneur and the product they sell, influences the performance of cybercrime components. Overall, our findings show that signals of a professional operational – product branding, refund policies and customer support, explain a significant portion of the variance in performance. Yet, the vendor of a product – rather than its nature or features – is the most powerful predictor for its success.

Chapter 6 explores the enablers of online anonymous markets and constructs a historical perspective of online anonymous market interventions. We take Operation Bayonet as a case and investigate the effects of the operation on all newly registered vendors on Dream Market by mapping their individual and historic characteristics to discern migration patterns and changes in vendor behavior. Our results suggest that Operation Bayonet is to be considered a game-changing police intervention – as users do not just move along after the Hansa Market shutdown. Few simply migrate, some take evasive measure - like changing their username and/or PGP-key - but many start with a clean slate on Dream Market. Inevitable, this means that vendors have to build their reputation and customer-base from scratch.

The previous chapters have shown how online anonymous markets play an important role in fulfilling resources in value chains of profit-driven cybercrime. We see that – although supply and demand are not completely matched – most forms of profit-driven cybercrime can make use of the potential for outsourcing that an online anonymous market provides. Leveraging measurements in the ecosystem, we have investigated commoditizaton of cybercrime components, the predictors for product success and how one particular law enforcement intervention impacted market users. Chapter 7 touches upon how measurements can help evaluate policing practice. We provide context to measurements in the online anonymous market ecosystem by synthesizing the state-of-the-art in intervention studies and identify common standards to measure interventions on online anonymous markets. We discuss in-depth how these measurements map to known aims and tactics of past interventions and present suggestions for novel measurements of future interventions.

Finally, Chapter 8 synthesizes the empirical findings and reports on the governance implications and presents the main take-aways. This dissertation adds a conceptual view on profit-driven cybercrime by analyzing value chains and the recurrence of cybercrime components among more than one form of profit-driven cybercrime. This way, we came to understand that many different forms of profit-driven cybercrime, do in fact use

congruent resources. Similar technical capabilities are therefore used in a wide range of profit-driven cybercrimes. Cybercrime policing can thus focus on more than just the phenomenon – that in the eye of an ordinary policemen might not have much in common with another type of cybercrime, but relies in full or in part on identical resources. In essence, this would allow law enforcement agencies to pinpoint where an intervention might be of best use.

Yet, we find there to be a large incongruity between industry and policing reports on the acclaimed profitability of criminal business models like financial fraud - e.g., phishing - or extortion - e.g., ransomware - and the size of markets for cybercrime commodities. The lucrativity of these business models should attract new entrepreneurs to build their value chain based on off the shelf cybercrime components. If this would be the case however, we should see more growth in the market for cybercrime commodities. The lack of proliferation in size and expansion of product portfolio of markets for cybercrime commodities, suggests that there are still bottlenecks in outsourcing crucial parts of criminal value chains. It seems that substantial entry barriers remain for aspiring cybercriminals.

SAMENVATTING

Veel wetenschappelijke studies en rapporten uit de beveiligingsindustrie hebben de opkomst van zogenaamde *cybercrime-as-a-service* beschreven. Het idee is dat gespecialiseerde leveranciers in de ondergrondse economie criminele ondernemers bedienen die behoefte hebben aan bepaalde capaciteiten - waarbij gespecialiseerde technische kennis wordt vervangen door weten wat te kopen. De impact van deze trend kan dramatisch zijn, omdat technische vaardigheid tot een onbeduidende drempel voor cybercriminaliteit verwordt.

Vormen van cybercriminaliteit die worden gemotiveerd door financieel gewin, maken gebruik van een unieke samenstelling van technische capaciteiten om succesvol te zijn. Winstgedreven cybercriminaliteit varieert van *carding* tot financiële malware en van afpersing tot *cryptojacking*. Gezien hun afhankelijkheid van technische capaciteiten, profiteren vooral deze vormen van cybercriminaliteit van een veranderend misdaadparadigma: de commoditisering van cybercriminaliteit. Dat wil zeggen: gestandaardiseerde aanbiedingen van technische capaciteiten die worden geleverd via gestructureerde markten door gespecialiseerde leveranciers die cybercriminelen kunnen contracteren om tools en technieken te gebruiken die in hun verdienmodel worden gebruikt. Commoditisering maakt de uitbesteding mogelijk van componenten die worden gebruikt bij cybercriminaliteit, bijvoorbeeld een *botnet* of witwasoplossing. Zo worden de drempels voor aspirant criminelen verlaagd en kan cybercriminaliteit mogelijk verder groeien.

Aangezien veel cybercriminele ondernemers niet over de vaardigheden beschikken om bepaalde onderdelen van hun verdienmodel zelf te organiseren, stimuleert dit hen deze onderdelen uit te besteden aan gespecialiseerde criminele verkopers. Met online anonieme markten - zoals Silk Road of AlphaBay - hebben deze ondernemers een nieuw platform gevonden om verkopers te contracteren en technische capaciteiten te verwerven voor een reeks cybercriminele verdienmodellen. Een samenstelling van technische capaciteiten die in een bedrijfsmodel wordt gebruikt, weerspiegelt de waardeketen. Hierin worden niet de criminele activiteiten zelf weergegeven, maar de technische capaciteiten die deze criminele activiteiten mogelijk maken.

Om een beter begrip te krijgen van hoe verdienmodellen in winstgedreven cybercriminaliteit worden beïnvloed door de commoditisatie van cybercriminaliteit, onderzoeken

we hoe uitbestede componenten kunnen voldoen aan technische capaciteiten die nodig zijn bij winstgedreven cybercriminaliteit. Hierbij gebruiken we een economische lens om een overzicht te geven van criminele activiteiten, middelen en strategieën in winstgedreven cybercriminaliteit. Weten hoe uitbesteding delen van de waardeketen vervult, kan op zijn beurt politiediensten helpen om knelpunten te benutten - het exploiteren van de zwakste schakel in de waardeketen waar criminelen kwetsbaar lijken te zijn.

Daarom bestudeert dit proefschrift het fenomeen van winstgedreven cybercriminaliteit middels waardeketens en verdienmodellen. Het onderzoekt hoe uitbesteding mogelijk wordt gemaakt door online anonieme markten en hoe uitbesteding kan worden verstoord. Dit leidt tot de volgende, centrale onderzoeksvraag:

Hoe vergemakkelijken online anonieme markten de uitbesteding van cybercrime componenten in winstgedreven cybercrime waardeketens?

In de volgende hoofdstukken wordt verslag gedaan van de onderzoeksrichtingen die gezamenlijk de bovenstaande onderzoeksvraag beantwoorden.

Hoofdstuk 2 heeft tot doel waardeketens in financiële malware te onderscheiden. Deze waardeketens zijn op dezelfde manier geconstrueerd als andere onderzoekers de spamwaardeketen hebben gereconstrueerd. We stellen vast dat er voor financiële malware aanvallen die gebruik maken van *man-in-the-browser* aanvalsvectoren, een duidelijke prikkel is om cybercrime componenten uit te besteden via ondergrondse markten. We hebben aangetoond dat een waardeketenbenadering nuttig is bij het onderzoeken van financiële malware en andere cybercriminele verdienmodellen. Dit biedt vervolgens de mogelijkheid om de belangrijke interacties tussen de strategieën van aanvallers enerzijds en de eigenschappen van cybercrime componenten en politietactieken anderzijds te bestuderen.

Hoofdstuk 3 bestudeert het fenomeen van commoditisering van cybercriminaliteit. Met behulp van zes jaar aan gegevens van acht online anonieme markten meten we de evolutie van gecommoditiseerde aanbiedingen voor dominante waardeketens van cybercriminelen en rapporteren we over het aantal transacties en de inkomsten van daadwerkelijke verkopen aan criminele ondernemers. Onze resultaten laten zien dat commoditisering op online anonieme marktplaatsen een minder duidelijk fenomeen is dan eerder werd aangenomen. We vinden geen aanbod voor veel van de capaciteiten die in prominente waardeketens worden waargenomen. Evenzo is er geen bewijs van een snelle groei, en dus van een sterke drang naar commoditisering, in tegenstelling tot de ietwat alarmerende taal die wordt aangetroffen in de rapportages van de beveiligingsindustrie.

Hoofdstuk 4 beschrijft hoe de opbrengsten van cybercrime kunnen worden witgewassen. Voortbouwend op eerder werk, richten we ons op bitcoin mixers in de ondergrondse economie. We onderzoeken bitcoin mixers en ondergrondse bitcoin wisseldiensten in een witwasexperiment. We onderzoeken de bruikbaarheid van deze diensten door dienstverleningskosten en reputatiemechanismen te analyseren. De resultaten van ons experiment suggereren dat het witwassen van de opbrengsten van cybercriminaliteit door middel van bitcoin mixers een gebruiksvriendelijk en operationeel crimineel dienstverleningsmodel is. De mogelijkheid om de kosten van witwassen te verlagen en tegelijkertijd meer anonimiteit te bieden, maakt het een praktisch denkbaar concept en daarmee waarschijnlijk geïmplementeerd in criminele verdienmodellen.

Hoofdstuk 5 onderzoekt de prestaties van producten in de *business-to-business* cybercrime marktsegmenten op AlphaBay. Omdat we weten dat niet alle producten en verkopers even succesvol zijn op de markt, willen we voorspellen welke kenmerken van zowel de criminele ondernemer als het product dat ze verkopen, de prestaties van cybercrime componenten beïnvloeden. Al met al laten onze bevindingen zien dat signalen van een professionele bedrijfsvoering - product marketing, terugbetalingsbeleid en klantenondersteuning, een aanzienlijk deel van de variantie in prestaties verklaren. Toch is de verkoper van een product - in plaats van de aard of kenmerken ervan - de krachtigste voorspeller van diens succes.

Hoofdstuk 6 verkent de mogelijkheden van online anonieme markten en construeert een historisch perspectief van online anonieme marktinterventies. We nemen *Operation Bayonet* als casus en onderzoeken de effecten van de operatie op alle nieuw geregistreerde verkopers op Dream Market door hun individuele en historische kenmerken in kaart te brengen om migratiepatronen en veranderingen in het gedrag van verkopers te onderscheiden. Onze resultaten suggereren dat *Operation Bayonet* moet worden beschouwd als een ingrijpende politie-interventie - aangezien gebruikers niet zomaar verder gaan na de sluiting van de Hansa Market. Weinig gebruikers migreren, sommigen nemen beschermende maatregelen - zoals het wijzigen van hun gebruikersnaam en / of PGP-sleutel - maar velen beginnen met een schone lei op Dream Market. Dit betekent onvermijdelijk dat verkopers hun reputatie en klantenbestand vanaf nul moeten opbouwen.

De vorige hoofdstukken hebben laten zien hoe online anonieme markten een belangrijke rol spelen bij het vervullen van capaciteiten in waardeketens van winstgedreven cybercriminaliteit. Hoewel vraag en aanbod niet volledig op elkaar zijn afgestemd, zien we dat de meeste vormen van winstgedreven cybercriminaliteit gebruik kunnen maken van de mogelijkheden voor uitbesteding die een online anonieme markt biedt. Door gebruik te maken van metingen in het ecosysteem, hebben we de commoditisatie van cybercrime

componenten, de voorspellers van productsucces en de impact van een bepaalde politie-interventies op de marktgebruikers onderzocht. Hoofdstuk 7 bespreekt hoe metingen kunnen helpen bij het evalueren van de politiepraktijk. We bieden context aan metingen in het online anonieme marktecosysteem door de *state-of-the-art* in interventiestudies te synthetiseren en gemeenschappelijke standaarden te identificeren om interventies op online anoniem te meten. We bespreken diepgaand hoe deze metingen overeenkomen met bekende doelen en tactieken van eerdere interventies en presenteren suggesties voor nieuwe metingen van toekomstige interventies.

Hoofdstuk 8 geeft een samenvatting van de empirische bevindingen en rapporteert over de *governance* implicaties en presenteert de belangrijkste *take-aways*. Dit proefschrift voegt een conceptuele kijk op winstgedreven cybercriminaliteit toe door waardeketens en de repetitie van cybercrime componenten te analyseren in meer dan één vorm van winstgedreven cybercriminaliteit. Zo kwamen we tot het inzicht dat veel verschillende vormen van op winst gerichte cybercriminaliteit in feite gebruik maken van soortgelijke technische middelen. Vergelijkbare technische capaciteiten worden gebruikt bij een breed scala aan op winst gerichte cybercriminaliteit. De opsporing van cybercrime kan zich dus richten op meer dan het fenomeen alleen - dat in de ogen van een gewone politieagent misschien niet veel gemeen heeft met een ander type cybercriminaliteit, maar volledig of gedeeltelijk afhankelijk is van identieke, technische middelen. In wezen zou dit politiediensten moeten helpen vaststellen waar een interventie het beste kan worden uitgevoerd.

Toch zien we dat de rapporten van de beveiligingsindustrie en de politie over de winstgevendheid van criminele verdienmodellen zoals financiële fraude en afpersing - bijvoorbeeld *phishing* of *ransomware* - niet rijmen met de omvang van de markten voor cybercriminaliteit capaciteiten. De lucrativiteit van deze verdienmodellen zou nieuwe ondernemers moeten aantrekken om hun waardeketen op te bouwen op basis van standaard cybercrime componenten. Als dit echter het geval zou zijn, zouden we meer groei moeten zien in de markt voor cybercrime capaciteiten. Het gebrek aan groei in omvang en uitbreiding van het productportfolio van markten voor cybercrime capaciteiten, suggereert dat er nog steeds knelpunten zijn bij het uitbesteden van cruciale onderdelen van criminele waardeketens. Het lijkt erop dat er nog steeds aanzienlijke drempels bestaan voor aspirant cybercriminelen.

AUTHORSHIP CONTRIBUTIONS

The dissertation is based on five peer-reviewed studies stemming from collaborative work with several co-authors. While all of these studies were led by myself, I was fortunate enough to receive valuable feedback and varying contributions of my co-authors in each of these studies. I will outline their respective contributions per study below.

For the first study (see Chapter 2), my co-authors, Bram Klievink and Michel van Eeten, have helped with improving the draft, refining its arguments, proof reading and polishing of the text. Collection and analysis of the underlying data used for this study was performed by me.

For the second study (see Chapter 3), my co-authors Samaneh Tajalizadehkhoob, Kyle Soska, Ugur Akyazi, Carlos Hernandez Ganan, Bram Klievink, Nicolas Christin and Michel van Eeten, have greatly helped to improve the text, place it within the concepts of related work and focus the main research question of the paper. Collection of the underlying data was performed by Kyle Soska and Nicolas Christin. Code to model, analyze and plot the data was developed by Samaneh Tajalizadehkhoob, Carlos Hernandez Ganan and Ugur Akyazi jointly.

For the third study (see Chapter 4), my co-authors, Jan-Jaap Oerlemans and Oskar van Deventer, have contributed by improving structure and text as well as sharpen the paper's argumentation. Collection and analysis of the data used for this study was performed by me primarily, but received valuable input from both Jan-Jaap Oerlemans as Oskar van Deventer.

For the fourth study (see Chapter 5), my co-authors Fieke Miedema, Ugur Akyazi, Arman Noroozian, Bram Klievink, and Michel van Eeten have all helped with improving my drafts and sharpening the paper's structure, methodology and logic. Code to model, analyze and plot the data was developed by Fieke Miedema, Arman Noroozian and Ugur Akyazi jointly.

Finally, the fifth study (see Chapter 6), builds on greatly valued support from my co-author, Thijmen Verburgh. He has helped me with polishing the text and reflect on the results of the analysis. The data collection and analysis were handled by us jointly.

In all these studies, I have been lucky enough to receive the support and vital contributions of my co-authors. To varying degrees and with their diverse backgrounds, they

have not only contributed to my studies in terms of ideas, feedback and writing, but have also contributed to the scientist I now am. I would like to especially emphasize the unwavering support of my *promotores* Bram Klievink and Michel van Eeten who have patiently guided a stubborn social scientist to the point of submitting papers to top-tier computer science conferences.

LIST OF PUBLICATIONS

- **Van Wegberg, R.S.**, Miedema, F., Akyazi, U., Noroozian, A., Klievink, B., & Van Eeten, M. (2020). "Go see a specialist? Predicting cybercrime sales on online anonymous markets from vendor and product characteristics". In *Proceedings of The Web Conference (WWW '20)* (pp. 816-826).

- Oerlemans, J.J., & **Van Wegberg, R.S.** (2019). "Opsporing en bestrijding van online drugsmarkten". *Strafblad, 17*(5), 25-31.

- Hartel, P., & **Van Wegberg, R.S.** (2019). "Crime and Online Anonymous Markets". In *International and Transnational Crime and Justice.* Natarajan, M. (ed.), 67-72.

- **Van Wegberg, R.S.**, Tajalizadehkhoob, S., Soska, K., Akyazi, U., Ganan, C. H., Klievink, B., Christin, N., & Van Eeten, M. (2018). "Plug and Prey? Measuring the commoditization of cybercrime via online anonymous markets". In *Proceedings of the USENIX Security Symposium (USENIX Security 18)* (pp. 1009-1026).

- Verburgh, T., Smits, E., & **Van Wegberg, R.S.** (2018). "Uit de schaduw: Perspectieven voor wetenschappelijk onderzoek naar dark markets". *Justitiële Verkenningen, 44*(5), 68-82.

- **Van Wegberg, R.S.**, & Verburgh, T. (2018). "Lost in the Dream? Measuring the effects of Operation Bayonet on vendors migrating to Dream Market". In *Proceedings of the Evolution of the Darknet Workshop* (pp. 1-5).

- **Van Wegberg, R.S.**, Oerlemans, J.J., & Van Deventer, O. (2018). "Bitcoin money laundering: mixed results? An explorative study on money laundering of cybercrime proceeds using bitcoin". *Journal of Financial Crime, 25*(2), 419-435.

- **Van Wegberg, R.S.**, Klievink, B., & Van Eeten, M. (2017). "Discerning novel value chains in financial malware".*European Journal on Criminal Policy and Research, 23*(4), 575-594.

- Klievink, B., **Van Wegberg, R.S.**, & van Eeten, M. (2017). "Een gezamenlijke rekening? Over digitale innovatie en samenwerking in een institutional void". *Bestuurskunde, 26*(1).

About the author

Rolf van Wegberg (1988) was born in Voorburg, The Netherlands. Originally trained as a criminologist, he received his MSc-degree (cum laude) from Leiden University in 2011 with a thesis on money laundering and the funding of terrorism in the Netherlands. After graduation, he joined the Department of Criminal Law and Criminology at Leiden University as a researcher and lecturer. In 2013 he moved to TNO, where he worked as a research scientist on the economics of cybercrime. Predominantly involved in TNO's Dark Web research program, he investigated new and evolving criminal business models.

Early 2015 he joined Delft University of Technology as a PhD candidate. His research was embedded in the MALPAY project - a stakeholder funded project focusing on cybercrime and the financial sector. In this project, he studied the criminal strategies used by cybercriminals for outsourcing key elements in their criminal business model, and the interactions between these strategies and the (security) policies of financial service providers and law enforcement.

During his PhD, he was elected president of the PhD Candidate Network of The Netherlands (PNN). In his two year term (2016-2018), he acted as the primary spokesperson and lobbyist for the nearly 10,000 PhD-candidates in The Netherlands. Right thereafter, he joined a committee tasked to write a new Good Governance-code for the Association of Universities in The Netherlands - which was adopted by the end of 2019.

More closely related to his scientific work, Rolf frequently serves as an expert-witness in cybercrime related court cases, trains law enforcement and judicial professionals in the various aspects of cybercrime policing and is regularly invited for international keynotes and interviews in (inter)national media outlets. In 2017, he spent early summer as a visiting researcher at the criminological research institute Transcrime (Milan, Italy) and was named one of 50 most influential cyber security experts in The Netherlands.

Currently, Rolf works as an assistant professor at the Faculty of Technology, Policy & Management of Delft University of Technology on a range of topics closely related to his PhD research within the field of cybercrime governance.


8x1EF5672C
4x2135BA3E
Dx1265F9E0
6x06078868
Dx26D2C331