

Go See a Specialist? Predicting Cybercrime Sales on Online Anonymous Markets from Vendor and Product Characteristics

van Wegberg, Rolf; Miedema, Fieke; Akyazi, Ugur; Noroozian, Arman; Klievink, Bram; van Eeten, Michel

DOI

[10.1145/3366423.3380162](https://doi.org/10.1145/3366423.3380162)

Publication date

2020

Document Version

Final published version

Published in

Proceedings of The Web Conference (WWW)

Citation (APA)

van Wegberg, R., 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)* (pp. 816-826). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3366423.3380162>

Important note

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

Copyright

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

Takedown policy

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

Go See a Specialist? Predicting Cybercrime Sales on Online Anonymous Markets from Vendor and Product Characteristics

Rolf van Wegberg
Delft University of Technology

Fieke Miedema
Delft University of Technology

Ugur Akyazi
Delft University of Technology

Arman Noroozian
Delft University of Technology

Bram Klievink
Leiden University

Michel van Eeten
Delft University of Technology

ABSTRACT

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 paper, 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 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. Consistent with earlier insights into carding forums, we identify prevalent product differentiators to be influencing the relative success of a product. While all these product differentiators do correlate significantly with product performance, their explanatory power is lower than that of vendor profiles. When outsourcing, the vendor seems to be of more importance to the buyers than product differentiators.

CCS CONCEPTS

• Security and privacy → Economics of security and privacy;

KEYWORDS

Criminal performance, Online anonymous markets, Cybercrime

ACM Reference Format:

Rolf van Wegberg, Fieke Miedema, Ugur Akyazi, Arman Noroozian, Bram Klievink, and Michel van Eeten. 2020. Go See a Specialist? Predicting Cybercrime Sales on Online Anonymous Markets from Vendor and Product Characteristics. In *Proceedings of The Web Conference 2020 (WWW '20)*, April 20–24, 2020, Taipei, Taiwan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3366423.3380162>

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '20, April 20–24, 2020, Taipei, Taiwan

© 2020 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-7023-3/20/04.

<https://doi.org/10.1145/3366423.3380162>

1 INTRODUCTION

Many cybercriminals can be described as freelancers. Specialized in specific tasks, like malware development or cash-out solutions, they trade self-made products and services to other cybercriminals [3, 4, 26]. 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 [24].

Within the B2B cybercrime segments on online anonymous markets, there are significant differences in the types of product offered, sales volume and vendor performance [20, 21, 24]. 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 [7]. 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 [8, 12, 13].

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 entrepreneurs 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 2 discusses the structure of and product differentiators in the market for B2B cybercrime solutions. Section 3 explains our methodology and presents our approach. Section 4 grasps the different product characteristics in newly constructed variables. Section 5 lays down our approach to cluster vendors into distinct profiles, and section 6 identifies predictors for B2B cybercrime sales. Section 7 discusses our findings both in terms of its limitations as well as our public policy take-aways. Section 8 connects our work to earlier contributions and section 9 concludes.

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.

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 [22, 24]. 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 [22]. Apart from drugs, products and services range from physical goods, like passports, to digital goods, like stolen credit cards or malware packages [22, 24]. 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 [2, 24].

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.

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 [10]. 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 [8, 13].

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 [8, 13].

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 [8, 13]. 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 [20]. We will further elaborate on this in Section 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.

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.

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 [22]. 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 [25].

We leverage the parsed and analyzed data set of Van Wegberg et al. [24] 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 [5, 22, 24]. 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. [24] 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.

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 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. [24] that large differences between categories of B2B offerings exist. Table 1 reports the number of listings, vendors, feedbacks and the total revenue per category of B2B cybercrime solutions on AlphaBay, as well as the average price, the revenue and lifespan (in days) of the listings in that category.

In line with previous research [24], 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 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 [6].


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.

Table 1: Listings per category on AlphaBay

Category	Total per category				Average per listing					
	# Listings	# Vendors	# Feedbacks	Revenue	Price	<i>min-max; SD</i>	Revenue	<i>min-max; SD</i>	Lifespan	<i>min-max; SD</i>
App	75	48	571	\$7,420.53	\$18.97	0-207; 34.56	\$98.94	0-840; 161.67	98.45	1-650; 152.28
Botnet	51	37	334	\$9,279.99	\$116.57	1-1,778; 336.63	\$181.96	1-1,778; 395.99	92.20	1-697; 168.43
Cash-out	6,221	1,226	113,897	\$3,341,405.44	\$71.07	0-6,974; 223.41	\$537.12	0-209,124; 3,842.52	85.00	1-798; 140.26
E-mail	377	151	3,412	\$41,191.79	\$34.44	0-1,100; 108.92	\$109.26	0-3109; 319.47	75.48	1-796; 125.24
Exploit	54	37	329	\$3,922.20	\$37.26	0-500; 103.45	\$72.63	0-1000; 159.50	117.30	1-708; 182.47
Hosting	7	6	47	\$423.56	\$21.06	3-50; 17.64	\$60.51	18-136; 45.64	68.14	1-173; 78.57
Malware	149	88	1,140	\$34,921.71	\$48.51	0-500; 95.21	\$234.37	0-5,346; 601.16	91.64	1-762; 149.61
Phone	135	80	1,259	\$52,457.95	\$67.98	0-3,200; 303.64	\$388.57	0-20,910; 1,896.05	92.98	1-745; 146.38
RAT	60	41	425	\$7,035.13	\$42.61	1-648; 122.74	\$117.25	0-1,256; 271.12	112.90	1-706; 184.00
Website	414	178	5,520	\$118,861.15	\$60.49	0-1,695; 158.53	\$ 287.10	0-11,088; 919.96	88.75	1-675; 131.70
Total	7,543	1,339	126,934	\$3,616,919.45						

**Figure 1: Cumulative distribution function of feedbacks and revenues across listings**

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 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 2.

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 such as hacked accounts, spam accounts or spyware. Finally, using the aforementioned list, we find that 849 listings ($\approx 11\%$) 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 ($\approx 14\%$) 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 ($\approx 6\%$) 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 [15]. 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 [15]. 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 ($\approx 64\%$).

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 ($\approx 3\%$) 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 ($\approx 98\%$) 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 ($\approx 2.5\%$) and that 1896 listings ($\approx 25\%$) 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 feedbacks.

Table 2: 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 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 [20]), its time on the market (experience [8, 13, 20]), the relative pricing of its products (price deviation [13]), having listings in one or multiple categories (diversity [20]), the amount of sales (performance [20]) and the sentiment of the feedback (reputation [8]). From earlier research we know that these vendor characteristics, evaluated separately, influence the performance of products on anonymous cybercrime markets [8, 13].

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. [20] 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).

5.1 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 [27], homicide [28] or sex offenders [9]. 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 [29]. 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 [30], 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 [9, 11, 27].

Table 3: 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

Table 3 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.

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.

5.2 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 4). 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). 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 6). 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 7). 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 8). When examining their listings, Loafers appear to sell mainly ‘make money’ and cash-out guides.

Table 4: 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: 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 6: 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 7: 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 8: 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

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 [5, 22, 24]. 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 1.

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 9 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

Table 9: 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** (0.0001)	0.01** (0.0001)	0.01** (0.0001)	0.01** (0.0001)
	Category botnet	-0.31 (0.25)	-0.38 (0.23)	-0.37 (0.24)	-0.42 (0.23)
	Category cash-out	0.62** (0.16)	0.23 (0.15)	0.46** (0.15)	0.21 (0.14)
	Category e-mail	-0.06 (0.17)	-0.44** (0.17)	0.07 (0.16)	-0.28 (0.16)
	Category exploits	-0.77** (0.25)	-0.72** (0.23)	-0.64** (0.24)	-0.62** (0.23)
	Category hosting	-0.28 (0.54)	-0.19 (0.51)	-0.14 (0.52)	-0.17 (0.50)
	Category malware	-0.21 (0.19)	-0.30 (0.18)	-0.16 (0.18)	-0.25 (0.18)
	Category phone	-0.21 (0.20)	-0.45* (0.19)	-0.25 (0.19)	-0.43* (0.18)
	Category RAT	-0.84** (0.24)	-0.96** (0.23)	-0.61** (0.22)	-0.76** (0.22)
	Category website	0.07 (0.17)	-0.18 (0.16)	0.09 (0.16)	-0.11 (0.16)
Vendor Variables	Vendor profile Generalist		0.15** (0.05)		0.16** (0.05)
	Vendor profile Loafer		-0.61** (0.07)		-0.53** (0.07)
	Vendor profile Professional		1.05** (0.05)		0.93** (0.05)
	Vendor profile Specialist		-0.48** (0.07)		-0.06 (0.08)
	Refund policy			0.65** (0.04)	0.43** (0.04)
Product Variables	Terms & Conditions			0.02 (0.07)	0.003 (0.06)
	Price deviation			-0.26** (0.02)	-0.27** (0.02)
	Customer support			0.50** (0.05)	0.36** (0.05)
	Mean sentiment			0.12** (0.05)	0.20** (0.04)
	Ratio special characters			-0.46* (0.20)	-0.86** (0.20)
	Ratio capitals			1.23** (0.06)	0.85** (0.06)
	Use of vendor name			0.76** (0.09)	0.57** (0.09)
	Constant	1.40** (0.16)	1.37** (0.15)	0.80** (0.15)	0.88** (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

Note: *p<0.05; **p<0.01

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 9 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 [7, 19] 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.

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.

7.1 Limitations

First and foremost, given the fact that we use scrape 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 [5, 22, 24]. 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 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' [8, 13].

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.

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.

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 [22]. Most existing studies include, or even focus on, drugs and physical goods, which represent a large share of the products offered on the markets [1, 2, 23]. In contrast, and most closely connected to our work, Van Wegberg et al. [24] 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 [8, 13, 20]. Both Decary-Hetu & Leppanen [8] and Holt et al. [13] 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. [20] 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. [16] yielded relevant insights into the offline contacts of online criminals. This offline angle in collaboration was also investigated by Lusthaus [17, 18] who interviewed over one hundred cybercriminals and unraveled how and where collaborations start. Next, Hutchings [14] studied the sharing of techniques amongst cybercriminals and identified distinct collaboration types, ranging from one-time partners to sustainable partnerships.

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 [8, 13]. 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 outperform ‘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 [2, 22, 25]. 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.

REFERENCES

- [1] 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>
- [2] Judith Aldridge and David Décary-Héту. 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>
 - [3] 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
 - [4] Juan Caballero, Chris Grier, Christian Kreibich, and Vern Paxson. 2011. Measuring Pay-per-Install: The Commoditization of Malware Distribution. In *Usenix Security Symposium*.
 - [5] 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> arXiv:1207.7139
 - [6] 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-anonymous-online-marketplace.pdf>
 - [7] David Décary-Héту 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>
 - [8] David Décary-Héту 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>
 - [9] Nadine Deslauriers-Varin and Eric Beaugregard. 2010. Victims' routine activities and sex offenders' target selection scripts: A latent class analysis. *Sexual Abuse* 22, 3 (2010), 315–342.
 - [10] Peter R Dickson and James L Ginter. 1987. Market segmentation, product differentiation, and marketing strategy. *Journal of marketing* 51, 2 (1987), 1–10.
 - [11] 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>
 - [12] Andreas Haslebacher, Jeremiah Onalapo, 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.
 - [13] 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>
 - [14] 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.
 - [15] 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.
 - [16] 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>
 - [17] Jonathan Lusthaus. 2018. *Industry of Anonymity: Inside the Business of Cybercrime*. Harvard University Press.
 - [18] 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>
 - [19] Carlo Morselli, David Décary-Héту, Masarah Paquet-Clouston, and Judith Aldridge. 2017. Conflict management in illicit drug cryptomarkets. *International Criminal Justice Review* 27, 4 (2017), 237–254.
 - [20] Masarah Paquet-Clouston, David Décary-Héту, and Carlo Morselli. 2018. Assessing market competition and vendors' size and scope on AlphaBay. *International Journal of Drug Policy* (2018). <https://doi.org/10.1016/j.drugpo.2018.01.003>
 - [21] 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.
 - [22] 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>
 - [23] 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>
 - [24] Rolf van Wegberg, Samaneh Tajalizadehkhooob, Kyle Soska, Ugur Akyazi, Carlos Hernandez Ganán, 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*.
 - [25] R.S. 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). <https://www.narcis.nl/publication/RecordID/oiatudelft.nl:uuid:8c080055-37fb-4f53-a949-099110f91659>
 - [26] 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>
 - [27] 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>
 - [28] 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>
 - [29] 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.
 - [30] J.K. Vermunt and J. Magidson. 2016. *Guide for Latent GOLD 5.1: Basic, Advanced, and Syntax*. Technical Report. Statistical Innovations Inc., Belmont, MA.