

# Economic Factors of Vulnerability Trade and Exploitation

Luca Allodi

Eindhoven University of Technology  
P.O. Box 513, Eindhoven, The Netherlands  
l.allodi@tue.nl

## ABSTRACT

Cybercrime markets support the development and diffusion of new attack technologies, vulnerability exploits, and malware. Whereas the revenue streams of cyber attackers have been studied multiple times in the literature, no quantitative account currently exists on the economics of attack acquisition and deployment. Yet, this understanding is critical to characterize the production of (traded) exploits, the economy that drives it, and its effects on the overall attack scenario. In this paper we provide an empirical investigation of the economics of vulnerability exploitation, and the effects of market factors on likelihood of exploit. Our data is collected first-handedly from a prominent Russian cybercrime market where the trading of the most active attack tools reported by the security industry happens. Our findings reveal that exploits in the underground are priced similarly or above vulnerabilities in legitimate bug-hunting programs, and that the refresh cycle of exploits is slower than currently often assumed. On the other hand, cybercriminals are becoming faster at introducing selected vulnerabilities, and the market is in clear expansion both in terms of players, traded exploits, and exploit pricing. We then evaluate the effects of these market variables on likelihood of attack realization, and find strong evidence of the correlation between market activity and exploit deployment. We discuss implications on vulnerability metrics, economics, and exploit measurement.

## KEYWORDS

Cybercrime; security economics; exploit and vulnerability trade

## 1 INTRODUCTION

The rapid expansion of the cyber-threat scenario reported in the recent literature is fostered by the presence of an ‘underground’ economy that supports the development, deployment, and monetization of cyber-attacks [35, 36]. A few studies analyze the dynamics of the underground and the markets that drive it [8, 36, 37], focusing on either the economic mechanisms that enable the market activity [8, 37], or the ‘after the fact’ analysis of its effects in the real world [15, 36]. However, it remains impossible to fully characterize the underground production of cyber-attacks without a quantitative account of its economic aspects. For example, several

allegations currently exist on the costs of an exploit in the underground markets. Similarly, the ‘economy of malware’ is thought to have significant repercussions on the realization of real-world attacks, yet no scientific account of this relation is currently present in the literature. Likewise, legitimate vulnerability markets [63, 77] have been designed to ‘compete’ with cybercrime markets, but the two systems remain hardly comparable without a more precise account of their economic aspects. In this paper we fill this gap by focusing on the economic aspects of exploit acquisition and deployment, hence providing an additional piece in the otherwise incomplete cybercrime puzzle.

Part of the reason why such analyses are scarce in the literature is the difficulty of identifying and studying ‘good’ or ‘influential’ underground markets. Criminal markets are known to be fraught with economic problems that hinder fair trade, and consequently market development [37, 69]. Moreover, markets trading attack technologies tend to be strongly segregated [8, 68, 69, 75], making their access and study more difficult to accomplish. For example, common segregation mechanisms include the implementation of *pull-in* mechanisms, language barriers (especially Russian/Portuguese/Chinese), and ingress monitoring [8]. On the other hand, non-English speaking attackers reportedly generate a significant fraction of attacks [43, 78]; this may be partly due to the still loose international regulation of the cyber-space [44], as well as economic and social aspects on welfare and higher education in developing countries [25, 49, 50]. For example, Russian cybercriminals are known to produce malware that does not attack ex-Soviet nations (ex-CIS), in an attempt to not catch the attention of the local authorities. On the other hand, a significant fraction of the malware detected at scale, as well as attack vectors such as exploit kits and booter services [35, 42, 48] are suspected to have been engineered by Russian attackers [43, 48, 70].

In this paper we characterize the economic aspects of vulnerability exploits as traded in a prominent Russian cybercrime market (RuMarket<sup>1</sup>) for user infections at scale (as opposed to targeted or 0-day attacks), and of their effect on risk of exploit in the wild. Through market infiltration we collect information on trade of vulnerability exploits spanning from 2010 to 2017, and correlate this data with Symantec data on exploits detected at scale. Our contribution can be summarized as follows:

- (1) The time between vulnerability disclosure and appearance of exploit in RuMarket is shortening, showing that attackers are becoming more reactive in delivering selected exploits. At the same time, number of actors and of traded exploits is increasing; exploit prices are inflating, and exploit-as-a-service models appear to allow for drastic cuts in exploitation costs. This is particularly relevant for the development of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CCS '17, October 30-November 3, 2017, Dallas, TX, USA

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4946-8/17/10...\$15.00

<https://doi.org/10.1145/3133956.3133960>

<sup>1</sup>We do not disclose the real name of the market to minimize threats to our anonymity.

economic models of the underground, and impacts attacker and risk modeling.

- (2) We find strong evidence of the relation between market activity and likelihood of exploitation in the wild. We find that exploits that spawn higher levels of discussion in RuMarket are associated with higher odds of exploitation in the wild, and that high market prices hinder exploit deployment at scale. This provides a quantitative link between attacker economics and attack realization, and can directly contribute in the development of (more) realistic attacker models.
- (3) Exploit prices in the underground markets are aligned with or above those of analogous 'legitimate' markets for vulnerabilities and vary between 150 and 8000 USD, whereas the arrival of new exploits is significantly slower than otherwise often assumed. This provides insights on the incentives to participate in the underground economy, and on the dynamics of exploit introduction.

*Scope of work.* The goal of this paper is to provide insights on exploit economics by analyzing one prominent Russian cybercrime market; this work *does not* aim at providing a full enumeration of *all* exploits traded in the underground: other cybercrime markets may feature different sets of exploits and/or foster different cybercriminal activities, as well as enforcing different market regulation mechanisms. Similarly, RuMarket *does not* focus on 0day exploits, whose employment for attacks at scale is reportedly very limited [18], and are outside of the scope of this work.

This paper proceeds as follows. The following section discusses related work. Methodology, data collection, and analysis procedure are discussed in Section 3. Section 4 reports our analysis, and Section 5 discusses our results.

## 2 RELATED WORK

The economics and development of underground markets have perhaps been first tackled by Franklin et al. [33]. On the other hand, Herley et al. [37] showed that cybercrime economics are distinctively problematic in that the lack of effective rule enforcement mechanisms may hinder fair trading, and as a consequence the existence of the market itself. A few studies analyzed the evolution of cybercrime markets [8, 23, 35, 53, 68, 75], and provided estimates of malware development [22] and attack likelihood [11], but no quantitative account of economic factors such as exploit pricing and adoption are currently reported in the literature [22, 63]. In this paper we provide the first empirical quantification of these economic aspects by analyzing data collected first-handedly from a prominent cybercrime market.

Recent work studied the services and monetization schemes of cyber criminals, e.g. to launder money through acquisition of expensive goods [36], or renting infected systems [35, 42]. The provision of the technological means by which these attacks are perpetrated remain however relatively unexplored [63], with the exception of a few technical insights from industrial reports [43, 70]. Similarly, a few studies estimated the economic effects of cybercrime activities on the real-world economy, for example by analyzing the monetization of stolen credit cards and banking information [15], the realization of profits from spam campaigns [45], the registration of fake online accounts [71], and the provision of booter services

for distributed denial of service attacks [47]. However, a characterization of the costs of the technology (as opposed to the earnings it generates), and the relation of trade factors on the realization of an attack is still missing. This work provides a first insight on the value of vulnerability exploits in the underground markets, and the effects of market factors on presence of attacks in the wild.

The presence of a cybercrime economy that absorbs vulnerabilities and generates attacks motivated the security community to study the deviation of 'legitimate' vulnerability markets that attract security researchers away from the illegal marketplaces [77]. Whereas several market mechanisms have been proposed [46, 57], their effectiveness in deterring attacks is not clear [46, 52, 61]. The so-called responsible vulnerability disclosure is incentivized by the presence of multiple bug-hunting programs by several providers such as Google, Facebook, and Microsoft, or 'umbrella' organizations that coordinate vulnerability reporting and disclosure [30, 63, 77]. It is however unclear how do these compare against the cybercrime economy, as several key parameters such as exploit pricing in the underground are currently unknown. Further, it remains uncertain whether the adoption of vulnerability disclosure mechanisms has a clear effect on risk of attack in the wild [61]. This study fills this gap by providing an empirical analysis of exploit pricing in the underground, and evaluating the effect of cybercrime market factors on the actual realization of attacks in the wild.

## 3 METHODOLOGY

Sections 3.1 to 3.3 present our methodology and provide a detailed description of our data. In Sec. 3.4 we outline the analysis procedure, assumptions, and data handling. Sec. 3.5 discusses observational biases of this study, and Sec. 3.6 addresses ethical aspects.

### 3.1 Market infiltration and evaluation

RuMarket is a forum-based market that can be reached from the open Internet. Access to the market requires explicit admission by the market administrators, who validate the access request by performing background checks on the requester. The main criterion for admission is the ability to demonstrate that the requester has control over other identities in affiliated Russian hacking forums, and that he/she has been active in the community.<sup>2</sup> Gaining access to RuMarket required approximately six months to build a credible profile, identify affiliated markets, and letting our alter-ego gain reputation within the hacking community. These activities required some level of proficiency in Russian. Section 3.6 provides a discussion on ethical aspects. As members of the market, we have access to all the (history of) product information, trades, and prices available to active participants. This analysis spans seven years of market activity (July 2010 - April 2017).

*Criteria for market evaluation.* It is important to first evaluate whether the selected market is a credible candidate for analysis, or is yet another example of many 'scam-for-scammers' underground forums [33, 37]. Following indications in the literature on the poor implementation of cybercrime markets [37, 75], in previous work [8] we performed an analysis of the markets' economic

<sup>2</sup>Admission criteria were initially enforced upon the 2013 arrest, performed by the Russian authorities, of a prominent market member. The enforcement of the admission criteria, albeit still present at the time of writing, has now loosened up.

**Table 1: Exploit identification and market activity**

Type	Ads	Vulns		Replies		Trade evidence		
		Tot.	Avg.	Tot.	Avg.	Yes	No	%
CVE	9	30	3.3	518	57.5	5	4	55%
Knwn	4	4	1	55	13.8	2	2	50%
0day	1	1	1	44	44	0	1	0%
Und.	5	-	-	65	13	1	4	25%

mechanisms (e.g. addressing information asymmetry [6], adverse selection, and moral hazard [28, 37]), traded goods, and participation. We here provide a summary of the considered criteria. A complete account of these aspects is given in [8].

*Cr.1 Enforcement of market regulation mechanisms;* market mechanisms enforcing market rules, such as punishment for rippers or presence of trade guarantors or escrows are known to be central to address foundational problems that cripple the economics of cybercrime markets [6, 28, 32, 37, 75] and hinder product quality. We found that in RuMarket, rippers are systematically punished, most sellers use the market escrow services to guarantee transactions, and that the high costs of market entry hinder unfair behavior.

*Cr.2 Evidence of trade.* We evaluated face evidence of actual trading activity in the market. Accounting for indications from economic literature [62, 67], we investigate trade-related feedback from market participants, discussions in the market threads, product evolution, and type of market interactions; all evidence points toward effective trade mechanisms that foster trading activity.

*Cr.3 Presence of prominent attack tools reported by the industry.* The relevance of RuMarket in the threat scenario is supported by the presence of traders for the most prominent attacks reported by the industry. Among those, we find several exploit kits [35] (e.g. Blackhole, RIG, Eleonore [3, 70]) and malware that led numerous infection campaigns (e.g. Zeus, Citadel [17, 19, 60]).

### 3.2 Sampling exploits in the underground

The unstructured nature of forum-based markets calls for a few additional considerations on data sampling: whereas most (criminal) goods such as drugs, weapons, and illegal pornography can be easily identified and described or demoed by vendors (and therefore measured by investigators [69]), the disclosure of too much information on an exploit would destroy its value [52], whereas revealing too little eventually leads to market death (as buyers cannot assess what they buy) [37]. In order to meaningfully sample data points, it is therefore critical to identify the exploit reporting criteria adopted by vendors.

To this aim, we randomly sample 50 posts from RuMarket generally referring to “эксплоит” (*exploit*) and (slang) variants thereof, and evaluate the type of reporting and received market response (i.e. number of replies, and trade evidence [69]). In our sample we find 19 ads selling 35 exploits overall. Four reporting mechanisms emerge: using the standard Common Vulnerabilities and Exposures (CVE) identifier [2]; describing an exploit as affecting a disclosed vulnerability (Knwn); describing it as a 0day; not describing it at all (Und.). Table 1 summarizes the results. Overall, we find nine adverts reporting 30 CVEs, one reporting a single 0day, and four reporting

one Knwn vulnerability each. Five additional posts (Und.) advert an undefined number of vulnerabilities without further details on affected software or type of exploit. The first observation is that posts reporting CVEs trade on average significantly more vulnerabilities than other posts ( $p = 0.03$  for a Wilcoxon rank sum test), indicating that this category likely represents the great majority of marketed exploits. Adverts reporting CVEs also show greater market activity, measured in terms of received replies ( $p = 0.05$ ), than other posts. Similarly, poorly described vulnerabilities are unlikely to show any evidence of trade, whereas CVE and Knwn vulnerabilities display similar rates. Overall, we find that vulnerabilities reported by CVE represent the significant majority of traded vulnerabilities (by almost an order of magnitude), and receive significantly more attention from the RuMarket community than the aggregate.

Further supporting the relevance of CVE reporting, we find that market participants actively look for CVE information when not immediately available; for example, an interested buyer of the KTR package asks (translated from Russian): “*Which exploits are bundled in the pack at the moment? If possible, specify the CVE*”; the seller complies. Numerous other examples go in this same direction. Critically, this mechanism allows buyers to perform a first assessment of the exploit, and to verify that the characteristics of the vulnerability it exploits match the vendor’s claims (e.g. allow for remote code execution or privilege escalation); this, alongside other market mechanisms described in previous work [8], directly addresses the problem of adverse selection, foundational to all markets of this type, and first underlined in [37]. Indeed, vulnerability identification is part of the regulation of the market itself: for example, a vendor was blocked by the forum administrators when trying to sell (for 3000 USD) an identifiable Windows PoC (CVE-2012-0002); the admin explains: “*This [exploit] is public (if not today, tomorrow). The DOS proof-of-concept is already public. Such sales are prohibited*”.

For these reasons in this study we use CVE-IDs as a sampling mechanism for traded exploits. This has also the advantage of allowing us to precisely measure additional characteristics of the vulnerability, including date of disclosure, technical severity, affected software, and presence in the wild, all of which would be impossible without a rigorous definition of published exploits. Importantly, this also rules out errors caused by double counting vulnerabilities, while accounting for the vast majority of published exploits (*ref.* Tab 1). The remaining bias is discussed in Sec. 3.5.

### 3.3 Data collection

In this analysis we employ three datasets. The collected data fields are reported in Table 2; in the Appendix we report an extended description of each field.

- (1) RuMarket. We query RuMarket and analyze results by reading discussion topics and extrapolating relevant information. Unfortunately the nature of the data limits the applicability of fully-automated data extraction procedures (e.g. product updates and multiple products per advert, see also [58]). We therefore employ semi-automated pattern matching and manual analysis to extract the information. We identify traded CVEs by querying RuMarket for matches to the case-insensitive regular expression `cve(-id)?(?i)` in the *Virus, attacks, and malware* commercial section of

**Table 2: Summary of collected variables used in the analysis**

Description and summary statistics of the collected data fields. *Unit* indicates the type of data field. *Lvls* indicates, for categorical variables, the number of factor levels. Descriptive statistics are provided for cardinal/ordinal variables, and categorical variables with only two factors (encoded as 1: presence of condition; 0: absence of condition). From the product descriptions naturally emerge the following package categorization (vulnerability descriptions from NVD [2]):

a) STANDALONE: packages traded as stand-alone exploits that are then personalized by the buyer. For example, CVE-2016-0189 “allow[s] remote attackers to execute arbitrary code [...] via a crafted web site”, and is traded in RuMarket as a STANDALONE exploit to which the customer can add his/her own ‘private’ (sic.) shellcode.

b) MALWARE: exploits embedded in malware packaging services. Exploits in these packages typically allow for privilege escalation. For example, CVE-2015-1701 “allows local users to gain privileges via a crafted application”, and in RuMarket is bundled in a MALWARE dropper that, when executed on the target machine, escalates to higher privileges and executes the custom code.

c) EKIT: exploit packages typically rented (as opposed to traded) as exploit kits, namely web servers that deliver exploits and custom payloads to victims that are redirected to the kit [48]. The rental period in our sample goes from a week to a month. EKITs operation requires the execution of arbitrary code on the victim system to remotely drop the malware. For example, CVE-2016-1019 “allows remote attackers to [...] execute arbitrary code”, and in RuMarket it is embedded in the notorious exploit kit RIG.

Variable	Dataset	Unit	Description	Lvls	Min	Mean	Max	sd
CVE	RuMarket, NVD, SYM	Cat.	The unique identifier of the vulnerability.	57				
CVEPub	NVD	Date	Date of vuln disclosure in NVD.		2006-04-11	2012-12-06	2016-11-10	970.84
Exp1Ven	RuMarket	Cat.	The identifier of the product vendor.	23				
Exp1VenReg	RuMarket	Date	Date of vendor registration in the market.		2008-05-25	2012-10-31	2016-03-27	873.05
Pack	RuMarket	Cat.	Bundle of exploits traded in the market.	38				
PubDate	RuMarket	Date	Date of exploit introduction in a package.		2010-07-29	2014-06-25	2017-01-19	640.02
PackType	RuMarket	Cat.	Pack classification in one of the categories STANDALONE; EKIT; MALWARE.	3				
PackPrice	RuMarket	USD	Acquisition cost of the package.		100	2417	8000	2408.28
PackActiv	RuMarket	Messages	Number of responses to package advert.		0	43.85	300	69.29
PackDeath	RuMarket	Date	Date of last reply for the package.		2010-12-24	2015-03-16	2017-03-27	642.63
Exp1Price	RuMarket	USD	Price estimate of single exploit.		13.64	969.00	8000	1708.76
SwVen	NVD	Cat.	Vendor of the vulnerable software.	3				
Sw	NVD	Cat.	Name of the affected software.	7				
CVSS	NVD	Ord.	Vulnerability severity measured by the Common Vulnerability Scoring System.		5	8.76	10	1.33
SYM	SYM	Cat.	Presence of exploit at scale.	2	0	0.84	1	0.38

the forum. This procedure returned 194 discussion threads and approximately 3000 posts to examine in April 2017. To minimize the chances of reporting ‘fake’ exploit products, we consider only vendors that have not been reported as ‘rippers’ or banned from the community. This leaves us with a sample of 89 traded exploits over 57 unique vulnerabilities embedded in 38 packages for STANDALONE, MALWARE and ExploitKit products, and attacking Microsoft, Oracle,<sup>3</sup> and Adobe software. This is quantitatively in line with previous studies on marketed exploits [3, 4, 9, 48].

- (2) NVD. The National Vulnerability Database [2], is the NIST-maintained vulnerability database reporting vulnerability characteristics, affected software, and severity.
- (3) SYM. Vulnerabilities for which Symantec’s threat explorer and attack signature databases report an exploit in the wild [26]. Vulnerabilities outside of SYM may still be actively exploited, but are unlikely to be exploited at scale [9, 26]. This allows us to correlate technical and market characteristics of vulnerabilities to the actual (mass) realization of an observed exploit in the wild.

We join the three datasets on the CVE-ID of the vulnerability.

<sup>3</sup>All vulnerabilities labeled as Oracle are relative to the Java platform. Some of those were disclosed while Java was Sun’s.

### 3.4 Analysis procedure

**3.4.1 Estimation of exploit prices.** When a package contains more than one exploit, the cost of a single exploit can only be estimated. From the literature on exploit development and deployment [9, 24, 29, 30, 48, 76] two aspects of vulnerabilities emerge as drivers of exploitation effort: 1) vulnerability type (e.g. memory corruption vs cross-site-scripting) [29, 48, 76]; 2) exploitation complexity (e.g. to evade attack mitigation techniques) [20, 24, 29].

**Vulnerability type.** The MITRE corporation maintains a community-developed standard (Common Weakness Enumeration, CWE in short) for the enumeration of software weaknesses [2]. This has the purpose of identifying the type of technical issue that generates the vulnerability. Table 12 in the Appendix provides a detailed breakdown of vulnerability CWE types by exploit package in RuMarket. We find that packages typically embed vulnerabilities of the same type (e.g. either remote code execution vulnerabilities or privilege escalation vulnerabilities), which suggests that significantly skewed distributions of exploitation efforts by vulnerability type within a package are unlikely.

**Exploit complexity.** The Common Vulnerability Scoring System (CVSS) [31] defines *Access Complexity* as a measure of whether a reliable exploit can be ‘easily’ obtained or additional measures or attacks are required to, for example, avoid attack mitigation techniques (memory randomization, canaries, etc.), or address specific software/system architectures [29, 31]. CVSSv2 assesses attack complexity in three categories: High, Medium, Low [65]. AC:High conditions have been shown to represent a threshold for exploit

adoption [9], whereas Medium and Low vulnerabilities require only limited exploitation efforts [65] and are commonly detected in the wild [9, 20]. Acknowledging this, the newer version of CVSS (v3) considers only High (existence of conditions outside of the attacker’s control) or Low (absence of conditions) values [31]. Out of 57 unique CVEs in our sample, we find 2 vulnerabilities characterized by a High CVSS attack complexity, whereas the remaining 55 include only limited or no exploitation complexities for the attacker. A further breakdown of attack complexity by package (Table 11 in the Appendix) shows that most packages prevalently include vulnerabilities with the same AC assessment. This once again suggests that exploit development efforts are not significantly skewed among vulnerabilities bundled in a package.

In light of these considerations, in this study we estimate unitary cost of exploit by assuming a uniform distribution of costs among exploits in a package.<sup>4</sup>

**3.4.2 Bootstrapped analysis of exploit prices.** In an effort to provide a more precise estimate of exploit costs, we employ a block bootstrap analysis ( $N = 10000$ ) of exploit prices. The bootstrap procedure randomly re-samples ( $N$  times), with replacement, exploit packages (i.e. our ‘blocks’, or sampling units) from RuMarket, and approximates the true unknown distribution of the population of traded exploits (of which we observe a sample) [27]. This allows us to infer the parameters of the true distribution and to build robust confidence intervals of price estimates.

**3.4.3 Regression analysis.** The nature of the sample requires a few additional precautions to be taken for a formal analysis. In particular, our exploit observations depend not only on the exploit, but also on the specific vendor who publishes the package where the exploit is bundled in. For example, qualified vendors may publish more reliable exploits that are more likely to generate attacks in the wild. Hence, the measure of an exploit implicitly depends on the vendor who publishes it (i.e. our data has an hierarchical structure [5]).<sup>5</sup> This ‘mixed effect’ should be captured to assure an unbiased quantitative analysis. We denote  $\mu_a$  as the (univariate) random effect for the vendor  $a$  such that the expected value of the observation for the  $i^{th}$  measurement is  $E(Y_{ai}|\mu_a) = \mu_{ai}$ , i.e. the expected value for the  $i^{th}$  observation is conditional on  $\mu_a$ . The general regression form of our analysis is derived from [5] and is:

$$f(\mu_{ai}) = z_{ai}\mu_a + \beta\mathbf{x}_{ai} \quad (1)$$

where  $f$  is the link function,  $z_{ai}\mu_a$  quantifies the random effect at the intercept, and  $\beta\mathbf{x}_{ai}$  is the vector of fixed effects and respective coefficients. Standard model diagnostics are run for all regressions. We report model Log-likelihoods for model comparison. The calculation of  $p$ -values and model power are not straightforward for mixed effects models. We report *pseudo* –  $R^2$  and  $p$  – values as approximations provided by the R packages lmerTest and MuMIn, alongside the standard deviation of coefficient estimation.

<sup>4</sup>Sec. 4.2 gives a detailed account of how this relates to different software packages.

<sup>5</sup>Exp1Ven and Package are both meaningful levels in the hierarchy of our sample. As it is the vendor of the exploit that publishes the exploits, fixes exploit cost, and determines exploit quality, we here consider the vendor as the main source of variance.

### 3.5 Limitations

The adopted CVE sampling mechanism may exclude some potentially relevant vulnerability that we cannot measure precisely. Results in Tab 1 indicate that this sampling bias is likely minimal. It is however worth noting that this is, unfortunately, an inherent limitation of all studies on this type of markets: without engaging in the trading activity, it is impossible to reliably measure what lies behind a market post. For example the excellent work in [69] conservatively estimates market size by assuming that user feedback relates to separate, single trade lots, as it is not possible to measure multiple trades in a single transaction. Similarly, as we cannot measure unidentified exploits, our analysis should be considered a conservative estimate of traded exploits in RuMarket.

The data collection in SYM reports exploits deployed *en-mass* against consumer (typically Windows) systems, and does not directly extend to targeted attacks and 0-day vulnerabilities.

### 3.6 Ethical aspects and data sharing

The market infiltration was performed while the author was at the University of Trento, Italy. All data collection happened at the Eindhoven University of Technology, the Netherlands. No activity involved the deception of market participants other than for our ‘identity’. We only engaged in discussion on non-hacking topics not to facilitate illegal activities. We use the anonymous network TOR to conceal our identities. To preserve our anonymity in the market we do not disclose the real name of the community. The collected data is available for sharing.<sup>6</sup>

## 4 DATA ANALYSIS

This analysis is structured in three parts. In the first (Sec. 4.1) we describe the market by analyzing the activity of market participants and the characteristics of the traded exploit packages. In the second part (Sec. 4.2) we analyze market factors driving exploit prices, and in the third (Sec. 4.3) the adoption of exploits in the wild.

### 4.1 Overview of RuMarket

**4.1.1 Exploit vendors.** It is first important to provide an overview of the exploit vendors that participate in RuMarket activities. Our RuMarket sample contains 22 uniquely identified vendors that trade CVEs in MALWARE, STANDALONE, and EKIT packages. The market mechanism generates strong disincentives for the creation of multiple accounts [8]. Following the approach adopted for similar applications in related work [69], in the following we consider vendor aliases as unique seller identifiers. Figure 1 shows the appearance of vendors trading exploits in RuMarket. The solid line reports the count of new vendors appearing in the market (i.e. vendors that did not publish a CVE exploit in the preceding years in RuMarket under the same alias); the dotted line reports the cumulative count of vendors. The number of vendors increases at a steady linear rate of approximately three new vendors per year during the observation period. This suggests that exploit trading in RuMarket is growing. Figure 2 offers a breakdown of vendor activity by product type by plotting the CDF of number of exploit packages introduced by each vendor and the number of exploits they embed in their products.

<sup>6</sup>Access procedure available at <http://security1.win.tue.nl>.




Figure 1: Count of vendors trading CVEs in RuMarket




Figure 2: Distribution of packages and exploits by vendor

EKIT vendors typically publish only one product, whereas MALWARE and particularly STANDALONE vendors appear to trade significantly more packages. This is interesting to observe as EKIT products (and to a lesser degree MALWARE products [60]) are traded under the ‘exploit-as-a-service’ model, whereby the seller maintains a service for a period of time during which customers rent the kit to deliver their own attacks. The maintenance operations include delivering vulnerable traffic to the customers, updating the exploit portfolio, and packing existing exploits to minimize detection in the wild (to generate the so-called *FUD*, ‘Fully UnDetectable’ exploits) [48]. The implied prolonged contractual form explains the prevalence of vendors with only one exploit package in their portfolio for EKIT and MALWARE vendors. On the other hand, EKIT vendors are by far the more ‘productive’ in terms of number of exploited vulnerabilities, with 50% of EKIT vendors contributing more than 10 exploits. STANDALONE vendors typically focus on a few exploits only, trading on average below three exploits, and only a small fraction of vendors trades overall more than 5 exploits. MALWARE vendors are the least productive in terms of exploited CVEs: whereas historically Internet worms and malware such as Slammer or Conficker exploited software vulnerabilities to replicate, in recent years infections happen mostly through Malware Distribution Networks [34, 35, 59] that implement the target exploitation by other means (e.g. exploit kits or ‘malvertising’), and allow for the malware to be ‘dropped’ on the attacked system, with only a few exceptions.

Foundational studies in economics [67] as well as more recent research on online marketplaces [21] put the emphasis on the relation between (expectation of) product quality and placement of the vendor in the market. Due to the unreliability of user feedback on online forums, criminal online markets often employ as a proxy for trustworthiness criteria such as time-on-market or number of




Figure 3: Distribution of exploit vendor age




Figure 4: Release of exploit packages by type per year

messages/specialty [8, 39]. These are costlier for malicious vendors to replicate than simply posting positive feedback on their own products. We use as a proxy measure of seller presence in the market the number of days the vendor have been registered on RuMarket at the time of package publication, and calculate it as  $ExpVenAge = PubDate - ExpVenReg$ . Figure 3 reports this distribution by package type. Exploit vendor age varies considerably by type of package. STANDALONE vendors are those with the highest average time on market at time of product publication. 50% of STANDALONE vendors have been registered on the market for at least six months, whereas only the top 30% of EKIT and MALWARE vendors are above this threshold. Overall, we find that only 18% of vendors publish their first package on the day of registration. 55% of vendors have been registered for at least a month, and 32% for at least a year. This indicates that RuMarket mechanisms encourage prolonged market activity, which may determine higher levels of trust among market participants [39].

**4.1.2 Exploit packages.** Our RuMarket sample reports data on 38 unique exploit packages; the breakdown is as follows: six EKIT, six MALWARE, and twenty-six STANDALONE packages. We consider the addition of new exploits in a pack as an update to an existing package. Figure 4 reports the number of updates to the exploit package portfolio by year in the market. In general, we can observe that the number of released products steadily increases every year. This trend appears to be mainly driven by STANDALONE packs, whereas EKIT and MALWARE packages are essentially stable in time, with the latter being the lowest on average. This is in line with the exploit authors’ activity described above, and suggests that these packages may enjoy a longer activity in the market. To evaluate this, we consider the days and volume of active discussion since publication in the market as a proxy measure of RuMarket’s interest in the product. We compute days of active discussion as






Figure 5: Market activity by exploit package

DaysActive=PackDeath-PubDate, and report PackActivity for volume of active discussion. Figure 5 plots the two distributions. The left plot reports longevity of RuMarket activity around an exploit package. Inspection of the market message board reveals that the length of activity around a package is not artificially inflated by the package vendor by continuously adding comments to the advert. As expected, we observe that longevity of discussion around EKIT packages is significantly higher than for STANDALONE and MALWARE packages ( $p = 0.015$  and  $p < 0.01$  respectively for a Wilcoxon rank sum test). RuMarket discussion around EKIT packages remains active for more than 500 days (approximately a year and a half) for 50% of packages, with the top 10% products remaining active in the market board for more than 3 years. Differently, 50% of STANDALONE and MALWARE packages remain active for up to approximately 220 days, and only less than 25% remain active for more than 500 days. Overall, we find that the average package remains active for a year since time of publication. The right plot in Figure 5 plots the distribution of replies by package type. EKIT packages receive on average significantly more replies than other pack types, which is in line with previous figures. Conversely, RuMarket interest around STANDALONE and MALWARE packages is significantly lower, with only a handful of packages receiving a comparable volume of discussion as the average EKIT. The lower interest of the RuMarket community may be driven by the higher difficulty of use of STANDALONE and MALWARE products, that require additional effort to deploy and deliver the attack compared to EKIT products [35]. Further, different price-tags, investigated below, may explain the overall market interest.

*Exploit pack prices.* Table 3 reports descriptive statistics of exploit pack prices and number of bundled exploits by package type. STANDALONE packages are traded at a mean price around 3000 USD up to 8000 USD, and bundle in between 1 and 4 exploits. The small standard deviation indicates that most STANDALONE packages bundle 1 exploit only. MALWARE packages are traded at a price range between 400 USD and 4000 USD, with most package prices set at around the 1000-2000 USD mark. Similarly to STANDALONE packages, MALWARE bundles typically include only one exploit, and up to three exploits. Finally, the lower 50% of EKITs are priced (accounting for an average rent of 2-3 weeks [41]) in the range 150-400 dollars, whereas the upper 50% are in the range 400-2000 USD. EKIT packages embed significantly more exploits than other package types. This is in line with previous findings in the literature [48] and, following Grier et al. [35], this allows for a greater flexibility




Figure 6: Log package price by year by product type

in terms of the range of selectable targets [10]. We give an account of the specific exploits in the next section.

Following [63], we further investigate possible outliers in our data to mitigate pricing noise. We find that the only four STANDALONE packages that received no trade reply from the RuMarket community were also traded at below average prices (in between 100/300 dollars each in year 2016). Similarly, we find only one EKIT that, despite embedding twelve exploits, is priced at 150 dollars, significantly below the EKIT mean of 560 USD. Figure 6 plots the results in terms of trend in price per product type by year with these outliers removed. MALWARE and STANDALONE packages show an increasing trend whereas EKIT product prices are steadily decreasing. Regression coefficients for the linear model displayed in Figure 6 are significant at the 5% level for MALWARE ( $\beta = 0.43, p = 0.05$ ) and EKIT ( $\beta = -0.28, p = 0.02$ ) but not significant for STANDALONE packages ( $\beta = 0.23, p = 0.14$ ). We find that ‘consumer’ services such as EKIT products are becoming more easily available to the users, a figure compatible with the increasing trend of ‘commodified’ attacks delivered in the wild [7, 35, 42, 56], whereas the remaining more ‘specialized’ sector of the market seems to be inflating. We do not find any significant association between number of exploits in the package and package price. This lack of correlation may indicate that the business model behind exploit trading, as well as other contextual considerations on market status, presence of similar exploit, and affected software should be considered in the analysis, as previously suggested by several authors [10, 14, 16, 35, 63, 72]. We give an extended account of this in the next section.

## 4.2 Analysis of exploits

*4.2.1 Exploit demographics.* Embedded in the packages we find 89 exploits targeting 57 unique CVEs in Microsoft, Adobe, and Oracle products. Table 4 reports the counts of exploited software for each product type. Microsoft vulnerabilities alone make up for more than half the exploits traded as STANDALONE products (56%); unsurprisingly, vulnerabilities in Oracle and Adobe products, as well as Internet Explorer vulnerabilities, are prevalent in EKIT bundles, as these products are by design exposed to Internet requests [35, 48]. Exploits bundled in MALWARE are for Windows and Adobe Acrobat. A Fisher Exact test rejects the null hypothesis of count uniformity ( $p = 0.0012$ ), suggesting that exploited software varies by package type.


Figure 7 plots the occurrences of exploit publications by year and by software vendor. We observe that during the first years of

**Table 3: Descriptive statistics of package prices and bundled exploits**

Type	n	Package price (USD)						no. bundled exploits							
		Min	0.025p	Mean	Median	0.975p	Max	sd	Min	0.025p	Mean	Median	0.975p	Max	sd
EKIT	6	150	157.92	693.89	400	1875	2000	708.94	2	2.12	6.83	7	11	11	4.26
MALWARE	6	420	428.75	1735	1250	3875	4000	1456.38	1	1	1.5	1	2.88	3	0.84
STANDALONE	26	100	100	2972.69	3000	8000	8000	2629.39	1	1	1.5	1	4	4	0.86
All	38	100	100	2417.46	1500	8000	8000	2408.28	1	1	2.34	1	11	11	2.63

**Table 4: Breakdown of traded vulnerabilities**


SwVendor	Software	MALWARE	STANDALONE	EKIT	Sum
Adobe		2	12	17	31
	flash	0	8	10	18
	acrobat	2	4	7	13
Microsoft		7	22	14	43
	office	0	11	2	13
	int. expl.	0	4	7	11
	windows	7	6	5	18
	silverlight	0	1	0	1
Oracle		0	5	10	15
	java	0	5	10	15
<b>Sum</b>		<b>9</b>	<b>39</b>	<b>41</b>	<b>89</b>



**Figure 7: Occurrences of exploit publication by year**

RuMarket operation there is a spike in number of published exploits for all platforms. Oracle products result as the most affected in that year, followed by Adobe and, closely, Microsoft products. This observation matches the surge around 2010-2013 of ‘cybercrime as a service’, thoroughly reported in the scientific literature and industry in that time-frame [35, 48, 70]. Interestingly, Oracle exploits seem to plunge after 2013; this coincides with the introduction in major web browsers of plugin-blocking features,<sup>7</sup> and a Java update (released in January 2013) that increases the default security settings of the plugin<sup>8</sup> (e.g. triggering certificate errors as exploited by several exploit kits [48]). This also independently supports previous findings on exploitation of Java vulnerabilities [40]. Following 2013, Microsoft and Adobe exploits are published at similar, steady rates. Anecdotaly, we observe that the shape of the described curve resembles the *Gartner Hype Cycle*<sup>9</sup> curve, whereby after a first spike at the beginning of a new product cycle (the

<sup>7</sup>[https://www.theregister.co.uk/2013/12/10/firefox\\_26\\_blocks\\_java/](https://www.theregister.co.uk/2013/12/10/firefox_26_blocks_java/), last visit Aug 2017.  
<sup>8</sup><http://bit.ly/2r8MLz1>, last visit Aug 2017.  
<sup>9</sup><http://gtnr.it/1g1Nnw0>, last visit Aug 2017.



**Figure 8: Number of repackaged exploits by year**

‘Peak of Inflated Expectations’) the market experiences a relative drop (‘Trough of Disillusionment’) followed by a ‘plateau’ where the technology reaches maturity (‘Plateau of Productivity’) [8].

**4.2.2 Exploit arrival.** Figure 8 reports the number of newly released (yellow) and repackaged (black) exploits in each package type by year. The introduction of new exploits in RuMarket is primarily driven by STANDALONE and EKIT packages, with MALWARE packages mainly re-introducing already published exploits. In particular, STANDALONE products seem to propose new exploits at a yearly rate of approximately 80% for each package. EKIT products introduced a significant number of exploits in 2011 (their ‘debut’ year on the markets [35, 70]), whereas newer exploit kits appear to embed a lower number of exploits. This confirms previous figures whereby exploit kits are specializing to use fewer, more reliable exploits than at their original introduction [48]. Table 13 in the Appendix reports the evolution of repackaged exploits by PackType. Most exploits first appear in STANDALONE and EKIT packages and re-appear in a pack of the same type, with a few exceptions. Among these, STANDALONE exploits seem to reappear in both MALWARE and EKIT packs, whereas EKIT exploits are prevalently re-packed in other kits. STANDALONE exploits seem therefore to play a role in the ‘innovation’ process in RuMarket; this may indicate the presence of an ‘exploit chain’ whereby the most reliable and effective STANDALONE exploits are selected for future inclusion in EKIT products for deployment at scale.

It is interesting to evaluate the *rate* at which exploit introduction happens. A few recent studies suggest that the rate of appearance of new exploits may be much lower than previously thought [10, 18], but no account of exploit timing in the cybercrime markets currently exists. Figure 9 shows a breakdown by software vendor of the distribution of days between the introduction of new exploits (left) and re-packed exploits (right). We exclude from the analysis six EKIT vulnerabilities that have been added to the respective






Figure 9: Distribution of days between exploit introduction

Table 5: Exploit age (days) at time of first publication

Type	Min	0.025p	Mean	Med.	0.975p	Max	sd	n
EKIT	1	4	372.48	294	1659.8	1745	470.16	25
MAL	185	185	185	185	185	185	-	1
STD	1	8	147.34	75	549.7	934	189.66	29
All	1	2.75	250.36	93	1368.85	1745	359.97	55

packages as updates, but whose date of addition is not reported in the market. This leaves us with 83 exploits, of which  $n = 55$  are introduced for the first time in the market. New exploits are introduced at similar rates for all software vendors, with 50% of exploits being introduced at approximately six months intervals (175 days). The ‘fastest’ 25% is introduced two months (62 days) after the appearance of an exploit for the same software platform, whereas the ‘slowest’ 25% appears after more than a year (401 days). These figures are in sharp contrast with current assumptions made in the literature, whereby essentially all ‘severe’ vulnerabilities are potentially exploited at scale by attackers [54, 66] (and must therefore be fixed immediately [73]). On the contrary, these findings support recent evidence pointing in the opposite direction: most vulnerabilities are not of ‘economic’ interest for an attacker, as the addition of a new vulnerability may not lead to a significant increase in targeted systems [12, 55]; this results in significantly skewed distributions of risk per vulnerability (as empirically shown in [7, 56], and analytically modeled in [12]). Exploit re-packaging (right plot in Fig. 9) happens at significantly faster rates: 75% of exploits are re-introduced within 184 days from first publication, indicating that their commercial interest is short-lived.

A different question is how ‘old’ are exploits when they first appear on the market. We compute exploit age as the difference in days between exploit publication in the market and publication of the corresponding CVE on the NVD, i.e.  $ExpAge = PubDate - CVEPub$ . Table 5 reports the distribution of exploit age for newly introduced exploits. The mean and median exploit age varies considerably by product type. STANDALONE exploits are on average significantly younger at time of publication than other exploits ( $p = 0.05$  for a Wilcoxon rank sum test). 50% of STANDALONE exploits are published in the market within 2.5 months (75 days) from the public disclosure date. The top 25% (not reported in Tab. 5) are published within 40 days, and the top 2.5% within approximately a week. The difference in exploit age for the EKIT and MALWARE categories is not statistically significant. Whereas some exploits do appear quickly after




Figure 10: Exploit age vs time of publication in RuMarket

disclosure in RuMarket, most exploits take around four months from disclosure date to be published. This may indicate that other factors such as effectiveness of older exploits [10], or delays in user system updates [55, 64], may affect timing of appearance of a marketed exploit. To evaluate the rate of change in time of arrival of new exploits, Figure 10 reports the exploit age distribution by year of publication. We observe that the mean exploit age decreases steadily for more recent publication dates ( $\beta = -0.32, p = 0.001$ ), indicating that exploit vendors are becoming faster in releasing exploits for newly disclosed vulnerabilities. The coefficient of the linear regression indicates that exploits appear at an approximately 30% faster rate every year.

4.2.3 Estimate of exploit prices. Conceptually, the lower bound cost of a pack can be summarized as follows:

$$PackPrice_{p \in P_T} = \sum_{i=1}^n CExpl_{pi} + CDev_p + CDepl_p \quad (2)$$

where  $CExpl_{pi}$  is the cost of the  $i^{th}$  exploit in package  $p$ , and  $CDev_p$  and  $CDepl_p$  are the development and deployment (including maintenance) costs of the pack. For example, on top of the sole exploits EKITs provide a web interface to control infections, as well as additional deployment services such as remote servers where the kit is hosted, or the redirection of vulnerable traffic to the kit [35, 48]. Similarly, MALWARE packages provide additional malware functionalities on top of the sole exploit. Hence, we have  $CDev_p > 0, CDepl_p > 0, \forall p \in P_T$ , with  $t \in \{EKIT, MALWARE\}$ . Unfortunately, an estimation of these costs would require an analysis of the source code of these packages [22], which is not publicly available. On the other hand, STANDALONE exploits are provided as-is, i.e. only the vulnerability exploit is traded, without further embellishments or services. This sets  $CDev_p \approx CDepl_p \approx 0$  for this category. This leaves us with only the term  $\sum_{i=1}^n CExpl_{pi}$  which, assuming a uniform distribution of exploit costs per package, (see discussion in Sec. 3.4), yields  $CExpl_{p \in P_{STD}, i} = 1/n \cdot PackPrice_p$ . We therefore only report STANDALONE exploit estimates.<sup>10</sup>

Table 6 reports price estimates for exploits against different software. In parenthesis we report the bootstrapped estimation of exploit prices. We report mean, median, standard deviation and 95% confidence intervals. Price estimates in the bootstrapped sample appear to diverge at the tails of the distribution w.r.t the observed sample, suggesting that outliers in the sample may bias sample statistics.

<sup>10</sup>The estimation for all packages is reported in the Appendix, Table 14.

**Table 6: (Bootstrapped) descriptive statistics of STANDALONE exploit price estimates in USD**

STANDALONE exploit prices are estimated on a uniform distribution by package. To approximate the true (unknown) distribution of exploits, we perform a bootstrap of our data ( $N = 10000$ ), reported in parenthesis. The column  $n$  reports number of exploits for that software. The bootstrapped data does not deviate substantially from our observations on the average. Fatter distribution tails indicate that RuMarket outliers tend to bias sample statistics. Exploits are priced between 150 and 8000USD with significant differences by software.

SwVendor	Software	Min	0.025p	Mean	Median	0.975p	Max	sd	n
Adobe		75	75	879.17	1250	1500	1500	693.54	12
		(75)	(100)	(1000.06)	(1040)	(1500)	(1500)	(521.91)	
	flash	75	75	568.75	150	1500	1500	652.3	8
		(75)	(87.5)	(562.05)	(545.45)	(1300)	(1500)	(316.52)	
	acrobat	1500	1500	1500	1500	1500	1500	0	4
		(1500)	(1500)	(1500)	(1500)	(1500)	(1500)	(0)	
Microsoft		150	150	2801.82	2250	8000	8000	2393.09	22
		(150)	(150)	(2442.13)	(2450)	(5600)	(8000)	(1601.69)	
	office	150	362.5	3195.45	4000	7250	8000	2504.04	11
		(150)	(1605.1)	(3407.31)	(3262.5)	(5750)	(8000)	(1112.54)	
	int. expl.	440	459.5	3035	1850	7625	8000	3504.22	4
		(440)	(440)	(3051.89)	(3000)	(8000)	(8000)	(1727.18)	
windows	700	800	2366.67	2250	4687.5	5000	1458.31	6	
	(700)	(1100)	(2349.27)	(2327.27)	(3750)	(5000)	(658.77)		
silverlight	150	150	150	150	150	150	150	1	
	(150)	(150)	(150)	(150)	(150)	(150)	(0)		
Oracle		25	25	1020	25	4502.5	5000	2224.89	5
		(25)	(25)	(1847.02)	(1020)	(5000)	(5000)	(1981.08)	
	java	25	25	1020	25	4502.5	5000	2224.89	5
	(25)	(25)	(1847.02)	(1020)	(5000)	(5000)	(1981.08)		

Looking at exploits by software, we find that the most expensive exploits in RuMarket are for Microsoft software, and are priced at 2500USD on the average. Among software from this vendor, Office and Windows exploits appear to be the most expensive with 50% of exploits above 2000 USD, and the top 2.5% quoted at about 7000 and 5000 USD respectively. As vulnerability patching and mitigation hinder the effectiveness of an exploit in the wild [55], we further evaluate whether exploit age affects exploit price estimates. We find a negative correlation between ExplAge and ExplPrice (albeit not significant when looking only at the exploit), suggesting that exploits lose value as they age ( $cor = -0.16$ ,  $p = 0.3$ ). We do not find evidence of dependence between exploit price and CVSS vulnerability severity.

**4.2.4 Regression analysis of exploit price estimates.** To evaluate the factors driving exploit price, we employ a set of mixed effect linear regression models over the response variable  $\log(ExplPrice)$ . We report regression results for the following three nested models:

$$M1: \log(ExplPrice_i) = \beta_0 + \beta_1 \log(ExplAge_i) + \epsilon_i$$

$$M2: \log(ExplPrice_i) = \dots + \beta_2 SwVen_i + \epsilon_i$$

$$M3: \log(ExplPrice_i) = \dots + \beta_3 SwVen_i \times \log(ExplAge_i) + \epsilon_i$$

Table 7 reports the regression results. Coefficient estimates are consistent among models. A Variance Inflation Factors (VIF) analysis does not indicate any significant collinearity between the model predictors. Log-log relationships can be interpreted as the *elasticity* between dependent and independent variables. For example, in M3 the coefficient for  $\log(ExplAge)$  ( $-1.357$ ) indicates that for a 1% increase in the variable  $ExplAge$  we can expect an average 1.4% decrease ( $1.01^{-1.357} = 0.986$ ) in  $ExplPrice$ . A rough quantitative


**Table 7: Regression on STANDALONE exploit pricing**

Variables:  $ExplPrice$  = price estimate of exploit;  $ExplAge$  = age of exploit when advertised;  $SwVen$  = software vendor. Exploit age is negatively correlated with price. Depreciation rate depends on the software vendor.

$\log(ExplPrice)$	Model 1	Model 2	Model 3
c	8.080 <sup>***</sup>	5.592 <sup>***</sup>	10.943 <sup>***</sup>
	(0.746)	(1.458)	(1.735)
$\log(ExplAge)$	-0.330 <sup>*</sup>	-0.268 <sup>*</sup>	-1.357 <sup>***</sup>
	(0.129)	(0.135)	(0.234)
Adobe		1.993	-4.846 <sup>*</sup>
		(1.395)	(2.068)
Microsoft		2.662 <sup>†</sup>	-3.483 <sup>†</sup>
		(1.375)	(1.891)
$\log(ExplAge) \times AD$			1.398 <sup>***</sup>
			(0.322)
$\log(ExplAge) \times MS$			1.276 <sup>***</sup>
			(0.272)
$Var(c ExplVen)$	2.209	1.520	1.598
Pseudo- $R^2$	0.05	0.28	0.38
Log-likelihood	-64.2	-60.6	-51.8
N	39	39	39

(<sup>†</sup>)  $p < 0.1$ ; (<sup>\*</sup>)  $p < 0.05$ ; (<sup>\*\*</sup>)  $p < 0.01$ ; (<sup>\*\*\*</sup>)  $p < 0.001$ ;

approximation of the effect can generally be obtained by simply looking at the regression coefficients (e.g.  $\beta_3 = -1.357$  indicates a decrease of approximately 1.4%). coefficients in M3 can only be interpreted simultaneously with the coefficients of the interaction effect  $\log(ExplAge) \times SwVen$ . We find that baseline prices for exploits vary widely by software vendor, and are negatively correlated



**Figure 11: Vulnerabilities exploited in the wild versus (left) package price; (right) replies received on the market.**

with the age of exploit; Adobe and Microsoft exploits retain their value significantly longer than Oracle exploits. This may indicate a prolonged economic interest in the exploitation of Microsoft and Adobe vulnerabilities, a finding consistent with related work on the persistence of vulnerabilities on end-user systems [55].

Exploit vendors are a significant source of variance in price of exploit. Pseudo- $R^2$  values indicate that the models' power is adequate in explaining the observed effect. In particular, Model 3 explains approximately 40% of the variance in exploit price estimates. Importantly, all variables of the model can be easily assessed with the sole knowledge of the vulnerability at any point in time.

### 4.3 Exploitation in the wild

In this section we evaluate the effects of the identified market variables on the exploitation in the wild of traded vulnerabilities. A consideration first: certain exploits may evade detection for some time, for example by means of frequent exploit repackaging. On the other hand, it is unlikely for an exploit to remain completely under-cover for a long time, while in the meanwhile delivering hundreds of thousands or millions of attacks [7, 56]. To lower uncertainty around exploit detection, we restrict our analysis to exploits published in RuMarket at least a year ahead of the SYM data collection (i.e. before the 1st of April 2016). This coincides with the median lifetime of an exploit package in RuMarket (see Fig. 5), and allows SYM a full calendar year to report an exploit at scale. We consider effects before that time to be unlikely to be caused by type II errors (i.e. no inclusion in SYM despite high attack volumes). This leaves us with  $n = 78$  exploits.

*Package price and market activity vs exploitation.* As we are considering the effect of the *acquisition* and *deployment* of an exploit by the attacker, we consider cost of package (as opposed to cost of exploit) because this reflects the upfront price the attacker needs to pay in order to deploy the attack. Figure 11 reports the distribution of package prices (left) and replies on the market (right) against exploitation in the wild. Overall, we find that exploits in SYM are traded at a lower upfront price than exploits not included in SYM. Similar conclusions can be drawn for the effect of market interest (as estimated by *PackActivity*) on exploitation: exploits to which the RuMarket community dedicated greater attention have a higher chance of exploitation in the wild than exploits around which developed less market activity. A break down by package types does not reveal any significant interaction between the variables. These

**Table 8: CVSS category vs SYM**

	CVSS Category		
	C	NC	Sum
Not exploited	4	8	12
Exploited	53	13	66
Sum	57	21	78

results also support recent findings underlying the importance of the economics of the attack process in the analysis of cyber-attack scenarios [14, 20, 56, 72], and in the development of sensible cybercrime-based measures for risk of cyber-attack [13, 39].

*Vulnerability severity.* Previous studies revealed the lack of correlation between technical vulnerability characteristics and exploitation [20]. The consideration of additional factors, such as presence of exploit in the black markets [9], is often advised by experts and best practices to obtain more significant tests for actual exploitation [38]. Following [9], we categorize vulnerabilities in two categories defined by the respective CVSS severity score: critical (C) ( $CVSS \geq 9$ ) and non-critical (NC) vulnerabilities ( $CVSS < 9$ ). Table 8 reports the corresponding distributions against SYM. Supporting previous research findings on Exploit Kits alone [9], we find that critical vulnerabilities traded in the cybercrime markets have a higher chance of exploitation in the wild (93% in our sample) than non-critical vulnerabilities (62%,  $p = 0.0021$ ).

*4.3.1 Regression analysis of exploitation in the wild.* To more rigorously evaluate the correlation between the identified market and vulnerability variables and exploitation in the wild, we select three logit regression models of the following form over the binary response variable SYM:

$$\begin{aligned}
 M1: SYM_i &= \beta_0 + \beta_1 \log(PackActivity_i) + \epsilon_i \\
 M2: SYM_i &= \dots + \beta_2 \log(PackPrice_i) + \epsilon_i \\
 M3: SYM_i &= \dots + \beta_3 CVSS:NC_i + \epsilon_i
 \end{aligned}$$

where SYM indicates presence or absence of exploit at scale; *PackActivity* is the number of replies received by the product advert; *PackPrice* is the upfront price to pay to obtain the exploit; and *CVSS : NC* is the CVSS categorization of the vulnerability severity as non-critical. Regression results are reported in Table 9.<sup>11</sup> Coefficients should be interpreted as the change in the odds ratio of exploit in the wild. For example, Model 3 indicates that for every increase in one unit of  $\log(PackActivity)$  there is a three-fold increase in odds of exploitation in the wild ( $\exp(1.101) = 3.01$ ); the coefficient significance indicates that exploits bundled in packages around which more market activity is developed are more likely to be detected in the wild than exploits with less activity. Interestingly, we find that package prices also have a significant and negative impact on odds of exploitation. This indicates that, everything else being equal, exploits bundled in more expensive packages are less likely to be detected in the wild than comparable exploits bundled in less expensive packages. These findings weigh favorably on the existence of a relation between market activity and exploit deployment in the

<sup>11</sup> An OLS robustness check run on average values of regressors for each CVE (reported in the Appendix) yields equivalent results.

**Table 9: Regression on odds of exploitation**

Variables: *PackActivity* = replies received in the market; *PackPrice* = price of package; *CVSSCat* = CVSS category. Market, economic and vulnerability factors are correlated with odds of exploit at scale.

<i>SYM</i>	Model 1	Model 2	Model 3
<i>c</i>	0.245 (1.526)	6.056 <sup>†</sup> (3.302)	6.754 <sup>*</sup> (3.086)
$\log(\text{PackActivity})$	0.673 <sup>†</sup> (0.383)	0.938 <sup>*</sup> (0.400)	1.101 <sup>*</sup> (0.451)
$\log(\text{PackPrice})$		-0.982 <sup>*</sup> (0.460)	-1.013 <sup>*</sup> (0.444)
CVSS:NC			-2.409 <sup>**</sup> (0.830)
$\text{Var}(c \text{ExplVendor})$	3.071	0.617	0.000
Pseudo - $R^2$	0.11	0.51	0.65
Log-likelihood	-28.4	-25.3	-20.4
N	78	78	78

(<sup>†</sup>)  $p < 0.1$ ; (<sup>\*</sup>)  $p < 0.05$ ; (<sup>\*\*</sup>)  $p < 0.01$ ; (<sup>\*\*\*</sup>)  $p < 0.001$ ;

wild [9, 15, 56]. Vulnerability severity has also a negative impact on likelihood of exploit, indicating that risk of exploitation for vulnerabilities in RuMarket increases for critical vulnerabilities. All models show satisfactory Pseudo- $R^2$  values, with Model 3 explaining most of the variance.

## 5 DISCUSSION

*Exploit measures.* An important aspect of threat assessment is the consideration of exploit metrics. Current approaches often implement these by looking at the technical requirements of an attack, including the evasion of attack mitigation measures [74] and complexity of the attack [20]. While a technical assessment of the ‘operational’ requirements of an attack can shed light on the relative ordering of attack preferences, it is hard to quantify absolute likelihoods. For example, a utility-maximizing attacker may decide to not perform (or delay) an attack because they do not believe that there is a positive payoff, given cost of exploit acquisition. Our analysis gives the first pointers in this direction by quantifying the relation between time, software, and exploit pricing. Importantly, this estimate only requires readily available information on the vulnerability, and the elastic relationship between age of exploit and price of exploit can be used to evaluate relative changes in exploit price as time passes. This directly affects current estimates of attack cost used in risk assessment practices [20, 38, 74]. Similarly, the *update* of attackers’ exploit portfolios is an important step driving the variance in risk of attacks [10, 18]. We find that new exploits are introduced at rates in between two and six months, and are approximately equal for all software. The process driving this update remains however to uncover: follow up studies may look at the factors driving appearance of exploit in the markets (e.g. by considering pre-existent exploits or software updates [10, 55]).

The dynamics of the underground markets have often been pointed at as an important block of overall risk of attacks, but a clear link between the two is currently missing. Whereas the problem of attack *attribution* remains open (i.e. we cannot establish a causality link), this paper provides important indications on the correlation between market operations and realization of attacks.

This weighs in favour of the importance of economic aspects of vulnerability exploitation to well-informed security practices (e.g. vulnerability assessment and prioritization [38, 73]). For example, our analysis of market activity and odds of exploitation in the wild reveals a significant and positive relationship between the two. Similarly, exploits that are more expensive to acquire have lower odds of exploitation than ‘cheaper’ exploits. This information is often ignored in risk-assessment studies [38], and condensed metrics for vulnerability assessment are used instead [54, 73]. Whereas existing vulnerability metrics are known to not correlate to attacks in the wild [9, 20], we find that they do once the effect of market inclusion is considered as well. Importantly, this provides a useful tool for a first evaluation of risk of exploit without insights from the cybercrime markets other than whether the vulnerability is present [9]. A more precise estimation can then be obtained by measuring market activity around the packages embedding the exploit. These results can be factored in current best practices for vulnerability risk management and exploit mitigation [31, 51, 74].

*Vulnerability economics.* Previous studies in the literature highlighted the operations of criminal markets for drugs, arms and pornography [69], and for the monetization of stolen information resulting from an attack [36]. However, little insight exists on the markets that, as opposed to (re)selling the result of an attack, trade the technology that enables the attack in the first place. A few estimates exist [1], but are mostly based on 0-day price allegations, vary widely, and their relevance for the overall risk of attack remains unclear [18, 63]. The scientific community long discussed on the idea of building ‘legitimate vulnerability markets’ [14, 46, 57], and the result is the institution of a few legitimate exploit markets [63] and of several ‘bug bounty’ initiatives that reward security researchers for the disclosure of new vulnerabilities [30], and discourage the participation in the underground economy. In this vein it is interesting to observe that the prices of modern bug-bounty programs are in line or below those we identify on RuMarket. For example, [30, Tab. 4 pp281] reports that the majority of vulnerability prizes awarded by Google in their Chrome Vulnerability Reward Program (VRP) are at or below 1000 USD, and that most external contributors to the program (i.e. vulnerability researchers) earn between 500 and 1000 US dollars. The median price of an exploit in RuMarket, showed in Table 6, is at approximately 2000 dollars, a higher but not distant figure from those indicated in [30]. It is however unclear whether the resulting balance weighs in favour of the legitimate or underground vulnerability markets: the dynamics balancing vulnerability finding (a notoriously demanding process [52]) and exploitation trade have not been fully investigated in the literature yet. For example, at the above rates it does not seem unlikely that vendors who sell their exploits multiple times to different buyers may still be better off participating in the cybercrime economy than moving to the ‘legitimate’ vulnerability markets, as in the latter vulnerabilities can realistically be traded only once (as the trade creates an association between the ‘0-day’ vulnerability and identity of whom discovered it). The results in this paper represent a first building block in the evaluation and enhancement of current legitimate vulnerability markets, to foster the responsible disclosure of vulnerabilities and attract skillful researchers away from criminal markets.

Our exploit price estimates provide additional insights on the effect of different criminal business models on exploit pricing and, therefore, accessibility of attack. With reference to Tab. 14 in the Appendix, and despite the upwards bias of the estimate (Eq. 2), EKIT exploits are priced, per unit, significantly below exploits in other package types. This effect is driven by the higher number of exploits bundled in EKIT (ref. Table 3), and underlines how the different business model employed by EKIT services may allow exploit vendors to drastically reduce exploit development and deployment costs. Lower prices may make these tools more accessible to ‘wannabe-criminals’, and therefore generate more attacks overall. This suggests that the criminal business model may play a central role in the diffusion of cyber-attacks, and calls for additional studies characterizing this effect. Further, we find that the studied market shows clear signs of expansion, with a growing number of vendors, exploits, attack products, and generally inflating package prices. This indicates that market activity is unlikely to stop in the near future, and that attacker economics will likely play an increasingly more relevant role in the cybersecurity scenario.

## 6 CONCLUSIONS

In this paper we presented the first quantitative account of exploit pricing and market effects on exploitation in the wild. Our findings quantify a strong correlation between market activities and likelihood of exploit. We find that the analyzed market shows signs of expansion, and that exploit-as-a-service models may allow for drastic cuts in exploit development costs. Further, we find that exploit prices are aligned with or above those of ‘legitimate’ vulnerability markets, supporting work on the identification of incentives for responsible vulnerability disclosure and attack economics.

## 7 ACKNOWLEDGMENTS

This line of work started in collaboration with Prof. Fabio Massacci at the University of Trento, Italy, and Prof. Julian Williams at Durham Business School, UK, whom the author thanks for their invaluable insights and feedback. This work has been supported by the NWO through the SpySpot project no. 628.001.004.

## REFERENCES

- [1] 2012. Shopping For Zero-Days: A Price List For Hackers’ Secret Software Exploits. (2012). [online] <http://www.forbes.com/sites/andygreenberg/2012/03/23/shopping-for-zero-days-an-price-list-for-hackers-secret-software-exploits/>.
- [2] 2015. NIST National Vulnerability Database (NVD). (2015). <http://nvd.nist.gov> [online] <http://nvd.nist.gov>.
- [3] 2015. An Overview of Exploit Packs (Update 22) Jan 2015. (2015). <http://contagiodump.blogspot.it/2010/06/overview-of-exploit-packs-update.html> [online] <http://contagiodump.blogspot.it/2010/06/overview-of-exploit-packs-update.html>.
- [4] Lillian Ablon, Martin C Libicki, and Andrea A Golay. 2014. *Markets for cybercrime tools and stolen data: Hackers’ bazaar*. Rand Corporation.
- [5] Alan Agresti and Maria Kateri. 2011. *Categorical data analysis*. Springer.
- [6] George A. Akerlof. 1970. The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Jour. of Econ.* 84 (1970), pp. 488–500.
- [7] Luca Allodi. 2015. The Heavy Tails of Vulnerability Exploitation. In *Proc. of ESSoS’15*.
- [8] Luca Allodi, Marco Corradin, and Fabio Massacci. 2015. Then and Now: On the Maturity of the Cybercrime Markets. *IEEE Transactions on Emerging Topics in Computing* (2015).
- [9] Luca Allodi and Fabio Massacci. 2014. Comparing vulnerability severity and exploits using case-control studies. *ACM Transaction on Information and System Security (TISSEC)* 17, 1 (8 2014).
- [10] Luca Allodi and Fabio Massacci. 2015. The work-averse attacker model. In *In the Proceedings of the 2015 European Conference on Information Systems (ECIS)*.
- [11] Luca Allodi and Fabio Massacci. 2017. Security Events and Vulnerability Data for Cybersecurity Risk Estimation. *Risk Analysis* 37, 8 (2017), 1606–1627. <https://doi.org/10.1111/risa.12864>
- [12] Luca Allodi, Fabio Massacci, and Julian Williams. 2017. The Work-Averse Cyber Attacker Model. Evidence from two million attack signatures. In *WEIS 2017*. Available at <https://ssrn.com/abstract=2862299>.
- [13] Luca Allodi, Shim Woohyun, and Fabio Massacci. 2013. Quantitative assessment of risk reduction with cybercrime black market monitoring. In *In Proc. of IWCC’13*.
- [14] Ross Anderson. 2008. Information Security Economics - and Beyond. In *Proc. of DEON’08 (DEON ’08)*, 49–49.
- [15] R. Anderson, C. Barton, R. Böhme, R. Clayton, M.J.G. van Eeten, M. Levi, T. Moore, and S. Savage. 2012. Measuring the Cost of Cybercrime. In *Proc. of WEIS’12*.
- [16] Hadi Asghari, Michel Van Eeten, Axel Arnab, and Nico Van Eijk. 2013. Security economics in the HTTPS value chain. Available at *SSRN 2277806* (2013).
- [17] Jonell Baltazar. 2011. *More traffic, more money: Koobface draws more blood*. Technical Report. TrendLabs.
- [18] Leyla Bilge and Tudor Dumitras. 2012. Before we knew it: an empirical study of zero-day attacks in the real world. In *Proc. of CCS’12*. ACM, 833–844.
- [19] Hamad Binsalleh, Thomas Ormerod, Amine Boukhtouta, Prosenjit Sinha, Amr Yousef, Mourad Debbabi, and Lingyu Wang. 2010. On the analysis of the Zeus botnet crimeware toolkit. In *Privacy Security and Trust (PST), 2010 Eighth Annual International Conference on*. IEEE, 31–38.
- [20] Mehran Bozorgi, Lawrence K. Saul, Stefan Savage, and Geoffrey M. Voelker. 2010. Beyond Heuristics: Learning to Classify Vulnerabilities and Predict Exploits. In *Proc. of SIGKDD’10*.
- [21] Luís Cabral and Ali Hortaçsu. 2010. THE DYNAMICS OF SELLER REPUTATION: EVIDENCE FROM EBAY. *The Journal of Industrial Economics* 58, 1 (2010), 54–78. <http://www.jstor.org/stable/40661997>
- [22] Alejandro Calleja, Juan Tapiador, and Juan Caballero. 2016. A Look into 30 Years of Malware Development from a Software Metrics Perspective. In *International Symposium on Research in Attacks, Intrusions, and Defenses*. Springer, 325–345.
- [23] A.A. Cárdenas, S. Radosavac, J. Grossklags, J. Chuang, and C. Hoofnagle. 2009. An economic map of cybercrime. In *Proc. of TPRC’09*.
- [24] Jared DeMott. 2015. Bypassing EMET 4.1. *IEEE Security & Privacy* 13, 4 (2015), 66–72.
- [25] Irina Dezhina and Loren Graham. 1999. Science and higher education in Russia. *Science* 286, 5443 (1999), 1303–1304.
- [26] Tudor Dumitras and Darren Shou. 2011. Toward a standard benchmark for computer security research: The Worldwide Intelligence Network Environment (WINE). In *Proc. of BADEGRS’11*. ACM, 89–96.
- [27] Bradley Efron and Robert J Tibshirani. 1994. *An introduction to the bootstrap*. Vol. 57. CRC press.
- [28] Kathleen M Eisenhardt. 1989. Agency theory: An assessment and review. *Academy of management review* 14, 1 (1989), 57–74.
- [29] Jon Erickson. 2008. *Hacking: the art of exploitation*. No Starch Press.
- [30] Matthew Finifter, Devdatta Akhawe, and David Wagner. 2013. An Empirical Study of Vulnerability Rewards Programs. In *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*. USENIX, Washington, D.C., 273–288.
- [31] First.org. 2015. *Common Vulnerability Scoring System v3.0: Specification Document*. Technical Report. FIRST, Available at <http://www.first.org/cvss>.
- [32] D. Florêncio and C. Herley. 2011. Sex, Lies and Cyber-crime Surveys. In *Proc. of WEIS’11*.
- [33] J. Franklin, V. Paxson, A. Perrig, and S. Savage. 2007. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proc. of CCS’07*. 375–388.
- [34] Maxim Goncharov. 2011. *Traffic Direction Systems as Malware Distribution Tools*. Technical Report. TrendLabs.
- [35] 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 *Proc. of CCS’12*. ACM, 821–832.
- [36] Shuang Hao, Kevin Borgolte, Nick Nikiforakis, Gianluca Stringhini, Manuel Egele, Michael Eubanks, Brian Krebs, and Giovanni Vigna. 2015. Drops for stuff: An analysis of reshipping mule scams. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1081–1092.
- [37] C. Herley and D. Florencio. 2010. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. *Springer Econ. of Inf. Sec. and Priv.* (2010).
- [38] Hannes Holm and Khalid Khan Afridi. 2015. An expert-based investigation of the common vulnerability scoring system. *Computers & Security* 53 (2015), 18–30.
- [39] Thomas J Holt, Olga Smirnova, and Alice Hutchings. 2016. Examining signals of trust in criminal markets online. *Journal of Cybersecurity* (2016), tyw007.
- [40] Philipp Holzinger, Stefan Triller, Alexandre Bartel, and Eric Bodden. 2016. An In-Depth Study of More Than Ten Years of Java Exploitation. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 779–790.

- [41] Ting-Kai Huang, Bruno Ribeiro, Harsha V Madhyastha, and Michalis Faloutsos. 2014. The socio-monetary incentives of online social network malware campaigns. In *Proceedings of the second ACM conference on Online social networks*. ACM, 259–270.
- [42] Alice Hutchings and Richard Clayton. 2016. Exploring the provision of online booter services. *Deviant Behavior* 37, 10 (2016), 1163–1178.
- [43] Group IB. 2011. *State and Trends of the Russian digital crime market*. Technical Report. Group IB.
- [44] David R Johnson and David Post. 1996. Law and borders: The rise of law in cyberspace. *Stanford Law Review* (1996), 1367–1402.
- [45] Chris Kanich, Christian Kreibich, Kirill Levchenko, Brandon Enright, Geoffrey M. Voelker, Vern Paxson, and Stefan Savage. 2008. Spamalytics: an empirical analysis of spam marketing conversion. In *Proc. of CCS'08 (CCS '08)*. ACM, 3–14.
- [46] Karthik Kannan and Rahul Telang. 2005. Market for software vulnerabilities? Think again. *Management Science* 51, 5 (2005), 726–740.
- [47] Mohammad Karami and Damon McCoy. 2013. Understanding the Emerging Threat of DDoS-as-a-Service.. In *LEET*.
- [48] Vadim Kotov and Fabio Massacci. 2013. Anatomy of Exploit Kits. Preliminary Analysis of Exploit Kits as Software Artefacts. In *Proc. of ESSoS 2013*.
- [49] Nir Kshetri. 2010. Diffusion and effects of cyber-crime in developing economies. *Third World Quarterly* 31, 7 (2010), 1057–1079.
- [50] Prashanth Loyalka, Martin Carnoy, Isak Froumin, Rafiq Dossani, and Jandhyala Tilak. 2012. Getting the Quality Right: Engineering Education in the BRIC Countries. (2012).
- [51] Pratyusa K. Manadhata and Jeannette M. Wing. 2011. An Attack Surface Metric. *TSE* 37 (2011), 371–386. <https://doi.org/10.1109/TSE.2010.60>
- [52] C. Miller. 2007. The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. In *Proc. of WEIS'07*.
- [53] Marti Motoyama, Damon McCoy, Stefan Savage, and Geoffrey M. Voelker. 2011. An Analysis of Underground Forums. In *Proc. of LMC'11*.
- [54] Mendes Naaliel, Duraes Joao, and Madeira Henrique. 2014. Security Benchmarks for Web Serving Systems. In *Proc. of ISSRE'14*.
- [55] A. Nappa, R. Johnson, L. Bilge, J. Caballero, and T. Dumitras. 2015. The Attack of the Clones: A Study of the Impact of Shared Code on Vulnerability Patching. In *Proc. of the 36th IEEE Symp. on Sec. & Privacy*. 692–708.
- [56] Kartik Nayak, Daniel Marino, Petros Efsthopoulos, and Tudor Dumitras. 2014. Some Vulnerabilities Are Different Than Others. In *Proc. of RAID'14*. Springer, 426–446.
- [57] Andy Ozment. 2004. Bug auctions: Vulnerability markets reconsidered. In *Third Workshop on the Economics of Information Security*.
- [58] Rebecca S Portnoff, Sadia Afroz, Greg Durrett, Jonathan K Kummerfeld, Taylor Berg-Kirkpatrick, Damon McCoy, Kirill Levchenko, and Vern Paxson. 2017. Tools for Automated Analysis of Cybercriminal Markets. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 657–666.
- [59] Niels Provos, Panayiotis Mavrommatis, Moheeb Abu Rajab, and Fabian Monrose. 2008. All your iFRAMEs point to Us. In *Proc. of USENIX'08*. 1–15.
- [60] Ashkan Rahimian, Raha Ziarati, Stere Preda, and Mourad Debbabi. 2014. On the reverse engineering of the citadel botnet. In *Foundations and Practice of Security*. Springer, 408–425.
- [61] S. Ransbotham, S. Mitra, and J. Ramsey. 2012. Are markets for vulnerabilities effective? *MISQ* 36, 1 (2012), 43.
- [62] Paul Resnick and Richard Zeckhauser. 2002. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. *Advances in applied microeconomics* 11 (2002), 127–157.
- [63] Jukka Ruohonen, Sami Hyrynsalmi, and Ville Leppänen. 2016. Trading exploits online: A preliminary case study. In *Research Challenges in Information Science (RCIS), 2016 IEEE Tenth International Conference on*. IEEE, 1–12.
- [64] Armin Sarabi, Ziyun Zhu, Chaowei Xiao, Mingyan Liu, and Tudor Dumitras. 2017. Patch Me If You Can: A Study on the Effects of Individual User Behavior on the End-Host Vulnerability State. In *International Conference on Passive and Active Network Measurement*. Springer, 113–125.
- [65] Karen Scarfone and Peter Mell. 2009. An analysis of CVSS version 2 vulnerability scoring. In *Proc. of ESEM'09*. 516–525.
- [66] Muhammad Shahzad, Muhammad Zubair Shafiq, and Alex X. Liu. 2012. A large scale exploratory analysis of software vulnerability life cycles. In *Proc. of ICSE'12*. IEEE Press, 771–781.
- [67] Carl Shapiro. 1982. Consumer Information, Product Quality, and Seller Reputation. *The Bell Journal of Economics* 13, 1 (1982), 20–35. <http://www.jstor.org/stable/3003427>
- [68] Aditya K. Sood, Rohit Bansal, and Richard J. Enbody. 2013. Cybercrime: Dissecting the State of Underground Enterprise. *IEEE Internet Computing* 17, 1 (jan.-feb. 2013), 60–68.
- [69] Kyle Soska and Nicolas Christin. 2015. Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem.. In *USENIX Security*, Vol. 15.
- [70] Symantec. 2011. *Analysis of Malicious Web Activity by Attack Toolkits* (online ed.). Symantec, Available on the web at [http://www.symantec.com/threatreport/topic.jsp?id=threat\\_activity\\_trends&aid=analysis\\_of\\_malicious\\_web\\_activity](http://www.symantec.com/threatreport/topic.jsp?id=threat_activity_trends&aid=analysis_of_malicious_web_activity). Accessed on June 1012.
- [71] Kurt Thomas, Damon McCoy, Chris Grier, Alek Kolcz, and Vern Paxson. 2013. Trafficking Fraudulent Accounts: The Role of the Underground Market in Twitter Spam and Abuse.. In *Usenix security*, Vol. 13. 195–210.
- [72] Michel Van Eeten and Johannes Bauer. 2008. *Economics of malware: Security decisions, incentives and externalities*. Technical Report. OECD.
- [73] Verizon. 2015. *PCI Compliance Report*. Technical Report. Verizon Enterprise.
- [74] Lingyu Wang, Tania Islam, Tao Long, Anoop Singhal, and Sushil Jajodia. 2008. An Attack Graph-Based Probabilistic Security Metric. In *Proc. of DAS'08*. LNCS, Vol. 5094. Springer, 283–296.
- [75] Michael Yip, Nigel Shadbolt, and Craig Webber. 2013. Why forums? An empirical analysis into the facilitating factors of carding forums. (2013).
- [76] Mingyi Zhao, Jens Grossklags, and Kai Chen. 2014. An Exploratory Study of White Hat Behaviors in a Web Vulnerability Disclosure Program. In *Proceedings of the 2014 ACM Workshop on Security Information Workers (SIW '14)*. ACM, New York, NY, USA, 51–58. <https://doi.org/10.1145/2663887.2663906>
- [77] Mingyi Zhao, Jens Grossklags, and Peng Liu. 2015. An empirical study of web vulnerability discovery ecosystems. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1105–1117.
- [78] 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*. Springer, 225–244.



**Table 11: CVSS Access Complexity by package**

	H	L	M	Sum
ADB	0	1	1	2
ADB2	0	0	1	1
ADB3	0	1	1	2
BANN	0	1	1	2
BOT	0	1	0	1
DOC	0	0	1	1
DOCPPT	0	0	1	1
DROP	0	2	0	2
DROP2	0	2	1	3
ELEN	1	2	8	11
FFLA	0	0	2	2
FLASH	0	1	2	3
GLUE	0	0	1	1
GLUE2	0	0	1	1
IE11	0	0	1	1
IE311	0	0	1	1
IEG8	1	0	0	1
JAVA	0	3	1	4
JAVA2	0	1	0	1
KATR	1	5	5	11
LPE	0	1	0	1
MSW	0	0	1	1
MSW2	0	0	2	2
NEUT	0	2	0	2
OFF	0	0	1	1
PAC	0	0	1	1
PDF	0	1	1	2
PRIV	1	2	1	4
PRIV1	0	1	0	1
R0	0	2	0	2
RIG	0	5	5	10
SILV	0	0	1	1
VIS	0	1	0	1
WMI4	0	0	4	4
XP	0	0	1	1
ZOMB	0	0	3	3
Sum	4	35	50	89

**APPENDIX**

**EXTENDED DESCRIPTION OF DATA FIELDS**

In the following we report a detailed list of the collected data fields.

**CVE.** The Common Vulnerability and Exposures ID assigned by the MITRE corporation to disclosed software vulnerabilities.

**CVEPub.** The date of vulnerability publication on NVD.

**ExplVen.** The anonymized market identifier of the user who posts the advert on the market.

**ExplVenReg.** The date of vendor registration on the forum.

**Pack.** The exploit package advertised by the seller on the market. Packages bundle one or more vulnerability exploits.

**PubDate.** The date of package publication in RuMarket.

**PackType.** The type of exploit package. The categorization emerges from the ExplVen package description.

**PackPrice.** The upfront price that the customer has to pay to acquire or rent the package. For rental packages we compute the mean rental price for a rental period of 3 weeks, in line with previous work on the duration of malware delivery campaigns and user infections [41]. All prices in RuMarket are reported in USD.

**PackActivity.** The RuMarket activity around an exploit package expressed as the number of replies received by the product advert on the market.

**PackDeath.** The date when RuMarket activity around the package stops.

**ExplPrice.** Estimation of cost of a single exploit. When only one exploit is bundled in a package, this corresponds to the package price, otherwise we provide an estimation. Further considerations are given in Sec. 3.4.

**SwVen.** The vendor of the vulnerable software affected by the published exploit.

**Sw.** The vulnerable software or platform.

**CVSS.** The severity of the exploited vulnerability as expressed by the Common Vulnerability Scoring System, the standard framework for evaluation of software vulnerability characteristics [31].

**SYM.** Presence or absence of an exploit in the wild at scale.

**SYM REGRESSION**

Table 10 below reports standard OLS coefficients using the average value of the fixed effects for each CVE as regressors. Results are equivalent to those shown by the mixed model in Table 9.

**Table 10: OLS regression of exploitation in the wild**

SYM	Model 1	Model 2	Model 3
c	-2.062 (1.479)	5.652 (3.580)	6.846 <sup>†</sup> (3.959)
log(PackActivity)	1.108* (0.460)	1.436** (0.549)	1.593* (0.627)
log(PackPrice)		-1.184* (0.546)	-1.299* (0.589)
CVSS:NC			-2.184 <sup>†</sup> (1.124)
McFadden R-sq.	0.2	0.4	0.5
Log-likelihood	-17.7	-14.1	-11.9
N	51	51	51

**Table 12: CWE counts by package**

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	Sum
ADB	0	0	0	0	1	0	0	1	0	0	0	0	0	0	2
ADB2	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1
ADB3	0	0	0	0	1	0	0	1	0	0	0	0	0	0	2
BANN	0	0	0	0	0	0	0	1	0	1	0	0	0	0	2
BOT	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1
DOC	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1
DOCPPT	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
DROP	0	0	0	0	0	0	0	0	0	0	0	2	0	0	2
DROP2	1	0	0	1	0	0	0	0	0	0	0	1	0	0	3
ELEN	2	2	0	0	3	0	1	1	1	0	0	0	1	0	11
FFLA	0	0	0	0	2	0	0	0	0	0	0	0	0	0	2
FLASH	2	0	0	0	0	0	0	1	0	0	0	0	0	0	3
GLUE	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1
GLUE2	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1
IE11	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1
IE311	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1
IEG8	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1
JAVA	0	0	0	0	3	0	0	0	1	0	0	0	0	0	4
JAVA2	0	0	0	0	1	0	0	0	0	0	0	0	0	0	1
KATR	1	2	0	1	5	0	0	1	0	0	0	1	0	1	11
LPE	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1
MSW	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1
MSW2	0	0	1	0	0	0	0	0	0	1	0	0	0	0	2
NEUT	0	0	0	0	2	0	0	0	0	0	0	0	0	0	2
OFF	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1
PAC	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1
PDF	0	0	0	0	1	0	0	1	0	0	0	0	0	0	2
PRIV	0	1	0	0	3	0	0	0	0	0	0	0	0	0	4
PRIV1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1
R0	0	0	0	0	0	0	0	0	0	0	0	2	0	0	2
RIG	2	1	0	0	4	0	0	1	0	1	0	0	0	1	10
SILV	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1
VIS	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1
WMI4	2	2	0	0	0	0	0	0	0	0	0	0	0	0	4
XP	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1
ZOMB	0	1	0	0	0	0	0	1	1	0	0	0	0	0	3
Sum	14	13	1	7	26	0	1	9	4	3	0	7	3	1	89

CWEs are uniformly distributed among packages, i.e. each pack contains vulnerabilities of the same or comparable type. This indicates that within a single package, exploit development cost are comparable and significantly skewed distributions of costs are unlikely.

1. Buf. Err. (CWE-119)
2. Code Inj. (CWE-94)
3. Data Handl. (CWE-19)
4. Input Val. (CWE-20)
5. Ins. Inf. (NVD-CWE-noinfo)
6. Link Following (CWE-59)
7. NONE
8. Num. Err. (CWE-189)
9. OS Comm. Inj. (CWE-78)
10. (NVD-CWE-Other)
11. Path Trav. (CWE-22)
12. Perm., Priv., Acc. Cntrl (CWE-264)
13. Res. Mngmt Err (CWE-399)
14. Use Af. Free (CWE-416)

**Table 13: History of re-packaged vulnerabilities**

CVE	no.	Sw	SwVendor	Type 1	Type 2	Type 3	Type 4
2015-8651	4	flash	adobe	STANDALONE	STANDALONE	STANDALONE	EKIT
2010-0188	4	acrobat	adobe	EKIT	EKIT	MALWARE	MALWARE
2012-1864	3	windows	microsoft	STANDALONE	MALWARE	MALWARE	
2015-1701	3	windows	microsoft	STANDALONE	MALWARE	MALWARE	
2010-4452	3	java	oracle	EKIT	EKIT	STANDALONE	
2006-0003	3	int_expl.	microsoft	EKIT	EKIT	EKIT	
2010-1885	3	windows	microsoft	EKIT	EKIT	EKIT	
2016-1019	2	flash	adobe	STANDALONE	EKIT		
2013-2729	2	acrobat	adobe	STANDALONE	STANDALONE		
2013-0640	2	acrobat	adobe	STANDALONE	STANDALONE		
2015-2545	2	office	microsoft	STANDALONE	STANDALONE		
2015-0057	2	windows	microsoft	STANDALONE	MALWARE		
2013-3660	2	windows	microsoft	STANDALONE	MALWARE		
2011-0611	2	acrobat	adobe	EKIT	EKIT		
2010-0886	2	java	oracle	EKIT	STANDALONE		
2008-2463	2	office	microsoft	EKIT	EKIT		
2015-0336	2	flash	adobe	EKIT	STANDALONE		
2013-3918	2	int_expl.	microsoft	STANDALONE	STANDALONE		
2015-2419	2	int_expl.	microsoft	STANDALONE	EKIT		
2014-6332	2	windows	microsoft	EKIT	STANDALONE		
2010-0840	2	java	oracle	EKIT	STANDALONE		
2011-3544	2	java	oracle	STANDALONE	EKIT		
2012-0158	2	office	microsoft	STANDALONE	STANDALONE		

**Table 14: (Bootstrapped) descriptive statistics of exploit price estimates in USD by software**

We compute expected exploit prices by considering a uniform distribution of exploit costs by package. To approximate the true (unknown) distribution of exploits per package and software, we perform a bootstrap of our data ( $N = 10000$ ). Estimates for MALWARE and EKIT exploits are only indicative as development costs of the package can not be accounted for. We report descriptive statistics of the original and of the bootstrapped sample means in parenthesis. The column  $n$  reports number of exploits for that software in the respective package type. The bootstrapped data does not deviate substantially from our observations on the average. Fatter distribution tails indicate that RuMarket outliers tend to bias sample statistics. Microsoft exploits are on average the most valuable in the market irrespective of package. MS Windows and Office exploits are consistently the most expensive. Adobe and Oracle are closely second and third. EKIT prices by exploit are substantially lower than for STANDALONE and MALWARE exploits.

Pack Type	SwVendor	Software	Min	0.025p	Mean	Median	0.975p	Max	sd	n	
EKIT	Adobe		13.64	13.64	98.24	31.67	333.33	333.33	111.28	17	
			(13.64)	(13.64)	(135.79)	(106.86)	(333.33)	(333.33)	(90.38)		
		flash	31.67	31.67	91.86	31.67	299.24	333.33	105.33	10	
				(31.67)	(31.67)	(132.99)	(98.7)	(333.33)	(333.33)	(89.9)	
	acrobat		13.64	13.64	107.36	13.64	310.61	333.33	127.34	7	
			(13.64)	(13.64)	(138.59)	(114.55)	(333.33)	(333.33)	(90.77)		
	Microsoft		13.64	13.64	115.71	120.83	284.09	333.33	96.76	14	
			(13.64)	(13.64)	(113.81)	(111.44)	(232.32)	(333.33)	(62.23)		
		office	13.64	17.84	97.73	97.73	177.61	181.82	118.92	2	
				(13.64)	(13.64)	(98.44)	(97.73)	(181.82)	(181.82)	(63.88)	
	int. expl.		13.64	16.34	97.47	120.83	181.82	181.82	71.89	7	
			(13.64)	(13.64)	(94.92)	(97.47)	(181.82)	(181.82)	(38.63)		
windows		13.64	15.44	148.45	181.82	318.18	333.33	130.6	5		
		(13.64)	(13.64)	(146.27)	(148.45)	(333.33)	(333.33)	(66.71)			
	Oracle	13.64	13.64	87.32	106.67	181.82	181.82	68.78	10		
			(13.64)	(13.64)	(94.75)	(92.26)	(181.82)	(181.82)	(39.1)		
java		13.64	13.64	87.32	106.67	181.82	181.82	68.78	10		
		(13.64)	(13.64)	(94.75)	(92.26)	(181.82)	(181.82)	(39.1)			
				(420)	(420)	(455.35)	(455)	(490)	(490)	(26.71)	
MALWARE	acrobat		420	421.75	455	455	488.25	490	49.5	2	
			(420)	(420)	(455.35)	(455)	(490)	(490)	(26.71)		
	Microsoft		333.33	333.33	1357.14	1500	3625	4000	1303.23	7	
			(333.33)	(333.33)	(1518.3)	(1375)	(4000)	(4000)	(745.71)		
	windows		333.33	333.33	1357.14	1500	3625	4000	1303.23	7	
			(333.33)	(333.33)	(1518.3)	(1375)	(4000)	(4000)	(745.71)		
Adobe		420	421.75	455	455	488.25	490	49.5	2		
STANDALONE	Adobe		75	75	879.17	1250	1500	1500	693.54	12	
			(75)	(100)	(1000.06)	(1040)	(1500)	(1500)	(521.91)		
		flash	75	75	568.75	150	1500	1500	652.3	8	
				(75)	(87.5)	(562.05)	(545.45)	(1300)	(1500)	(316.52)	
	acrobat		1500	1500	1500	1500	1500	1500	0	4	
			(1500)	(1500)	(1500)	(1500)	(1500)	(1500)	(0)		
	Microsoft		150	150	2801.82	2250	8000	8000	2393.09	22	
			(150)	(150)	(2442.13)	(2450)	(5600)	(8000)	(1601.69)		
		office	150	362.5	3195.45	4000	7250	8000	2504.04	11	
				(150)	(1605.1)	(3407.31)	(3262.5)	(5750)	(8000)	(1112.54)	
	int. expl.		440	459.5	3035	1850	7625	8000	3504.22	4	
			(440)	(440)	(3051.89)	(3000)	(8000)	(8000)	(1727.18)		
	windows		700	800	2366.67	2250	4687.5	5000	1458.31	6	
			(700)	(1100)	(2349.27)	(2327.27)	(3750)	(5000)	(658.77)		
		silverlight	150	150	150	150	150	150	0	1	
			(150)	(150)	(150)	(150)	(150)	(0)			
Oracle		25	25	1020	25	4502.5	5000	2224.89	5		
		(25)	(25)	(1847.02)	(1020)	(5000)	(5000)	(1981.08)			
	java	25	25	1020	25	4502.5	5000	2224.89	5		
			(25)	(25)	(1847.02)	(1020)	(5000)	(1981.08)			