

Exploit Brokers and Offensive Cyber Operations

Matthias Dellago
Andrew C. Simpson
Daniel W. Woods

ABSTRACT

A necessary step in conducting offensive cyber operations is developing or acquiring an exploit, i.e., a means for taking advantage of a software vulnerability or security deficiency. While these can be developed within government agencies, they can also be procured from private actors. Studying these private markets present an opportunity to understand offensive cyber operations, especially as markets break from the secretive culture of intelligence agencies. This article provides novel evidence of such opportunities by collecting data in the form of the prices quoted by an exploit broker who claims to sell to governments. We find exploit price inflation of 44% per annum, and higher prices for exploits targeting mobile devices relative to desktop devices. Exploits requiring additional capabilities like physical access to the device are quoted at a discount, and no-click remote access vulnerabilities carry a heavy premium. The broker does not quote prices for any exploits that specifically target industrial control systems or IoT devices. We conclude by discussing how these results inform the future of offensive cyber.

INTRODUCTION

The emergence of offensive cyber operations (OCO) – “the adversarial manipulation of digital services or networks”^[1] – creates new considerations in military strategy and government policy. The resulting debates consider issues like the nature of cyber weapons,^[2,3] the possibility of cyber war^[4,5] the role of norms of responsible behavior^[6,7] and, most importantly for this paper, the role of private actors in developing and deploying offensive cyber technology.^[8,9] Such issues are even spilling over into the public sphere as evidenced by Nicole Perlroth's New York Times bestseller^[10] arguing that

© 2022 Matthias Dellago, Andrew C. Simpson, Daniel W. Woods



Matthias Dellago is a computer science master's student at the University of Innsbruck. He earned his bachelor's degree in physics from the University of Vienna. His research interests include security and privacy, especially from an economic perspective.



Andrew Simpson holds a BSc in Computer Science from Swansea University and an MSc and a DPhil from the University of Oxford. He is currently an Associate Professor in Software Engineering in the Department of Computer Science at the University of Oxford.

private actors who supply offensive cyber technology are facilitating repressive regimes in targeting opposition politicians and journalists.

To incorporate such actors into national cyber strategy and to ensure responsible behavior,^[11] it is important first to understand the market structures through which they operate. We apply the tools of security economics^[12] to understand the business processes and price structure surrounding the supply of offensive cyber technology. Doing so provides a rare opportunity to collect empirical data on offensive cyber operations, as such private actors break from the secretive culture of intelligence agencies. While prior work has focused on bug bounty programs^[13] and illegal underground forums,^[14] we provide a longitudinal analysis of a zero-day exploit broker whose customers are “government organizations (mainly from Europe and North America).”^[15] Our empirical results show that the mean exploit price is increasing by \$234 per day or 44% per annum. Exploits of both Apple operating systems and mobile devices have a higher average price in our dataset. In terms of the application targeted, exploits targeting communications (e.g., emails and messengers) have the highest average price. Further, we found no evidence that this broker procures exploits of technologies specifically targeting industrial control systems. These findings may not generalize beyond the idiosyncratic broker we study, especially given prices are based on the maximum price advertised for each exploit rather than actual payouts.

Turning to the question of this CDR Special Edition, these insights can inform the future of offensive cyber. First, exploit price inflation represents a growing constraint on offensive operations. Importantly, this constraint binds ex-ante unlike imposing costs via deterrence. Increased rewards for exploits in private markets function to increase the staffing costs for states maintaining internal offensive cyber capabilities and may motivate export controls and other policy interventions.



Daniel Woods is a Lecturer in Cybersecurity at the University of Edinburgh. His post is jointly appointed by the British University in Dubai. He received his Ph.D. in Cybersecurity from the University of Oxford's Computer Science Department.

RELATED WORK

Selling exploits to a broker is but one of many ways for an independent security researcher to share information. The options available include:^[16] privately reporting the information to the vendor (possibly in exchange for a bug bounty) or to a legitimate third party; selling the information on the black market; and sharing the information publicly. Before we turn to economic incentives, it is worth noting that many researchers share information without any financial reward. For example, the CERT Coordination Center (CERT/CC) have been running coordinated vulnerability disclosure for over 30 years without offering any financial reward, and have exchanged over 430K emails in the process.^[17] Similar institutions exist outside the US.^[18]

Vulnerability Markets

Multiple sales channels exist for researchers seeking monetary compensation. Bug bounty programs, in which researchers are rewarded for reporting directly to the vendor,^[19] sit at the legitimate end of the spectrum. Black markets, in which criminals offer financial rewards for exploits, sit at the illicit end of the spectrum.^[20,21,14] Exploit brokers can be considered gray markets existing somewhere between bug bounties and black markets, with legitimacy varying based on who the broker sells to.

These institutions display many properties of traditional markets. For example, bug bounties display upwards sloping supply curves.^[22,13] Perhaps more surprisingly, black markets have developed enforcement mechanisms that prevent dishonest practices^[14] and freelancers have declined from 80% to 20% of total participants (as of 2014) as criminal organizations form.^[20] In terms of outcomes, empirical works show that bug-bounty programs are effective^[23,24] and efficient^[25,26] security interventions. We also see that exploits procured in black markets are used by threat actors.^[27]

Although different sales channels exist, the viability and rewards of each channel will vary depending on the particular exploit. On the supply side, Luna et al.,^[24] find that experienced researchers display different work patterns to entry-level researchers. In terms of equilibrium price, exploit kits are priced in thousands of dollars,^[20] and the average bug bounty on the HackerOne platform was just \$318,^[26] whereas zero-day exploits can be priced in the millions.^[11] This motivates considering markets for zero-days separately.

Zero-Day Markets

Zero-day exploits take advantage of a security vulnerability that is not known to the software vendor or the wider security community. Such exploits are powerful because two important tools are not available to defenders, namely applying software patches designed to fix the underlying vulnerability and scanning for “signatures,” the behavioral patterns and code of past exploits. This means zero-day exploits can target more devices and be detected less easily than N-day exploits, where N is the number of days since the exploit or vulnerability was public. This section does not exhaustively examine the technical or policy aspects of zero-days but does try to do so for empirical studies of market structure. In terms of technical analysis, Stone^[28] analyzed the 24 zero-days detected in the wild in 2020, nine of which were variants on “previously disclosed vulnerabilities” or incompletely patched. This raises the question of how markets deal with zero-day variants. Turning to policy, Fidler^[11] outlines the national and international policy apparatus surrounding zero-days considering issues like export bans that likely impact market participants and structure.

The nature and ethics of zero-day markets were probed at a 2013 workshop,^[29] which documented how zero-day markets operated largely in the shadows. A year later, Ablon et al.^[20] assembled a “sparse and inconsistent” table of prices for zero-days and note that whether prices are increasing or decreasing is an open question. Table 2 of Meakins’ work^[30] provided a snapshot of pricing for a limited number of vulnerabilities across four different brokers. Interestingly, they show the high-end prices at Western brokers are an order of magnitude higher than at the broker operating in Russia. Table 2 does not differentiate between the properties of an exploit, such as whether physical access or user interaction is required.^[30]

While the previous papers^[20,30] opted for a comparative study of multiple brokers, we provide an in-depth study of just one broker. This allows us to identify the longitudinal development of prices and answer the open question of whether prices are increasing.^[20,26] Although we have only studied the maximum prices quoted by one broker, which is an imperfect proxy of the actual fee paid to researchers. Further, we also collect information about not only the systems targeted in an exploit (as in ^[20,30]), but the capabilities required to use that exploit. The next section describes the process by which exploits are sold, as this sheds light on some of the open questions.

RESEARCHER-BROKER RELATIONS

The supply side of zero-day markets consist of a researcher selling an exploit to the broker. We describe the process using a mixture of testimonies from researchers,^[31,32] the websites of brokers^[15,33] and research articles^[29,34]. The seller contacts the broker, whether through connections or directly, and shares the exploit's specifications. Important criteria are:

- ◆ The targeted software, OS and architecture.
- ◆ The type of vulnerability (e.g., use-after-free).
- ◆ Attack vector (website, document, etc).
- ◆ Reliability (typically probability of success needs to be > 90%).
- ◆ Speed of exploitation (on the order of seconds).
- ◆ Does the exploit crash running processes?
- ◆ Is user interaction required?
- ◆ Does the exploit work with default settings?
- ◆ Any other relevant limitations.

The broker responds with a non-binding preliminary offer, usually less than the publicly advertised maximum payout, after taking limitations into account. The seller may then submit their exploit for evaluation by the broker. It is customary to allow for an assessment period of up to two weeks.^[31,15] In this time the broker tests the zero-day and compares their result to the specifications provided by the seller. Given no contract has been signed before verification, the seller generally has to trust the broker not to share the information about the exploit,^[35] although some brokers sign a contract with the seller before the submission.^[31,33]

Whether the contract is signed before or after the validation period, the contract specifies: payment terms (warranty), intellectual property rights, exclusivity and support requirements.^[31,32] The payment is usually spread out over the course of a few months to a year. The contract is contingent on no patch being developed that purposely (or accidentally) fixes the vulnerability underlying the exploit. Depending on the terms of the contract, the seller may be required to either to provide a replacement exploit or forfeit all outstanding payments. This also serves as an incentive to honor possible exclusivity agreements. Previous research indicates that exploits are quite likely to survive this period^[34] providing they are used responsibly.^[31] This contractual structure is sometimes referred to as “split the risk.”

Compromising exclusivity, coined “double dipping” by Schwartz,^[31] by selling the same exploit to multiple parties is risky due to the small size of the market. With estimates of active researchers ranging from 400^[31] to 1500^[35] and a much smaller number of brokers and buyers, the discovery of dishonesty becomes quite likely. The seller would thus incur legal troubles and reputational loss.^[31]

Until 2015, brokers did not publicly advertise prices.^[29] Instead, market participants needed to navigate informal professional networks, a sign of an immature market. Since then, at least two brokers have begun to publicly advertise prices paid to researchers (notably omitting what buyers pay).^[35,33] Our research studies one of these brokers.

RESEARCH DESIGN

Our goal is to capture the development of exploit prices offered by Zerodium. This broker buys zero-day exploits from freelance developers and provides them to government agencies.^[15] Their website lists prices for different exploits, in a graphic designed to resemble a periodic table, an example of which is depicted in Figure 1. It should be noted that the prices listed are the maximal amount, and actual payouts depend on “quality of the submitted exploit (full or partial chain, supported versions/systems/architectures, reliability, bypassed exploit mitigations, default vs. non-default components, process continuation, etc.).”^[15]

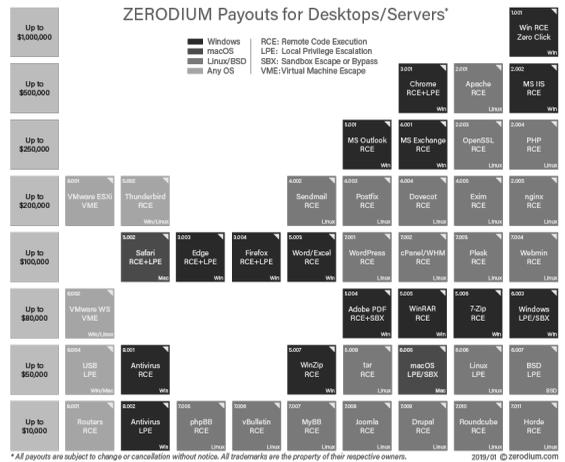


Figure 1: The Figure lists different tiers of prices on the left-most column and a brief description of an exploit in each square, for example “Chrome RCE+LPE on Windows OS” and a price on the left-hand side is bought for “up to \$500,000.”

Data Collection

We collected the longitudinal data via the Internet Archive's Wayback Machine^[36]. Using their CDX API, we determined that the price table has changed only seven times since the program's inception in 2015. The Internet Archive samples are published much more frequently (243 times since 2015) than new prices, which increases confidence that our sample does not miss data.

The tables are available only as images, which we transcribed manually. We extracted each tile from the seven price tables. We then combined exploits that have the same name but are listed in different tables (i.e., at different times) to see how prices vary over time. Thus, we recover a sample of 205 unique types of exploits. This leads to a sparse and irregular panel data set from 2015 until the present, in which the quoted price is the dependent variable.

We further classify the exploits in order to run linear regressions, we chose this functional form for interpretability.¹ For each advertised price, we calculate the number of days since our first observation.

1 Future work will explore models tailored to our irregular, auto-correlated panel data that contains outliers.

Statistic	N	Mean	St. Dev.	Min	Max
price (\$)	543	183.8k	293,063.8	5k	2.5m
days (since start of sample period)	543	846.4	439.3	0	1,387
osandroid	543	0.14	0.3	0	1
osbsd	543	0.01	0.43	0	1
osios	543	0.2	0.4	0	1
oslinux	543	0.2	0.4	0	1
osmac	543	0.04	0.2	0	1
osunspecified	543	0.2	0.4	0	1
oswindows	543	0.2	0.4	0	1
oswindows.phone	543	0.01	0.1	0	1
BrowserTrue	543	0.2	0.4	0	1
EmailTrue	543	0.1	0.3	0	1
MessengersTrue	543	0.1	0.3	0	1
Web.ServerTrue	543	0.1	0.3	0	1
antivirusTrue	543	0.02	0.1	0	1
Requires.Local.AccessTrue	543	0.05	0.2	0	1
Local.Privilege.EscalationTrue	543	0.3	0.5	0	1
Mitigation.BypassTrue	543	0.01	0.1	0	1
Remote.Code.ExecutionTrue	543	0.7	0.5	0	1
Full.Chain.with.PersistenceTrue	543	0.03	0.2	0	1

Table 1: Descriptive data. All rows apart from price and days are dummy variables. The mean value column describes the proportion of the 543 vulnerabilities for which that property is true. For example, 14% of the exploits target Android devices and 70% provide remote code execution functionality.

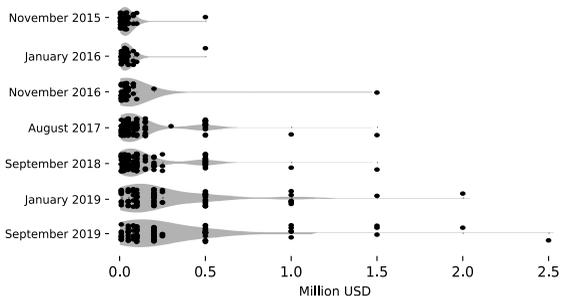


Figure 2: Univariate scatter plot of exploit bounties offered at the different captured snapshots. One dot represents one exploit. Violin plots are added to visualize the concentration of exploit types around certain price bands.

Model 1 shows that around 20% of the variance can be explained by temporal fixed effects. Models 2-4 isolate the explanatory power of the dummy variables based on targeted OS, targeted application, and exploit type respectively. The targeted OS has the least explanatory power (Model 2), likely because each OS contains a range of applications with varying levels of interest and security. For example, an exploit of WhatsApp messenger has the same price for both the iOS or Android version. The targeted application (Model 3) and type of exploit (Model 4)

We then build several dummy variables.² We extracted³ explanatory variables like the vendor and type of product (e.g., messenger, browser, etc.), as well as the kind of exploit (remote code execution, local privilege escalation, etc.). These can be seen in Table 1.

RESULTS

Figure 2 shows how exploit prices are distributed. The majority of prices are \$100k or less, especially in the early years of our sample. The most expensive exploits inflate in price rapidly from 2016, growing by 500%. Prices cluster around salient values, such as the cluster at \$500K that emerged from 2017 onward. Negotiated prices may not display such clustering.

We then ran a number of regressions to understand what explains this variance. Table 2 contains a number of log-linear models (1 through 7, with column heads at the top of the table) with the exploit's dollar price as the dependent variable. We opted for log-linear over linear models after inspecting QQ plots, but for comparison we include the equivalent linear regressions in Table 3 in the appendix.

2 A dummy variable is equal to 1 if the property is true, and 0 otherwise. For example, the dummy variable BrowserTrue is equal to 1 if the exploit targets a browser application and 0 otherwise.

3 We used the OS labels from the price table, using unspecified when none was provided, which was a minority of cases. To create the other categories (vendor, product type and exploit type) we automatically searched the exploit names for certain keywords. For instance, an exploit whose name contained any of the keywords "messenger," "signal," "telegram," "whatsapp" or so on was categorized under "Messenger." We chose these categories with the intent of grouping similar exploits, to allow for descriptive modelling.

EXPLOIT BROKERS AND OFFENSIVE CYBER OPERATIONS

Model	1	2	3	4	5	6	7
android		1.740***			0.828***		0.493**
		-0.257			-0.229		-0.167
bsd		0.502			0.095		-0.233
		-0.649			-0.52		-0.502
ios		2.063***			0.991***		0.657***
		-0.256			-0.232		-0.169
linux		0.655*			0.388		0.107
		-0.259			-0.212		-0.155
mac		0.674*			0.147		-0.15
		-0.316			-0.271		-0.234
windows		1.188***			0.783***		0.499**
		-0.255			-0.21		-0.156
windows phone		1.444*			-0.768		-0.937
		-0.61			-0.518		-0.521
BrowserTrue			0.249*		0.266*		0.236*
			-0.12		-0.106		-0.108
EmailTrue			0.036		0.124		0.125
			-0.14		-0.126		-0.128
MessengersTrue			1.821***		1.180***		1.161***
			-0.137		-0.14		-0.142
Web.ServerTrue			-0.876***		-0.593***		-0.589***
			-0.136		-0.131		-0.134
antivirusTrue			-0.947**		-1.019***		-1.055***
			-0.297		-0.252		-0.255
Requires.Local.AccessTrue		-0.495*	-0.377*		-0.358*		
				-0.197	-0.167		-0.17
Local.Privilege.EscalationTrue		1.054***	0.482***		0.542***		
				-0.091	-0.095		-0.095
Mitigation.BypassTrue		1.183*	1.074*		1.077*		
				-0.563	-0.466		-0.474
Remote.Code.ExecutionTrue		0.633***	0.591**		0.587***		
				-0.097	-0.096		-0.097
Full.Chain.with.PersistenceTrue		3.164***	2.758***		2.930***		
				-0.25	-0.241		-0.234
days						0.001***	0.001***
						-0.0001	-0.0001
Constant	10.167***	9.931***	10.296***	9.275***	9.384***	10.191***	9.381***
	-0.191	-0.174	-0.164	-0.178	-0.159	-0.111	-0.123
Observations	543	543	543	543	543	543	543
R2	0.203	0.388	0.473	0.487	0.662	0.186	0.647
Adjusted R2	0.194	0.373	0.462	0.477	0.647	0.185	0.635

Note: *p<0.05; **p<0.01; ***p<0.001

Table 2: Linear regressions with log-transformed price (\$) as the dependent variable. Time-based fixed effects included for all but Model 6 and 7.

effect size, which is particularly striking given Figure 6 shows the broker did not trade such exploits when it was launched in 2015. Exploits of web servers (*Web.ServerTrue*) and anti-virus products (*antivirusTrue*) are comparably cheaper, as seen in Figure 6.

Turning to properties of the exploit, we find a number of reassuringly obvious results. The variable for full-chain-persistence has the largest effect size, which is unsurprising given such an exploit can be used to compromise any other application on the device. Conversely, the least powerful exploits – those that require local access (e.g., to insert a USB driver) – are comparably cheaper, which can be seen in the regression coefficient *Requires.Local.AccessTrue*. Figure 7 shows the average price for each type of exploit over time.

have more explanatory power, with individual dummies contributing a lot—removing the messenger and full-chain-persistence dummies lead to 38% and 33% reductions in R^2 in Model 3 and Model 4 respectively.

Comparing the R^2 of Model 5 and Models 2-4 shows that the variables have additional predictive power when taken together, and so we proceed by analyzing this model. The coefficients for targeted OS in Model 5 suggest that exploits targeting mobile devices (e.g., Android and iOS) are more expensive than those targeting desktops, which can also be seen in Figures 4 and 5. Figure 3 shows the total bounties available has been consistently high for Apple, but Google has recently overtaken Microsoft, likely due to the increasing cost and availability of exploits targeting mobile devices.

Turning to specific products, Model 5 shows that *MessengerTrue* (e.g., WhatsApp, iMessage, Signal and so on) has the second largest

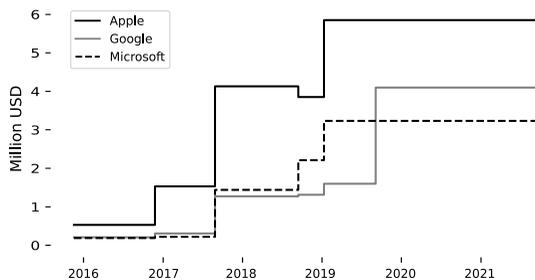


Figure 3: The sum of all exploit bounties for certain vendors, from 2015 until present.

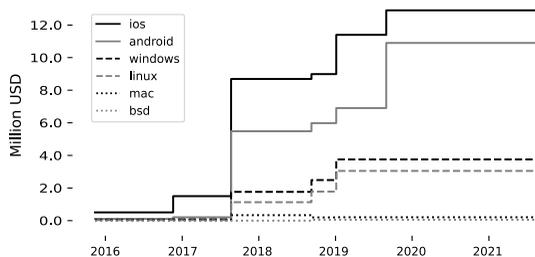


Figure 4: The sum of all exploit bounties by OS, from 2015 until present.

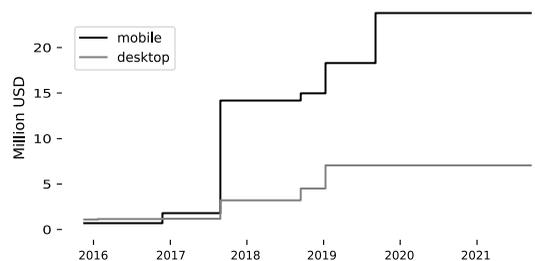


Figure 5: A comparison of total bounties offered for mobile versus desktop/server systems (2015 until present).

justified by the broker receiving multiple exploit submissions.^[38] This suggests that bursts in the supply of exploits can drive down prices, although we doubt buyers see an equivalent reduction in price. We cannot observe whether price fluctuations influence researcher attention. Ultimately, we can only caution against reading too much into prices and call for a more sophisticated economic analysis in future work.

We can, however, make a number of reliable observations. The monetary cost of exploiting certain systems is a consideration in itself. For example, exploits of modern messenger applications can now cost over a million dollars. These costs no doubt drive law enforcement's calls for "exceptional access,"^[39] whereby technology companies would be required to build-in backdoor vulnerabilities that governments can access in response to an incident and/or investigation.

High exploit prices feed into the challenge of retaining security researchers, who can leave and sell their expertise to the highest bidder, this impacts both government agencies and the

All the longitudinal figures show that exploits generally become more expensive over time. This can also be observed in the days variable in Model 7 – here we impose a linear relationship between the number of days since the first set of prices and that exploit's price. The linear model in the generally appendix (Table 3) shows a mean increase of \$234 per day in our observation period, which translates into 44% growth per annum. Also inspecting the fixed effects on each time period, we find larger effects for the later periods.

DISCUSSION

We first consider what these results tell us about wider debates, and then reflect on using exploit brokers as a data source.

Interpreting prices

Interpreting prices is notoriously difficult.^[37] Are prices high because many governments target that system (demand driven) or because that product is particularly secure (supply driven)? In 2020, the broker announced that purchases of iOS exploits were temporarily suspended, which was

vendors of these products.^[10] These problems will persist given that the average price of an exploit increases by \$234 per day (see Table 3) or 44% per year. This speaks to the open question of whether exploit prices are increasing.^[20]

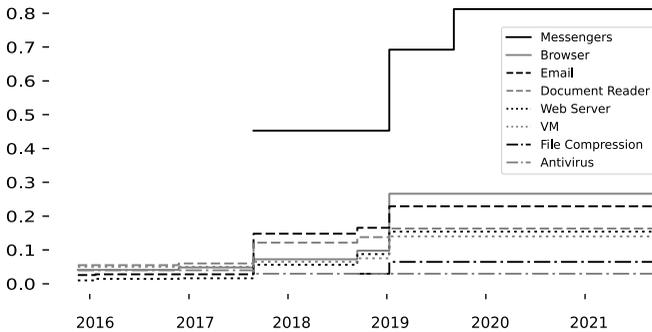


Figure 6: Comparing the average prices of exploits for different product categories, from 2015 until present.

Omissions

We observed a lack of exploits that could conceivably cause physical damage targeting Industrial Control Systems or even IoT devices. This is likely because ICS systems display a different security model from mobile applications, not to mention the increased barriers to conducting ICS research. The majority (64%) of ICS vulnerability advisories had no patch, and instead relied on network segmentation to avoid compromise.^[40] Thus, only exploits that provide “access to a control system network”^[40] are valuable. Similarly, IoT devices are often insecure by default.^[41] Exploit markets are unlikely to exist where the barrier to compromise is low enough for internal expertise.

An alternative explanation for this omission is that other brokers or criminal groups trade in exploits providing such access, or even that this broker trades in them without announcing prices. More generally, that Zerodium^[15] and Crowd-Fense^[33] offer public prices suggests that trading in these specific exploits is not deemed to incur prohibitive reputation or legal risk. We return to the question of cyber norms in the final section.

Data Sources

Building an empirical picture of offensive cyber operations runs against the interests of those conducting such operations. While circumstances can exist under which belligerents claim credit for cyber operations,^[42] secrecy is the default.^[43] The dynamics of offensive cyber will outpace time-lagged sources used by traditional intelligence studies like declassified documents^[44,45] or officers retiring and then revealing details.^[46] To address this, cyber strategy scholars have turned to novel data sources.

The ease of duplicating operational computer code leads to publications by third parties like governments,^[47] private firms,^[48] and academics.^[49,50] Each samples in a different way, leading to very different pictures of cyber operations. Egloff^[47] argues that beyond establishing facts (sense-making), state-led attributions also aim to influence public and elite opinion—this sampling bias would lead one to believe cyber operations are primarily conducted by a handful of governments against the West.^[51,Table V] In contrast, the Citizen Lab's analyses of malware campaigns^[49,50] suggest journalists are the primary target. Stretching the definition of offensive

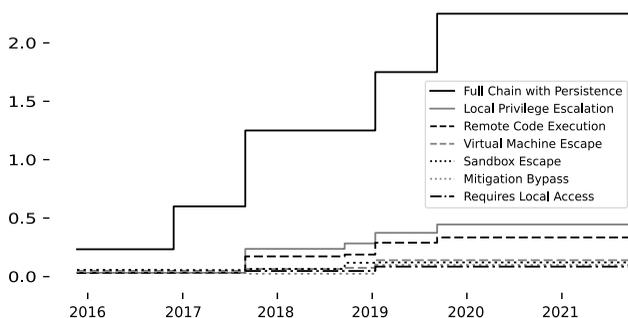


Figure 7: Comparing the average prices of different kinds of exploits, from 2015 until present.

cyber to include information operations, empirical sources like Twitter^[52,53] and message boards^[54] suggest the public is the primary target. Thus, the choice of evidence base leads to a different characterization of offensive cyber operations.

We argue that our data supplements existing data sources. In particular, relying on analyses of operational code leads to a reporting bias in that it necessarily ignores undetected operations. An additional strength of our data is that markets are updated in real time, whereas code analysis takes time and declassification takes even longer.

Limitations

Our findings are based on maximum payouts, whereas real payouts are negotiated and likely much lower. Thus, our data over-inflates the demand and supply of exploits. Zerodium’s publishing of disclosures likely serves the commercial interest of generating publicity, and so Zerodium’s prices may fundamentally differ from those of other brokers. For example, Meakins^[30] showed that brokers who operate in Russia quote far lower prices. Beyond limitations in the data, our modelling was crude and requires further refinement.

FUTURE OF OFFENSIVE CYBER

Exploit prices quoted by brokers provides insights into offensive cyber operations. We discovered that:

- 1) Exploits of iOS and windows are the most expensive for mobile and desktop respectively with mobile exploits higher on average.
- 2) Exploits of messengers and browsers are more expensive than those of web servers and anti-viruses.
- 3) Exploits requiring local access are comparably inexpensive.

We also observed a general trend towards exploits becoming more expensive over time. So, what does this mean for the future of offensive cyber operations? The rest of this section speculates on three aspects of this question.

Could offensive cyber operations be constrained by exploit markets?

Continued exploit price inflation represents an increasing economic constraint on offensive cyber operations. Scholars of security economics have long argued that increasing the cost to attackers is viable route forward given that perfect security is not achievable.^[55,56] The resulting

barrier to entry provides a hard ex-ante limit on offensive cyber operations, whereas cyber deterrence imposes ex-post disincentives that rely on victims detecting, attributing, and authorizing a response to offensive cyber operations, each of which is uncertain.

Such constraints vary according to the targeted system and the capabilities of the offensive actor as our regressions show. The future implications are unclear, but it is clear that no-click, remote access exploits are inflating fastest (see Figure 7). As a result, exploits requiring local access are relative bargains. One could speculate that cyber operations targeting domestic actors are becoming relatively more cost-effective because capabilities like physical access to the targeted device are more realistic for domestic actors. Note this also assumes offensive actors are rational, which may not always be the case.

Could offensive cyber operations be constrained by vulnerability researchers?

First, exploit brokers offer incentives for individuals to leave government agencies, although admittedly higher private-sector salaries are nothing new. These incentives are relatively higher for more talented researchers thus creating staffing problems. This could motivate export controls and other legal limitations on the sale of zero-days.^[11] Perhaps more interestingly, researchers are paid based on how long the zero-day remains un-patched (see Section 3). Could researchers exert pressure against wanton use of the exploit that increases the likelihood of detection and hence a patch that disrupts the payment plan? This turns on how much market power researchers have.

The number of independent active sellers (between 400^[31] and 1500^[35] individuals) relative to buyers (a small number of states conducting OCO) suggest the power is limited. Further, looking at Zerodium's total payouts, \$50 million for exploits since their founding in 2015,^[35] suggests an annual pay of \$5.5k - 20.8k per researcher. Such estimates should be interpreted in light of researchers having multiple income streams (e.g., multiple brokers, bug bounties, and other security work) and the reality that superstar effects mean a minority collect the majority of payments.^[13]

How do exploit markets interact with cyber norms?

Another cost incurred by offensive operations is reputation damage, such as that mediated by norms of responsible state behavior.^[6] Norms constrain what can be publicly advertised as brokers seek to avoid scandal. Market actors selling offensive cyber appear to have created outrage among journalists who focus on their use by repressive regimes.^[10,57] Broker's demonstrate their understanding of such reputation risk by establishing “due diligence and vetting process”^[15], although we have no further details on what exactly this entails.

Alternatively, one could imagine how brokers quoting a price for a given exploit could legitimize using such exploits, acting as private norm entrepreneurs in doing so.^[7,58] It could be that these brokers are normalizing the use of exploits for espionage, given that exploits specifically targeting communications (e.g., messaging and email) are among the most common and also have a higher average price. Looking forward, this motivates ongoing analysis of brokers' offerings to understand which systems it is “normal” to target with offensive cyber operation.♥

ACKNOWLEDGEMENTS

We would like to thank Rainer Bohme, Max Smeets, Simon Rock, Andrew Dwyer and Amy Ertan for their insightful comments and useful feedback. We are also grateful to the other participants for the fruitful discussion during the author workshop. This research was funded by the Air Force Office of Scientific Research.

APPENDIX

Price (\$)	1	2	3	4	5
period					234 ^{***}
					-25
android	249,135 ^{***}			-143,682 ^{***}	-97,025 ^{***}
	-63,895			-49,100	-35,647
bsd	-6,651			-279,139 ^{**}	-232,700 ^{**}
	-161,136			-111,202	-106,807
ios	351,662 ^{***}			-103,540 ^{**}	-54,922
	-63,657			-49,655	-35,924
linux	62,229			-191,314 ^{***}	-142,278 ^{***}
	-64,342			-45,433	-32,991
mac	40,448			-192,090 ^{***}	-149,388 ^{**}
	-78,372			-58,031	-49,898
windows	110,775 [*]			-122,970 ^{***}	-77,292 ^{**}
	-63,218			-45,033	-33,211
windows phone	67,613			-999,072 ^{***}	-981,444 ^{***}
	-151,481			-110,890	-111,026
BrowserTrue		-12,468		17,671	18,295
		-30,267		-22,740	-22,907
EmailTrue		2,950		52,266 [*]	53,673 [*]
		-35,382		-27,060	-27,361
MessengersTrue		426,659 ^{***}		365,398 ^{***}	364,179 ^{***}
		-34,564		-29,869	-30,189
Web.ServerTrue		-60,188 [*]		19,193	20,056
		-34,370		-28,114	-28,429
antivirusTrue		-112,511		-113,953 ^{**}	-110,883 ^{**}
		-74,923		-53,849	-54,361
Local.Privilege.EscalationTrue		192,633 ^{***}	92,681 ^{**}	96,774 ^{***}	
			-20,138	-20,328	-20,257
Mitigation.BypassTrue		197,569	94,563	97,809	
			-124,779	-99,817	-100,937
Remote.Code.ExecutionTrue		139,874 ^{***}	60,792 ^{***}	65,631 ^{***}	
			-21,465	-20,594	-20,737
Requires.Local.AccessTrue		-112,396 ^{**}	-76,307 ^{**}	-68,726 [*]	
			-43,641	-35,798	-36,147
Full.Chain.with.PersistenceTrue		1,015,542 ^{***}	1,153,635 ^{***}	1,139,252 ^{***}	
			-55,464	-51,477	-49,843
Constant	21,510	63,916	-168,446 ^{***}	65,440 [*]	-91,838 ^{***}
	-43,185	-41,380	-39,452	-33,971	-26,148
Observations	543	543	543	543	543
R2	0.241	0.325	0.494	0.689	0.679
Adjusted R2	0.223	0.311	0.483	0.675	0.668

Note: *p<0.1; **p<0.05; ***p<0.01

Table 3: Equivalent linear regressions with log-transformed price (\$) as the dependent variable.

NOTES

1. Florian J. Egloff and James Shires, The better angels of our digital nature? offensive cyber capabilities and state violence, *European Journal of International Security*, 1-20, 2021.
2. Jacqueline Eggenschwiler and Jantje Silomon, Challenges and opportunities in cyber weapon norm construction *Computer Fraud & Security*, 2018(12):11{18, 2018.
3. Andrew Dwyer, The NHS cyber-attack: A look at the complex environmental conditions of WannaCry, *RAD Magazine*, 44, 2018.
4. John Arquilla and David Ronfeldt, Cyberwar is coming! *Comparative Strategy*, 12(2):141-165, 1993.
5. Thomas Rid, Cyber war will not take place, *Journal of Strategic Studies*, 35(1):5-32, 2012.
6. Martha Finnemore and Duncan B Hollis, Constructing norms for global cybersecurity. *American Journal of International Law*, 110(3):425-479, 2016.
7. Louise Marie Hurel and Luisa Cruz Lobato, Unpacking cyber norms: private companies as norm entrepreneurs, *Journal of Cyber Policy*, 3(1):61-76, 2018.
8. Myriam Dunn Cavely, Cyber-security and private actors. In *Routledge handbook of private security studies*, 89-99. Routledge, 2015.
9. Jamie Collier, Cyber security assemblages: a framework for understanding the dynamic and contested nature of security provision, *Politics and Governance*, 6(2):13-21, 2018.
10. Nicole Perlrot, *This Is How They Tell Me the World Ends: The Cyber-weapons Arms Race*, chapter 10, Bloomsbury Publishing, 2021.
11. Maylin Fidler, Anarchy or regulation: Controlling the global trade in zeroday vulnerabilities, *PhD diss.*, Freeman Spogli Institute for International Studies, Stanford University, 2014.
12. Ross Anderson and Tyler Moore, The economics of information security, *science*, 314(5799):610-613, 2006.
13. Kiran Sridhar and Ming Ng, Hacking for good: Leveraging hacker one data to develop an economic model of bug bounties, *Journal of Cybersecurity*, 7(1):tyab 007, 2021.
14. Luca Allodi, Marco Corradin, and Fabio Massacci, Then and now: On the maturity of the cybercrime markets the lesson that black-hat marketers learned, *IEEE Transactions on Emerging Topics in Computing*, 4(1):35-46, 2015.
15. Zerodium, Frequently Asked Questions, www.zerodium.com/faq.html, 2021.
16. Charlie Miller, The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales, In *Workshop on the Economics of Information Security*, 2007.
17. Kiran Sridhar, Allen Householder, Jonathan M. Spring, and Daniel W. Woods, Cybersecurity information sharing: Analysing an email corpus of coordinated vulnerability disclosure, In *Workshop on the Economics of Information Security*, 2021.
18. Leonie Maria Tanczer, Irina Brass, and Madeline Carr, CSIRTs and global cybersecurity: How technical experts support science diplomacy, *Global Policy*, 9:60-66, 2018.
19. Rainer Bohme, A comparison of market approaches to software vulnerability disclosure, *International Conference on Emerging Trends in Information and Communication Security*, 298-311, Springer, 2006.
20. Lillian Ablon, Martin C Libicki, and Andrea A Golay, *Markets for cyber-crime tools and stolen data: Hackers' bazaar*, RAND Corporation, 2014.
21. Luca Allodi and Fabio Massacci, Comparing vulnerability severity and exploits using case-control studies, *ACM Transactions on Information and System Security (TISSEC)*, 17(1):1, 2014.
22. Mingyi Zhao, Jens Grossklags, and Peng Liu, An empirical study of web vulnerability discovery ecosystems, *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 1105-1117, ACM, 2015.
23. Sam Ransbotham, Sabyaschi Mitra, and Jon Ramsey, Are markets for vulnerabilities effective? *MIS Quarterly*, 43-64, 2012.
24. Donatello Luna, Luca Allodi, and Marco Cremonini, Productivity and patterns of activity in bug bounty programs: Analysis of hackerone and google vulnerability research, *Proceedings of the 14th International Conference on Availability, Reliability and Security*, 67, ACM, 2019.
25. Matthew Finifter, Devdatta Akhawe, and David Wagner, An empirical study of vulnerability rewards programs, *USENIX Security Symposium*, 2732-88, 2013.
26. Thomas Walshe and Andrew Simpson, An empirical study of bug bounty programs, *2020 IEEE 2nd International*.

NOTES

27. Chris Grier, Lucas Ballard, Juan Caballero, Neha Chachra, Christian J. Dietrich, Kirill Levchenko, Panayiotis Mavrommatis, Damon McCoy, Antonio Nappa, Andreas Pitsillidis, et al. Manufacturing compromise: the emergence of exploit-as-a-service, *Proceedings of the 2012 ACM conference on Computer and communications security*, 821-832, ACM, 2012.
28. Maddie Stone, The state of 0-day in-the-wild exploitation, USENIX Association, February 2021.
29. Serge Egelman, Cormac Herley, and Paul C. van Oorschot. Markets for zero-day exploits: Ethics and implications, *Proceedings of the 2013 New Security Paradigms Workshop*, NSPW '13, 41-46, New York, 2013, Association for Computing Machinery.
30. Joss Meakins, A zero-sum game: the zero-day market in 2018, *Journal of Cyber Policy*, 4(1):60-71, 2019.
31. Maor Shwartz, Selling 0-days to governments and offensive security companies, Blackhat, 2019.
32. Vlad Tsyrlkevich, Hacking team: a zero-day market case study, www.tsyrlkevich.net, 2015.
33. Crowdfense, Bug Bounty Program, www.crowdfense.com/bug-bounty-program.html, 2021.
34. Lillian Ablon and Andy Bogart, Zero Days, Thousands of Nights, RAND Corporation, Santa Monica, CA, 2017.
35. Zerodium, www.zerodium.com, 2022.
36. The Internet Archive, *The Wayback Machine*, www.web.archive.org, 2021.
37. Scott Sumner, Never reason from a price change, www.themoneyillusion.com/never-reason-from-a-price-change/, 2010.
38. @Zerodium. We will NOT be acquiring any new Apple iOS LPE, Safari RCE, or sandbox escapes for the next 2 to 3 months due to a high number of submissions related to these vectors. Prices for iOS one-click chains (e.g. via Safari) without persistence will likely drop in the near future. twitter.com/Zerodium/status/1260541578747064326, May 2020.
39. Harold Abelson, Ross Anderson, Steven M Bellovin, Josh Benaloh, Matt Blaze, Whitfield Diffie, John Gilmore, Matthew Green, Susan Landau, Peter G Neumann, et al. Keys under doormats: mandating insecurity by requiring government access to all data and communications. *Journal of Cybersecurity*, 1(1):69-79, 2015.
40. Dragos, Inc. ICS cybersecurity year in review 2020, 2021.
41. Irina Brass, Leonie Tanczer, Madeline Carr, and Jason Blackstock. Regulating iot: enabling or disabling the capacity of the internet of things? *Risk&Regulation*, 33:12-15, 2017.
42. Michael Poznansky and Evan Perkoski. Rethinking secrecy in cyberspace: The politics of voluntary attribution. *Journal of Global Security Studies*, 3(4):402-416, 2018.
43. Dakota S Rudesill. Cyber operations, legal secrecy, and civil-military relations. In *Reconsidering American Civil-Military Relations*, pages 245-262. Oxford University Press, 2020.
44. Len Scott and Peter Jackson. The study of intelligence in theory and practice. *Intelligence & National Security*, 19(2):139-169, 2004.
45. Richard Aldrich. 'grow your own': cold war intelligence and history supermarkets. *Intelligence and National Security*, 17(1):135-152, 2002.
46. Nigel West. Fiction, faction and intelligence. *Intelligence & National Security*, 19(2):275-289, 2004.
47. Florian J Egloff. Public attribution of cyber intrusions. *Journal of Cybersecurity*, 6(1):tyaa012, 2020.
48. Juan Andrés Guerrero-Saade. The ethics and perils of apt research: an unexpected transition into intelligence brokerage. In *Proceedings of the 25th Virus Bulletin International Conference*, 2015.
49. Bill Marczak, Adam Hulcoop, Etienne Maynier, Bahr Abdul Razzak, Masashi Crete-Nishihata, John Scott-Railton, and Ron Deibert. Missing link: Tibetan groups targeted with 1-click mobile exploits. Citizen Lab Research Report No. 123, 2019.
50. Bill Marczak, John Scott-Railton, Noura Al-Jizawi, Siena Anstis, and Ron Deibert. The great ipwn: Journalists hacked with suspected nso group imessage 'zero-click' exploit. Citizen Lab Research Report No. 135, 2020.
51. Brandon Valeriano and Ryan C Maness. The dynamics of cyber conflict between rival antagonists, 2001-11. *Journal of Peace Research*, 51(3):347-360, 2014.
52. Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. Disinformation warfare: Understanding state-sponsored trolls on twitter and their influence on the web. In *Proceedings of the 2019 World Wide Web Conference*, 218-226, 2019.
53. Darren L Linvill and Patrick L Warren. Troll factories: Manufacturing specialized disinformation on twitter. *Political Communication*, 37(4):447-467, 2020.

NOTES

54. Savvas Zannettou, Tristan Caulfield, William Setzer, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. Who let the trolls out? towards understanding state-sponsored trolls. In Proceedings of the 10th ACM Conference on Web Science, 353–362, 2019.
55. Marco Cremonini and Patrizia Martini. Evaluating information security investments from attackers perspective: the return-on-attack (ROA). In Workshop on the Economics of Information Security, 2005.
56. Ross Anderson. Security engineering: a guide to building dependable distributed systems. John Wiley & Sons, 2020.
57. Cory Doctorow. Cyber-mercenaries helped saudis hack an NYT reporter. <https://pluralistic.net/2021/10/24/breaking-the-news/#kingdom>, 2021.
58. Nancy Ayer Fairbank. The state of microsoft?: the role of corporations in international norm creation. *Journal of Cyber Policy*, 4(3):380–403, 2019.

