

# Booting the Booters: Evaluating the Effects of Police Interventions in the Market for Denial-of-Service Attacks

Ben Collier

Ben.Collier@cl.cam.ac.uk

Department of Computer Science & Technology,  
University of Cambridge, Cambridge, CB3 0FD, UK

Richard Clayton

Richard.Clayton@cl.cam.ac.uk

Department of Computer Science & Technology,  
University of Cambridge, Cambridge, CB3 0FD, UK

Daniel R. Thomas

Daniel.Thomas@cl.cam.ac.uk

Computer & Information Sciences,  
University of Strathclyde, Glasgow, G1 1XH, UK

Alice Hutchings

Alice.Hutchings@cl.cam.ac.uk

Department of Computer Science & Technology,  
University of Cambridge, Cambridge, CB3 0FD, UK

## ABSTRACT

Illegal *booter services* offer denial of service (DoS) attacks for a fee of a few tens of dollars a month. Internationally, police have implemented a range of different types of intervention aimed at those using and offering booter services, including arrests and website takedown. In order to measure the impact of these interventions we look at the usage reports that booters themselves provide and at measurements of reflected UDP DoS attacks, leveraging a five year measurement dataset that has been statistically demonstrated to have very high coverage. We analysed time series data (using a negative binomial regression model) to show that several interventions have had a statistically significant impact on the number of attacks. We show that, while there is no consistent effect of highly-publicised court cases, takedowns of individual booters precede significant, but short-lived, reductions in recorded attack numbers. However, more wide-ranging disruptions have much longer effects. The closure of HackForums' booter market reduced attacks for 13 weeks globally (and for longer in particular countries) and the FBI's coordinated operation in December 2018, which involved both takedowns and arrests, reduced attacks by a third for at least 10 weeks and resulted in lasting change to the structure of the booter market.

## CCS CONCEPTS

• **Networks** → Denial-of-service attacks; • **Social and professional topics** → *Computer crime*; • **Security and privacy** → Social aspects of security and privacy; • **Mathematics of computing** → Time series analysis.

---

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).

*IMC'19, October 21–23, 2019, Amsterdam, Netherlands*

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6948-0/19/10...\$15.00

<https://doi.org/10.1145/3355369.3355592>

## KEYWORDS

denial of service attacks; DDoS; UDP-reflection; booter; stresser; cybercrime; police interventions

## ACM Reference Format:

Ben Collier, Daniel R. Thomas, Richard Clayton, and Alice Hutchings. 2019. Booting the Booters: Evaluating the Effects of Police Interventions in the Market for Denial-of-Service Attacks. In *Internet Measurement Conference (IMC '19), October 21–23, 2019, Amsterdam, Netherlands*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3355369.3355592>

## 1 INTRODUCTION

'Booter', or 'stresser', services provide Denial of Service (DoS) attacks as-a-service. DoS attacks generate large amounts of traffic which overwhelm end-users or web services, taking them offline or making legitimate access impossible [26]. Booter operators advertise customer-facing websites, where individuals can set up accounts and order attacks [54], with payments accepted using digital services such as PayPal or through transfers of cryptocurrency [22, 28]. A range of different packages and membership options are available, with \$10 to \$20 being typical for a month's worth of DoS attacks of sufficient size to disrupt an end-user connection or a website which does not have specialist DoS protection.

Operating a booter service or purchasing a DoS attack is illegal in most jurisdictions. However, when asked during research undertaken in 2014, booter operators were unconcerned about the possibility of law enforcement taking action against them [22]. Nevertheless, a number of police actions have taken place in recent years, and this research aims to measure whether there has been any impact on booter provision and booter usage.

There has been much research into the effects of police action on offline illicit markets and services, however there is little understanding of what 'best practice' looks like for cybercrime, which tends to be more geographically dispersed and organised around online communities and markets which are particularly resilient [1, 43]. In practice, law enforcement interventions in online illicit markets are not straightforward. Displacement is a problem for both online and offline

markets [23, 44, 63] and participants may seek new vendors or locations following the shutdown of established markets [58, 59]. In this paper, we evaluate the effects of a range of different kinds of interventions – high profile court cases and sentencing of booter providers, arrests, takedowns of booter websites, and messaging campaigns targeted at users.

We measure the impact of interventions using two datasets. The first is a dataset of victims of reflected UDP amplification attacks, a technique widely used by booters, which covers a five-year period from 2014. This dataset is known to have good coverage of many widely abused protocols [61]. The second dataset is of self-reported DoS attack numbers collected from booter websites. It covers 75% or more of active booters over the 18-month period in which the majority of interventions have occurred.

We start (§2) by considering major police interventions against booter providers and users and then discuss (§3) the attack datasets we use. We present (§4) our negative binomial regression model and show several linkages between interventions and deviations from seasonal trends. We do further analysis (§4.1) of the data at a country level and show the impact of a UK-specific advertising campaign. After reviewing related work (§5), we discuss what makes for an effective police intervention (§6) and finally draw conclusions (§7). Ethical considerations are outlined in the Appendix.

## 2 THE INTERVENTIONS

Law enforcement can respond to cybercrime not only by arresting criminals, but by seeking to prevent people becoming involved, by disrupting harmful criminal activity, and protecting those at risk of becoming victims. For example, the UK Home Office sets out their ‘Pursue, Prevent, and Protect’ approach in their *Serious and Organised Crime Strategy* [21].

Here, we outline some of the main interventions against booter services, their providers, and their users in the 2014–2019 period for which we have DoS attack data. The events come from an analysis of Brian Krebs’ popular blog [30–38], but were also widely covered in the press. We have added one further event (§2.7) which was not widely reported but was very visible within booter communities.

### 2.1 LizardStresser

In early 2015, the backend database of users of the *LizardStresser* booter service, which had only been in operation for a few weeks, was leaked. On 28 August 2015 six UK individuals who had purchased attacks were arrested in ‘Operation Vivarium’. Approximately 50 others who had registered with the site received a ‘cease and desist’ home visit from UK police. Although merely registering at a booter is not an offence in the UK the individuals were told that DoS attacks are “illegal, can prevent individuals from accessing vital online services, and can cause significant financial and reputational damage to businesses” [45].

Only one related court case has been reported. On 22 December 2015 a 17-year-old pleaded guilty to a DoS attack (and another offence), receiving a 12-month sentence in a

young offenders’ institution. He also had to pay over £1 000 compensation to the DoS victim [46].

On 6 October 2016 two 19-year-olds were arrested in the US and the Netherlands for allegedly running the *LizardStresser* booter service [32]. On 27 March 2018 the American received a three month prison sentence and was ordered to pay \$350 000 in restitution after pleading guilty and cooperating with authorities [49].

### 2.2 Netspoof, etc.

On 8 April 2016, a 20-year-old man was sentenced in the UK after pleading guilty to 6 charges under the Computer Misuse Act and 4 under the Serious Crime Act. He had operated four booter services, including one called *Netspoof*, and received a 2-year youth detention sentence, suspended for 18 months, 100 hours unpaid work and £800 costs [36].

In December 2016, 12 people were arrested for purchasing DoS attacks on *Netspoof*, 30 cease and desist notices were issued, computers were seized from 11 of the suspects, one protective visit was made, and two cautions were issued [47]. This was part of a coordinated international action against users of booter services which took place between 5 and 9 December 2016. It involved Europol as well as law enforcement authorities from Australia, Belgium, France, Hungary, Lithuania, the Netherlands, Norway, Portugal, Romania, Spain, Sweden, the UK, and the US. Overall, there were 34 arrests, and 101 suspects were interviewed and cautioned [16].

### 2.3 HackForums & Mirai

On 28 October 2016, the “Server Stress Testing” (SST) section of *HackForums*, a large and long-running English-language underground forum, was removed, and advertisements for booter services were banned [33]. The proximate cause of this shuttering was a telephone conversation between *HackForums*’ owner and an FBI agent (personal communication, 2018), but it was the culmination of a series of events.

On 8 September 2016, Brian Krebs posted an article [34] about a leaked backend database from the *vDOS* booter, which contained entries for tens of thousands of users and more than 150 000 attacks. He claimed the operators made more than \$600 000 over two years, and that the service was operated by two men in Israel, with support from others in the US. Hours later, the two Israeli men were arrested [30]. This was followed a DoS attack against Krebs’ website that was reported as being the largest ever recorded at the time [35].

On 30 September the source code for *Mirai*, the botnet responsible for this DoS attack, was released on *HackForums*, apparently because the author feared arrest and wanted an excuse for having a copy of the code. On 21 October, a *Mirai* botnet launched an even larger attack against the Internet infrastructure firm Dyn causing outages for many popular sites, including Twitter, PayPal, Reddit, and Netflix [31].

The *HackForums* ban left a gap in the market for advertising and discussing booters and a few months later *stresserforums.net* was set up. There was some displacement to this new forum, and it had almost 800 members and over 7 000

posts when it closed in April 2018, around the time of the Webstresser takedown (detailed below). A number of forums with the same name have subsequently failed to gain traction.

The Israelis were charged, but it appears the matter is still proceeding. In the UK, a 19-year-old pleaded guilty to a number of offences relating to vDOS. On 19 December 2017, he was sentenced to 16 months in a young offenders institution, suspended for two years, and was ordered to complete 20 days of rehabilitation activity [56].

On 18 September 2018, three men aged 21 to 22 years were sentenced for authoring and using the Mirai botnet. They each received five years probation, 2 500 hours of community service, and were ordered to pay \$127 000 restitution. This light sentence (for the US) was due to the cooperation of the defendants in helping the investigation, identifying other suspects, and assisting industry to combat incidents [37].

On 26 October 2018, one of the men received a further sentence for using Mirai for DoS attacks against Rutgers University. He was sentenced to 2 500 hours of community service, six months home confinement, and was ordered to pay \$8.6 million in restitution [38].

## 2.4 Titaniumstresser

In the UK a 19-year-old received a 24-month sentence on 25 April 2017 (reduced on appeal to 21 months) [2] for operating the *Titaniumstresser* booter and personally committing 594 DoS attacks against 181 targets. The service was reportedly used for 1.7M attacks. His case was also widely reported on 22 November 2016 when he had pleaded guilty to two offences under the Computer Misuse Act and one money laundering offence. At a confiscation hearing in March 2018, he was ordered to repay £69 629 in compensation or face an additional two years in prison [20].

## 2.5 Webstresser

On 24 April 2018 the domain for the *Webstresser* booter was seized and its alleged administrators were arrested in the UK, Croatia, Canada, and Serbia [17]. Europol subsequently reported that 250 UK users of Webstresser were to receive police visits or warnings [18] and that users in the Netherlands, Italy, Spain, Croatia, Australia, Canada, and Hong Kong would also be “targeted”.

## 2.6 FBI action: Xmas2018

On 19 December 2018 the FBI announced that they had arrested three booter operators, running *DownThem*, *Amprnode*, and *Quantum Stresser* and that they had seized 15 domain names for booter websites [13]. Dataset analysis shows this immediately took seven booter services offline (the other domain names were either duplicates for the same service or the booter was not operational at that time). The timing, just before Christmas, was intended to disrupt a pattern of increased DoS activity over the holiday period [8].

## 2.7 NCA Google search advert warnings

The HackForums ban on adverts for booter services meant providers had to find other ways of attracting customers. One booter service purchased Google search adverts (displayed when people searched for relevant keywords) [9]. However, advertising can also be used by the police.

From late December 2017 to June 2018 the UK National Crime Agency (NCA) bought search adverts from Google which warned young users on UK IP addresses of the illegality of DoS attacks when they searched for booter-related terms. As with the in-person warnings delivered in Operation Vivarium, the aim was to divert people away from cybercrime, by informing them of potential legal consequences.

## 3 DATASETS

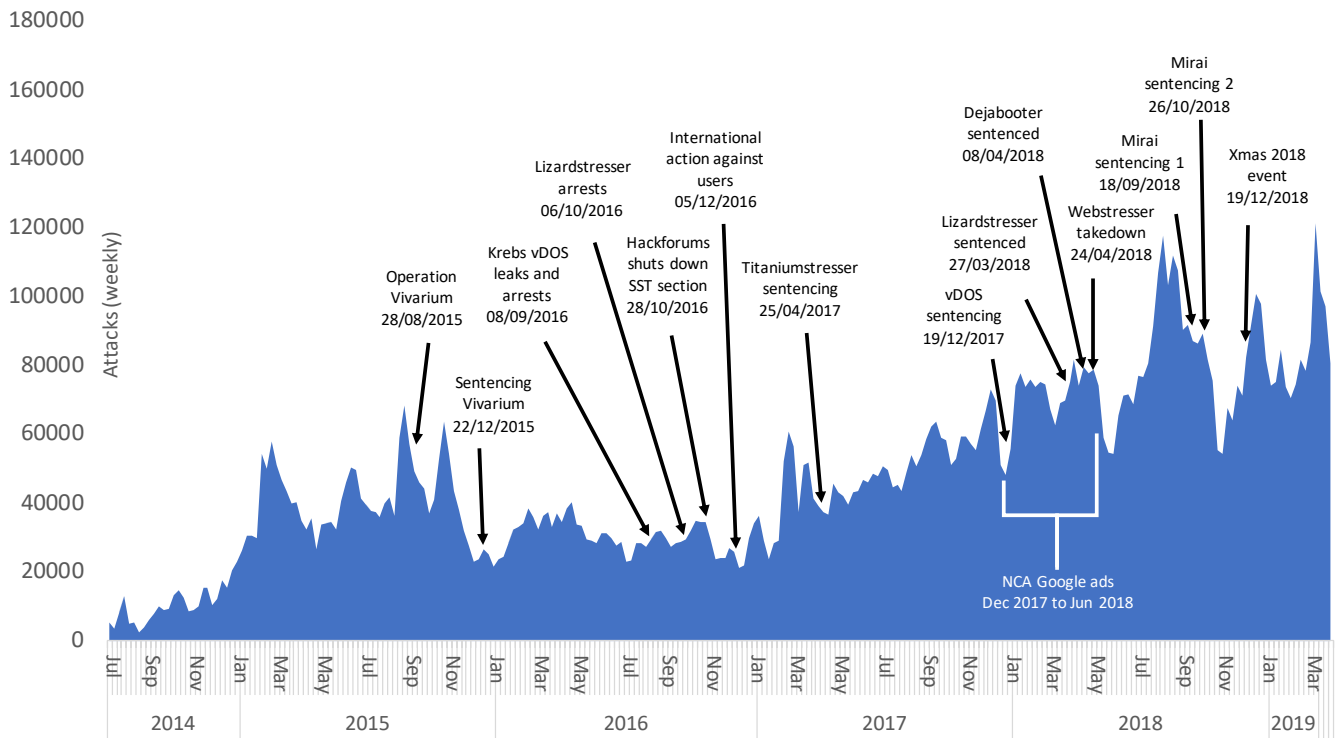
We use two datasets, provided to us by the Cambridge Cybercrime Centre, to measure whether the number of DoS attacks by booter services is affected by police interventions.

The first dataset is of reflected amplified UDP DoS attacks in which a small incoming UDP packet generates a much larger response – and if the source IP of the original packet is spoofed then substantial amounts of traffic can be directed at a victim. The dataset is of victim IPs seen by a large number of honeypot machines roped into attacks using the protocols QOTD, CHARGEN, time, DNS, PORTMAP, NTP, LDAP, MSSQL Monitor, MDNS, and SSDP. Full details of the dataset are provided by Thomas et al. along with a statistical analysis to show high levels of coverage for many of these UDP protocols [61]. For analysis we group flows of packets to the same victim IP or prefix for the same protocol until there is a gap of at least 15 minutes with no packets being received by any sensor. We then check to see if any sensor received more than 5 packets. If so then we deem it an attack, if not then we classify the event as a scan.

While this dataset counts traffic volume we cannot reliably translate this into the traffic volume which victims would experience and so we focus on the number of attacks rather than their size. The problem is that we do not know how many real reflectors booters are using and so we are unable to scale our observed volumes appropriately.

There are limitations to this dataset because not all reflected attacks are associated with booters and booters can perform other types of attack. However, we believe it is broadly representative of booter activity. The attack logs of three prominent booters<sup>1</sup> from 2014, 2016, and 2017/2018 all

<sup>1</sup>Of the 285 414 attacks for booter.io recorded between 2014-03-24 and 2014-09-07, 261 968 (91%) are for method names indicating UDP reflection (UDP, CHARGEN, UDPLAG). Of the 169 845 attacks recorded between 2016-05-01 and 2016-07-23 for vDOS 123 751 (72%) are for method names indicating UDP reflection (DNS, NTP, SNMP, PORTMAP) [61]. Of the 412 059 attacks recorded for Webstresser (§2.5) between 2017-10-18 and 2018-02-26, 339 181 (82%) were probably UDP reflection attacks based on their name. The UDP honeypot dataset contains 97% of the LDAP, NTP and PORTMAP attacks but only 111 306 (33%) of the attacks overall, mostly because of only 9% coverage for ‘SUDP’. It is hard to know what method names mean without either performing self attacks [55] or observing the attacks with sensors. Coverage for named methods that might be UDP based is: ‘LDAP’ 41 097/42 136 (98%), ‘NTP’ 32 094/33 171



**Figure 1: Timeline of intervention events and the number of reflected UDP DoS attacks per week.**

show over 70% of attacks were likely UDP reflection attacks. Krupp et al. were able to attribute 26% of DNS and 13% of NTP attacks to specific booters which they had bought attacks from [40]. Noroozian et al. found that over 48% of UDP reflection attacks and 62% of victims are on IP addresses in access networks [48], Sharma found that 89% of US, 98% of UK, 71% of FR, and 89% of DE victims were home users [57]. Sharma also found that over 50% of attacks were less than 5 minutes. This is a pattern of behaviour which we would expect from booters. Nevertheless there will be other sources of UDP reflection attacks included within this dataset.

The second dataset comes from the booters themselves which, presumably to assist in marketing, report a running total of the attacks they have performed. This data has been collected on a weekly basis since November 2017.

We know that booters use SQL databases to hold details of users and attacks because large numbers of these databases have been leaked. The source of many booters has also been leaked and we invariably find PHP code such as: `$TotalUsers = $odb->query("SELECT COUNT(*) FROM 'users'")->fetchColumn(o); $TotalAttacks = $odb->query("SELECT COUNT(*) FROM 'logs'")->fetchColumn(o);` i.e.

(97%), 'PORTMAP' 11 497/11 858 (97%), 'SUDP' 21 922/235 905 (9%), 'UDP KILL' 3 058/10 507 (29%), and 'UDPRAND' 1 638/5 604 (29%). Of the 72 878 attacks whose names suggest they are not entirely UDP based (TS3KILL, TS3, VOX, FRAG, ZAP, ICMP, DOMINATE, ACK, VSE, SYN, COD, RST), we observed 21 598 (30%) with coverage rates between 20% and 50%.

fetching counts directly from the SQL database, followed by display code such as `<li>Users: <?php echo $TotalUsers; ?> Attacks: <?php echo $TotalAttacks; ?></li>`.

A handful of booters have clearly inflated their counts (one counted from 150 000 rather than zero – trivially done in the PHP code) and some wipe their databases (and hence zero their counts) from time to time. One booter reported values which were regularly multiples of 1000 and we exclude it. However, we can see no other obvious artificial patterns within the dataset, but for good measure we performed some statistical tests on the weekly totals to determine if they might have been algorithmically generated.

Count data tends to be heteroskedastic i.e. as numbers go up the variance in the series will be found to increase as well. Many of the smaller or shorter booter series show too high a degree of variance or nonlinearity to perform meaningful tests for this effect. However, we conducted linear regression analysis and performed White's heteroskedasticity test on those booters where this was valid. We also performed skewness kurtosis tests for normality on these series, as real-world data are often normally distributed, and faking with random data would produce uniform distributions. Our analysis indicated that the top ten most active booters' attack series were normally distributed or heteroskedastic (with most being both) at 95% confidence. We further checked if simple multipliers were being applied to otherwise genuine data, but no sequences of any length had values which were all divisible by any prime less than 50.

We conclude that if booters were generating fake data to feed their live attack counters, they would have to have considerable statistical acumen to reproduce the distributions we observe. Of course we cannot completely rule out forgery by booter operators with a deep knowledge of statistics, but this does not seem especially likely. Furthermore, the booter self-reported dataset shows moderate correlation with our own reflected attack dataset (a correlation coefficient of 0.47) and, most importantly, shows large drops in attack numbers in the same places and for the same durations as the significant drops we observe in our attack time series, which we believe to correspond to law enforcement interventions.

This second dataset can also be used to determine when booters first appear, how many booters are taken down each week, and how many subsequently reappear. Unfortunately, the ‘birth’ data is irredeemably biased by the data collection process in that new booters were only searched for at somewhat irregular intervals. However, the ‘death’ and ‘resurrection’ data can be usefully analysed to determine if interventions affect users (they choose to do fewer attacks) or booter operators (they choose to enter or leave the market).

## 4 MODELLING THE DATA

It’s extremely difficult to measure directly the effect on crime of law enforcement interventions. Empirical associations, causal effects, and the presence of extraneous variables are all hard to quantify, and thus mechanisms are hard to demonstrate through ‘true experiments’. In a forthcoming paper, we attempt to trace some of these mechanisms empirically through mixed-methods qualitative work, however here we focus on an in-depth quantitative approach. Where practical or ethical issues make it impossible to carry out classic experiments with treatment and control groups and randomisation (as in this case), it is well-established within criminology and the social sciences that quasi-experimental designs are an appropriate way of making tentative claims about the effects of particular interventions [6][52][3].

Where large-scale interventions are attempted which affect entire populations, establishing a suitable control group can be impossible. We adopt what Cook and Campbell [10] classify as a time series design, now often referred to as an interrupted time series approach. This is appropriate where data takes the form of a time series of observations, with interventions occurring at specific points in time, assumed to have an immediate effect, with a clear pre-intervention functional form, a suitable number of pre-intervention observations, and a reasonable assumption that no unaccounted-for variable is responsible for the change in the time series. Denial of service attacks constitute event count data, which often have skewed outcomes in practice, and our time series is indeed non-normalised in distribution. Therefore, a maximum likelihood estimation approach, rather than an ARIMA approach (which relies on normally-distributed data), is indicated. We use a negative binomial rather than poisson regression model, as the events (denial of service attacks) are not independent, rather there is a simple trend to the data [50][4][11]. We

restricted our modelling to the period June 2016 to April 2019 as there is a clear and fairly constant linear trend over this period. Weekly totals were used as daily attack counts showed a high degree of volatility.

Negative binomial regression is an established technique for modelling count data, and can account for seasonal patterns and non-stochastic slope components [19]. It is well-suited to intervention analysis, and has been used to measure the effects of interventions on criminal offending [5, 60].

Our aim was to analyse the effects of different interventions on the booter market, once seasonal variation and the underlying trend of the data were accounted for, fitting for optimum log-pseudolikelihood. Thus for all periods in the time series which drop significantly below the modelled series, we added dummy ‘intervention’ variables to model the effect sizes of these disruptions. We found five such interventions that were statistically significant and one of the key conclusions of this paper is they correspond closely to events discussed in §2 above.

The model parameters are displayed in Table 1 and Figure 2 shows the correspondence between the model and measured attacks. We model seasonality over twelve one-month periods, for which we need eleven seasonal variables in the model. We included a component in the model to account for the changing date of Easter in the seasonal analysis, as the patterns of booting are strongly linked to school holidays. We found no evidence of multicollinearity in the regression components used. Statistical data about the interventions is shown in Table 2. We not only present the overall impact (in the final column) but also whether these interventions are significant when we apply the overall model solely to the attacks against particular countries.

The intervention with the biggest impact (at a 95% confidence level) was the FBI’s Xmas2018 intervention which lasted for 10 weeks, during which there was a reduction of between 37% and 27% in overall recorded attacks. However, in some countries the effect lasted for only three weeks, and for France the impact was not statistically significant. The shutdown of HackForums’ SST section (§2.3) was also long-lived with the market being suppressed for 13 weeks and for longer in some countries.

The other interventions can be seen to result in smaller but still significant (in most countries) drops in attack numbers. The reporting of high-profile court cases and sentencing corresponds with short, immediate drops in attack numbers. For takedowns, the effect is delayed, with the Webstresser takedown (§2.5) taking effect after a fortnight and lasting 3 weeks. This may be because the totals are distorted by attacks directed at the Netherlands which went up by 146% (i.e. more than doubled), presumably caused by reprisal attacks against the Dutch police who had spearheaded the operation.

### 4.1 Analysing by country

Having seen how the various interventions affected countries differently we now take a step back and ask to what extent countries have seen similar patterns of growth in attacks.

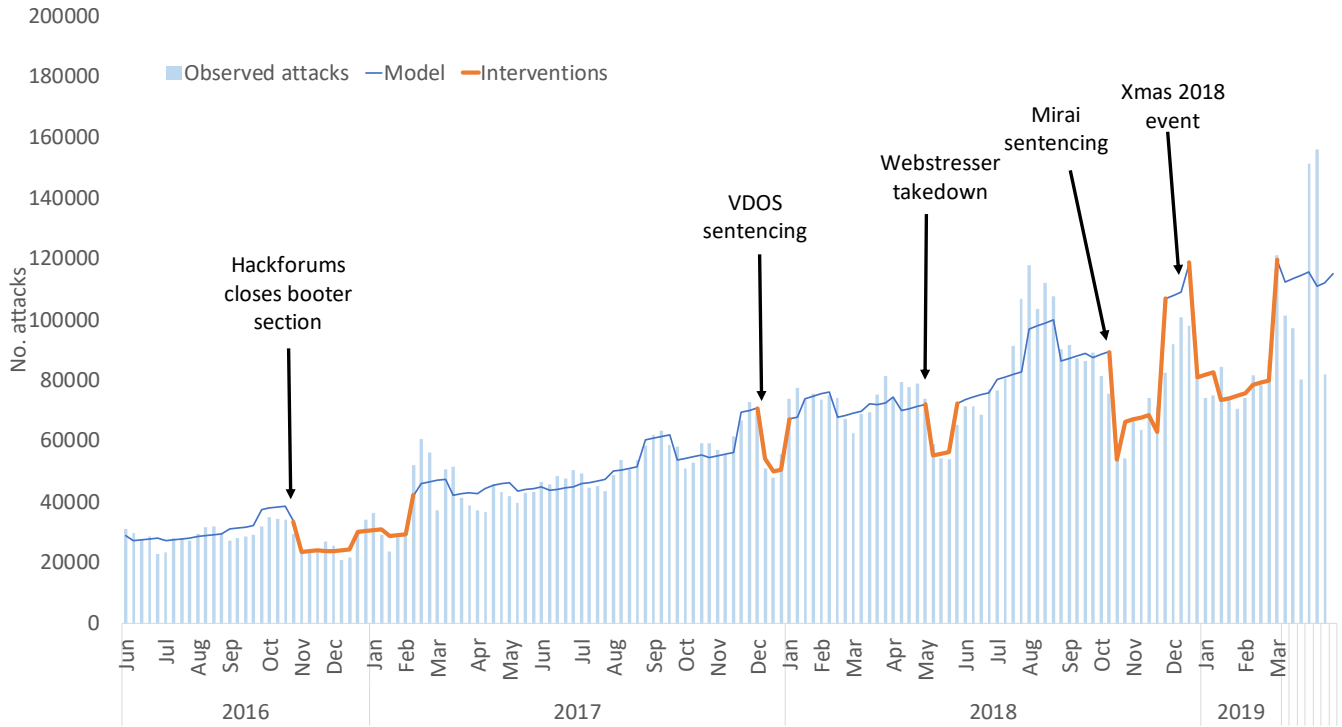


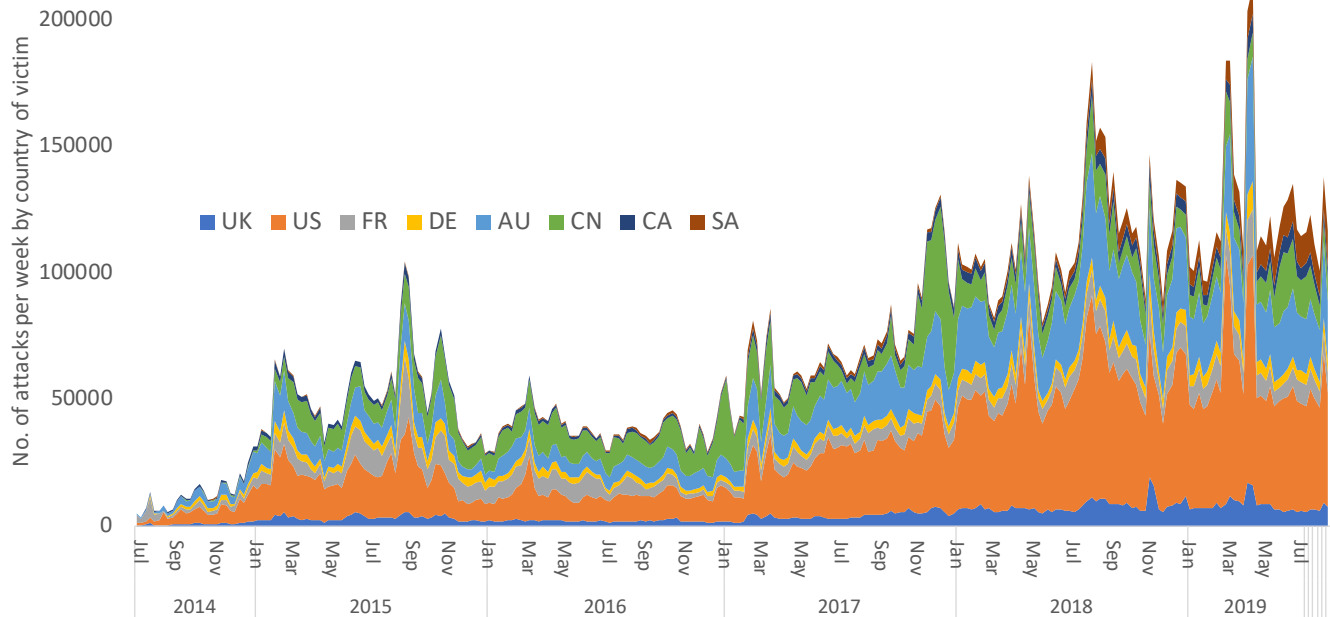
Figure 2: Total attack numbers over time (light blue bars) with negative binomial model (dark blue line) overlaid. Labels indicate the statistically significant interventions (modelled over periods shown by the dark orange line).

	Date	Coef.	Std.error	z	P >  z	95% CI	
						Lower	Upper
Xmas2018	19/12/2018	-0.393	0.039	-10.05	0.000**	-0.469	-0.316
Webstresser	24/04/2018	-0.238	0.0574	-4.15	0.000**	-0.351	-0.126
Mirai sentencing and arrests	26/10/2018	-0.516	0.049	-10.46	0.000**	-0.613	-0.420
HackForums SST forum closed	28/10/2016	-0.360	0.039	-9.16	0.000**	-0.437	-0.283
vDOS sentencing	19/12/2017	-0.275	0.057	-4.83	0.000**	-0.387	-0.164
Easter		-0.016	0.094	-0.17	0.864	-0.200	0.168
seasonal_2		0.076	0.066	1.15	0.25	-0.053	0.205
seasonal_3		-0.051	0.060	-0.86	0.390	-0.168	0.066
seasonal_4		-0.025	0.057	-0.44	0.660	-0.137	0.087
seasonal_5		-0.098	0.062	-1.59	0.110	-0.220	0.023
seasonal_6		-0.134	0.069	-1.95	0.050*	-0.269	0.001
seasonal_7		-0.125	0.054	-2.32	0.020*	-0.230	-0.019
seasonal_8		-0.078	0.060	-1.3	0.190	-0.196	0.040
seasonal_9		0.069	0.058	1.19	0.240	-0.045	0.184
seasonal_10		-0.086	0.048	-1.77	0.080	-0.181	0.009
seasonal_11		-0.111	0.051	-2.16	0.030*	-0.211	-0.010
seasonal_12		0.091	0.047	1.93	0.050	-0.001	0.182
time		0.010	0.000	27.04	0.000**	0.009	0.011
_cons		10.289	0.060	170.88	0.000**	10.171	10.407

Table 1: Negative binomial regression model showing model composition, including key interventions, seasonal components, first order trend, and constant with significance and effect size. Asterisks indicate if inclusion of an intervention made a significant (\*) or strongly significant (\*\*) contribution to the model. The seasonal variables model the month-by-month seasonality of the data. We also included a separate component for Easter as school holidays are linked to rises in attacks and the date of Easter is not fixed.

Intervention		UK	US	RU	FR	DE	PL	NL	Overall
Xmas2018 Intervention 19/12/2018	Mean	-27%	-49%	-33%	-1%	-28%	-23%	-16%	-32%
	L95/U95	-43/-28%	-55/-42%	-43/-22%	-13/11%	-36/-20%	-37/-5%	-27/-3%	-37/-27%
	Duration	9 weeks	9 weeks	9 weeks	N/A	8 weeks	3 weeks	8 weeks	10 weeks
	Signif.	0.000**	0.000**	0.000**	0.828	0.000**	0.014*	0.018*	0.000**
Mirai sentencing and other actions 24/10/2018	Mean	-27%	-31%	-5%	-9%	-32%	-47%	-19%	-40%
	L95/U95	-42/-9%	-41/-20%	-16/7%	-31/21%	-40/-23%	-56/-36%	-35/0%	-46/-34%
	Duration	2 weeks	7 weeks	2 weeks	N/A	6 weeks	2 weeks	6 weeks	8 weeks
	Signif.	0.006**	0.000**	0.41	0.533	0.000**	0.000**	0.053	0.000**
Webstresser takedown 24/04/2018	Mean	-10%	-24%	-16%	-22%	-29%	-29%	146%	-21%
	L95/U95	-21%/3%	-40/-4%	-33/6%	-35/-7%	-36/-22%	-42/-14%	94/211%	-30/-12%
	Duration	N/A	4 weeks	2 weeks	4 weeks	9 weeks	6 weeks	4 weeks	3 weeks
	Signif.	0.120	0.022*	0.14	0.006*	0.000**	0.001**	0.000**	0.000**
vDOS sentencing 16/12/2017	Mean	-20%	-4%	-37%	-30%	-4%	16%	-24%	-24%
	L95/U95	-33/-5%	-18/12%	-47/-24%	-37/-23%	-17/10%	-17/62%	-33/-13%	-32/-25%
	Duration	3 weeks	3 weeks	2 weeks	2 weeks	N/A	N/A	3 weeks	3 weeks
	Signif.	0.011*	0.563	0.000**	0.000**	0.532	0.373	0.000*	0.000**
HackForums 28/10/2016	Mean	-48%	-30%	-13%	-52%	-32%	2%	-35%	-30%
	L95/U95	-53/-42%	-37/-21%	-23/-3%	-59/-43%	-41/-23%	-19/28%	-42/-27%	-33/-25%
	Duration	15 weeks	7 weeks	14 weeks	15 weeks	7 weeks	N/A	15 weeks	13 weeks
	Signif.	0.000**	0.000**	0.02*	0.000**	0.000*	0.86	0.000*	0.000**

**Table 2:** Estimated effect sizes of statistically significant (at the global scale) interventions by country, showing the effects of each intervention component in separate negative binomial models of attack numbers over time in each country. Effects in red cells are not significant and the green cells are a significant increase rather than decrease. Asterisks indicate inclusion of intervention in the model made a significant (\*) or strongly significant (\*\*) contribution to the model. Countries were chosen by prominence in number of attacks, or factors which made them of interest (such as NL retaliation for Webstresser takedown)



**Figure 3:** Stacked area graph showing total attack numbers split by country of victim – top 8 countries by number of attacks.



	Feb-15	Feb-16	Feb-17	Feb-18	Feb-19
US	45%	25%	31%	45%	47%
FR	11%	19%	8%	5%	10%
DE	9%	6%	5%	5%	6%
CN	8%	16%	55%	12%	7%
UK	7%	6%	4%	7%	8%
PL	3%	3%	1%	3%	9%
RU	3%	3%	3%	2%	2%
NL	3%	3%	2%	2%	3%
Total	88%	81%	108%	82%	92%

**Table 3: Share of attacks by country of victim over time (includes double-counting when attacks are attributed to more than one country).**

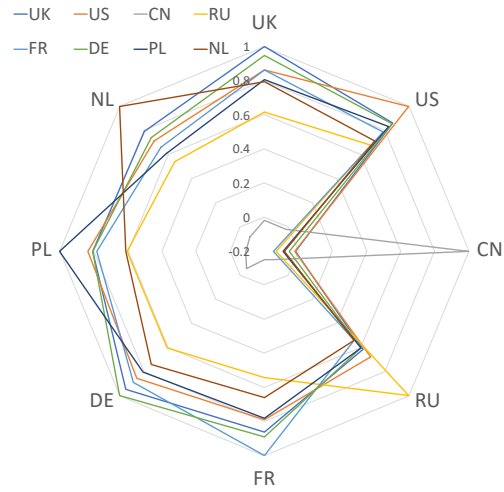
Although DoS attacks can be directed anywhere and be for any purpose, prior studies have shown that most attacks are on end-users (assumed to be games players) and on gaming related websites [7, 48]. Sharma showed that timezone determines when peak attacks occur [57] and we take the view that there will be a significant correlation between the country of the attacker and the country of the victim.

Table 3 sets out the eight countries which receive the highest number of DoS attacks and their percentage share over the past five years. The US now accounts for the largest number of UDP reflection attacks which we observe – 47%. Note that this table, and the time series in Figure 3 include some double-counting as an artefact of how attacks are conservatively assigned to countries.

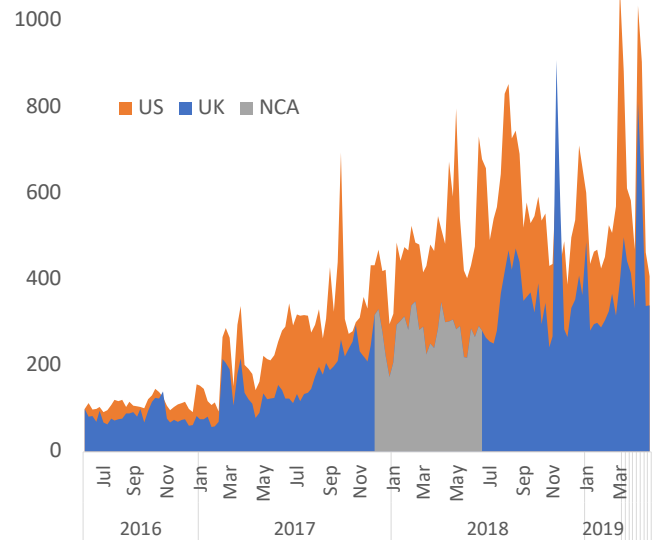
There is strong correlation between the attack time series for the UK, US, France, Germany and Poland (see Figure 4). These countries all show a flat series until the beginning of 2017, then steady growth across 2017 and 2018 with a strong seasonal pattern. The Netherlands is fairly similar, with a slightly lower degree of correlation. Russia is lower still, with less growth over time and smaller effects from interventions (though still showing a reasonable degree of correlation with the other series). China stands apart, showing no correlation to the other nations or impact from interventions, with a largely flat pattern over time and this lack of similarity led us to exclude it from the analysis we presented in Table 2.

In the US, France, Germany, Netherlands, Russia, and Poland, we observe a continuing upward trend from the beginning of 2017 up to the Webstresser takedown in April 2018. In the UK, however, this upward trend flattened off entirely from December 2017 until June 2018. This flat trend continues until August, whereupon there was a large spike in attacks and the series begins to grow again.

In Table 2, we present a comparison of the effects of the different interventions we observe as significant at the global level in individual countries. We selected seven nations which had both large numbers of booter attacks and whose booter markets we believed might credibly be expected to experience disruption as a result of these interventions. We ran negative binomial models (as described for the overall attack series) for the series of attacks over time for each of



**Figure 4: Correlation between numbers of attacks over time between countries, where 1 indicates complete correlation, 0 indicates no correlation, and -1 indicates negative correlation.**



**Figure 5: US and UK attack counts comparison. Non-stacked graph with totals scaled so both start at 100 in June 2016, with 200 representing a doubling. The NCA advertising intervention period which affects the UK data is highlighted in grey.**

these nations, enabling us to judge whether significant effects could be observed in different nations, and to compare the size of these effects. Where the per-country data showed large spikes localised around particular events (such as retaliation against the Dutch police for the Webstresser arrests), these were included as a component in the model. For reasons



of space, we do not present the details of the individual per-country model parameters. The variation in effect sizes between nations provide important evidence for a causal link between interventions and the drops in attack numbers we observe – for example, the Webstresser arrests show strong negative effects in the US and Europe, but not in Russia (as might be expected given their approach to extradition), and correspond to a statistically significant increase in attack numbers in the Netherlands (which we believe to be due to reprisals against the Dutch police who led the action). France is not significantly affected by a number of the interventions, possibly accounted for by the existence of a large Francophone booter market (where other European nations tend to use English-language services) which may insulate it from news reporting of major sentencing decisions and takedowns.

Figure 5 shows the figures for the US and the UK and highlights the period during which the NCA were purchasing search adverts (§2.7). As can be seen there is a strong case to be made that the adverts, which were only displayed within the UK, have caused a reduction in the number of attacks on UK IP addresses. For comparison, the UK and US linear trends from the period Jan 2017 until Dec 2017 had slopes of 3.2 and 5.3 respectively, while for the period of the NCA intervention, the US has a linear trend with a slope of 6.8, while the UK trend has reduced to a nearly-flat slope of -0.1. We believe that the adverts have led to a clear and lasting reduction in the number of attacks. The US and UK have a shared language and culture around booting with participants often sharing similar online communities and so correlation between growth rates would be expected. Attacks on the UK are considerably lower in number (around 16% as many per day on average) than attacks on the US, so a suppression in UK growth due to the NCA intervention would be unlikely to impact US growth even if UK targets are hosted in the US.

#### 4.2 Analysing by UDP protocol

Breaking attacks down by the UDP protocol used shows a number of underlying patterns (Figure 6). Protocols appear to go in and out of vogue at different times and there are short term spikes. From analysis of booter attack logs we know that booters experiment with switching to different protocols, or perhaps choose not to reflect packets off the honeypots providing our dataset [61], which may explain these spikes. The steady rise in attack numbers from the beginning of 2017 to the end of 2018 appears to be largely driven by an increase in attacks using the LDAP protocol which is the only protocol with consistent growth over time. It has a large amplification factor which has driven its popularity, but there are not many real LDAP reflectors and so the honeypots are likely to be used and continue to be used, so the data will be very representative of overall traffic.

Many of the drops in attacks seen after interventions are caused by drops in attacks for a particular protocol. This is most likely due to the protocol only being used by a particular booter. For the Webstresser takedown, we observe a small

drop in LDAP and a large drop in DNS. Some protocols are perceived by users to be ‘best’ to use at that time, so the drop seen following the shutdown of the HackForums SST section was largely in the CHARGEN and NTP protocols, whereas for the Xmas2018 intervention, the drop appears to largely occur in the LDAP protocol, and to a lesser extent, DNS. The recovery following this latter intervention sees increases in DNS, LDAP and NTP attacks.

While in other countries the rise in LDAP largely drives the overall upward trend across 2017 and 2018, in China this is not the case, with LDAP largely replacing NTP attacks, and the overall number of attacks remaining static. The rise in LDAP takes place six months later in China than in the rest of the world – near the end of 2017, rather than near the beginning. Attacks against China use a much smaller range of protocols than against the US, largely focusing on NTP and SSDP, with LDAP increasingly prominent since the start of 2018. The US, conversely, additionally sees substantial use of DNS and PORTMAP (and, historically, CHARGEN). We hypothesise that this is because the ‘Great Firewall of China’ blocks DNS traffic. Attacks targeting the UK appear to be almost entirely LDAP since mid-2017.

#### 4.3 Self-reported booter dataset

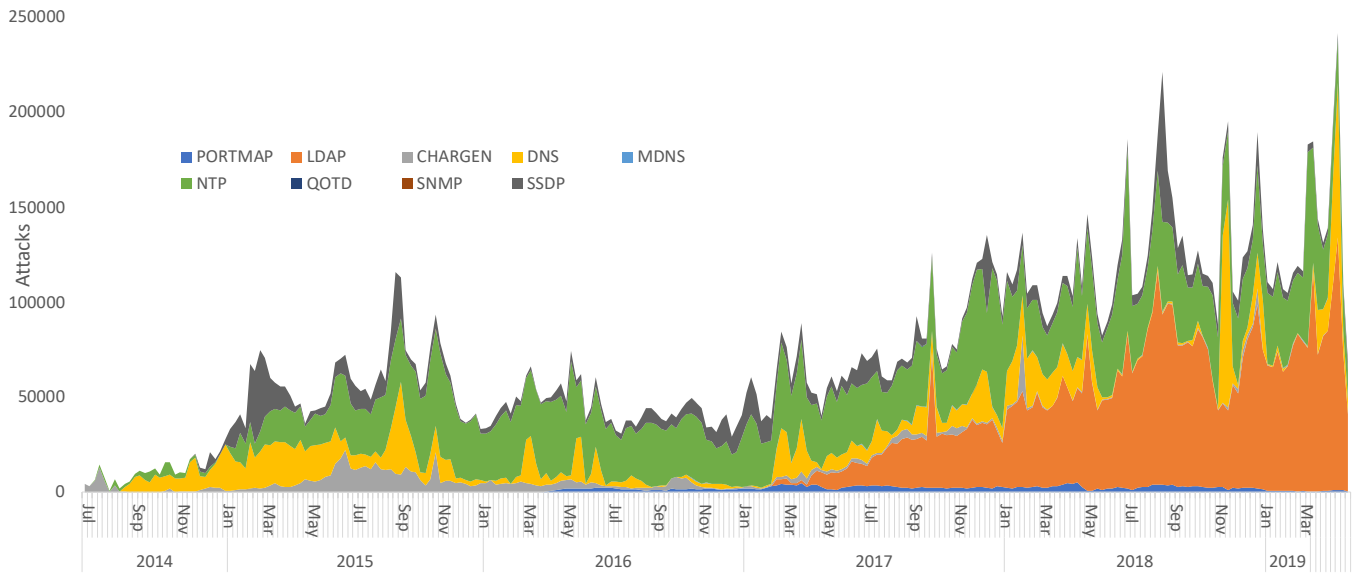
The self-reported attack data is shown in Figure 7. The data is unlabelled because there are 150 different booters involved, some of which are still active. We see a generally increasing level of attacks over time, with clear changes following particular interventions. Booters, especially those of medium size, tend to be fairly unstable, and the data reflects this, with outages clearly visible throughout. The effects of these outages often appear to be ‘absorbed’ by displacement to other booters (as seen in March 2018), so the overall attack numbers remain steady overall.

In Figure 8 we plot the number of booters leaving (‘deaths’) or re-entering (‘resurrections’) the market each week. The spikes in new booters (‘births’) are an artefact of the data collection process (aperiodic searches for new booters) and we do not analyse those. Most weeks there is little change, with two exceptions.

Webstresser was taken down in April 2018 (§2.5)<sup>2</sup> and we believe this disrupted a number of smaller booters that had subcontracted their attacks to it and thus we see a spike in ‘deaths’. The attack data shows that the two biggest booters are largely unaffected (though decreases in their attacks show that there is some discouragement of users), and after a couple of weeks new booters begin to appear.

The effects of the Xmas2018 intervention (§2.6) are highly visible in the data and again there is a spike in ‘deaths’. There is an initial large drop in attack numbers, which builds back up slightly to a plateau of reduced attack numbers (relative to the pre-intervention trend), which lasts until March. Prior

<sup>2</sup>Webstresser did not self-report attack numbers, so is not included in Figure 7. Analysis of partial attack logs from Webstresser indicate that it was of a similar size to the largest current booter before it was taken down. Unfortunately, this data is not comparable with the self-reported data due to its incompleteness and so is not presented.



**Figure 6: Stacked area graph of total attacks per week over time, split by protocol used by attacker. Most of the growth comes from LDAP.**

to the intervention, the market was distributed between three major players and numerous smaller providers. The intervention closes two of the three with the remaining one maintaining a substantial share (about 60%) of an eventually recovering market. Some of the smaller booters benefit from this displacement as well, but not to the same extent.

Immediately after the intervention some short-lived booters enter the market, then leave after a few weeks. More new booters start appearing after a month, but with little impact on the overall total number of attacks, so it appears that the publicity has dissuaded users from attacking, or they cannot locate working booters, or a search leads them to visit one of the domains seized by the FBI and the splash page has dissuaded them from further action. Growth in attack numbers does not occur until March when one of the booters taken down in December returns under a similar name.

Note that following the Xmas2018 intervention, while the UDP reflection data shows a general flattening of the previously rising trend, the self-reported data continues to grow from March 2019 onwards. From our separate qualitative research on botter communities, we believe that this is a move away from UDP reflection attacks to using botnets that send traffic directly, often as Layer 7 (TCP) attacks.

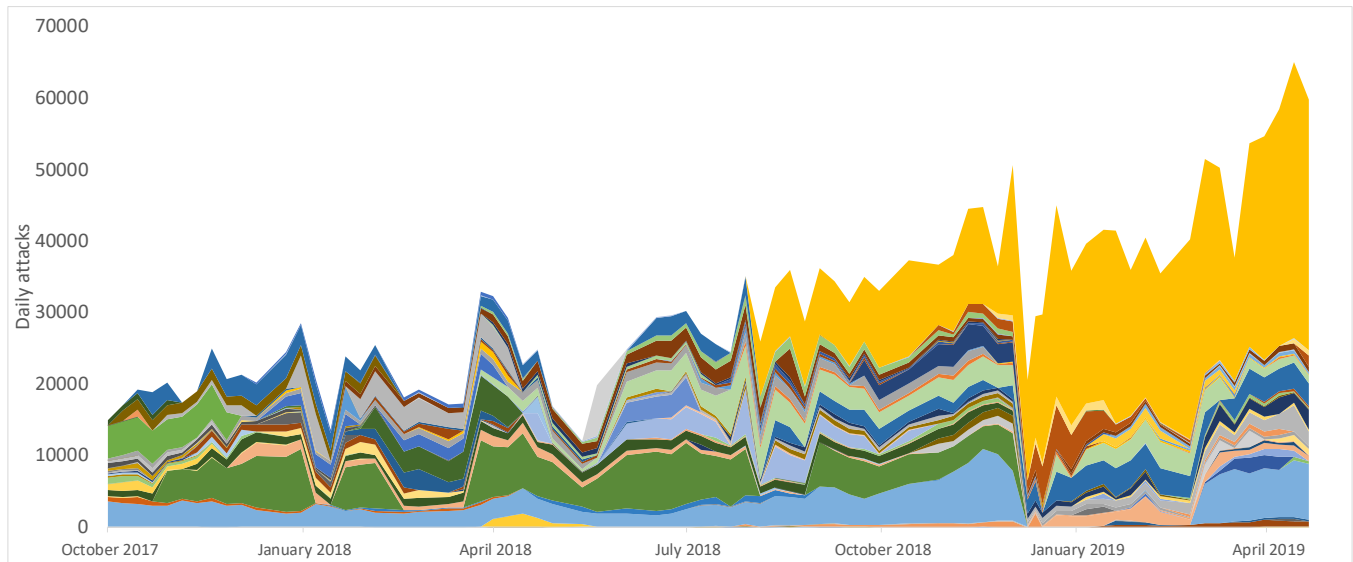
## 5 RELATED WORK

The closest work to this is the simultaneously published paper by Kopp et al. which uses IXP and ISP flow data to study the Xmas2018 intervention by the FBI [29]. They also found that there was reduction in attacks as a result, but found it to be smaller, possibly because they only model attacks over the period Oct 2018 to Jan 2019, thereby ignoring seasonal effects.

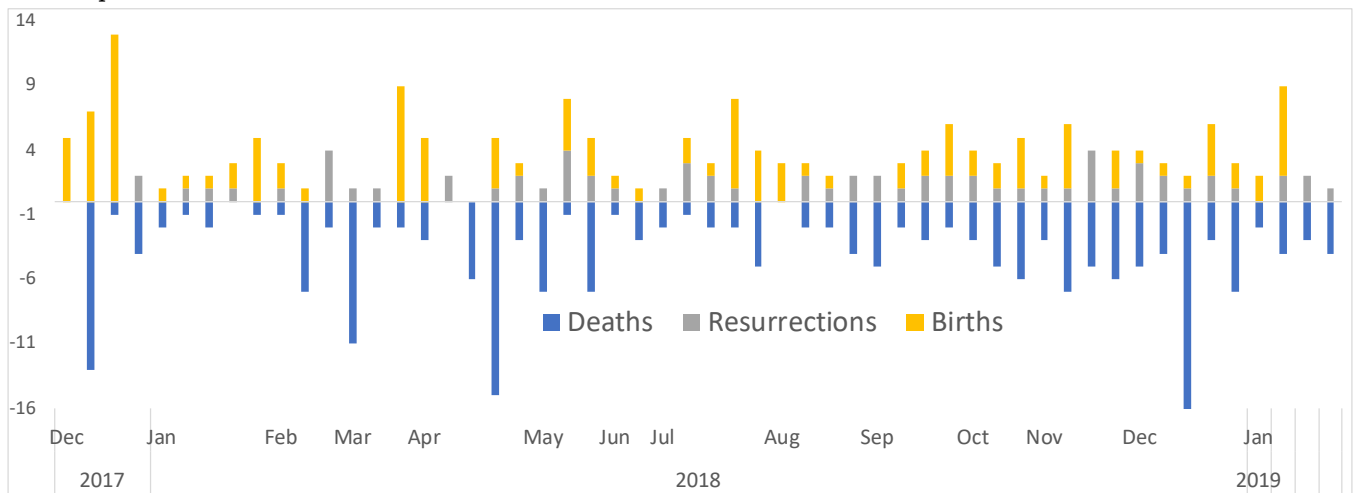
Analysis of botter databases began with Karami and McCoy’s analysis of twBooter [27], then Santanna et al. analysed databases from 15 booters [53]. Karami et al. [28] evaluated the disruptive effects of PayPal shutting down accounts linked with booters. They found that despite operators changing to accept bitcoin payments, the intervention had a negative effect on revenue. Brunt et al. [7] followed this up by scraping the vDOS website and analysing its dumped database to show that only 11% of the customers who had previously paid by PayPal switched to Bitcoin. However, vDOS probably still remained profitable until the operators were arrested [30].

The other major strategy for measuring botter attacks is using UDP honeypots such as the hopscotch [61] or AmpPot [41]. Noroozian et al. [48] used AmpPot and analysed the victims of DDoS finding that most were on access networks. In contrast Jonker et al. [25] combined AmpPot data with UCSD’s network telescope and found that most DoS attacks were on web servers. Several research groups have bought attacks from booters in order to analyse the attacks [29]. Santanna et al. analysed attacks from 14 booters [55]. Krupp et al. [40] were able to use purchased attacks to link other attacks to individual botter services with a precision of 99% and recall of 69% using a *k*-NN classifier using the set of honeypots used in the attack, the TTL values, and the victim port entropy. Krupp et al. [39] used a different approach for attribution and used selective reply from AmpPot honeypots to give scanners a fingerprint which enabled them to attribute 58% of attacks to a scanner IP with 99.9% certainty.

Dupont [15] examines attempts to deal with botnets and identifies three main categories of response: incapacitation, disruption and harm reduction. More widely, current law enforcement approaches to online crime focus on disrupting markets through investigation and high-profile prosecution of



**Figure 7:** Stacked area graph showing total number of DoS attacks per week over time as self-reported by booter provider websites.



**Figure 8:** Number of booters entering and leaving the market per week. “Deaths” (booters not responding) and “resurrections” (running again after a death) are recorded weekly for all booters which we are aware of. “Births” relate to intermittent sweeps to detect new booters being set up and should be viewed cautiously.

prominent ‘key players’, aided by intelligence gathering and standard policing approaches such as cultivating informants and infiltration [24].

Criminological studies have analysed the effects of police crackdowns on online drug markets. These ‘crackdowns’ are time-limited intensive law enforcement efforts which aim at deterrence. The assumption is that increasing the likelihood and severity of punishment will increase the perceived risk of operating in these markets. This increased risk should raise the cost of the end product, reduce consumption, and lead to the exit of key market players [51]. However, the literature suggests that the effect of these interventions is limited, with

drug markets showing more resilience than would be expected if their members’ participation were determined on a purely calculative, economic basis [12]. In addition to ‘crackdowns’, there are a range of additional law enforcement strategies for tackling online marketplaces [24], but there is limited data on the effectiveness of these strategies in practice.

Comparing the effects of different kinds of interventions in the market for booter services suggests that it is particularly susceptible to law enforcement action. Criminological work on other online illicit markets shows that law enforcement interventions are often limited in their effects. Décarry-Hétu’s research on drug cryptomarkets shows that the wide-ranging

Operation Onymous police action was successful in reducing numbers of active dealers, and to a smaller extent, consumption of drugs, for around two months but there was no long-term observable effect on trade in drug cryptomarkets [12]. The impact we see from the Xmas2018 intervention is longer term and more pronounced. Also, we believe that the structure of the booter market has been changed with a single booter now predominating. Research by Ladegaard shows that, contrary to our findings for booter services, high-profile sentencing judgements actually appear to increase revenue for drug cryptomarkets [42]. This accords with a long history of criminological research which suggests that harsh sentencing has little deterrent effect on crime [42]. This may indicate that involvement in the booter market is influenced by rather different factors than involvement in drug cryptomarkets.

Our use of time series data to evaluate the effects of interventions intended to reduce crime has a long history. For example, Dugan et al. [14] use this approach to evaluate the effects of a range of interventions designed to reduce aeroplane hijackings from 1931 to 2003. They found the introduction of metal detectors and increased enforcement significantly reduced hijackings, while tighter baggage and customer screening did not.

## 6 DISCUSSION

We are well-aware of flaws in the datasets that we have used for our analysis. The datasets are incomplete, they are very noisy and, as shown by the rise in attacks in the Netherlands when everywhere else was seeing a decrease, local effects can substantially affect global trends. However our datasets are large and cover long periods of time and so despite these limitations we believe they provide a reasonable approximation to ‘ground truth’ about levels of attacks and the effects which different interventions have on them.

We are also cautious in making claims about cause and effect. While we have been able to locate interventions or disruptions which correspond to major drops in the time series of attacks, they may actually be due to factors of which we are unaware.

That said, we see that there is consistency between the type of intervention and the effect it produces both globally and at a country level. We now discuss this by considering the different types of intervention and the extent to which they produce three primary outcomes – dissuading providers (reducing supply), dissuading users (reducing demand), and producing structural changes to the market.

### 6.1 High-profile court cases

Media coverage of the prosecution or sentencing of booter providers appears to have no consistent effect on the number of attacks we observe. The reporting of the two Mirai court cases shows a clear and significant reduction in attack numbers but with a substantial variation between different countries, presumably because events in foreign countries are seen as less salient than those closer to home. Additionally, these took place at a time during which other events were

occurring which may have disrupted the market for booter services. We argue that this indicates that the reporting of booter sentences has no consistent negative effect on attacks. Given that these are not linked to increased shutdowns of booter services, we argue that any effect we do observe is to reduce the demand for attacks from users.

### 6.2 Taking down individual booters

The Webstresser takedown had a deep but short-term effect on the market for booter services. Webstresser was the biggest booter at the time of its shutdown, but there was only limited structural change to the market. A number of smaller booters disappeared, but they made little contribution to overall attack totals. There was a short-term drop in attacks by medium-sized providers, but mainly in mainland Europe and the US. Although there may have been a deterrent effect, the reduction may in fact be due to reselling, and Webstresser may have been providing the actual attack infrastructure and other booters were merely a shop-front. There was no lasting effect on the overall trend of attack numbers or the structure of the market, and the market had recovered to previous levels after a few weeks.

### 6.3 Wide-ranging interventions

Compared to the limited effects of the takedown of Webstresser, more wide-ranging takedowns had a much longer-lasting effect. The Xmas2018 intervention was by some margin the most effective, preceding a 10 week decrease of 27% to 37% fewer attacks than would be expected. The shutdown of the SST section of HackForums led to a shallower effect, but one which lasted for 13 weeks. We view this as a form of ‘takedown’ because it was the closure of a series of shop-fronts for booters which directly affected how easy it was to find a booter. It also removed a space for discussions where users could compare the effectiveness of booters, share practices and generally reinforce the booter culture.

We believe these wide-ranging takedowns affected the structure of the market, causing a number of booters to leave the market permanently, along with a move away from multiple mid-range providers towards a market dominated by a single booter. We also see clear evidence of a suppression of user demand for services, with lower overall numbers of attacks for a sustained period.

### 6.4 Targeted messaging campaigns

The NCA’s search adverts campaign targeting potential booter users in the UK appears to be correlated with a striking change in the time series of UK attacks. Where the other major booter-using nations continue their upward trend, the UK deviates at this point, flattening throughout the period of the campaign and only resuming an upward trend a few months after it has ended. This suggests that the campaign may have had the effect of dissuading new users from becoming involved, halting the rising demand for attacks for a period of seven or eight months. It further suggests that the rise in attacks (at least in the UK) comes

from increased demand for these services linked to new users entering the market, rather than extra activity by existing booter users.

## 6.5 Displacement and deterrence

While we can observe displacement to alternative booter providers when takedowns occur, this is often time-limited for smaller providers as the influx of users can overwhelm them (ironically this can be seen as a ‘denial of service’) and lead to their services stopping working effectively.

Although there is much commonality, we observe differences in the effect sizes of interventions between countries. These may be language effects in that news fails to spread, particularly to China. Additionally, Russia and France saw no significant drop from the Xmas2018 event, suggesting that their booter communities may have been using a different set of providers. Conversely, the Webstresser takedown had an effect across several countries, possibly because the market was so concentrated around Webstresser. The current concentration of the market round a single booter means that if it were to be taken down then we would expect to see an international effect.

There is little evidence in the literature of any deterrent effect of media reporting of sentencing for other kinds of criminal online activity [42], so our findings for booting merit further consideration. Equally, although adverts may not affect the behaviour of those already involved in booting, the NCA campaign appears to have halted the rise in DoS attacks in the UK for as long as it was running.

We argue, given previous research on the booter community [22], that this we have not observed a ‘classical’ deterrent effect where interventions affect the risk calculus of actors involved in crime. Instead, we believe that the effect is explained by cultural factors in the booter community, which is particularly reliant on a widespread, persistent narrative that booting is not serious crime, involves low levels of harm, and is effectively legal. The effectiveness of the NCA campaign implies that this narrative is a key factor for new users entering the booter market. We therefore argue, on the evidence from the UK, that messaging campaigns could be a key part of an effective strategy against booting.

## 7 CONCLUSIONS

This is the first academic research, of which we are aware, to study the effects of police interventions on the market for booter services and it is also one of the first evaluations of cybercrime interventions more generally.

We have modelled the number of reflected UDP DoS attacks using a negative binomial regression with intervention components modelled where the series drops below that expected from the seasonally adjusted, upward trend. We then linked these drops to widely reported intervention events: discussion boards closing, shutting down booters and the sentencing of booter operators.

We have used self-reported data of the number of attacks carried out by individual booters to determine whether the

impact of interventions has been to cause booters to withdraw from the market (voluntarily or otherwise) or to affect the demand for their services. We find some evidence for displacement – when one booter closes there is an uptick in the attacks performed by others. This does however make the market more ‘brittle’, meaning that the long term impact of the Xmas2018 intervention may make any future action against the booter which now has 60% of the market especially disruptive.

The main impact we see is that interventions against booters can successfully cause a reduction in attack numbers. We see a strong effect from the targeted messaging of the NCA search adverts campaign in the UK which appeared to be particularly effective at keeping new users out of the market.

The most successful interventions appear to be mass takedowns – the Xmas2018 intervention saw the closure of large numbers of booter sites, and the Hackforums intervention led to the de-facto closure of several major shopfronts. Both these interventions made it harder for users to find working booters. Arrests (and subsequent sentencing) do have an effect, but it is more short-lived. It is an open question whether arrests are essential to reinforce the impact of a takedown, although they should of course prevent the same booter operator just starting up again the following day.

We argue that there are three mechanisms underlying the effects we see in our data. Firstly, messaging campaigns appear to suppress user demand for services by undermining the widespread perception in the booter community that their activity is low-harm and essentially legal. A further advantage of messaging approaches is that they are relatively cheap, do not pull people into the criminal justice system, and avoid the criminogenic effects and harms of harsher enforcement action. Secondly, there appears to be a destabilising effect of website takedowns, which dissuade booter providers and reduce the accessibility of these services. Finally, wide-ranging website takedowns appear to have a structural effect on the market for booter services, concentrating them around particular providers (and potentially making them more susceptible to further intervention).

## ACKNOWLEDGMENTS

We wish to thank the Cambridge Cybercrime Centre for access to the booter datasets, the UK National Crime Agency and US Federal Bureau of Investigation for providing information about their activities, and vigilante.pw for the booter.io database. Finally, we thank our colleagues, particularly our shepherd Alberto Dainotti, Ildikó Pete, Alexander Vetterl, and the anonymous reviewers for their invaluable feedback and support. This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) [grant number EP/M020320/1].

## REFERENCES

- [1] Angus Bancroft and Peter Scott Reid. 2017. Challenging the techno-politics of anonymity: The case of cryptomarket users. *Information, Communication & Society* 20, 4 (2017), 497–512. <https://doi.org/10.1080/1369118X.2016.1187643>

- [2] BBC News. 2017. Teenage cyber hacker Adam Mudd gets jail term reduced. <https://www.bbc.co.uk/news/uk-england-beds-bucks-herts-40744373>
- [3] Richard Berk, Geoffrey Barnes, Lindsay Ahlman, and Ellen Kurtz. 2010. When second best is good enough: A comparison between a true experiment and a regression discontinuity quasi-experiment. *Journal of Experimental Criminology* 6, 2 (2010), 191–208.
- [4] Richard Berk and John M. MacDonald. 2008. Overdispersion and Poisson regression. *Journal of Quantitative Criminology* 24 (2008), 269–284.
- [5] Anthony A. Braga and Brenda J. Bond. 2008. Policing crime and disorder hot spots: A randomized controlled trial. *Criminology* 46, 3 (2008), 577–607.
- [6] Gerben J. Bruinsma and David Weisburd. 2007. Experimental and quasi-experimental criminological research in the Netherlands. *Journal of Experimental Criminology* 3, 2 (2007), 83–88.
- [7] Ryan Brunt, Prakhar Pandey, and Damon McCoy. 2017. Booted: An analysis of a payment intervention on a DDoS-for-hire service. In *Workshop on the Economics of Information Security (2017-06-26) (WEIS)*. 12.
- [8] Catalin Cimpanu. 2018. Law enforcement shut down DDoS booters ahead of annual Christmas DDoS attacks. <https://www.zdnet.com/article/law-enforcement-shut-down-ddos-booters-ahead-of-annual-christmas-ddos-attacks/>
- [9] Richard Clayton. 2018. Google doesn't seem to believe booters are illegal. <https://www.lightbluetouchpaper.org/2018/08/28/google-doesnt-seem-to-believe-booters-are-illegal/>
- [10] Thomas D. Cook and D. T. Campbell. 1979. . Goodyear Publishing Company, Santa Monica, CA, USA, Chapter The design and conduct of true experiments and quasi-experiments in field settings.
- [11] Nicholas Corsaro. 2018. Interrupted Time Series Analysis Using STATA. Justice Research Statistics Association (JRSA) Conference. <http://www.jrsa.org/events/presentations/western-2018/corsaro.pdf>
- [12] David Décarry-Hétu and Luca Giommoni. 2017. Do police crackdowns disrupt drug cryptomarkets? A longitudinal analysis of the effects of Operation Onymous. *Crime, Law and Social Change* 67, 1 (2017), 55–75. <https://doi.org/10.1007/s10611-016-9644-4>
- [13] Department of Justice. 2018. Criminal charges filed in Los Angeles and Alaska in conjunction with seizures of 15 websites offering DDoS-For-Hire services. <https://www.justice.gov/opa/pr/criminal-charges-filed-los-angeles-and-alaska-conjunction-seizures-15-websites-offering-ddos>
- [14] Laura Dugan, Gary LaFree, and Alex R. Piquero. 2005. Testing a rational choice model of airline hijackings. *Criminology* 43, 4 (2005), 1031–1065.
- [15] Benoit Dupont. 2017. Bots, cops, and corporations: On the limits of enforcement and the promise of polycentric regulation as a way to control large-scale cybercrime. *Crime, Law and Social Change* 67, 1 (2017), 97–116. <https://doi.org/10.1007/s10611-016-9649-z>
- [16] Europol. 2016. Joint international operation targets young users of DDoS cyber-attack tools. <https://www.europol.europa.eu/newsroom/news/joint-international-operation-targets-young-users-of-ddos-cyber-attack-tools>
- [17] Europol. 2018. World's biggest marketplace selling internet paralysing DDoS attacks taken down. <https://www.europol.europa.eu/newsroom/news/world%E2%80%99s-biggest-marketplace-selling-internet-paralysing-ddos-attacks-taken-down>
- [18] Europol. 2019. Authorities across the world going after users of biggest DDoS-for-hire website. <https://www.europol.europa.eu/newsroom/news/authorities-across-world-going-after-users-of-biggest-ddos-for-hire-website>
- [19] A. C. Harvey and C. Fernandes. 1989. Time series models for count or qualitative observations. *Journal of Business & Economic Statistics* 7, 4 (1989), 407–417. <https://doi.org/10.1080/07350015.1989.10509750>
- [20] Hertfordshire Constabulary. 2018. Computer hacker Adam Mudd ordered to pay back £70,000. <https://www.herts.police.uk/news-and-appeals/Computer-hacker-Adam-Mudd-ordered-to-pay-back-70,000-C>
- [21] Home Office. 2018. Serious and Organised Crime Strategy. <https://www.gov.uk/government/publications/serious-and-organised-crime-strategy-2018>
- [22] Alice Hutchings and Richard Clayton. 2016. Exploring the provision of online booter services. *Deviant Behavior* 37, 10 (2016), 1163–1178. <https://doi.org/10.1080/01639625.2016.1169829>
- [23] Alice Hutchings, Richard Clayton, and Ross Anderson. 2016. Taking down websites to prevent crime. In *APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, Toronto, ON, Canada, 1–10. <https://doi.org/10.1109/ECRIME.2016.7487947>
- [24] Alice Hutchings and Thomas J. Holt. 2017. The online stolen data market: Disruption and intervention approaches. *Global Crime* 18, 1 (2017), 11–30. <https://doi.org/10.1080/17440572.2016.1197123>
- [25] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. 2017. Millions of Targets Under Attack: A Macroscopic Characterization of the DoS Ecosystem. In *Proceedings of the 2017 Internet Measurement Conference (IMC '17)*. ACM, 100–113. <https://doi.org/10.1145/3131365.3131383>
- [26] Mohammad Karami and Damon McCoy. 2013. Rent to pwn: Analyzing commodity booter DDoS services. *Usenix login* 38, 6 (2013), 20–23.
- [27] Mohammad Karami and Damon McCoy. 2013. Understanding the emerging threat of DDoS-as-a-Service. In *LEET '13*. USENIX, 4.
- [28] Mohammad Karami, Youngs Park, and Damon McCoy. 2016. Stress testing the booters: Understanding and undermining the business of DDoS services. In *Proceedings of the 25th International Conference on World Wide Web (WWW)*. International World Wide Web Conferences Steering Committee, Montréal, Québec, Canada, 1033–1043. <https://doi.org/10.1145/2872427.2883004>
- [29] Daniel Kopp, Matthias Wichtlhuber, Ingmar Poesse, José Jair Santanna, Oliver Hohfeld, and Christoph Dietzel. 2017. DDoS Hide & Seek: On the Effectiveness of a Booter Services Takedown. In *Internet Measurement Conference, October 21–23, 2019 (IMC '19)*. ACM.
- [30] Brian Krebs. 2016. Alleged vDOS proprietors arrested in Israel. <https://krebsonsecurity.com/2016/09/alleged-vdos-prorietors-arrested-in-israel/>
- [31] Brian Krebs. 2016. DDoS on Dyn impacts Twitter, Spotify, Reddit. <https://krebsonsecurity.com/2016/10/ddos-on-dyn-impacts-twitter-spotify-reddit/>
- [32] Brian Krebs. 2016. Feds charge two in Lizard Squad investigation. <https://krebsonsecurity.com/2016/10/feds-charge-two-in-lizard-squad-investigation/>
- [33] Brian Krebs. 2016. Hackforums shuts booter service bazaar. <https://krebsonsecurity.com/2016/10/hackforums-shutters-booter-service-bazaar/>
- [34] Brian Krebs. 2016. Israeli online attack service 'vDOS' earned \$600,000 in two years. <https://krebsonsecurity.com/2016/09/israeli-online-attack-service-vdos-earned-600000-in-two-years/>
- [35] Brian Krebs. 2016. KrebsOnSecurity hit with record DDoS. <https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/>
- [36] Brian Krebs. 2016. 'Operation Tarpit' targets customers of online attack-for-hire services. <https://krebsonsecurity.com/2016/12/operation-tarpit-targets-customers-of-online-attack-for-hire-services/>
- [37] Brian Krebs. 2018. Mirai botnet authors avoid jail time. <https://krebsonsecurity.com/2018/09/mirai-botnet-authors-avoid-jail-time/>
- [38] Brian Krebs. 2018. Mirai co-author gets 6 months confinement, \$8.6M in fines for Rutgers attacks. <https://krebsonsecurity.com/2018/10/mirai-co-author-gets-6-months-confinement-8-6m-in-fines-for-rutgers-attacks/>
- [39] Johannes Krupp, Michael Backes, and Christian Rossow. 2016. Identifying the scan and attack infrastructures behind amplification DDoS attacks. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 1426–1437. <https://doi.org/10.1145/2976749.2978293>
- [40] Johannes Krupp, Mohammad Karami, Christian Rossow, Damon McCoy, and Michael Backes. 2017. Linking amplification DDoS attacks to booter services. In *Research in Attacks, Intrusions, and Defenses (RAID)*, Vol. LNCS 10453. Springer, 427–449.
- [41] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Research in Attacks, Intrusions, and Defenses (RAID) (2015-11)*, Vol. 9404 LNCS. Springer, 615–636. [https://doi.org/10.1007/978-3-319-26362-5\\_28](https://doi.org/10.1007/978-3-319-26362-5_28) ISSN: 0302-9743.
- [42] Isak Ladegaard. 2018. We know where you are, what you are doing and we will catch you: Testing deterrence theory in digital drug markets. *The British Journal of Criminology* 58, 2 (2018), 414–433. <https://doi.org/10.1093/bjc/azx021>

- [43] Alexia Maddox, Monica J. Barratt, Matthew Allen, and Simon Lenton. 2016. Constructive activism in the dark web: Cryptomarkets and illicit drugs in the digital ‘demimonde’. *Information, Communication & Society* 19, 1 (2016), 111–126. <https://doi.org/10.1080/1369118X.2015.1093531>
- [44] Ashich V. Naik, Alok Baveja, Rajan Batta, and Jonathan P. Caulkins. 1996. Scheduling crackdowns on illicit drug markets. *European Journal of Operational Research* 88, 2 (1996), 231–250. [https://doi.org/10.1016/0377-2217\(94\)00201-0](https://doi.org/10.1016/0377-2217(94)00201-0)
- [45] National Crime Agency. 2015. Operation Vivarium targets users of Lizard Squad’s website attack tool. <https://perma.cc/gVMC-SVNE>
- [46] National Crime Agency. 2015. Teenager admitted trying to buy gun on the dark web. <https://perma.cc/98GH-G2BD>
- [47] National Crime Agency. 2016. Operation Vulcanalia targets users of netspoof website attack tool. <https://perma.cc/CXB6-LKX6>
- [48] Arman Noroozian, Maciej Korczynski, Carlos Hernandez Ganan, Daisuke Makita, Katsunari Yoshioka, and Michel Van Eeten. 2016. Who Gets the Boot? Analyzing Victimization by DDoS-as-a-Service. In *Research in Attacks, Intrusions and Defenses (RAID)* (2016-09), Vol. LNCS 9854. Springer, 368–389. <https://doi.org/10.1007/978-3-319-45719-217>
- [49] Patrick Howell O’Neill. 2018. Lizard Squad’s ‘@fbirelosers’ hacker gets smaller sentence for helping FBI arrest his friends. <https://www.cyberscoop.com/zachary-buchta-lizard-squad-sentence/>
- [50] D. Wayne Osgood. 2000. Poisson-based regression analysis of aggregate crime rates. *Journal of Quantitative Criminology* 16, 1 (01 Mar 2000), 21–43. <https://doi.org/10.1023/A:1007521427059>
- [51] Peter Reuter and Mark A.R. Kleiman. 1986. Risks and prices: An economic analysis of drug enforcement. *Crime and Justice* 7 (1986), 289–340. <https://doi.org/10.1086/449116>
- [52] Robert J. Sampson. 2010. Gold standard myths: Observations on the experimental turn in quantitative criminology. *Journal of quantitative criminology* 26, 4 (2010), 489–500.
- [53] José Jair Santanna, Romain Durban, Anna Sperotto, and Aiko Pras. 2015. Inside booters: An analysis on operational databases. In *Proceedings of the 2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*. IEEE, 432–440. <https://doi.org/10.1109/INM.2015.7140320>
- [54] José Jair Santanna, Ricardo de O. Schmidt, Daphne Tuncer, Joey de Vries, Lisandro Z. Granville, and Aiko Pras. 2016. Booter blacklist: Unveiling DDoS-for-hire websites. In *12th International Conference on Network and Service Management (CNSM)*. IEEE, Montréal, Québec, Canada, 144–152.
- [55] José Jair Santanna, Roland Van Rijswijk-Deij, Rick Hofstede, Anna Sperotto, Mark Wierbosch, Lisandro Zambenedetti Granville, and Aiko Pras. 2015. Booters – An analysis of DDoS-as-a-service attacks. In *IFIP/IEEE International Symposium on Integrated Network Management*. IEEE, 243–251. <https://doi.org/10.1109/INM.2015.7140298>
- [56] Alex Scapens. 2017. Cyber attack teenager who helped gang target Netflix, Amazon and NatWest avoids jail. <https://www.manchestereveningnews.co.uk/news/greater-manchester-news/cyber-attack-teenager-who-helped-14058060>
- [57] Yashovardhan Sharma. 2018. Characterising the Victims of DDoS Attacks. MPhil Thesis, University of Cambridge.
- [58] Russell G. Smith, Nicholas Wolanin, and Glenn Worthington. 2003. E-crime solutions and crime displacement. *Trends & Issues in Crime and Criminal Justice* 243 (jan 2003), 1–6.
- [59] Kyle Soska and Nicolas Christin. 2015. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In *24th USENIX Security Symposium (USENIX Security 15)*. USENIX, Washington, DC, USA, 33–48.
- [60] Rebecca Steinbach, Chloe Perkins, Lisa Tompson, Shane Johnson, Ben Armstrong, Judith Green, Chris Grundy, Paul Wilkinson, and Phil Edwards. 2015. The effect of reduced street lighting on road casualties and crime in England and Wales: Controlled interrupted time series analysis. *J Epidemiol Community Health* 69, 11 (2015), 1118–1124. <https://doi.org/10.1136/jech-2015-206012>
- [61] Daniel R. Thomas, Richard Clayton, and Alastair R. Beresford. 2017. 1000 days of UDP amplification DDoS attacks. In *2017 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, Scottsdale, AZ, USA, 79–84. <https://doi.org/10.1109/ECRIME.2017.7945057>
- [62] Daniel R. Thomas, Sergio Pastrana, Alice Hutchings, Richard Clayton, and Alastair R. Beresford. 2017. Ethical issues in research using datasets of illicit origin. In *Proceedings of the Internet Measurement Conference (IMC)*. ACM, 445–462. <https://doi.org/10.1145/3131365.3131389>
- [63] Evan Wood, Patricia M. Spittal, Will Small, Thomas Kerr, Kathy Li, Robert S. Hogg, Mark W. Tyndall, Julio S.G. Montaner, and Martin T. Schechter. 2004. Displacement of Canada’s largest public illicit drug market in response to a police crackdown. *Canadian Medical Association Journal* 170, 10 (2004), 1551–1556. <https://doi.org/10.1503/cmaj.1031928>

## A ETHICAL CONSIDERATIONS

The ethical case for the UDP reflection honeypots operated by the Cambridge Cybercrime Centre is described in Thomas et al. [61] and summarised below.

Reflecting UDP packets could assist criminals in performing DoS attacks. However the hopscotch reflector limits the number of packets it reflects to any IP address. It attempts to only reflect to the criminals’ scanners (so that they use the honeypots) but not (at any scale) to victims. Furthermore, when any hopscotch sensor identifies a victim this is reported to a central server which informs all the other sensors of the attack, so that they all refuse to reflect any packets at all to the victim.

The result is that if the hopscotch sensors are used for illegal attacks then after a very short period there will be rather *less* traffic delivered to the victim than if the sensors did not exist, because some of the attack traffic generated by the criminal is being absorbed by the sensors rather than going to real reflectors and being reflected and amplified.

‘White-hat’ scanners identify reflectors and report them to the relevant ISP or hosting company who will then contact the owner of the machine. The hopscotch sensors do not reply to identified ‘white-hat’ scanners, or to known researchers, in order to avoid wasting their time or affecting their results.

Data on self-reported attacks by booters was collected by logging in to the booter websites as a visitor. No data was collected that required buying a subscription or requesting an attack before data would be displayed.

Both datasets can be obtained by academic researchers from the Cambridge Cybercrime Centre so that other researchers can fully reproduce or build upon our results. <https://www.cambridgecybercrime.uk/process.html>

In line with previous research in this space, we have not published the identity of any booters which are still operational at the time of publication.

The ethical issues in the use of leaked or seized booter databases to support our results is discussed by Thomas et al. [61] and put in a broader context in a later paper on the use of leaked data for research [62]. Only the attack logs were analysed for this paper.

We followed our institution’s ethical review procedure throughout. We carefully designed our experiments to operate ethically and we had no human subjects. Our related qualitative research which supported this work was approved by our ethics committee (#621).