An Empirical Study of Bug Bounty Programs

Thomas Walshe and Andrew Simpson Department of Computer Science, University of Oxford
Wolfson Building, Parks Road, Oxford OX1 3QD, United Kingdom
Email: firstname.secondname@cs.ox.ac.uk

***Abstract*—The task of identifying vulnerabilities is commonly outsourced to hackers participating in bug bounty programs. As of July 2019, bug bounty platforms such as HackerOne have over 200 publicly listed programs, with programs listed on HackerOne being responsible for the discovery of tens of thousands of vulnerabilities since 2013. We report the results of an empirical analysis that was undertaken using the data available from two bug bounty platforms to understand the costs and benefits of bug bounty programs both to participants and to organisations. We consider the economics of bug bounty programs, investigating the costs and benefits to those running such programs and the hackers that participate in finding vulnerabilities. We find that the average cost of operating a bug bounty program for a year is now less than the cost of hiring two additional software engineers.**

***Index Terms*—Bug bounty programs, vulnerability disclosure, software security**

I. Introduction

An increasingly popular approach to identifying vulnerabilities in software is to offer rewards to security researchers that are external to an organisation ('hackers') to find and disclose vulnerabilities [1]. This approach is now seeing adoption in areas such as e-voting systems, government systems and self-driving cars [2]–[4].

As an example. the Swiss government launched a program offering €132,000 for hackers to find vulnerabilities in an e-voting system. Rewards of up to €44,000 were made available to hackers who discovered undetectable ways of manipulating votes [2]. As another example, the US Department of Defense (DoD) launched the 'Hack the Pentagon' pilot program in April 2016, with the aim of assessing the benefit of opening up vulnerability discovery to hackers. Within six hours 138 vulnerabilities were found and reported [5]. The success of the program has led to the DoD introducing a new vulnerability disclosure policy, opening up new domains to hackers [3], [6]. There have also been suggestions for US government departments to participate in searches for vulnerabilities in open-source projects [7].

The formal exchange of information in this context is typically facilitated by a bug bounty platform. The number of new bug bounty programs available on platforms such as HackerOne[1] has increased year-on-year since 2013 [8], with 50 new organisations launching a program in 2018. This increase is illustrated in Table I.

As of January 2019, the top 25 companies using HackerOne have used the platform to obtain reports for over 19,000 vulnerabilities, at an average of 0.71 vulnerabilities reported for each day the program is run — resulting in $11.9 million being paid out to hackers for successfully finding vulnerabilities [8].

An empirical study into the use of rewards programs has the potential for economic analysis to be performed, with a view to explaining why some companies are increasingly utilising the global network of security expertise offered by hackers, instead of hiring additional staff — and whether others should follow that pattern [9]. These questions drive this contribution.

In order to give appropriate consideration both to related work and to our questions of interest, we first need to establish appropriate criteria. One such criterion is the cost of employment of a software engineer, which we calculate as follows.

The average salary for a London-based software engineer is £41,700 [10], which is equivalent to $51,708 (exchange rate £1 to $1.24, as of 17^{th} of July 2019). Of course, the hiring process has additional initial costs when new staff are employed — Blatter et al. estimated the cost of hiring new skilled workers to be equivalent to 10–17 weeks salary [11]. This puts the total cost of hiring an additional software engineer between $61,652 and $68,613 for the first year. Therefore, an average value of $65,133 will be used to represent the cost of hiring an additional software engineer throughout this paper.

II. Background and Motivation

A. Background

Vulnerability management aims to improve the security of a system through the identification, remediation and mitigation of software vulnerabilities [12]. The latest version of the Building Security in Maturity Model (BSIMM9) identifies the operation of bug bounty programs (CMVM3.4) as a mature activity that addresses the need for vulnerability management [13]. Bug bounty programs are being integrated into secure software development lifecycle (SDLC) frameworks to aid security teams in the release and maintenance phases [14].

Many large organisations (including Google, Facebook and Microsoft) host their own bug bounty and vulnerability rewards programs [15]–[17]. Many smaller organisations choose to use bug bounty platforms such as HackerOne, BugCrowd and Cobalt to advertise their programs [18]–[20]. HackerOne offers a range of both free and paid for services to organisations wishing to run a bug bounty program. Free hosting allows organisations to advertise their program to hackers, with only a 5% surcharge on any bounty payouts. Monetisation of the platform comes in the form of live-support and the operation of a fully managed program — allowing for organisations without experience to benefit from the operation of a bug bounty

[1] https://www.hackerone.com

TABLE I
NUMBER OF NEW PROGRAMS ON HACKERONE FROM 2013 TO 2019

Year	New programs
2013–2014	11
2014–2015	26
2015–2016	33
2016–2017	37
2017–2018	43
2018–2019	50

program. One benefit offered by platforms is the increased visibility of programs, allowing for a large number of security researchers to search for vulnerabilities [21]. Raymond argued that increasing the number of people searching for vulnerabilities is beneficial with regards to minimising the number of vulnerabilities: "Given enough eyeballs, all bugs are shallow" [22].

White hat hackers are independent security researchers that are authorised by an organisation to identify security vulnerabilities in hardware, software or networks [23]. They must comply with the 'rules of engagement' set out by an organisation to prevent misuse or intrusion into areas that are out of scope [24]. Kranenbarg et al. recommend that new hackers should be taught about the rules around vulnerability disclosure as to avoid accident misuse [25].

A benefit to these ethical hackers is the lucrative rewards that can be paid out after being the first to discover and document a critical vulnerability, with the highest single payout (at the time of writing) standing at $100,000 [26]. This also legitimises the work of many hackers who might otherwise sell the vulnerability reports on cybercrime markets for a similar price [27]. This crowdsourced approach to security brings together hackers from many backgrounds and disciplines, resulting in a wide range of approaches to finding vulnerabilities that might not otherwise be achieved by an organisation's security team [28], [29].

B. Related work

Finifter et al. considered the vulnerability reward programs offered by Google and Mozilla for the Chrome and Firefox web browsers [9]. Vulnerability reports over three years (2010–2013) were analysed to find the payouts, severity and frequency at which reports are submitted. Comparing the two programs shows that Google has fixed three times as many vulnerabilities identified by hackers than Mozilla. According to Finifter et al. this is due to a larger number of white hat hackers (when compared to the Google program) taking part in the program, as well as a broader reward structure [9].

The daily cost to operate each program is reported as $485 for Google and $658 for Mozilla; over the course of a year, the total cost is $177,025 ($485 $\times$ 365 days) and $240,170 ($658 $\times$ 365 days). This is broadly comparable to the salary of three or four additional software engineers, with the current average salary of a software engineer being $65,133. The authors argue that it is economically viable to run bug bounty programs instead of hiring additional researchers. Further work is recommended by the authors to include economic models to identify phases during the operation of a program which, in part, motivates the present contribution.

A similar study by Zhao et al. analyses web vulnerability reports on the Wooyun and HackerOne bug bounty platforms [30]. Wooyun served as the predominant platform in China from 2010 until being shut down in 2016 [31]. Analysis of vulnerability reports is used to determine the number of white hat hackers participating in the search for vulnerabilities. HackerOne is shown to have a consistent growth of approximately 75 new hackers each month from 2014 to 2015. Wooyun received almost 200 hackers each month; however, as most of the programs do not give rewards, there is little incentive to participate. An interesting observation is that over half of all hackers have only submitted one report; the top 100 have an average of 147 submissions per person. Basic economic analysis by the authors shows that offering a monetary reward attracts more hackers than non-monetary rewards or acknowledgements in a hall of fame.

Unfortunately, [30] does not look into the use of the delayed full disclosure policy used by Wooyun. This controversial policy can be used to force an organisation to fix vulnerabilities more quickly. However, the policy creates conflict between organisations and hackers, which is one of the contributing factors that led to the shutdown of Wooyun in 2016, with several members being arrested [32].

C. Motivation

Since the 2015 study by Zhao et al. [30], the number of programs available on HackerOne has almost tripled from 82 to 212 [18]. In the last two years, an additional 2289 rewarded reports have been disclosed on the Chromium bug tracker — up from 501 in the years 2010 to 2013 [9], [33]. This significant increase in publicly available data gives strong motivation to verify the hypotheses presented in the previous two studies. For the oldest programs on HackerOne, six years of disclosure data is available. A temporal study can be conducted investigating the frequency and severity of reports over time. This allows for the long term impact of operating a bug bounty program to be assessed and compared to the impact of hiring additional security researchers.

The creation of new economic models for bug bounty programs has the potential to be useful to organisations wanting to set up a new program or maintain an existing one. Such models could be used to determine the optimal rewards structure over the lifetime of a program to attract the largest number of hackers.

The question *Is it beneficial for a company to make use of bug bounty programs, instead of hiring additional software engineers?* is addressed at least in part by answering the following subsidiary questions:

1) What is the expected cost per year of operating a bug bounty program, and what results can be expected?
2) What is the expected benefit per year?
3) How does the program activity vary across its lifetime?

4) Does the promise of higher bounty payouts result in an increased number of vulnerabilities being reported?

The motivation for Question 1 is to quantify the cost to an organisation of running a program. This running cost can be directly compared to the average cost of hiring a software engineer (which, as discussed, we have estimated at $65,133). If companies are to make use of bug bounty programs, the cost of operating a program should be less than the cost of hiring security researchers, while yielding a similar number of vulnerabilities found per year. The benefit of operating a bug bounty program is explored in Question 2. Question 3 considers the long-term effect of operating a program. Understanding how the rate of vulnerability discovery changes over time allows for the long-term costs of operating a program to be estimated. Finally, Question 4 considers the program rewards structure and the effect this has on hacker productivity. A company will want to keep the operating costs low and the rate of vulnerability discovery high. Investigating the link between bounty payouts and rate of vulnerability discovery allows for an optimal rewards structure to be proposed.

III. Data

HackerOne and BugCrowd were chosen because of the large quantity of publicly available data on their platforms. They are also the two most searched for bug bounty platforms [34]. Synack also provides a platform that is growing in popularity; however, there is currently no public data provided [35].

The lack of publicly available data and poor transparency in many programs is an issue that inevitably limits current research [36], [37]. Increased transparency may benefit both hackers and organisations wishing to learn about uncommon vulnerabilities that might be found [38].

For HackerOne, all of the program data available on the directory was collected. The data for each program was: program launch date, number of reports resolved, minimum bounty, and the average bounty. The Hacktivity section provides information on the latest reports submitted. Each entry contains the username of the hacker that submitted the report and the time of submission. Additional information is sometimes displayed, which includes: a description of the vulnerability, severity level, and the bounty awarded. Every entry in this section was collected.

The leaderboard provides information on the top 100 hackers: reputation is used as a measured of number of valid reports submitted, signal is the average reputation per report, and impact is the average reputation per bounty. This information was collected for user analysis.

Program data was also collected from BugCrowd; this information was not as extensive as the HackerOne directory. For each program, the number of vulnerabilities discovered and the average payout was collected. BugCrowd provides detailed user performance statistics for each hacker. This data was collected for users with the highest number of points in a program's Hall of Fame.

All data was collected from December 2018 to January 2019. A summary of all data collected is shown in Table II.

TABLE II
Summary of information collected from HackerOne and BugCrowd

	HackerOne	BugCrowd
Programs	212	99
Reports	5,832	Not available
User data	100	92

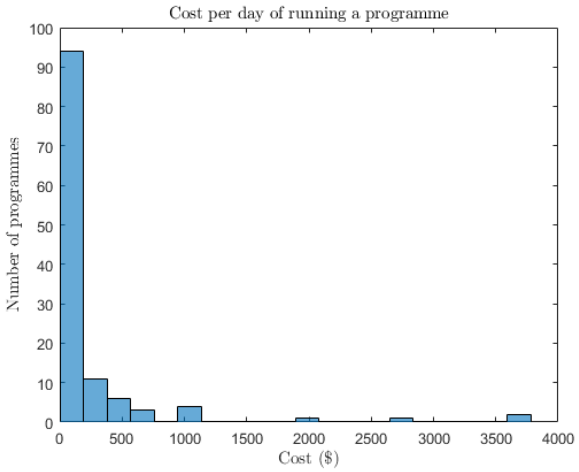

Fig. 1. Histogram of the cost per day to operate a bug bounty program.

IV. Results

This section presents results that help answer the questions introduced in Section II-C. Results pertaining to Question 1 are presented in Sections IV-A and IV-B. Questions 2 and 3 are explored in Sections IV-C and IV-D respectively. The remaining subsections give consideration to additional interesting data collected from HackerOne and BugCrowd.

A. What is the expected cost per year of operating a bug bounty program?

The total cost to run each program is given on HackerOne's directory. Taking this figure and dividing it by the total number of days since the program was launched gives the average daily running cost for each program.

The average daily cost of operating a bug bounty program is $230, with a 95% confidence interval (CI) of [$126, $334]. This gives an average yearly cost of $83,950 ($230 $\times$ 365 days). This is more than the $65,133 required to hire each additional software engineer.

The daily cost of operating each program is displayed in Figure 1. This shows that a large proportion of programs have an operating cost of $250 per day or less. There are several outliers found on HackerOne; we consider these below.

- PayPal launched a program in August of 2018 and has paid out over $700,000 to date. The high operating cost of $3,766 per day is due to the large number of critical vulnerabilities reported; these pay out up to $23,000 each.
- The newest available program was Postmates, and within one week after launch they had already spent $26,124 to find 75 vulnerabilities. The daily cost was $3,732; this

TABLE III
CONVERSION FROM CVSS 3.0 SCORE TO SEVERITY RATING.

Severity	Critical	High	Medium	Low	None
Score	10.0–9.0	8.9–7.0	6.9–4.0	3.9–0.1	0.0

TABLE IV
SEVERITY OF THE MOST RECENT 3,000 DISCLOSED REPORTS ON HACKERONE, ALONG WITH AVERAGE PAYOUTS.

Severity	Proportion (%)	Average cost ($)	95% CI
Critical	8.5	7,924	[$3,324, $12,524]
High	19.6	2,610	[$2,046, $3,174]
Medium	33.5	786	[$618, $954]
Low	28.8	225	[$177, $272]
None	9.6	0	[$0, $0]

high operating cost is expected to fall as the obvious vulnerabilities are reported.

- Oath's program, which was launched in 2014, has cost almost $1,000,000 per year to run. The daily cost was $2,730. This is the most active program on HackerOne with over 5,000 vulnerabilities reported.

Severity ratings are either set by the organisation for each type of vulnerability (e.g. remote code execution designated as a critical vulnerability) or calculated by the hacker using a Common Vulnerability Scoring System (CVSS) score [39]. CVSS scores are calculated using exploitability, impact, temporal and environmental metrics, and represent the overal impact of a given vulnerability. A typical conversion from a CVSS 3.0 score to severity rating is shown in Table III [40].

Table IV shows the average payout for a vulnerability of a given severity, with Figure 2 showing the payout distribution. Munaiah et al. found that there is a weak correlation between CVSS score and the bounty payout, with the CVSS score being more likely to underestimate the actual bounty payout [41]. This is also found in the skewed distribution shown in Figure 2.

As discussed by Finifter et al. [9], certain architectural decisions can lower the proportion of critical vulnerabilities present in software — which has the effect of lowering the operating cost. This is seen when, for example, comparing the privilege-separated architecture of the Chrome web browser to the Firefox web browser: the former has a lower proportion of critical vulnerabilities [9].

The payout distribution for each level of severity is shown in Figure 2. Critical vulnerabilities have the largest variation in payouts, ranging from $200 to $200,000 per report.

B. What is the expected benefit per year?

On average, a program can expect 0.429 new valid report submissions each day and 156 valid reports over the course of a year. Newly-released programs can expect to see as many as 20 new valid reports per day. This may be due to hackers switching to new programs, hoping to find the easy initial vulnerabilities.

With an average of 156 reports per year, an organisation can expect 13 critical vulnerabilities to be discovered. The relative proportion of each severity is shown in Table IV. Critical

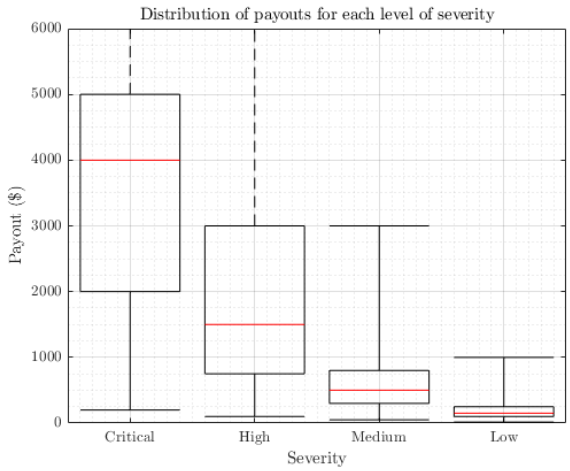

Fig. 2. Distribution of the payouts for each level of vulnerability severity. The maximum value for critical is $200,000 and for high is $15,000.

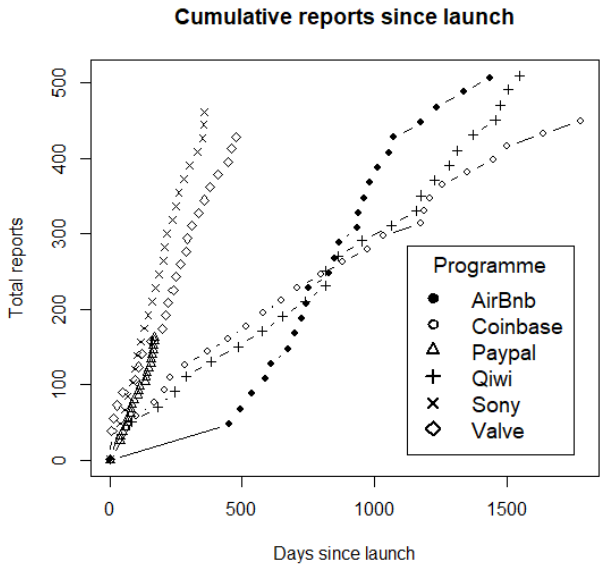

Fig. 3. Cumulative number of reports over the lifetime of six programs.

vulnerabilities are the most damaging to the organisation and often covers unauthorised access, remote code execution and privilege escalation [42].

C. How does the program activity vary across its lifetime?

Figure 3 shows the cumulative number of reports for six programs since launch. Over the lifetime of a program, the rate of vulnerability discovery remains relatively constant after initial release. Figure 4 shows the cumulative number of reports for the first 100 days after launch. Four of these companies (AirBnB, Sony, CoinBase and Paypal) see an initial high level of reports before the rate of discovery lowers.

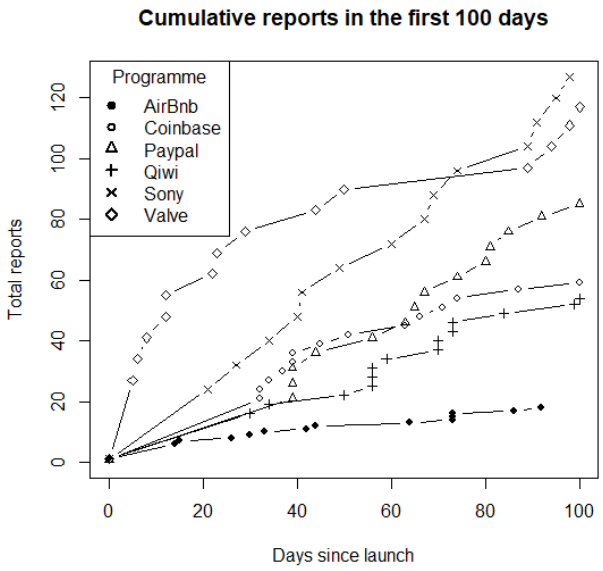

Fig. 4. Cumulative number of reports for six programs over the first 100 days.

D. Does the promise of higher bounty payouts result in an increased number of vulnerabilities being reported?

Figure 5 shows the relationship between the average program bounty and the number of reports submitted. Linear regression is used to plot a trendline with an R^2 value of 0.0527. This suggests that linearly increasing the average bounty will not result in proportionally more reports. There is a weak relationship between the average bounty of a program and the number of reports submitted, which is evidenced by the low correlation coefficient of 0.02295.

An interesting observation is that many hackers are willing to work without the promise of monetary rewards. They are instead motivated by reputation on the website, which is needed to be invited to private programs. Table IV shows that 9.6% of all reports submitted do not result in a bounty payout. Algarni et at. found that, among the top users on the site, monetary rewards are still the primary motivation [43].

E. Available programs and hacker participation

The number of new programs available on HackerOne is shown in Figure 6.

HackerOne uses reputation to measure the validity of reports submitted. Users with high reputation submit a large number of valid reports. Figure 7 shows the distribution of the reputation for the top 100 users, with these users being responsible for 40,690 vulnerability submissions out of almost 72,000. The distribution of reputation and number of vulnerabilities reported shows that a small number of users submit significantly more reports that the average user. The top five users of HackerOne have submitted a combined total of 7,121 vulnerability reports, 17.5% of vulnerabilities submitted by

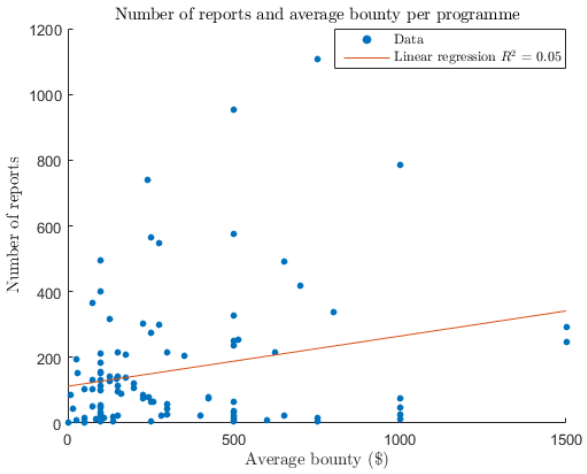

Fig. 5. The number of reports for each program plotted against the average bounty payout.

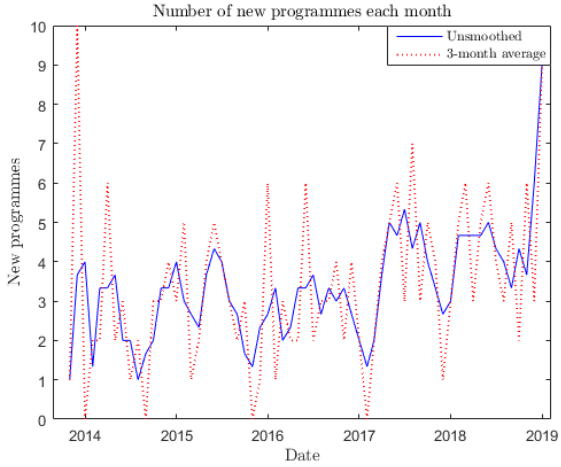

Fig. 6. Number of new programs on the site HackerOne from 2013 to 2018.

the top 100 and almost 10% of all vulnerabilities submitted to HackerOne.

BugCrowd provides detailed user activity information. It can be seen that the top 10 hackers have contributed to 118 programs on average.

HackerOne defines a standard response for organisations to adhere to when responding to hacker-submitted reports. This includes a limit of five days to respond to a report and 10 days to triage [44]. The distribution of responses meeting the standards is shown in Figure 8. On average, an organisation will meet the standard in 97% of responses. However five organisations meet this target less than 20% of the time.

F. Organisation classification

Analysing the range of businesses that use bug bounty platforms allows for any bias towards businesses of a certain size or sector to be identified.

Using the small and medium enterprise (SME) and Large business definitions [45], each organisation on HackerOne

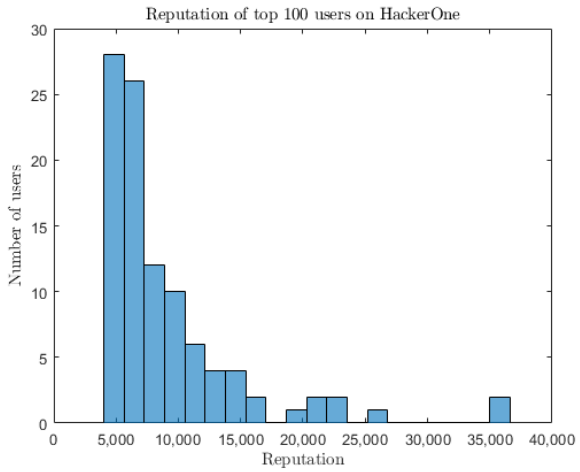

Fig. 7. Reputation of users on HackerOne.

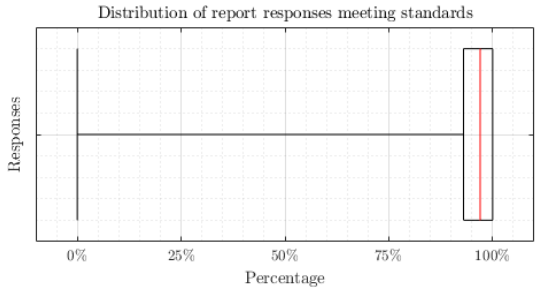

Fig. 8. Distribution of the percentage of report responses that meet the standards set by HackerOne.

and BugCrowd can be categorised into the following sizes of business: micro, small, medium and large. Data is publicly available to categorise 142 companies; the results of this are shown in Figure 9.

From the list of organisations, 193 can be placed into 18 predefined categories. This is shown in Figure 10. The remaining businesses could not be categorised due to insufficient information. Software development companies and Opensource projects account for 41% of programs listed on HackerOne and BugCrowd.

V. Analysis

This section provides analysis of the results presented in Section IV.

A. Q1: What is the expected cost per year of operating a bug bounty programs?

With an average annual cost of $83,815, the cost of running a full-time program is greater than the $65,133 required to hire an additional software engineer. The running cost is highly dependent on the success of hackers at finding vulnerabilities, as well as the rewards structure in place. Oath's program costs almost $1,000,000 per year to run, which is, in part, due to the high payout for critical vulnerabilities found over 31 web domains and within 21 Android and IOS applications.

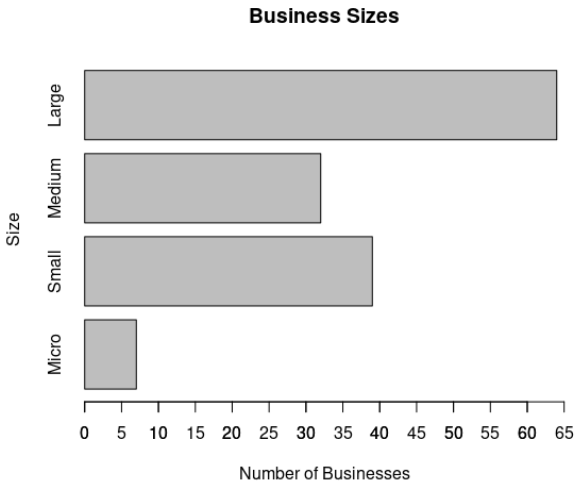

Fig. 9. Sizes of businesses on bug bounty platforms.

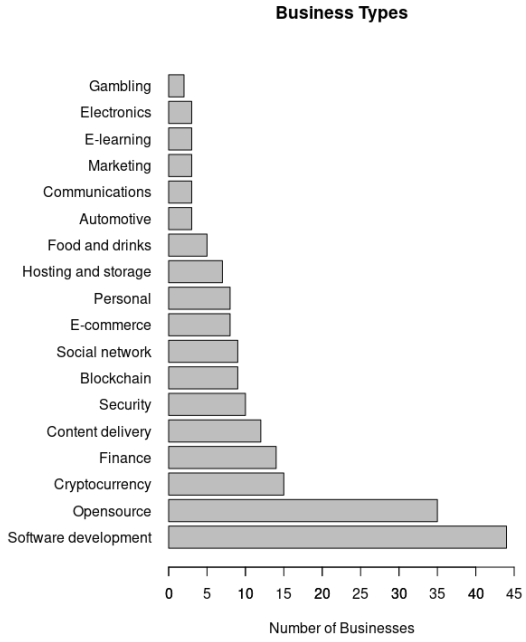

Fig. 10. Types of businesses on bug bounty platforms.

Finifter et al. found the daily operating cost to be $658 and $485 for Firefox and Chrome [9]. While these are above the current operating average of $230, they are not outliers in the current distribution of operating costs.

As the severity of vulnerabilities present in the software cannot be predicted before the program is launched, the relative proportion of critical vulnerabilities is unknown. Once a program has been launched, the rewards structure can be altered if large numbers of high-severity vulnerabilities are reported, this can be used to reduce the operating cost.

Zhao et al. gave consideration to the severity distribution

of vulnerabilities on Wooyun [30]. Only 10% of submitted vulnerabilities were classified as of low severity, compared to 38.4% of vulnerabilities on HackerOne being classified as of low or no severity. The general decrease in high severity vulnerabilities could be in part responsible for lowering the average operating costs, as the payouts are far less for low severity vulnerabilities.

Another trend is the increase in payouts for critical and high severity vulnerabilities. Firefox and Google typically paid out $3000 and $1000 for vulnerabilities of this severity [9]. The maximum bounty on HackerOne in 2015 was $7,560 by Twitter [30]. The current average payout for a critical vulnerability is $7,924 with a maximum of $200,000. Having an above average proportion of critical vulnerabilities can increase the operating cost significantly.

B. Q2: What is the expected benefit per year?

Finifter et al. found that Firefox and Google received 63 and 167 vulnerability reports respectively per year [9]. These rates are comparable to current rates seen on HackerOne and BugCrowd.

Zhao et al. reported that HackerOne saw 400 vulnerability reports per month in 2015 [30]. This has now increased to over 2,500 vulnerability reports per month across all programs. This increase could be due to an increase in the number of hackers participating in programs. HackerOne now boasts over 200,000 hackers [46].

Newly-released programs can expect an initial spike in the number of reports submitted as hackers use automated testing to find vulnerabilities caused by cross-site scripting or SQL injection. After this initial period, a stable level of 156 vulnerabilities per year can be expected.

As the number of programs available on bug bounty platforms increases, the number of hackers may continue to grow. This may result in the average number of vulnerabilities found per year increasing.

Al-Banna et al. found that a major concern of current program operators was the volume of low quality submissions received, as well as the associated high cost of processing them [47]. The use of the standardised reporting forms by organisations can encourage hackers to produce reports of higher quality [48].

C. Q3: How does the program activity vary across its lifetime?

After an initial high rate of reported vulnerabilities after the program is launched, most programs continue on a fairly constant rate of vulnerability discovery; this was also observed by Finifter et al. [9]. With an almost constant rate of discovery, the average cost of operating the program should not be expected to rise unless the average severity of vulnerabilities increases.

D. Q4: Does the promise of higher bounty payouts result in an increased number of vulnerabilities being reported?

There is a weak correlation between the success of a program and the average bounty payout. Hackers value the challenge or opportunity to learn almost as much as the bounty payouts when selecting programs in which to participate [49]. This may explain the success of programs with low or no bounty payouts. Finifter at al. also found similar results. Mozilla had a far greater average bounty payouts than Google, but found much fewer vulnerabilities [30].

The non-linear relationship between number of reports and average bounty makes it difficult to model the effect of increasing the average bounty of a program; however, the weak correlation suggests there is some benefit to increasing the average bounty.

The average bounty for all programs is $318. This can be used as a guideline for new programs when setting a rewards structure.

E. Further observations

1) Available programs: Since the study by Zhao et al. [30], HackerOne has seen continual growth in the number of programs available and the number of hackers participating. This demonstrates the viability of bug bounty programs to provide benefit to both organisation and hackers.

2) Hacker participation: The participation of hackers on BugCrowd and HackerOne is very skewed, with the top 100 users contributing a significant proportion of overall reports. This was also found by Zhao et al. [30], with each hacker in the top 100 submitting 147 reports on average. The top 100 have now submitted 406 reports on average.

3) Standard response: Organisations meet the standard response guidelines set out by HackerOne in 97% of responses. Adhering to these guidelines ensures that hackers receive a reward promptly and organisations are able to mitigate any vulnerabilities identified. Large volumes of poor quality or invalid reports will reduce the effectiveness of a organisation to meet the standards. A theoretical framework introduced by Laszka et al. aims to reduce the numbers of these reports [50].

4) Size of organisation: Although small and medium businesses do not typically have the security resources available to large organisations, SMEs make up 50% of businesses using bug bounty platforms (compared to 99.3% of businesses in the UK [51]). This demonstrates that operating a bug bounty program may be viable for a business, regardless of size.

5) Type of organisation: As shown in Figure 10, there is a wide range of sectors currently operating bug bounty programs. Open-source projects make up 18.1% of listed programs on bug bounty platforms, with many of these projects being funded by governments in order to improve the security of popular open-source software. The success of the European Commission's Directorate-General for Informatics (DIGIT) funded programs has the potential to lead to an increase in the number of open-source programs [52].

VI. Limitations

Without quantifying the productivity of software engineers, the usefulness of comparing the cost of employing additional software engineers to the cost of operating a bug bounty program is limited. However organisations will be able to

compare the efficiency of their software engineers with the results detailed within this paper. The productivity of software engineers can be measured by the number of vulnerabilities they find per year, the chance to detect vulnerabilities, or the average number of vulnerabilities they introduce per line of code written.

Of course, the previous experience and education of an employee will influence their effectiveness at finding or preventing the existence of certain types of vulnerability. As such, we have to be cautious when comparing the 'average' software engineer with the 'average' hacker.

In comparison, hackers from all technical backgrounds will participate in bug bounty programs. The types of vulnerabilities found are not limited to those relating to the skill set of the software engineers. For organisations with technical skill gaps, there is an additional benefit to the search for vulnerabilities via bug bounty programs; however, this is not quantified here.

VII. A Research Agenda

We now outline a research agenda for future investigations into bug bounty programs. Clearly, the research agenda is driven by the needs of those in managerial positions responsible for decision-making in the area of bug fixing. Further research areas are identified in Sections VII-A–VII-D.

A. Interviews with organisations

Not all programs currently use a CVSS score to assign a bounty payout to a given vulnerability. Furthermore, Munaiah et al. identified that there is a weak correlation between CVSS score and the bounty payout for a vulnerability [41]. Conducting interviews with organisations that currently operate bug bounty programs may provide insight into the process used by operators to set rewards structures — including how hackers might be incentivised to search for the often more complex high and critical severity vulnerabilities.

B. Behavioural study

The study by Zhao et al. proposed a simple method to classify hackers into three productivity groups: high, medium and low [30]. Using the hacker data readily available on BugCrowd, this classification could be extended to find the current number of high productivity hackers. Tracking this statistic over time could be used to determine changes in hackers attitude towards bug bounty programs. Decreasing numbers of high-productivity hackers might suggest that full-time participation in programs does not provide adequate rewards. This concern is raised by Pentland et al. [53].

A preliminary behavioural studied was undertaken by Hunag et al., looking into the number of programs each hacker would participate in [54]. Expanding upon this work could reveal if there have been any behavioural changes in recent years.

Research into the behaviour of hackers might involve questions such as *How does the productivity of hackers change over time?*, *How much does each productivity group contribute to each program?*, and *Does the change in average payouts affect the motivation of hackers?*

C. Program changes

Over the lifetime of a bug bounty program any changes to the rewards structure are visible on the program page. This includes any changes to minimum bounty, total bounty and top bounty range, as well as any changes to payouts. In order to build robust economic models, (at least) the following questions need to be considered: *How often do organisations alter the program rewards structure, and how do these changes affect the rate of vulnerability discovery?*, *With the increasing number of programs available on platforms such as HackerOne, do organisations need to increase the payouts in order to attract hackers?*, and *How do the average bounties change over time?*

D. Correlation of major vulnerabilities

When studying the activity on a bug bounty platform over time, it may be possible to correlate the disclosure of major vulnerabilities with an increased rate of vulnerability discovery and predict the effects of future disclosures. The importance of being able to model future vulnerability discovery trends was highlighted by Zhao et al. [55], [56].

The Spectre vulnerability in 2018 might have led to an increase in the number of remote exploitation vulnerabilities discovered [57]. This might have resulted in an increased number of critical vulnerabilities being reported.

The following questions could be the focus of a study into the correlation of major vulnerability disclosure and bug bounty activity: *Does the disclosure of a major vulnerability result in an increased number of accepted reports?* and *Can previous trends in vulnerability disclosure be used to predicted future increases in the running costs of bug bounty programs?*

VIII. Conclusion

This empirical study has looked into the benefits of operating a bug bounty program, and compared the operating costs to that of hiring additional software engineers. Analysis of existing programs has shown that the average cost of operating a program for a year is now less than the cost of hiring two additional software engineers.

The rate of vulnerability discovery was investigated over the lifetime of a program. It was found that hackers are able to discover vulnerabilities at a fairly constant rate over the lifetime of a program. Further work is needed to determine the impact of changing bounty payouts on the rate of discovery.

In conclusion, we would argue that bug bounty programs can be seen as a valuable complementary technique that can be deployed in an organisation's search for vulnerabilities.

Acknowledgements

The authors would like to thank the anonymous reviewers for their helpful comments.

The research described in this paper was undertaken as part of the "Data and models for secure software engineering" project, which is funded by the UK's National Cyber Security Centre.

REFERENCES

[1] H. Fryer and E. Simperl, "Web science challenges in researching bug bounties," in *Proceedings of the 2017 ACM on Web Science Conference*. ACM, 2017, pp. 273–277.

[2] Euro News. (2019, Feb) Euro News: Switzerland paying e-voting hackers. [Online]. Available: https://www.euronews.com/2019/02/13/switzerland-offers-cash-to-hackers-who-can-crack-its-e-voting-system

[3] Department of Defense. (2018, Oct) United States Department of Defense: Expanding hack the Pentagon. [Online]. Available: https://dod.defense.gov/News/News-Releases/News-Release-View/Article/1671231/department-of-defense-expands-hack-the-pentagon-crowdsourced-digital-defense-pr/

[4] AI Trends. (2019, Feb) AI trends: Bug bounties and AI systems: The case of AI self-driving cars. [Online]. Available: https://www.aitrends.com/ai-insider/bug-bounties-and-ai-systems-the-case-of-ai-self-driving-cars/

[5] HackerOne. (2016, Jul) HackerOne: Hacking the pentagon. [Online]. Available: https://www.hackerone.com/blog/hack-the-pentagon-results

[6] A. T. Chatfield and C. G. Reddick, "Cybersecurity innovation in government: A case study of US Pentagon's vulnerability reward program," in *Proceedings of the 18th Annual International Conference on Digital Government Research*. ACM, 2017, pp. 64–73.

[7] A. Schwartz, R. Knake, and Belfer Center for Science and International Affairs, *Government's Role in Vulnerability Disclosure: Creating a Permanent and Accountable Vulnerability Equities Process*. Harvard Kennedy School, Belfer Center for Science and International Affairs, 2016.

[8] HackerOne. (2019, Jan) HackerOne: Bug bounty program directory. [Online]. Available: https://hackerone.com/directory

[9] M. Finifter, D. Akhawe, and D. Wagner, "An empirical study of vulnerability rewards programs," in *USENIX Security Symposium*, 2013, pp. 273–288.

[10] Glassdoor. (2019, Jan) Glassdoor: Software engineer salaries in london. [Online]. Available: https://www.glassdoor.co.uk/Salaries/london-software-engineer-salary.htm

[11] M. Blatter, S. Muehlemann, and S. Schenker, "The costs of hiring skilled workers," *European Economic Review*, vol. 56, no. 1, pp. 20–35, 2012.

[12] P. Foreman, *Vulnerability management*. Auerbach Publications, 2009.

[13] BSIMM. (2018, Oct) BSIMM9: Building security in maturity model version 9. [Online]. Available: https://www.bsimm.com/download/

[14] BugCrowd. (2018, Aug) BugCrowd: Integrating crowdsourced security with the SDLC. [Online]. Available: https://www.bugcrowd.com/blog/integrating-crowdsourced-security-with-the-software-development-lifecycle/

[15] Google. (2019, Jan) Google: Vulnerability reward program. [Online]. Available: https://www.google.com/about/appsecurity/reward-program/index.html

[16] Facebook. (2018, Sep) Facebook: Whitehat program. [Online]. Available: https://www.facebook.com/whitehat

[17] Microsoft. (2018, Jul) Microsoft: Bug bounty programs. [Online]. Available: https://www.microsoft.com/en-us/msrc/bounty

[18] HackerOne. (2019, Jan) HackerOne: HackerOne bug bounty platform. [Online]. Available: https://www.hackerone.com/

[19] BugCrowd. (2019, Jan) BugCrowd: BugCrowd bug bounty platform. [Online]. Available: https://www.bugcrowd.com

[20] Cobalt. (2019, Jan) Cobalt: Bug bounty platform. [Online]. Available: https://cobalt.io/

[21] A. Kuehn and M. Mueller, "Analyzing bug bounty programs: An institutional perspective on the economics of software vulnerabilities," in *TPRC, the 42nd Research Conference on Communication, Information and Internet Policy*, 2014.

[22] E. Raymond, "The cathedral and the bazaar," *Knowledge, Technology & Policy*, vol. 12, no. 3, pp. 23–49, 1999.

[23] M. Rouse. (2007, Jun) Search Security: Whitehat hacker definition. [Online]. Available: https://searchsecurity.techtarget.com/definition/white-hat

[24] A. Laszka, M. Zhao, A. Malbari, and J. Grossklags, "The rules of engagement for bug bounty programs," in *International Conference on Financial Cryptography and Data Security*, 2018.

[25] M. W. Kranenbarg, T. J. Holt, and J. van der Ham, "Don't shoot the messenger! a criminological and computer science perspective on coordinated vulnerability disclosure," *Crime Science*, vol. 7, no. 1, pp. 16–17, 2018.

[26] E. Kovacs. (2018, Jul) SecurityWeek: Intel's record bounty. [Online]. Available: https://www.securityweek.com/intel-pays-100000-bounty-new-spectre-variants

[27] L. Allodi, "Economic factors of vulnerability trade and exploitation," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2017, pp. 1483–1499.

[28] D. Votipka, R. Stevens, E. Redmiles, J. Hu, and M. Mazurek, "Hackers vs. testers: A comparison of software vulnerability discovery processes," in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 374–391.

[29] T. D. LaToza and A. van der Hoek, "Crowdsourcing in software engineering: Models, motivations, and challenges," *IEEE software*, vol. 33, no. 1, pp. 74–80, 2016.

[30] M. Zhao, J. Grossklags, and P. Liu, "An empirical study of web vulnerability discovery ecosystems," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 1105–1117.

[31] Wooyun. (2016, Jun) Wooyun: Vulnerability disclosure platform. [Online]. Available: https://www.wooyun.org

[32] P. Paganini. (2016, Aug) Security Affairs: Wooyun group arrested. [Online]. Available: https://securityaffairs.co/wordpress/49916/hacking/wooyun-group-arrested.html

[33] Google. (2019, Jan) Google: Chromium bug tracker. [Online]. Available: https://www.chromium.org

[34] Google. (2019, Mar) Google Trends: Search popularity of bug bounty platforms. [Online]. Available: https://trends.google.com/trends/explore?q=HackerOne,bugcrowd,safehats,Intigriti,Synack

[35] Synack. (2019, Mar) Synack: Bug bounty platform. [Online]. Available: https://www.synack.com/solution/

[36] R. Böhme, "A comparison of market approaches to software vulnerability disclosure," in *Müller G. (eds) Emerging Trends in Information and Communication Security. ETRICS 2006. Lecture Notes in Computer Science, vol 3995*. Springer, Berlin, Heidelberg, 2006, pp. 298–311.

[37] J. L. Christian, "Bug bounty programs: Analyzing the future of vulnerability research," Ph.D. dissertation, Utica College, 2018.

[38] J. Ruohonen and L. Allodi, "A bug bounty perspective on the disclosure of web vulnerabilities," *arXiv preprint arXiv:1805.09850*, 2018.

[39] First. (2019, Mar) First: User guide to CVSS 3.0. [Online]. Available: https://www.first.org/cvss/user-guide

[40] HackerOne. (2019, Mar) HackerOne: Paypal program page. [Online]. Available: https://hackerone.com/paypal

[41] N. Munaiah and A. Meneely, "Vulnerability severity scoring and bounties: Why the disconnect?" in *Proceedings of the 2nd International Workshop on Software Analytics*. ACM, 2016, pp. 8–14.

[42] H. Homaei and H. R. Shahriari, "Seven years of software vulnerabilities: The ebb and flow," *IEEE Security & Privacy*, vol. 15, no. 1, pp. 58–65, 2017.

[43] A. M. Algarni and Y. K. Malaiya, "Most successful vulnerability discoverers: Motivation and methods," in *Proceedings of the International Conference on Security and Management (SAM)*. The Steering Committee of The World Congress in Computer Science, Computer , 2013, p. 1.

[44] HackerOne. (2019, Feb) HackerOne: Response target metrics. [Online]. Available: https://docs.hackerone.com/programs/response-target-metrics.html

[45] European Commission. (2015, Sep) Publications office of the European Union: User guide to the SME definition. [Online]. Available: http://www.innovation.public.lu/en/brochures-rapports/s/sme-definition/sme-definition-user-guide-2015.pdf

[46] HackerOne. (2019, Feb) HackerOne: 188 fascinating facts. [Online]. Available: https://www.hackerone.com/blog/118-Fascinating-Facts-HackerOnes-Hacker-Powered-Security-Report-2018

[47] M. Al-Banna, B. Benatallah, D. Schlagwein, M. C. Barukh, and E. Bertino, "Friendly hackers to the rescue: How organizations perceive crowdsourced vulnerability discovery," in *Pacific Asia Conference on Information Systems (PACIS)*, 2018.

[48] D. Mu, A. Cuevas, L. Yang, H. Hu, X. Xing, B. Mao, and G. Wang, "Understanding the reproducibility of crowd-reported security vulnerabilities," in *27th USENIX Security Symposium, USENIX Security 18)*, 2018, pp. 919–936.

[49] HackerOne. (2018, Nov) HackerOne: 2018 hacker report. [Online]. Available: https://www.hackerone.com/sites/default/files/2018-01/2018-Hacker-Report

[50] A. Laszka, M. Zhao, and J. Grossklags, "Banishing misaligned incentives for validating reports in bug-bounty platforms," in *European Symposium on Research in Computer Security*. Springer, 2016, pp. 161–178.

[51] BIS. (2018, Oct) Department for Business Innovation and Skills: Business population estimates for the uk and regions 2018. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/746599/OFFICIAL_SENSITIVE_-_BPE_2018_-_statistical_release_FINAL_FINAL.pdf

[52] DIGIT. (2019, Jan) Directorate-general for informatics: Commission launches open source bug bounties. [Online]. Available: https://ec.europa.eu/newsroom/informatics/item-detail.cfm

[53] A. Pentland, D. L. Shrier, and H. E. Shrobe, *New Solutions for Cybersecurity*. MIT Connection Science and Engineering, 2018.

[54] K. Huang, M. Siegel, S. Madnick, X. Li, and Z. Feng, "Diversity or concentration? hackers' strategy for working across multiple bug bounty programs," in *37th IEEE Symposium on Security and Privacy (S&P)*, 2016.

[55] M. Zhao, J. Grossklags, and K. Chen, "An exploratory study of white hat behaviors in a web vulnerability disclosure program," in *Proceedings of the 2014 ACM workshop on security information workers*. ACM, 2014, pp. 51–58.

[56] M. Zhao, A. Laszka, T. Maillart, and J. Grossklags, "Crowdsourced security vulnerability discovery: Modeling and organizing bug-bounty programs," in *The HCOMP Workshop on Mathematical Foundations of Human Computation, Austin, TX, USA*, 2016.

[57] K. Afifi-Sabet. (2019, Feb) IT Pro: Meltdown and spectre. [Online]. Available: https://www.itpro.co.uk/exploits/30478/what-are-meltdown-and-spectre-and-are-you-affected