The Hackers' Viewpoint: Exploring Challenges and Benefits of Bug-Bounty Programs

Omer Akgul
University of Maryland

Taha Eghtesad
University of Houston

Amit Elazari
UC Berkeley

Omprakash Gnawali
University of Houston

Jens Grossklags
Technical University of Munich

Daniel Votipka
University of Maryland

Aron Laszka
University of Houston

Abstract

In recent years, bug-bounty programs have garnered popularity and became a significant part of the security culture of many organizations. Bug-bounty programs enable these organizations to improve their security posture by harnessing the outside perspective of a diverse crowd of security experts (bug hunters). However, bug-bounty programs also suffer from inefficiencies, such as duplicate and invalid bug reports, which are resource consuming for organizations and bug hunters alike. To address these issues, it is crucial to understand how bug hunters make decisions, what motivates them, and what challenges they face. We present the results of an initial survey conducted among bug hunters to address these questions. We recruited 56 security experts who participate in bug-bounty programs to answer open-ended questions regarding various aspects of their participation in bug-bounty programs. Their responses provide a detailed overview of the motivations of security experts and the challenges that they face.

1 Introduction

Traditionally, organizations relied on the work of internal security experts (e.g., security testing conducted by red teams) and outsourced experts (e.g., pentesting) to discover vulnerabilities in their products and services. In contrast, bug-bounty programs—also known as vulnerability-reward programs or "crowd-sourced" security—provide incentives to external security experts to evaluate the security of an organization's products and services in scope, and to report vulnerabilities in exchange for rewards (financial or otherwise, such as recognition). Bug bounties differ from vulnerability disclosure programs (i.e., programs that promote non-incentive based disclosure), and their value has recently received both growing business and policy (regulatory proposals) recognition [5]. Spearheaded by Netscape as a forerunner in 1995, now many large technology companies such as Google, Intel, Facebook, and Microsoft run bug-bounty programs, which have garnered popularity and became a significant part of the security culture in many organizations due to their multifaceted benefits. For example, they may enable organizations to improve their security posture by harnessing the diverse expertise of large crowds of security experts and support recruiting.

However, due to their crowd-sourced nature, bug-bounty programs also suffer from inefficiencies. Organizations that run bug-bounty programs may receive many invalid or duplicate reports (i.e., multiple bug hunters reporting the same bug), which are time and resource consuming [18]. Further, there exists competition between programs to attract the attention of productive bug hunters [12], and the security impact and financial costs of bug-bounty programs are uncertain in advance. On the other hand, bug hunters face similar uncertainties regarding their findings and rewards, and they are often disappointed by the responses of programs due their reports being marked as having lower impact or being duplicates, for example. To mitigate some of these issues, bug-bounty platforms such as Bugcrowd[1] and HackerOne[2] have emerged, which act as a marketplace and facilitator connecting organizations and bug hunters. However, even on platforms, many issues persist.

To improve bug-bounty programs and platforms and to address these exigent issues, we must understand how bug hunters work, what motivates them, and what challenges they face [9]. Indeed, a number of research efforts have investigated bug-bounty programs (e.g., Finifter et al. [6], Zhao et al. [17], Maillart et al. [12], Laszka et al. [10], Luna et al. [11], Elazari [4], Walshe and Simpson [16]). Nonetheless, a common limitation of these prior efforts is that they consider only data collected from the perspective of bug-bounty programs (e.g., vulnerability reports, bounty payments, program policies and terms). As a result, while they may offer a clear picture of how organizations work, they provide only a limited view of bug hunters' work, which considers merely their output and neglects their motivations and the challenges that they face. In related work, Votipka et al. interviewed hackers to understand how they discover vulnerabilities; however, their study did not consider bug-bounty programs [15]. In light of this, there

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

is a clear and relevant research gap regarding bug hunters' views on bug-bounty programs. Bug-bounty platforms have reported on this gap ([2, 3, 7]), focusing on bug hunter motivations. However, these reports do not consider challenges faced by bug hunters, appear to be crafted mostly for marketing purposes, and are not verified by independent entities.

Research Questions As a first step toward addressing this gap, we present the results of an initial survey we conducted among bug hunters to understand their motivations and the challenges they face. We specifically investigate the bug-bounty ecosystem from the viewpoint of bug hunters based on the following questions.

- **RQ1:** How do bug hunters choose which specific bug-bounty program to work on?
- **RQ2:** What makes bug hunters stop working on a particular bug-bounty program?
- **RQ3:** What are the main benefits of working on bug bounty in general?
- **RQ4:** What are the main challenges that bug hunters face when working on bug-bounty programs?
- **RQ5:** What are the most useful features of bug-bounty platforms, such as HackerOne and Bugcrowd, for bug hunters?

To explore these questions, we first conducted a survey that asked participants to answer open-ended questions regarding various aspects of working on bug-bounty programs. We then analyzed their responses to identify the breadth of important factors and issues in "bug hunting." The results of the survey provide a detailed overview of security experts' motivations and the challenges that they face.

Organization The remainder of this paper is organized as follows. We present our research methodology in Section 2, followed by the description of our participants' demographics and experience levels in Section 3, and our results in Section 4. Finally, we briefly discuss our findings, limitations and future work in Section 5.

2 Methodology

To identify the factors that influence participation in bug bounties, we performed an online open-ended survey ($n = 56$). This section describes our survey, recruitment process, and data analysis methodology in detail. The study was collaboratively approved by our institutions' ethics review boards.

Survey The survey comprised of three main sections. The first section of the survey included open-ended questions asking participants to list factors they consider when deciding to participate in bug bounties. We chose to use an open-ended listing approach, called *free listing*, which is common in anthropological research when the domain is not well understood [1].

The following five questions each focused on one of our research questions:

- **Choosing a program** (RQ1): What are all the factors you consider when deciding in which bug-bounty program to participate?
- **Leaving a program** (RQ2): What are all the issues that could make you stop working on a particular bug-bounty program?
- **Benefits of bug hunting** (RQ3): What are all the benefits of working on bug-bounty programs for you? [3]
- **Challenges of bug hunting** (RQ4): What are all the challenges that you face working on bug bounty? What factors make working on bug bounty difficult for you?
- **Useful platform Features** (RQ5): What are the most useful features of bug-bounty platforms for you?

For each question, we stressed that survey participants should list all factors they may consider even if they are not always considered at every decision point. Additionally, we asked survey participants to spend time continuing to recall factors if they thought there might be more they could remember. This is a common prompt in listing exercises, with the goal of eliciting less common responses [1].

Next, we asked survey participants to self-report their bug-hunting skill and experience (e.g., estimated number of vulnerabilities discovered, revenue earned, and amount of time spent on bug-bounty programs) and concluded with several demographic questions (e.g., educational background, age group, country of residence) to understand our sample population's makeup.

Pilots Through personal connections, we recruited three security experts, who regularly work on bug-bounty programs, to participate in a pilot to evaluate the phrasing of our questions. First, we sent them the survey and gave them several days to complete it. Then, we discussed their responses with them in a focus group session, which we conducted via teleconferencing. We proceeded with our main data collection after confirming the questions were clear and provided good construct validity.

Recruitment We recruited participants through advertising on social media (through the authors' personal accounts) as well as mailing lists and private Slack channels used by members of the bug-bounty community. We continued recruitment until we had evidence to believe we had sufficiently saturated the factor space for each question. Specifically, we made sure the frequency of new unique items listed in responses had dropped significantly compared to earlier responses.

[3] Our goal here is to understand why bug hunters continue to participate in the overall bug-bounty marketplace.

In total, we received over 61 complete responses to the survey. We removed 5 responses due to poor quality (e.g., unintelligible answers, unreasonably fast completion times, or obviously duplicated responses), leaving us with 56 responses for analysis. The survey was active for 7 months (May – December 2019). The last 15 responses largely confirmed the factors identified in the first 41 responses, indicating saturation of factors, and thus we concluded our data collection.

Data Analysis We analyzed our open-ended survey responses with exploratory, inductive qualitative coding procedures. Despite the exploratory nature, we expect future work to directly use the factors we identified. Therefore, we chose to calculate inter-coder reliability metrics [13]. We report descriptive statistics for participant demographics and technical skills.

We developed the codebook and established reliability on the first 41 responses (73.2% of all responses, 64.0–78.2% of the total factors listed[4]). We use the same procedure described below for each open-ended question.

The initial codebook was developed by three researchers on 10 participant responses (∼25% of the responses at the time; 15.3–28.0% of all factors listed). Two of the three researchers then attempted to establish good reliability by independently coding batches of 10 responses at a time using this codebook. After each batch, we resolved differences, updated the codebook, and reiterated the assigned codes. We ran out of new responses without being able to establish good reliability. However, after some time had passed (∼30 days), the researchers revised the codebook and two researchers independently re-coded 16 out of 41 responses (∼40% of the responses at the time; 27.0–38.3% of the all factors listed) for each question and achieved "almost perfect" [8] reliability (measured with Cohen's K, given in Table 3). Finally, with reliability established, one researcher re-coded the rest of the responses and another sampled some of the responses for further confirmation. The final 15 responses were coded by one researcher with the established and validated codebook.

3 Participants

We received 56 valid responses from security researchers participating in bug-bounty programs. Participants were mainly from North America, South Asia, and Europe; young in age; and overwhelmingly male (see Table 1).

As observable in Table 2, survey participants generally described themselves as closer to "Expert" than "Beginner" in bug hunting (median 3 on a 1–5 scale, 4-5 being twice as frequent as 1-2); they reported finding more than 10 bugs that received bounties, generally had less than or equal to 3 years of experience, and typically spent 5-10 hours per week on bug-bounty hunting.

[4]Responses are *unitized* based on number of factors listed in a response (i.e., codes are assigned to individual factors, not entire responses). Responses may include multiple factors; thus, percentage of responses is not always the same as the percentage of all listed factors.

Gender	Male	51
	Female	2
	Other	1
Age	18-29	34
	30-39	16
	40-40	4
	50-59	2
Residence	North America	21
	South Asia	14
	Western Europe	10
	Eastern Europe	2
	Southeast Asia	2
	South America	2
	Africa	2
Education	Completed H.S. or below	18
	Some college, no degree	13
	Trade/technical/vocational training	1
	Associate's degree	3
	Bachelor's degree	15
	Professional degree	1
	Master's degree	4

Table 1: Participant demographics. Numbers might not add up to total participant number due to "other" and "prefer not to answer" options.

There is no ground truth for the demographics of bug hunters; however, the demographics of our survey participants are similar to bug-hunter demographics reported by popular bug-bounty platforms [2, 3, 7].

4 Results

In this section, we present the results of our analysis of the survey responses. We discuss the most prevalent factors mentioned by the survey participants. We further note how many times each factor was mentioned for context (see Table 3). We also analyzed the ordering of the factors given by the participants, and calculated their average rank (see Table 3). Note that these numbers by themselves might not accurately convey the perceived importance of factors.

RQ1: How do bug hunters choose which bug-bounty programs to work on? Our analysis indicates that bug hunters are generally motivated by the program's promised rewards ($n = 36$, 64.3%) and by the payout distribution based on the reported bug's severity ($n = 16$, 28.6%). Survey participants were also encouraged by wide program scopes ($n = 28$, 50.0%), familiarity with the technology of the assets in scope (iOS app, hardware, web, etc.) ($n = 22$, 39.3%), and technical

Skill level	5 (Expert)	12
	4	10
	3	23
	2	8
	1 (Beginner)	3
Number of vulnerabilities discovered	0	4
	1-10	8
	11-50	19
	51-100	7
	101-500	12
	> 500	4
	Other	2
Years of experience	≤ 1	19
	1-3	20
	3-5	10
	$5 <$	7
Hours spent per week on bug bounty	< 5	13
	5-10	18
	10-20	12
	20-30	8
	30-40	1
	> 40	4
Hours spent per week on technical tasks not related to bug bounty	< 5	14
	5-10	7
	10-20	3
	20-30	8
	30-40	7
	> 40	17

Table 2: Participant skill levels and experience with bug-bounty hunting. All sections refer to metrics on bug-bounty hunting except the last one. Data is self report and largely consists of estimates by survey participants. Numbers might not add up to the total number of participants due to "other" and "prefer not to answer" options. For "Number of vulnerabilities discovered" and "Years of experience", buckets were generated by the authors.

challenge promised by the assets in scope ($n = 6$, 10.7%). Further, survey participants valued the program's reputation ($n = 15$, 26.8%); they decided to participate in programs based on how responsive the program managers were ($n = 21$, 37.5%) and how familiar they were with the company or product they were investigating ($n = 15$, 26.8%). Some were also interested in the business domain of the company ($n = 2$, 3.6%) or the country it operated in ($n = 1$, 1.8%).

Although not as prevalent, some survey participants noted they were motivated by how long the programs were running ($n = 6$, 10.7%) and how saturated they seemed to be with bug reports ($n = 8$, 14.3%). Some chose to participate in bug-bounty programs depending on whether they were private (invite only) or public (open to all) ($n = 4$, 7.1%). Others cared about public disclosure policies ($n = 6$, 10.7%) and legal safe harbor guarantees ($n = 4$, 7.1%). Only a few mentioned explicitly that they sought career opportunities from bug-bounty programs ($n = 2$, 3.6%).

RQ2: What makes bug hunters stop working on a bug-bounty program? The most common reasons for leaving a bug-bounty program are all related to communication issues: slow (or lack of) responses to bug reports and messages ($n = 30$, 53.6%), dissatisfaction with their report's classification (different severity level than expected, disagreements with duplicates, etc.) ($n = 26$, 46.4%), and difficulty communicating and working with bug-bounty program managers ($n = 23$, 41.1%).

We also observe that survey participants sometimes get bored of hunting bugs on specific programs or switch to programs that might be more interesting ($n = 8$, 14.3%). Similarly, our survey participants noted that they would leave bug-bounty programs due to assets being too secure ($n = 5$, 8.9%) or due to assets being based on technologies which survey participants were not familiar with ($n = 1$, 1.8%) or interested in ($n = 3$, 5.6%).

Some bug hunters indicated they left programs which were more saturated with other active bug hunters ($n = 5$, 8.9%). Others noted too many duplicates ($n = 4$, 7.1%), limited scope ($n = 3$, 5.6%), and the age of the program ($n = 2$, 3.6%) under reasons for dropping out of programs.

A few of our survey participants had issues with legal threats ($n = 2$, 3.6%), disagreements with disclosure policies ($n = 2$, 3.6%), poor bug-bounty platform support ($n = 1$, 1.8%) among many other less frequent issues (see Table 3).

RQ3: How do bug hunters benefit from working on bug bounty? Consistent with our findings on choosing between programs, the most commonly listed benefit is receiving monetary rewards ($n = 42$, 75%). Interestingly, survey participants also largely saw bug-bounty hunting as a training environment ($n = 32$, 57.1%) that might help with career development ($n = 11$, 19.6%) and with building reputation in the community ($n = 14$, 25%). Some also mentioned the benefits of legal safe harbors when hacking ($n = 4$, 7.1%).

Compared to conventional work environments, survey participants benefited from the flexibility (i.e., hours and schedule) of bug-bounty work ($n = 16$, 28.6%). A significant fraction of survey participants also noted that they simply enjoy working on bug bounties ($n = 20$, 35.7%), while few said that they have altruistic reasons for participating in bug-bounty programs ($n = 5$, 8.9%).

Although the prevalence levels are different, our findings are consistent with those of HackerOne [7] and Bugcrowd [3] with respect to the benefits of bug hunting and the reasons why bug hunters choose particular programs.

Question	Code Name	Description	#	Rank
Choosing a program	**Reward:**	expected monetary or non-monetary rewards (e.g., SWAG, hardware, subscription).	36	1.94
	Scope:	number of domains or assets that are included in the program.	28	2.07
	Technology in scope:	familiarity with or interest in the technology of the assets (e.g., web, iOS).	22	3.36
	Responsiveness:	how fast and effectively program managers communicate with hackers.	21	2.67
	Bounty table:	reward rules and ranges set by the managers (e.g., $50 for low-criticality bugs, but $5000 for high-criticality bugs).	16	2.25
	Company familiarity:	company behind the program is widely known, or you or your peers use its products or services.	15	2.73
	Program repute:	program's reputation in the community for being pleasant to work with (i.e., what other hackers say about the program).	15	2.93
	Saturation:	number of reports received or number of hackers working on the program.	8	3.63
	Public disclosure:	public vulnerability disclosure is generally allowed following the resolution of the issue, permissive NDAs.	6	4.00
	Age:	for how long the program has been running.	6	4.00
	Technical challenge:	intellectually challenging or stimulating assets.	6	4.83
	Legal safe harbor:	program is committed to not pursue legal actions after hackers who follow the rules and/or explicitly authorizes testing in accordance with the rules.	4	4.50
	Private or public:	private programs (accessible only by invitation) vs. public programs (accessible by anyone).	4	3.25
	Career opportunities:	future career opportunities with the company.	3	4.33
	Business domain:	business domain of the company behind the program (e.g., social media, insurance, medical).	2	2.50
	Country:	where the company behind the program is located.	1	3.00
Leaving a program	**Responsiveness:**	how fast and effectively program managers communicate with hackers.	30	2.03
	Dissatisfaction with responses:	rewards are lower than promised by rules (e.g., downgraded severity, impact, disagreements about duplicates).	26	2.12
	Difficulty working with managers:	program managers are difficult to work with (e.g., disrespectful, requiring extra work).	23	2.39
	Boredom:	bored of working on the program or a more interesting program launches.	8	2.00
	Secure assets:	finding bugs is too difficult.	5	1.0
	Reward:	fair but unsatisfying monetary or non-monetary rewards.	5	1.40
	Saturation:	number of reports received or number of hackers working on the program.	5	1.80
	Duplicates:	too many reports marked as duplicates.	4	3.5
	Technology in scope:	familiarity with or interest in the technology of the assets (e.g., web, iOS).	3	1.00
	Scope:	number of domains or assets that are included in the program.	3	3.00
	Lacking communication or language skills:	communication difficulties because you feel that you lack language skills, experience anxiety in communication, etc.	2	3.0
	Legal threats:	fear of threats of legal implication (civil or criminal).	2	3.0
	Age:	for how long the program has been running.	2	1.00
	Limited vulnerability disclosure:	restrictive vulnerability disclosure policies and NDAs that may prevent you from publishing your work following the resolution/mitigation of the issue.	2	2.5
	Not enough time:	not having enough time for participating in bug bounty.	2	1.00
	Program repute:	program's reputation for being pleasant to work with (e.g., fair, responsive)	1	1.0
	Unrepresentative reputation system	hackers' reputation points do not reflect real experience and are not transferable between platforms.	1	1.0
	Poor platform support	dissatisfaction with how platforms handle issues, such as mediating between hackers and programs.	1	1.0
Benefits of bug hunting	**Monetary rewards:**	monetary compensation.	42	1.67
	Learning/improving:	learning or improving skills.	32	1.88
	Enjoyment:	enjoyment or challenge of white-hat hacking.	20	2.70
	Flexibility:	work schedule and place flexibility (compared to traditional employment).	16	1.38
	Reputation:	earning platform reputation points, building a following, etc.	14	2.71
	Career:	building relations and reputation with companies for employment and other work opportunities.	11	3.18
	Altruism:	improving cybersecurity for the sake of helping others, hacking to make the internet safer for everyone.	5	4.4
	Legal safe harbour:	hacking without the threat of legal actions if they obey the rules.	4	2.5
	Non-monetary rewards:	non-monetary compensation (e.g., SWAG, hardware, subscriptions).	4	4.25
	Community:	bug bounty creates a community of hackers.	3	2.67
Challenges of bug hunting	**Responsiveness:**	how fast and effectively program managers communicate with hackers.	16	1.93
	Too much labor work:	menial tasks (e.g., CAPTCHA, waiting for timeouts, obfuscation, setting up test accounts)	12	2.08
	Assets outside expertise	assets are outside area of expertise, lacking certain required skills.	11	1.5
	Difficulty working with managers:	program managers are difficult to work with (e.g., disrespectful, requiring extra work).	9	1.89
	Dissatisfaction with responses:	rewards are lower than promised by rules (e.g., downgraded severity, impact, disagreements with duplicates).	9	2.33
	Duplicates:	too many reports marked as duplicates.	8	1.63
	Secure assets:	finding bugs is too difficult.	7	1.86
	Not enough time:	not having enough time for participating in bug bounty.	6	1.17
	Unclear scope:	program scope is not defined clearly.	6	1.83
	Saturation:	number of reports received or number of hackers working on the program.	6	2.33
	Stress and uncertainty:	fear of burning out, social isolation during work, irregular income, etc.	5	2.0
	Limited scope:	number of domains or assets that are included in the program are limited.	3	1.0
	Limited vulnerability disclosure:	restrictive vulnerability disclosure policies and NDAs that may prevent you from publishing your work following the resolution/mitigation of the issue.	1	1.0
	Boredom:	bored of working on bug bounty programs.	1	2.0
	Reward:	fair but unsatisfying rewards.	1	3.0
	Lacking communication or language skills:	communication difficulties because you feel that you lack language skills, experience anxiety in communication, etc.	1	3.0
	Poor platform support:	dissatisfaction with how platforms handle issues, such as mediating between hackers and programs.	1	4.0
	Unrepresentative reputation system	hackers' reputation points do not reflect real experience and are not transferable between platforms.	1	4.0
	Program-platform conflict of interest:	conflict of interest due to platforms receiving payments from programs.	1	5.0
Useful platform features	**Program directory:**	listing many programs in one place, with statistics, details, etc. (being able to view Uber, Paypal, etc. programs on one page with statistics).	17	1.59
	Ease of reporting:	easy to generate, submit, and track reports and their status.	16	1.75
	Viewing disclosed vulnerabilities:	platform provided interface for viewing bugs found by others.	15	1.47
	Mediation:	platform resolving disputes between hackers and programs.	13	2.62
	Ease of payment:	receiving payments in a standardized, hassle-free way.	12	2.42
	Community:	platform making effort to create a community of hackers	11	2.09
	Platform managed disclosure:	platform provided tools/mechanisms to publicly disclose resolved bugs.	6	2.17
	Reputation system:	platform managed reputation system for hackers.	6	2.67
	Private program invitations:	access to private programs on the platform.	5	2.2
	Platform triage:	triaging managed by the platform (e.g., HackerOne triages your report instead of Uber).	5	2.4
	None:	there are no useful features that platforms provide	4	-
	Standardized rules:	platform standardizing how scopes, rewards, criticality, etc. are defined.	3	2.33
	Resources for learning:	platform providing free resources on how to hack (e.g., Bugcrowd University).	2	2.0
	Platform rewards:	e.g., platform SWAG, funded travel.	1	4.0

Table 3: Codebook for each question. We list how many times a code was mentioned by unique participants and the average rank in the listing order. Note that some of the same codes are identified under multiple questions (particularly with "leaving" and "challenges"). Cohen's K's are 0.81, 0.82, 0.91, 0.84, and 0.84, respectively.

RQ4: What are the main challenges that bug hunters face when working on bug bounty? As expected, challenges experienced with bug-bounty programs in general overlap with reasoning raised as to why bug hunters quit bug-bounty programs. For instance, communication issues are considered by many to be a challenge (responsiveness, dissatisfaction with responses, and difficulty working with managers were mentioned by 28.5%, 16.1%, and 16.1% of the study participants, respectively). Similarly, we see survey participants listing duplicates, secure or out-of-expertise assets, and sat-

uration of programs (many hackers working on a program, influx of submitted bug reports etc.).

Interestingly, some challenges were not listed under reasons why bug hunters quit bug-bounty programs. For instance, some mentioned the inconvenience of menial tasks ($n = 12$, 21.4%), such as setting up testing accounts and environments, dealing with CAPTCHAs, and timeouts. Further, we noticed complaints about bug-bounty programs' scopes not being clearly defined ($n = 6$, 10.7%) and the general uncertainty and stress bug hunters face when participating in bug-bounty programs ($n = 5$, 8.0%).

RQ5: Which bug-bounty platform features do bug hunters find the most useful? Survey participants mostly consider the fundamental feature of bug-bounty platforms to be the most useful; that is, being able to view and select from many active programs through one interface ($n = 17$, 30.4%). In addition, some survey participants appreciated private program invitations ($n = 6$, 10.7%). Some survey participants noted that interfaces which show disclosed (in some cases partially redacted) bug reports were useful ($n = 15$, 26.8%).

A large minority of survey participants enjoyed the streamlined process of generating bug reports and tracking their status ($n = 16$, 28.6%), receiving payments in a standardized way ($n = 12$, 21.4%), and easily disclosing bug reports to the public if the involved parties agree ($n = 6$, 10.7%). Survey participants also thought the platform's involvement in disputes between bug hunters and program managers ($n = 13$, 23.2%), the platform itself handling triage ($n = 5$, 8.9%), and platform standardized bug-bounty program rules were useful ($n = 3$, 5.4%).

Similar to the benefits of bug-bounty programs, survey participants found platform efforts to create a "community" of hackers to be useful ($n = 11$, 19.6%). Some also expressed that the platform-managed reputation system of bug hunters was a helpful feature ($n = 6$, 10.7%). A minority of our participants said that bug-bounty platforms had no useful features ($n = 4$, 7.1%)

5 Discussion and Future Work

Our results suggest that in order to create the most attractive programs bug-bounty programs should try to improve the responsiveness of their bug-report management systems, clearly communicate the scope and payment rules, and (if possible) keep a wide and dynamic scope. Further, our results suggest that emphasizing the learning potential of bug-bounty programs and providing (more) educational resources (e.g., by making more resolved bug reports public) might increase participation. Improving educational resources would not only benefit bug hunters but also provide a more skilled workforce to bug-bounty programs and the computer security industry in general which is known to be lacking the necessary workforce [14].

Our work might be affected by general qualitative research limitations such as sampling issues, satisficing, self-selection bias, social desirability, and demand effects. To ensure data quality, we recruited through multiple means and made sure responses to free-response questions made sense. We believe to have reached a diverse enough sample (see Section 3) to identify most of the answers bug hunters would give to our research questions. Further, the research outcomes should directly be beneficial to our survey participants which might reduce satisficing and demand effects. As with other qualitative research, for a more complete view of our results, we need validation with other experimental designs. Interviews conducted directly with bug hunters could be particularly useful in obtaining in-depth explanations as to why the factors are important and provide insight into how programs can emphasize benefits or work on addressing challenges. While the focus of this work is on bug-bounty program, some of the insights may shed light on managing vulnerability disclosure programs as well.

Building on this work, accurately knowing the importance of each factor from the bug hunters' perspective would allow researchers to give effective and concrete recommendations to both bug-bounty platforms and programs to shift the marketplace to be more productive (programs reaching the right bug hunter) and sustainable (reducing strain on the bug hunters) as well as providing program managers with insights as they render platform services and manage the program. Further, an analysis of the importance of factors based on the experience/skill levels of bug hunters might produce actionable recommendations on what bug-bounty programs/platforms should do to provide incentives for the bug hunters that best fit their programs. We attempted to convey the importance that bug hunters attribute to identified factors by noting how frequently factors were mentioned, as well as the average rank of the factors (see Table 3). While both metrics should give some context, they are likely limited to our sample (more experienced bug hunters) and require further validation. To this end, we are planning a larger-scale survey that directly aims to measure the importance of factors to provide accurate data and to allow for statistical inferences to be made. Additional factors may be added, and expansion to vulnerability disclosure programs may be considered.

Acknowledgments We thank our participants and the anonymous reviewers of our paper for their insightful comments and suggestions. This work was supported in part by the National Science Foundation under Grant CNS-1850510.

References

[1] H. Russell Bernard. *Research methods in anthropology: Qualitative and quantitative approaches*. Rowman & Littlefield, 2017.

[2] Bugcrowd. The state of bug bounty, 2018. https://www.bugcrowd.com/resources/reports/state-of-bug-bounty-2018/.

[3] Bugcrowd. Inside the mind of a hacker, 2020. https://itmoah.bugcrowd.com/.

[4] Amit Elazari. Private ordering shaping cybersecurity policy: The case of bug bounties. *An edited, final version of this paper in Rewired: Cybersecurity Governance, Ryan Ellis and Vivek Mohan eds. Wiley*, 2019.

[5] Amit Elazari. The evolving security policy landscape and how it impacts you. 2020 IoT Village (virtual conference), https://www.youtube.com/watch?v=qnusAKB2OU8, May 2020.

[6] Matthew Finifter, Devdatta Akhawe, and David Wagner. An empirical study of vulnerability rewards programs. In *22nd USENIX Security Symposium*, pages 273–288, 2013.

[7] HackerOne. The 2020 hacker report, 2020. https://www.hackerone.com/sites/default/files/2020-04/the-2020-hacker-report.pdf.

[8] J. Richard Landis and Gary G. Koch. The measurement of observer agreement for categorical data. *Biometrics*, pages 159–174, 1977.

[9] Aron Laszka, Mingyi Zhao, and Jens Grossklags. Banishing misaligned incentives for validating reports in bug-bounty platforms. In *Proceedings of the 21st European Symposium on Research in Computer Security (ESORICS)*, pages 161–178, September 2016.

[10] Aron Laszka, Mingyi Zhao, Akash Malbari, and Jens Grossklags. The rules of engagement for bug bounty programs. In *22nd International Conference on Financial Cryptography and Data Security (FC)*, pages 138–159. Springer, 2018.

[11] Donatello Luna, Luca Allodi, and Marco Cremonini. Productivity and patterns of activity in bug bounty programs: Analysis of HackerOne and Google vulnerability research. In *14th International Conference on Availability, Reliability and Security (ARES)*, pages 1–10, 2019.

[12] Thomas Maillart, Mingyi Zhao, Jens Grossklags, and John Chuang. Given enough eyeballs, all bugs are shallow? Revisiting Eric Raymond with bug bounty programs. *Journal of Cybersecurity*, 3(2):81–90, 2017.

[13] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW Issue):1–23, 2019.

[14] Steve Morgan. The 2019/2020 official annual cybersecurity jobs report, 2020. https://cybersecurityventures.com/jobs/.

[15] Daniel Votipka, Rock Stevens, Elissa Redmiles, Jeremy Hu, and Michelle Mazurek. Hackers vs. testers: A comparison of software vulnerability discovery processes. In *39th IEEE Symposium on Security and Privacy (S&P)*, pages 374–391, 2018.

[16] Thomas Walshe and Andrew Simpson. An empirical study of bug bounty programs. In *2020 IEEE 2nd International Workshop on Intelligent Bug Fixing (IBF)*, pages 35–44, 2020.

[17] Mingyi Zhao, Jens Grossklags, and Peng Liu. An empirical study of web vulnerability discovery ecosystems. In *22nd ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1105–1117, 2015.

[18] Mingyi Zhao, Aron Laszka, and Jens Grossklags. Devising effective policies for bug-bounty platforms and security vulnerability discovery. *Journal of Information Policy*, 7:372–418, 2017.