Unfairness in the Bug Bounty Ecosystem: Problems, Metrics, and Solutions

YANGHERAN PIAO and DANIEL W. WOODS, University of Edinburgh, UK

As bug bounty programs evolved since the 1990s, bug hunters have started to raise questions about "fairness". At a high-level, the discourse echos the complaints of other workers in the gig economy, such as rideshare drivers. However, closer inspection reveals vulnerability research has unique considerations resulting from the production of complex, information goods. This requires a dedicated study. This article identifies and classifies issues of unfairness into four categories: policy ambiguity; prosecution risk; responsiveness; and hacker relations. Turning to solutions, we discuss mechanisms to improve fairness: service level agreements; reputation systems; collective action; and appeals processes. We then identify metrics that assess real-world bug bounty platforms, tailored to capture different dimensions of (un)fairness. Finally, we conduct a case study on Chromium's bug bounty program, evaluating responsiveness across almost 10 years of operation against a benchmark SLA.

Keywords: bug bounty, fairness, vulnerability disclosure, gig economy

Reference Format:
Yangheran Piao and Daniel W. Woods. 2025. Unfairness in the Bug Bounty Ecosystem: Problems, Metrics, and Solutions. In *Proceedings of the 24th Workshop on the Economics of Information Security (WEIS'25).* , 23 pages. Tokyo, Japan. June 23–25, 2025.

1 INTRODUCTION

Bug bounty platforms and programs have capitalized on the gig economy's labor model to enhance cybersecurity efforts. These programs reward security researchers, known as bug hunters, who identify and report vulnerabilities [1, 18, 20]. Vendors utilize these programs to increase software security at a fraction of the cost of security engineers [5], providing hunters with recognition and financial compensation [44]. This offers a mutually beneficial business model for technology companies and ethical hackers [4, 21, 35].

However, bug hunters are the weaker party in the transaction relative to large technology companies, much like other types of work in the gig economy [16]. This imbalance often leaves them with less bargaining power and next to no job protection. Unlike traditional cybersecurity employees who may benefit from labor laws, minimum wage guarantees, and health benefits, bug hunters operate as independent contractors [84]. The nature of bug hunting also introduces a high level of uncertainty in earnings, as payments are contingent upon discovering bugs that meet criteria that can be ambiguous [19].

Bug hunters face more uncertainty than other gig economy workers. In food delivery, the task is well-defined (delivering meals from the restaurant to the customer's address), and workers can reliably estimate the time of arrival. In contrast, bug bounty programs involve producing novel information goods that cannot be described ex-ante, or have their production time estimated [17]. This differs from typical information tasks on sites like MTurk where task completion time is predictable, such as filling out surveys, object recognition, translation, data cleaning, and so on. Inexperienced hackers encounter even greater uncertainty about payout likelihood and the time investment [20, 21, 42].

Authors' Contact Information: Yangheran Piao, lawrence.piao@ed.ac.uk; Daniel W. Woods, daniel.woods@ed.ac.uk, University of Edinburgh, Edinburgh, UK.

Prior work has framed unfair treatment as a limitation of bug bounty platforms [10, 12, 22, 23], rather than a systemic issue that can be addressed. Discussions frequently mention issues such as mismatched reward payouts [9], inconsistency in severity assessments [24], and poor communication [25, 27]. However, these are often presented as isolated problems rather than part of a broader structural challenge. Addressing this gap, our study focus on three research questions to advance understanding of fairness in BBPs:

RQ1: What are the causes and specific manifestations of unfairness in the bug bounty ecosystem?

RQ2: What mechanisms are expected to improve bug bounty fairness?

RQ3: How can fairness in bug bounty industry be measured?

Contributions: We review literature on unfairness in bug bounty programs (**RQ1**), categorize the issues into four types, which can be mapped to existing notions of organizational justice:

(1) Prosecution risk concerns uncertainty about protection from legal actions against hackers.
(2) Policy ambiguity covers adverse decisions related to program scope, severity classification, and reward amount.
(3) Responsiveness relates to delays in communications, time-to-patch and payment.
(4) Hacker relations captures subjective feelings of disrespect or not being taken seriously.

We then discuss four potential mechanisms to improve fairness in bug bounty programs (**RQ2**):

(1) Service Level Agreements (SLAs) whereby vendors commit to response times, and communication standards.
(2) Reputation systems that share experiences to help hunters make informed choices about BBP participation.
(3) Collective action whereby hunters jointly negotiate and boycott unfair BBPs, similar to a labor union.
(4) Appeals processes that subject the BBPs' decisions to external scrutiny, and the possibility of being over-turned.

We finally derive metrics that evaluate fairness across all identified dimensions (**RQ3**). We also conduct the case study of Chromium, showing that the proposed historical metrics can capture the fairness of responsiveness in BBPs.

2 BACKGROUND

2.1 The Economics of Unfair Bug Bounty

Fairness in labor markets has been a longstanding concern in both ethical and economic discussions [11, 85]. Adams' equity theory [30] holds that workers evaluate fairness by comparing their inputs to the outcomes they receive relative to others. When workers perceive an imbalance, feelings of inequity arise, leading to dissatisfaction and disengagement. Fairness concerns are even more pronounced in the gig economy [31]. Platforms often create information asymmetries, limit worker bargaining power, and set ambiguous rules that favor platform owners rather than workers. Workers have little recourse to challenge decisions.

These insights broadly apply to the bug bounty ecosystem, which also has its own unique challenges [22]. Hunters face unfairness due to key factors: *incomplete contracts* and *sunk costs*. Both of these weaken hunters' bargaining power within the bug bounty ecosystem.

Incomplete contracts: As information goods [17], vulnerabilities lack a reliable and objective valuation standard [55]. In the absence of a external standards, vendors establish bespoke policies and processes, using subjective evaluation criteria to determine whether to pay a reward and how much to offer. This subjectivity results in incomplete contracts, leaving hunters uncertain about whether they will be compensated for their discoveries [35].

One of the most evident manifestations of incomplete contracts is the issue of duplication [36, 37, 51]. In cases of duplicate reports, bug hunters invest time and effort into discovering a vulnerability, only to be informed that someone

else has already reported it, rendering their work worthless. This is almost impossible in other parts of gig economies, because at the moment a driver or deliverer accepts an order, the same trip can only be completed once.

Uncertainty in payment also arises from subjective evaluation. For example, Wunder *et al.* [24] showed that practitioners do not reliably score vulnerabilities using an established standard (CVSS). Vendors who follow their own proprietary classification policies are unlikely to be any more consistent. As a result, hackers cannot reliably predict how their findings will be rated or compensated [38].

Sunk costs: In addition to incomplete contracts, bug hunters incur sunk costs that further restrict their ability to secure fair compensation. Sunk costs in bug hunting can be divided into: (1) time invested in finding a specific bug; and (2) time invested in understanding the system.

First, the time required to discover a vulnerability ranges from a single day to several weeks [8]. If a platform offers a reward that does not fairly reflect the time and effort invested, hunters must either accept an inadequate payment or reject the offer, effectively wasting weeks of work. This is particularly problematic since there are very few alternative legitimate buyers of specific security vulnerabilities, leaving hunters with limited options [56].

Second, researchers tend to specialize in specific domains, often investing considerable time in understanding the inner workings of a particular system. This specialization makes it difficult for hunters to switch between different bounty programs. For example, a researcher who has spent years mastering iOS security may find that abandoning Apple's bug bounty program and starting fresh in a different ecosystem would require discarding all the time invested in learning iOS security. This understanding of iOS cannot be easily re-purposed for alternatives to BBP participation.

While bug bounty programs offer a legitimate channel for vulnerability disclosure, they are not the only option available to hackers. Alternatives often offer higher rewards, such as selling bugs on black markets or to intelligence agencies [56, 57]. While these paths pose ethical and legal risks, such financial incentives could drive hackers toward underground forums or exploit brokers [41, 58], particularly when bounty platforms appear unfair. In this context, ensuring fairness is also important to ensure researchers submit vulnerability reports to the vendor.

To conclude, incomplete contracts and sunk costs create a fundamental power imbalance between hunters and programs. Vendors hold the authority to set terms, control payouts, and interpret policies. This gives vendors negotiating power over hunters, who often lack the leverage to push back. As a result, the current structure of the bug bounty economy systematically disadvantages hunters.

2.2 Related Works

The discussion of unfair treatment in bug bounty programs has largely been fragmented in prior research. Many studies have highlighted individual challenges faced by vulnerability hunters, which have often been framed as shortcomings or inefficiencies within bug bounty platforms rather than as manifestations of a broader systemic fairness issue.

Previous research on bug hunters has primarily examined factors influencing their participation in bug bounty programs, such as financial incentives, skill development opportunities, and the pursuit of recognition [9, 12, 34]. Additionally, studies have also identified challenges including poor communication with vendors, severity downgrades, unclear scope definitions, and inadequate platform support [9, 10, 12, 21, 23, 25]. However, these cons have not been explicitly linked to fairness considerations.

There have been efforts to measure various aspects of bug bounty platforms, primarily focusing on their effectiveness and efficiency. Some studies have analyzed the distribution of rewards [36, 45–47], participant engagement [8, 40, 42, 43], and program policies [59] to assess how different platforms operate. Other work has attempted to quantify the effectiveness of bounty programs by examining the number of reported vulnerabilities [5, 46], the time of fix or patch

deployment [40, 48, 49], and the overall return on investment for vendors [5, 44]. Nevertheless, such metrics are commonly used to evaluate program success, yet they do not explicitly capture the experiences of hunters or assess whether these programs provide equitable treatment.

Furthermore, these studies which focus on operational efficiency fail to address a critical factor—sustainability [27]. Fairness is not just about optimizing short-term efficiency or maximizing the number of bugs reported; it is also about ensuring that these programs remain viable and attractive to skilled hunters in the long run [61]. Sustainable programs require trust, transparency, and equitable treatment to maintain a steady influx of high-quality researchers willing to invest their time and expertise. Unfair treatment can erode hunter participation over time, leading to talent attrition and a declining quality of reports. Thus, sustainability is not only about the welfare of hunters but also about the long-term success of vendors in maintaining a resilient security posture.

Discussions about fairness also have been found within the hacker community, where hunters themselves have reported experiencing various forms of unfairness in bug reporting. These often take place on forums and social media platforms, where hunters share their concerns regarding inconsistent reward decisions [79, 82], unexpected ignore or scope exclusions [14, 82, 83, 93], the lack of transparency in vendor evaluations [80, 81], delayed or unclear response [14, 82], and legal risk [83].

Despite growing exploration around hunter perception and platform measurement, a systematic analysis of fairness dimensions in bug bounty ecosystems remains largely absent. The above works motivates our work, which seeks to define, categorize, and measure fairness issues in bug bounty programs to provide a structured understanding of such kind of challenges faced by hunters.

2.3 Concepts of Fairness

In the fields of management and psychology, *organizational justice* is a research paradigm that focuses on fairness in organizational processes, outcomes, and communications [30, 86, 87]. It is a widely used framework to understand how perceptions of fairness shape attitudes, motivation, and behavior. This overall justice is typically divided into three dimensions: distributive, procedural, and interactional fairness [88, 89].

Distributive justice refers to the perceived fairness of outcomes and resource distributions [94], typically whether rewards, benefits, or costs are allocated in a just manner across individuals [95]. This notion of fairness emphasizes that workplaces can be judged by principles of equity (rewards proportional to contributions), equality (everyone receives the same), or need. *Procedural justice* focuses on the fairness of processes and rules [29]. A process is perceived as procedurally just if it is consistent, accurate, transparent, unbiased, ethical and allows voice or input from those affected [28]. *Interactional justice* includes two sub-components: interpersonal and informational. The interpersonal aspect refers to the quality of treatment people receive, meanwhile the information aspect captures the adequacy of explanations provided when decisions are implemented [32, 33]. It captures whether individuals feel respected, valued, and informed. These three types of justice provide a useful perspective for understanding unfairness in bug bounty programs. We map key issues in the bug bounty ecosystem to each of these fairness types in Section 3.2.

A different perspective is provided by notions of individual versus group fairness that have been studied in the fields of philosophy, law, and social science. This notion has been applied in recent years with the rise of machine learning algorithms [66–68]. *Individual-based fairness* is the theory that each person is treated equitably relative to relevant criteria, focusing on case-by-case justice [69]. A common formulation is that like cases should be treated alike. The emphasis is on similarity and consistency in treatment at the individual level [70]. In contrast, *group-based fairness* assesses justice across groups of people, often defined by demographic or other categorical attributes [71].This notion

often imports the normative idea that outcomes should be proportionately balanced across groups to avoid systemic bias [72]. Group fairness is concerned with aggregate equality (e.g. no protected group is disadvantaged on average), whereas individual fairness is concerned with each person's treatment.

The bug bounty ecosystem can be studied through both an individual and group-based conception of fairness. At the individual level, researchers can compare their own effort, skill, or contribution to the outcomes they personally receive in the form of rewards, acknowledgments, or treatment by the BBP. The group-based perspective on fairness has been applied to consider whether the current ecosystem may disadvantage marginalized people and new comers [9, 12].

3 DIMENSIONS OF UNFAIRNESS IN BBPS

3.1 Definition

We synthesize prior work and real-world case-studies to derive four key dimensions of unfairness (see Table 1).

Table 1. Types of unfairness in bug bounties

Category	Issue	Example
Prosecution Risk	- Unclear or lack of legal protection - History of prosecution	[23, 26, 53, 54, 59]
Policy Ambiguity	- Unclear bug bounty policy - Misclassification or downgrade of severity - Lower than expected rewards - Rejecting in-scope vulnerabilities	[14, 37, 49, 79, 81, 83, 91]
Responsiveness	- Slow or delayed responses - Delayed reward payments - Denying duplicated reports	[9, 45, 49, 51, 79, 82]
Hacker Relations	- Dismissive or disrespectful communications	[9, 10, 12, 27]

Prosecution risk focuses on the threat of legal action after reporting the bug, and the unclear legal risks before finding the bug. Without explicit safe harbor provisions, hunters may be exposed to potential lawsuits, prosecution, or legal retaliation simply because they report a vulnerability in good faith [23, 53]. Some programs provide unclear legal boundaries [54], leaving hunters uncertain about what is permissible. The absence of clear legal protections discourages ethical hacking efforts, as hackers may hesitate to report security issues for fear of facing legal consequences. Vague or poorly defined policies regarding responsible disclosure and legal liability force researchers to operate in a precarious legal gray area [26, 59]. This lack of clarity increases the likelihood of unintentional violations, where hunters unknowingly breach terms of service, network access policies, or even cybersecurity laws.

Policy ambiguity concerns decisions about assessing whether the bug is in-scope, the severity and reward amount. Unclear policies create more uncertainty for hunters [37, 45, 81, 91], as ambiguous guidelines may result in valid reports being rejected. The verification of whether an issue reported by hunters qualifies as a valid bug is highly subjective, which can cause bugs to be classified as out of scope [49, 51, 98]. This leads hunters to feel they were unfairly denied compensation for their findings [14, 83].

Responsiveness concerns the timeliness and predictability of key processes in bug reporting. One of the most common issues is delayed responses from vendors, with some reports going unanswered for weeks or even months [49, 82]. In extreme cases, vendors become completely unresponsive, leaving hunters uncertain about the status of their submissions. Such delays waste the time and effort of researchers [9, 45]. Similarly, delayed reward payments exacerbate the problem,

with some hunters waiting weeks or even months to receive their bounties [13]. Such uncertainty discourages continued participation, particularly for those who rely on bug bounty programs as a source of income.

Delayed responses increase the likelihood of duplicate reports, in which multiple researchers independently discover the same vulnerability. Platforms often operate on a "first-to-report" basis, meaning that only the first submission is rewarded while others receive nothing. In some cases, certain bug bounty programs have started offering small compensation or shared rewards to later reporters, which is a positive step toward fairness [90].

Hacker Relations refers to the quality of communication and respect shown towards bug hunters. Studies have found that vendors display an unprofessional attitude when handling reports [9, 12, 27]. This can manifest in several ways, such as vague or unhelpful responses, a lack of acknowledgment of the hunter's efforts, or dismissiveness toward legitimate concerns. As poor communication fails to provide meaningful updates on report status, leaving hunters in limbo for extended periods.

3.2 Mapping to Organizational Justice

As shown in Figure 1, the unfairness perceived in bug bounty ecosystems can be mapped onto organizational justice dimensions. Distributive justice concerns emerge in how rewards are allocated, procedural fairness issues arise in the clarity of policies and timeliness of processes, and interactional fairness pertains to the tone and transparency of communications.

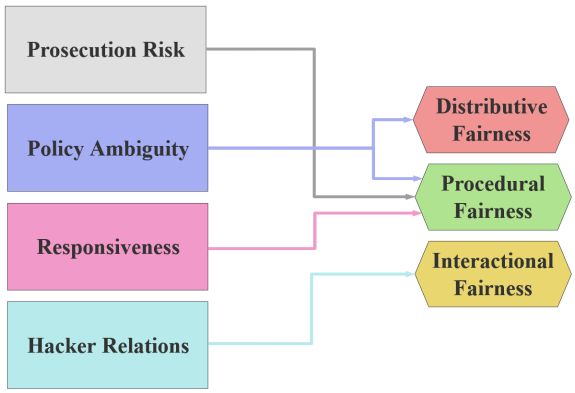

Fig. 1. Mapping unfairness in bug bounties to organizational justice

Prosecution risk can be mapped to *procedural justice*, which concerns whether decision-making processes are consistent, transparent, predictable, and fair. Since many BBPs lack clear safe harbor policies, creating uncertainty about the applicable rules and potential consequences. This absence leaves hunters feeling unprotected as the rules of the game are either unclear or perceived to be biased against them, which touches on notions of procedural injustice.

Policy ambiguity can be mapped to both *procedural justice* and *distributive justice*. Vague BBP criteria violate procedural fairness by depriving hunters of a transparent and unbiased process for evaluating their findings. Moreover, when hackers submit valid vulnerabilities, investing time, skill, and effort, yet are denied rewards due to subjective or ambiguous rules, this creates a mismatch between labor and return, constituting also a violation of distributive fairness.

Responsiveness is related to *procedural justice*. Unpredictable or slow responses violate the principles of consistency, timeliness, and transparency within procedural justice, leaving researchers unsure whether their report will be acted

upon at all. Hacker relations fall within the scope of *interactional justice*. Interactional fairness emphasizes that respectful dialogue and adequate explanation matter. In the context of bug bounties, this means providing clear and polite communication about decisions. When hackers are ignored or not acknowledged, it violates the interpersonal dimension of interactional justice. Meanwhile, poor communication from the organization violates the informational dimension.

4 MECHANISMS TO IMPROVE FAIRNESS

This section provides four different mechanisms to improve fairness in bug bounty programs, as shown in the overview in Figure 2, each could address different dimensions of unfairness.

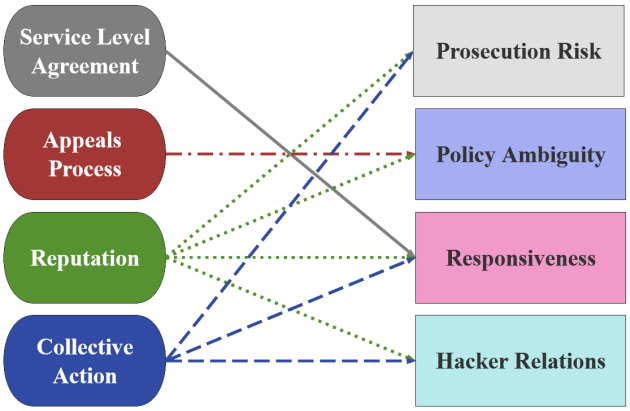

Fig. 2. Mechanisms and the corresponding types of unfairness they could address

4.1 Service Level Agreements

SLAs address responsiveness issues by setting expectations regarding response times, resolution times, and payment processing. In practice, this means vendors commit to predefined timelines for acknowledging reports, providing status updates, and processing payments. A well-structured SLA would specify, for instance, that all reports will receive an initial response within a set number of days, that high-severity vulnerabilities must be addressed within a certain time frame, and that rewards will be processed within a fixed period after report validation [27]. These agreements reduce ambiguity, helping hunters plan their efforts and avoid the indefinite waiting periods.

Bug bounty platforms can introduce an independent *third party* to monitor vendors' response times and assess their adherence to SLAs. The third party should make the times visible to hunters, allowing them to make informed decisions about where to invest their efforts. For example, HackerOne recommends four "response efficiency metrics" (namely time to first response, triage, bounty, and close), recommending "healthy" target response times [92]. However, there is no precedent for vendors offering compensation to hackers for failing to adhere to an SLA. This is perhaps unsurprising given the commitment creates unknown liability given there is no limit to how many bugs can be submitted.

Challenges. While SLAs can help ensure timely responses and payments, SLAs struggle to address subjective aspects of the review process like policy ambiguity or hacker relations. Additionally, vendors may avoid committing to strict deadlines, especially for complex vulnerabilities requiring extensive internal coordination.

4.2 Appeals Processes

Appeals provide hunters with an opportunity to challenge decisions made by BBPs. Hunters cannot take vendors to court when bug bounty policies waive the rights of participants. The most relevant clauses prevent hackers from taking legal action over disputes. Even when rights have not been waived, hackers may lack the resources and will to take legal action against a powerful vendor.

One alternative is for platforms that manage the bug bounty program to help resolve disputes between the vendor and reporter. For example, HackerOne offer a mediation service for disputes [96]. The platform promises to "facilitate discussions between hackers and customers to enable a more favorable outcome for everyone involved". However, this falls short of arbitration with the power to over-turn the vendor's decision, with hunters reporting that this mediation rarely results in a change in decision and the intermediary platform is inherently biased [12, 93].

More ambitiously, BBPs could introduce a structured peer review process to provide an additional layer of oversight and transparency [6, 25]. This mechanism reduces the risk of unilateral or biased decision-making by vendors by involving multiple stakeholders in the evaluation of submitted reports. Community-driven assessments, alongside vendor input and external expert reviews [52], can help ensure that reports are not dismissed via misclassification.

Challenges. Appeals processes can help address policy ambiguity by providing a route to challenge vendor decisions. However, this adversarial relationship is unlikely to improve prosecution ambiguity or hacker relations. Furthermore, it actively slows down responsiveness as appeals processes can be drawn out.

4.3 Reputation

One of the core issues in the current ecosystem is the asymmetry of information. New hunters often have little knowledge about how a vendor treats reports until they personally experience the process, which can lead to negative feelings and wasted effort [9]. Reputation provides transparency for hunters, potentially addressing all four dimensions of fairness. Vendors known for timely responses, fair payouts, and transparent policies naturally gain a positive reputation, attracting more skilled researchers. Conversely, vendors with a history of mistreatment may find it harder to engage top talent, as negative experiences are widely shared within the community.

Thus far, BBP reputations have been established through informal channels. This means vendors face minimal direct consequences for poor treatment of hunters [15], aside from scattered complaints on social media or within hacker communities. Nevertheless, there are isolated reports of hacker communities sharing informal ratings and anecdotal experiences about different bug bounty programs, using this information to decide which BBPs to participate in [7]. We are not aware of a formalized reputation system, such as ratings aggregated by a third-party. Such a system could aggregate data on vendors' response times, payment reliability, dispute resolution practices, and overall treatment of researchers.

Challenges. The subjective nature of informal reputation systems makes them difficult to standardize. Unlike SLAs, which establish concrete expectations and measurable outcomes, reputation systems rely on anecdotal experiences within the community. This can punish the most egregious examples of unfairness, in which experiences are documented and shared in public channels, and also punish vendors who are consistently unfair. It is perhaps most effective in setting community expectations rather than addressing individual grievances. Further, vendors can be unfairly punished if they are judged based on isolated incidents, or unreliable anecdotes.

4.4 Collective Action

Worker collectives have long been a powerful tool for addressing imbalances of power, particularly in industries where individuals face systemic disadvantages when negotiating with larger entities [62]. In the context of bug bounty programs, collective action offers a way for hunters, who typically operate as independent freelancers, to counteract power asymmetries with vendors [63]. Hacker collectives can force vendors to change in order to maintain engagement from top talent, by threatening to boycott a program or refuse to engage with vendors known for unfair treatment [9].

This leverage can address and improve fairness, which can be seen in the bug bounty ecosystem in China where many hackers form into teams. Hackers report better treatment and faster response times after joining teams, thereby improving *hacker relations* and *responsiveness*. Teams can even address *prosecution risk*, as evidenced by an anecdote in which a vendor was threatening to prosecute a hacker and dropped the case after a team intervened [9].

Collective action can also feed into the other mechanisms. For example, teams enhance reputation systems by sharing information and only submitting reports to BBPs with a history of fair decisions. There is also evidence that team leaders discuss historic decisions with vendors, providing a function similar to mediation on HackerOne [96]. Finally, teams in China have pushed to create an independent appeals process, which involves multiple teams forming an alliance and inviting programs to participate [6]. When a dispute arises between a member of the alliance and a program regarding a bug, a vote is initiated on how to resolve the dispute. Although the outcome is currently unknown, this is a important step by the hacker collective to enhance fairness.

Challenges. A major obstacle of implementation is maintaining engagement and coordination among a highly decentralized group of researchers. Unlike traditional labor unions, which operate within a structured environment, hackers are often scattered across different platforms, geographies, and communities. Their participation in bug bounty programs is usually voluntary and fluid, making it difficult to sustain long-term collective efforts. Without a formalized structure, maintaining active participation and commitment to collective initiatives can be difficult.

4.5 Outlook

There has been limited experimentation with mechanisms that address unfairness in bug bounty programs. HackerOne provide examples of mechanisms implemented by vendors. The platform gently recommends that vendors should commit to response SLAs [92], and also create a very weak appeals process [96]. However, vendors can ignore these recommendations while still hosting their BBP on HackerOne. As such, we believe there is still considerable room for vendors to experiment with SLAs and appeals processes. A cynic would say vendor-led initiatives will always be lagging, given vendors benefit from the current power asymmetries.

In terms of hacker-led mechanisms, reputation systems remain largely informal relying on anecdotal reports, but there has been successful experimentation with hacker collectives in China [9]. Collectives serve the twin goals of sharing information about unfair BBPs, and also coordinating hackers to exit unfair programs. Members of these collectives report to receive better treatment by vendors.

One commonality is that all mechanisms function better with more information. SLAs require metrics upon which promises can be made, appeals processes can be bench marked against industry standards, reputation systems can avoid unreliable anecdotes, and collectives can ground negotiations in objective realities. With this in mind, the next section derives metrics that quantify fairness in bug bounty programs.

5 TOWARDS FAIRNESS METRICS

Fairness metrics could help compare different platforms, assess the impact of proposed fairness-enhancing mechanisms, and identify outliers who need change. However, it is unclear how to actually measure the diffuse unfairness factors in bug bounty ecosystem. In this section, we identify three broad approaches to measurement:

(1) **Historical** metrics capture how BBPs have processed reports in the past.
(2) **Perception** metrics capture how hackers feel about specific BBPs.
(3) **Contractual** metrics analyze bug bounty policies via natural language processing techniques.

The measurement approach should be tailored to each aspect of fairness, as shown in the following Table 2. Additionally, Table 4 in Appendix A describes and classifies the proposed metrics.

Table 2. Metrics for measuring different aspects of fairness

Target	Metric	Reference
Prosecution Risk	Previous Prosecution Record	
	Perceived Researcher Protection	
	Clarity of Legal Risk Disclosure	[59]
Policy Ambiguity	Expectation Consistency	
	Hunter Comprehension	
	Reward-Severity Ratio	
	Clarity of Submission Guidelines	[50, 59]
	Clarity of Scope Definition	[50, 59]
	Policy Readability Test	[50, 59]
Responsiveness	Average First Response Time	[75, 76, 92]
	Average Triage Time	[92]
	Average Fix Time	[48, 49, 76, 77]
	Average Bounty Decision Time	
	Average Evaluation Time	
	Average Payment Time	[92]
Hacker Relations	Satisfaction Score	
	Average Dispute Resolution Time	
	Dispute Resolution Rate	

5.1 Metrics for Measuring Prosecution Risk

Prosecution risk can be measured using historical, perception, and contractual metrics. From a historical perspective, a key metric is the *previous prosecution record*, i.e. number of hackers prosecuted or legally threatened after submitting a report. For most BBPs, this number is likely zero, as vendors typically do not pursue legal action against hunters. However, a non-zero value carries high information content, as even a single case of prosecution can create deterrence among the hunter community. There is a difficult definitional question about what counts as prosecution, whether there needs to be a formal police report or legal action or whether sending a scary letter suffices. Additionally, legal threats are not always publicly documented, making it difficult to obtain reliable data.

Instead contractual metrics could analyze the bug bounty policy for the presence of safe harbors [78]. These metrics assess whether a bug bounty program includes legal safeguards such as safe harbor provisions and whether its rules and policies are clear enough to prevent legal ambiguity. *The clarity of legal risk disclosure* [59] measures whether a program

explicitly states whether researchers will face legal repercussions for submitting reports. It could also measure which laws or jurisdictions apply to the program [3]. NLP techniques can assist in analyzing policy documents, detecting vague or missing legal protections, and assessing whether vendors' policies align with industry legal standards.

Prosecution risk is not just about written policies, it is also about how hunters perceive their treatment. *Perceived researcher protection* serves as a perception-based metric, capturing how safe hunters feel when engaging with a bug bounty program. Even if a program has strong safe harbor provisions on paper, hunters may still feel vulnerable to legal threats if the vendor has a vague policy enforcement [23]. Surveys can help gauge hunter sentiment, identifying whether security researchers trust a platform's legal protections or feel at risk despite policy assurances.

5.2 Metrics for Measuring Policy Ambiguity

Policy ambiguity can be measured using all three types of metric. From a perception standpoint, *expectation consistency* serves as a metric for assessing can quantify the gap between perceived fairness and actual program behavior. If many hunters report that they expected higher rewards, different severity ratings, or clearer acceptance criteria, this suggests that the program's policies and past enforcement lack sufficient clarity or predictability. A consistent mismatch between expectations and reality may indicate that policy documentation does not accurately reflect the vendor's real-world decision-making process. Additionally, *hunter comprehension* is an assessment of hunters' understanding of program policies via a questionnaire to evaluate their clarity. By assessing comprehension levels across different experience groups, it can determine whether hunters clearly grasp the rules, processes, and expectations set by a program [37].

From a historical perspective, the reward-severity ratio can evaluate whether the rewards issued by the program align with the declared severity ratings of reported bugs. Ideally, bugs classified at similar severity levels should receive proportionally similar rewards. However, if hunters observe inconsistent or arbitrary reward allocations, it can indicate that policy enforcement is not only ambiguous but also discretionary [24, 102]. For example, if two reports in same type and classified as "critical" receive vastly different payouts, hunters may question whether the program has clear internal standards for bounty determination.

From contractual metrics, *the clarity of submission guidelines* measures whether the program provides explicit instructions on what a valid bug report must include [5]. Ambiguous or insufficiently detailed submission guidelines can result in unnecessary report rejections due to formatting errors, missing proof-of-concept (PoC) elements, or incomplete impact descriptions [65]. Programs with clear and structured submission templates help researchers better align their reports with vendor expectations [73, 74]. Similarly, *the clarity of scope definition* determines whether the program explicitly states what vulnerabilities are in-scope and out-of-scope. By analyzing how well policies outline specific in-scope vulnerabilities and include concrete examples, we can measure the extent to which ambiguity leads to unnecessary confusion or unfair decisions [5, 50].

Even when enough policies exist, they must be accessible and understandable to be effective. *The policy readability test* evaluates the linguistic complexity and readability of program policies, ensuring that hunters from diverse backgrounds can easily grasp the rules, such as the Flesch Reading Difficulty Scale [5, 50]. Some policies are laden with legal jargon or excessively long-winded, making them difficult to navigate. Programs that score poorly in readability tests are more likely to contribute to ambiguity because hunters may misinterpret unclear language or fail to find relevant information.

5.3 Metrics for Measuring Responsiveness

Measuring responsiveness requires historical metrics that capture how efficiently vendors handle bug reports at various stages. HackerOne and other platforms have established suggested standards for responsiveness [92]. We adopt and

expand on these existing metrics, focusing on several key time intervals that influence the fairness and effectiveness of BBPs. One of the responsiveness metrics is *average first response time*, which measures the duration from a report's submission to the vendor's initial response. Faster initial responses indicate a program that values hunter contributions, whereas prolonged silence signals poor communication. Beyond the first response, it must also ensure timely assignment of reports to relevant reviewers. This is captured by *average triage time*, which measures the duration from submission to the assignment of a reviewer. Delays in triage can slow down the overall resolution process, as reports remain in limbo without being assessed for validity or severity.

Once a report is triaged, its resolution speed is crucial for both the security of the vendor's product and the fairness to the hunter. *Average fix time* tracks the time taken from bug submission to be fixed [48, 49, 76, 77]. While some delays are expected for complex vulnerabilities requiring extensive changes, excessively long fix times can make hunters feel that their findings are undervalued if vendors do not prioritize patching security issues in a timely manner. *The average bounty decision time* calculates the time from submission to the final determination of the bounty amount. Many hunters depend on bounty programs as a source of income, and uncertainty about whether they will be compensated. Programs that take too long to make bounty decisions signal inefficiency and lack of commitment to fair compensation.

In addition to the core phases, reports often require multiple rounds of assessment. *Average evaluation time* captures the duration between the first assignment and either reassignment or report resolution. Some reports may bounce between different reviewers, leading to delays in processing. Finally, the *average payment time* measures how long it takes for rewards to be disbursed once a bounty decision has been made. Even after a bounty is awarded, hunters frequently experience delays in receiving payment.

5.4 Metrics for Measuring Hacker Relation

Hacker relations are best captured by studying perceptions and how vendors handle conflicts with hunters. The *satisfaction score* measures hunters' subjective experiences regarding communication professionalism and respect. Hunters who feel dismissed, ignored, or treated with hostility may perceive the program as unfair, even if other structural elements, such as rewards are reasonable. This metric can be quantified through post-interaction surveys where hunters rate their experiences based on responsiveness, tone, and helpfulness of vendor representatives.

Beyond perception, the historical metric, like *average dispute resolution time* tracks the time taken to resolve disputes or reassess cases where hunters contest a decision, which can be quantified by analyzing the average time taken from the moment a dispute is filed to its resolution. Excessive delays indicate inefficiencies in the dispute-handling process and signal that vendors deprioritize hunter concerns [9]. Similarly, *dispute resolution rate* measures how frequently disputes are resolved amicably versus those that remain unresolved or are escalated to higher levels. A low resolution rate suggests that vendors may be unresponsive to hunter concerns or unwilling to engage in constructive discussions. This metric could be calculated by tracking the number of disputes raised by hunters and categorizing them into "resolved," "unresolved," or "escalated" (e.g., requiring third-party mediation or public disclosure).

6 CASE STUDY – CHROMIUM

In this section, we conducted a case study using real-world data from Chromium's program as testing a hypothetical SLA. By applying a subset of the proposed time-based metrics to this well-documented, long-running BBP, we assessed how well historical metrics perform in measuring responsiveness. We use the standard and recommended values from HackerOne for comparison [92].

Table 3. Comparative overview of Chromium measurement research in our paper and literature

Study	Year	Bug Bounty[$]	Fairness[⚖]	Measured Stages
Meneely *et al.* [101]	2014	—	—	No specific
Munaiah *et al.* [102]	2016	✓	—	No specific
Aljedaanial *et al.* [103]	2020	—	—	Report Submission→Fix
Paul *et al.* [99]	2021	—	—	First Response→Final Decision
Rajbhandari *et al.* [75]	2022	—	—	Report Submission→Fix Report Submission→Assign
Atefi *et al.* [49]	2023	✓	–	Report Submission→Fix
Franken *et al.* [77]	2023	—	—	Report Submission→Fix
Aljedaani *et al.* [48]	2024	—	—	Report Submission→Fix
Us	2025	✓	✓	Report Submission→First Response→ Assign→Fix→Final Decision→Payment

[$]The study specifically focuses on or mentions bug bounties.
[⚖]The study specifically focuses on or mentions fairness.

We collected the publicly available bug review process data from the Chromium bug tracking platform [97] that successfully received a bounty between January 2014 and November 2023. We used the customized Python script to collect the data in early November 2024. In total, we scraped 1,727 reporting process records, covering the entire lifecycle: from report submission to the first response, assignment, fix, and reward decision.

To analyze the responsiveness aspect of bug bounty fairness, we measured and computed the following historical metrics: (1) Time from submission to first response; (2) Time from submission to assignment; (3) Time from submission to fix; (4) Time from submission to bounty decision. We define the first response as the reviewer's first reply or actions such as assign, and locate it based on the content of the reply and marks. For assignments and fixes, we directly determine the corresponding posts and times based on marks. For decision-making on rewards, we filter using keywords such as "congrats" to obtain the complete duration. All identified posts were manually checked by the authors to ensure they were at the correct stage.

Most prior studies measuring different response times in Chromium have not focused on fairness; instead, they have primarily examined aspects related to bug-handling productivity or platform efficiency (see Table 3). In contrast, our work captures multiple sub-stages of the vulnerability review process from a hunter-centered perspective. We then compare these findings against the Google's own commitment and benchmark SLA, specifically through the viewpoint of fairness.

For first response time, the overall average is 20 hours and 22 minutes, with response times remaining under 1 day for nearly all years, as shown in Figure 3. The only exception was 2022, where the average extended to 1 day and 4 hours. HackerOne sets the first response standard at 5 days, while the recommended is 1 day. Comparing Chromium's performance to these standards, in Figure 4, we find that 74.89% of all first responses occurred within 1 day, meeting the recommended best practice. Moreover, the compliance rate is more than 60% every year, 98.14% of responses were completed within 5 days, meeting HackerOne's standard requirement. This indicates that Chromium's initial response to bug reports is typically fast and efficient, aligning closely with the best practice.

For time to triage, as illustrated in Figure 5, the average response time has consistently remained within 1 to 2 days, with an overall mean of 1 day and 23 hours. The only exception occurred in 2017, where the average time extended to 3 days due to outlier cases where triage took more than a month. HackerOne defines 10 days as the standard response time

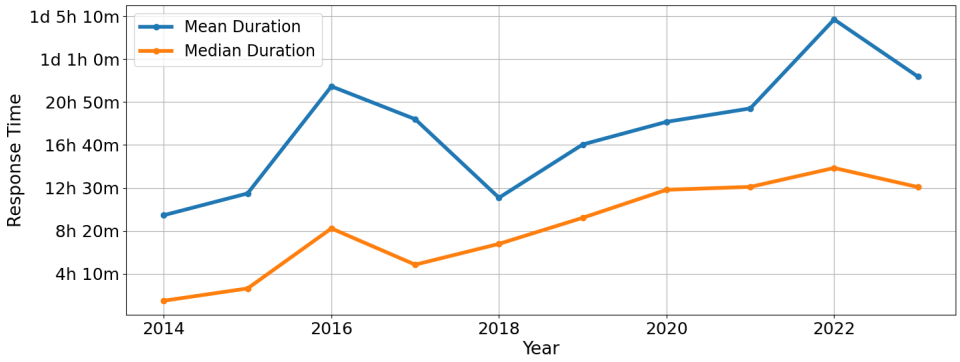

Fig. 3. Yearly average Submission-to-First Response duration, 2014-2023

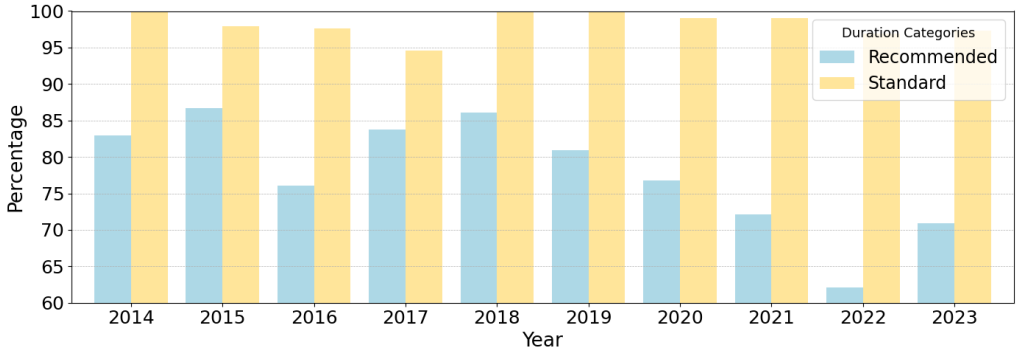

Fig. 4. Yearly percentages of first response durations within recommended and standard, 2014-2023

for triage, whereas the recommended is 2 days. When comparing Chromium's triage performance to these benchmarks, as illustrated in Figure 6, we find that 74.01% of all triage cases in the dataset were completed within 2 days, aligning with the recommended best practice. Moreover, 97.83% of all cases met HackerOne's 10-day triage standard, indicating that Chromium's bug bounty program performs well in ensuring timely triage decisions.

For fix times shown in Figure 7, the overall average duration is 54 days and 12 hours, but a clear downward trend has been observed in recent years, with 2023 seeing a significantly reduced average of 20 days and 9 hours. This improvement aligns with Chromium's previously stated goal: "For critical vulnerabilities, we aim to deploy the patch to all Chrome users in under 30 days." Our analysis shows that 68.51% of all reported bugs meet this standard, demonstrating progress in reducing fix times.

In contrast, HackerOne does not publicly provide a standardized response efficiency metric for fix times. This is likely due to the diverse range of vendors on the platform, whose internal security workflows, engineering resources, and patch deployment capabilities vary widely. Unlike Chromium, which operates as a unified entity with centralized patch deployment, HackerOne serves as a hosting platform for multiple organizations, each with its own vulnerability management process. Furthermore, HackerOne does not have direct visibility into whether or when a reported

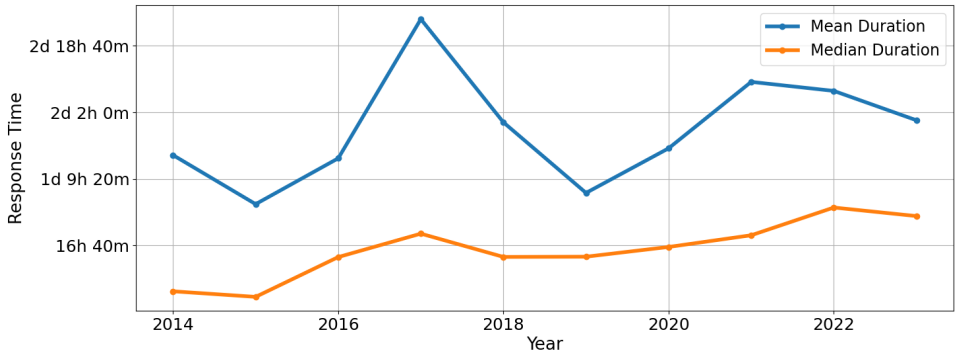

Fig. 5. Yearly average Submission-to-Triage duration, 2014-2023

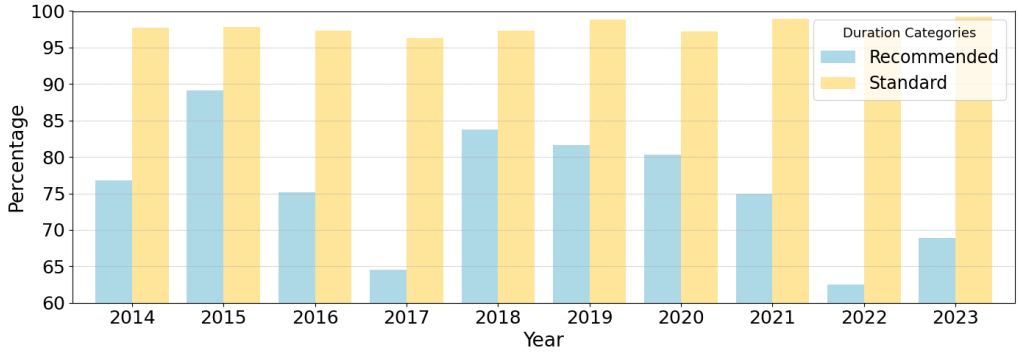

Fig. 6. Yearly percentages of triage durations within recommended and standard, 2014-2023

vulnerability is actually fixed, as vendors are responsible for implementing patches independently. This structural difference makes it challenging to establish a universal benchmark for fix times across HackerOne's ecosystem.

Due to Chromium's policy of determining bounty amounts only after a bug has been fixed or during the patching process, its decision-making time is consistently longer than the time taken to fix the issue, with an average of 73 days and 20 hours, as shown in Figure 8. Nearly half of the decision time is less than one month (46.33%). Like fix time, the time taken to reach bounty decisions has also shown a clear downward trend, with an average of 46 days and 8 hours in the past 3 years. On HackerOne, the timing of bounty payments varies depending on the specific program. Some programs issue payments before the vulnerability is fixed, and payment does not necessarily indicate that the issue has been resolved. As a result, HackerOne's recommended timeline for when a bounty should be paid is one day after triage is completed.

By comparison, Chromium takes significantly longer to make a bounty decision, even before considering the additional time required for payment processing. The two systems operate differently - HackerOne, as a platform, holds programs' funds in escrow, allowing for rapid disbursement of rewards. For Chromium, although we can not specifically measured payment processing time, we have observed a reduction in the time mentioned in bounty payment notifications, decreasing from around a month in the early years to approximately a week in recent cases. Additionally,

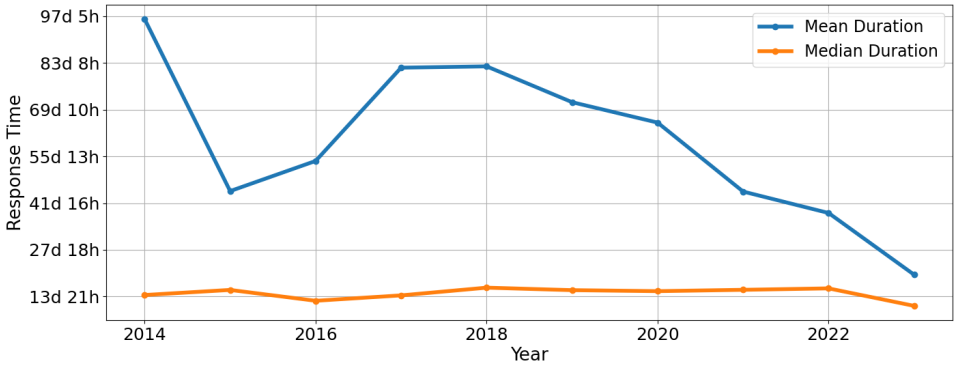

Fig. 7. Yearly average Submission-to-Fix duration, 2014-2023

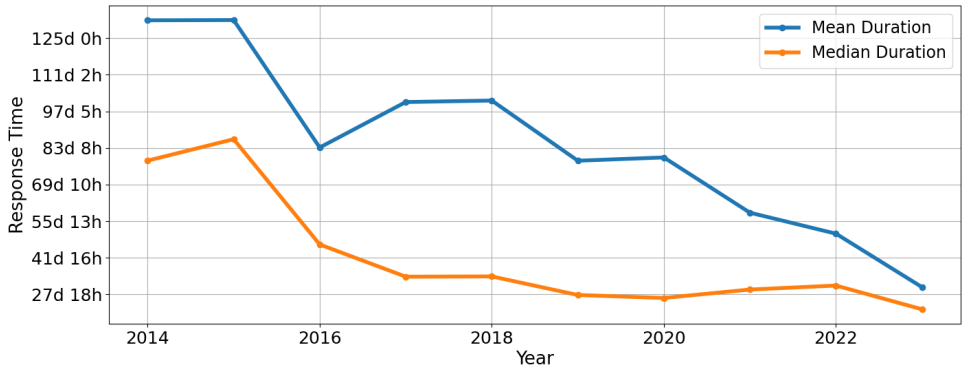

Fig. 8. Yearly average Submission-to-Reward Decision duration, 2014-2023

the number of follow-up inquiries from hunters regarding payment delays has also declined, suggesting improved efficiency in the payout process.

Overall, as a bug bounty program, Chromium generally meets or exceeds the benchmarks set by HackerOne in various aspects of responsiveness and efficiency. Our analysis indicates that first response times are consistently fast, and time to triage also remains well within the expectations. While Chromium's fix time remains long, but this has improved in recent years. In terms of bounty decision times, Chromium takes a noticeably longer time to determine rewards than HackerOne. Given that HackerOne's recommended standard for bounty payout is one day after triage, Chromium appears to lag in this area. However, the payment time itself has improved, with historical payout times decreasing from about one month to approximately one week, and inquiries from hunters about delayed payments have also all but disappeared.

The case study of Chromium demonstrates that the historical metrics we propose can effectively capture improvements in bug bounty programs. The case highlights how responsiveness indicators, such as first response time, triage duration, fix time, and bounty decision time, these can reflect refinements in program efficiency and fairness. By analyzing

these historical trends, we can observe efforts of Chromium's efforts to make its bounty process more predictable and transparent.

7 DISCUSSION

Given our work aims to improve fairness in bug bounty ecosystem, we discuss the steps that can be taken by different stakeholders, include vendors, internal roles, hackers, and third parties.

7.1 Vendors

Vendors can unilaterally make their bug bounty programs more fair for hackers. Concrete options, loosely ranked from easiest to hardest to implement, for vendors include: Establishing a clear safe harbor policy to protect security researchers from prosecution, ensuring they can report without legal risks. Additionally, vendors can improve the readability of their bug bounty policies, making expectations and guidelines clearer to hunters. Next, to enhance responsiveness, vendors can adopt the SLA and report their historical performance. The last improvement is creating an appeals procedure, allowing hunters to challenge decisions. This can range from an internal review of complaints to accepting input from third-party mediators, such as platforms, and even removing clauses that waive a hacker's right to bring a legal claim.

In addition, there are loftier goals that are difficult to implement and measure, such as clarifying ambiguity in policies, consistently enforcing the policy, and showing more respect to hackers. The first two recommendations are not new, and are only included to address vendors are yet to adopt. Creating a safe harbor is a basic act of corporate responsibility. Amateur security researchers already face significant legal uncertainty related to computer crime laws [3]. Organizations should create a responsible disclosure policy (RDP), even if the organization cannot afford to fund a bug bounty program. Further, the vendors should consider how complex and readable the document is given that one study found a third of policies are "somewhat difficult to read" [45].

Few vendors have adopted an SLA for their bug bounty program. Further, the significance depends entirely on the thresholds that have been set. For example, Section 6 showed that Chromium achieved close to perfect compliance with the standard response durations introduced by HackerOne, but compliance drops to 60% if the recommended response times are considered. Nevertheless, reporting on the average duration provides informational value to potential hunters, as well as creating accountability for internal teams

Finally, vendors could commit to an appeals process, even if they are only considered internally. This seems necessary given we know that prior work has shown experts inconsistently evaluate vulnerabilities [24]. The discussion resulting from appeals can help evaluators discuss discrepancies and converge on a more consistent policy. Appeals to third-parties like platforms and courts help because they help disseminate norms across the whole bug bounty ecosystem. While this appears to result in more pay-outs to hackers, the result may actually be that hackers may even accept lower payouts if they can be more easily predicted. Uncertainty comes at a higher price in other areas of the economy.

7.2 Hackers

Hackers lack agency alone, but appear to have more power as a collective. In the Chinese ecosystem, teams appear to have both implicit and explicit influence [9]. Explicit influence is evidenced by anecdotes where team leaders intervene to have a decision over-turned. Implicit influence is evidenced by the general feeling that team members are treated with more respect. In this way, visibility being part of a collective may improve hacker relations, at least for members of that collective. However, explicit demands need to be made to achieve more systemic change, such as by asking

vendors to commit to SLAs or appeals processes. Such moves would require a large collective, such as the alliance of teams that has formed in China in order to demand a new bug bounty decision structure that involves hackers as peer-reviewers [6].

However, collectivization requires hackers forgoing independence. This may be difficult given the free-lance nature of bug bounty research is attractive. Further, collectives will be undermined by churn in hacker activity, such as hackers leaving to start a full-time job. These challenges undermine the viability of hacker collectives.

7.3 Third-Parties

Bug bounty platforms like HackerOne and BugCrowd have a clear interest in ensuring fair bug bounty decisions. Indeed, HackerOne recommend minimum response times, akin to a service-level agreement, and a mediation service, which represents a soft appeal. However, these initiatives are merely advisory, likely reflecting the reality that vendors creating BBPs is holding back growth in the ecosystem rather than the supply of hackers. Until that changes, bug bounty platforms are unlikely to force behavioral change, unless the unfairness is egregious.

An interesting direction for researchers is to create tools to audit RDPs and BBPs at scale. This might involve shaming organizations who do not provide safe harbors for legitimate research, similar to academic works trying hold ISPs accountable foir hosting abuse [60]. While bug bounty policies can be measured, particularly with advanced NLP methods, it is more difficult to study historic bug bounty reports and decisions. Chromium show remarkable transparency in publishing redacted email logs of their successful reports, which enabled our case-study. However, the study is exposed to a sample bias as we do not observe response times for unsuccessful bug reports. It remains to be seen how researchers can better collect bug reports.

Finally, governments have a role in clarifying computer crime laws [2]. The core problem is that laws apply to the broadly defined concept of unauthorized access regardless of intent, which allows vendors to prosecute researchers unless the policy affirmatively creates a safe harbor. Beyond this criminal issue, there is a potential question of whether bug hunters should be considered as vulnerable researchers in the gig economy. Although we have emphasized unfairness in this article, we concede that the long-term life prospects of bug hunters is likely very good compared to say a food delivery worker. As such, bug hunters should not be considered as vulnerable gig economy workers.

7.4 Fairness Within Internal Roles

While much of our work focuses on unfairness experienced by bug hunters, it is equally important to acknowledge the challenges faced by individual triagers and platform operators. These employees are responsible for managing a large volume of incoming reports, many of which may be invalid, low-quality, or duplicates [100]. This creates substantial operational burdens, requiring triage teams to allocate time and resources efficiently, while ensuring that only valid and impactful reports are rewarded.

One could argue that fairness is also meaningful for internal operational actors, such as triagers and platform operators. Ensuring fairness for these individuals involves not only distributing workloads equitably and providing sufficient compensation, but also upholding interactional fairness. These individuals deserve respectful, reasonable interactions from both vendors and hackers, rather than being treated as faceless gatekeepers.. This raises the important design question of how to build bug bounty systems that safeguard fairness at multiple layers: not only between companies and external participants but also within the internal organizational roles that uphold the system. This challenge could become a promising direction for future research.

8 CONCLUSION

Bug bounty programs have historically been structured against the interests of hunters. Vendors unilaterally interpret ambiguous bounty policies, which typically waive the hackers' rights. This leaves hackers unable to challenge decisions or seek redress for delays, let alone ask to be treated with respect in day-to-day interactions.

Studying the literature suggests unfairness can be decomposed into: (1) Prosecution risk: uncertainty about protection from legal actions. (2) Policy ambiguity: adverse decisions related to program scope, severity classification, and reward amount. (3) (Un)responsiveness: delays in communications, time-to-patch and payment. (4) Hacker relations: subjective feelings of disrespect. Some of these dimensions have been measured by prior work [9, 49], whereas others have not. We then map organizational justice theory to these unfairness factors observed in practice.

Vendors can unilaterally improve the situation by introducing SLAs that commit to responding within certain time frames. We conducted a case-study showing Google's Chromium adheres in 98%+ of cases to the standard response times outlined by HackerOne, whereas around 74% adhere to "recommended best practice" response times [92]. We recommend other bug bounty platforms report this information, especially given we could do run the calculations with open-source, unstructured data.

Solving the ambiguity in bug bounty policies will be more challenging. Unfortunately, this dimension of unfairness is harder to measure. Further, vendors are unwilling to subject decisions to independent appeals processes, as evidenced by the weakness of HackerOne's mediation policy [96]. However, hackers can respond by collectivizing. Chinese bug hunters report that joining teams has resulted in more fair treatment along all four fairness dimensions [9], even with relatively small teams of typically 10 members and 500 maximum. It remains to be seen what hackers could achieve were they to organize into far larger groupings that could effectively boycott vendors who do not commit to addressing fairness.

REFERENCES

[1] Anderson, R., and Moore, T.: "Information security economics–and beyond," in *Annual International Cryptology Conference*, Berlin, Heidelberg, 2007, pp. 68-91.

[2] Graves, J.T., Acquisti, A., Anderson, R.: "Perception versus punishment in cybercrime," *Journal of Criminal Law and Criminology*, vol.109, no.2, 2019, pp. 313–36.

[3] Hrle, T., Milad, M., Li, J., and Woods, D.: ""Just a tool, until you stab someone with it": Exploring Reddit Users' Questions and Advice on the Legality of Port Scans," in *European Symposium on Usable Security(EuroUSEC)*, 2024, pp. 322-336.

[4] Perlroth, N.: "This is how they tell me the world ends: The cyberweapons arms race," Bloomsbury Publishing USA, 2021.

[5] Walshe, T., and Simpson, A.: "An empirical study of bug bounty programs," in *2020 IEEE 2nd International Workshop on Intelligent Bug Fixing (IBF)*, IEEE, 2020.

[6] White hat hackers self-discipline committee code of conduct V1.0 (in Chinese), Available: https://www.0xu.cn/article/intelnet/safe/2985.html

[7] ChaMd5 Team's Public Statement on Terminating Vulnerability Reporting with Ctrip Security Response Center (in Chinese), Available: https://mp.weixin.qq.com/s/bFImTdovQiJ71PPnMDgVWw

[8] Walshe, T., and Simpson, A.: "A longitudinal study of hacker behaviour," in *37th ACM/SIGAPP Symposium on Applied Computing*, 2022, pp. 1465–1474.

[9] Piao, Y., Hrle, T., Woods, D., and Anderson, R.: "Study Club, Labor Union or Start-Up? Characterizing Teams and Collaboration in the Bug Bounty Ecosystem," in *2025 IEEE Symposium on Security and Privacy (SP)*, IEEE Computer Society, 2025, pp. 20-20.

[10] Fulton, K. R., Katcher, S., Song, K., Chetty, M., *et al.*: "Vulnerability discovery for all: Experiences of marginalization in vulnerability discovery," in *44th IEEE Symposium on Security and Privacy (SP)*, 2023, pp. 1997–2014.

[11] Leventhal, G. S.: "What Should be Done with Equity Theory? New Approaches to the Study of Fairness in Social Relationships," in *Social Exchange: Advances in Theory and Research*, Springer US, Boston, MA, 1980, pp. 27-55.

[12] Akgul, O., Eghtesad, T., Elazari, A., Gnawali, O., *et al.*: "Bug hunters' perspectives on the challenges and benefits of the bug bounty ecosystem," in *32nd USENIX Security Symposium (USENIX Security)*, 2023, pp. 2275–2291.

[13] "Hacking capitalism—A new profit-sharing workforce platform," Available: https://www.lutasecurity.com/post/hacking-capitalism-a-new-profit-sharing-workforce-platform

[14] "The price of security: How bug bounty platforms fail ethical hackers," [LinkedIn Post], 2024. Available: https://www.linkedin.com/pulse/price-security-how-bug-bounty-platforms-fail-ethical-hackers-isufi-terae/

[15] Blakley, B., and Cranor, L.: "Katie Moussouris: Vulnerability disclosure and security workforce development," *IEEE Security & Privacy*, vol. 21, no. 1, 2023, pp. 11–18.

[16] Ellis, R., Huang, K., Siegel, M., Moussouris, K., *et al.*: "Fixing a hole: The labor market for bugs," in Shrobe, H., Shrier, D. L., Pentland, A. (eds.), *New Solutions for Cybersecurity*, MIT Press, 2018, pp. 130–159.

[17] Shapiro, C.: "Information Rules: A Strategic Guide to the Network Economy," Harvard Business School Press, 1999.

[18] Böhme, R.: "A comparison of market approaches to software vulnerability disclosure," in *International Conference on Emerging Trends in Information and Communication Security*, Berlin, Heidelberg: Springer Berlin Heidelberg, 2006.

[19] Zhao, M., Laszka, A., and Grossklags, J.: "Devising effective policies for bug-bounty platforms and security vulnerability discovery," *Journal of Information Policy*, vol. 7, pp. 372-418, 2017.

[20] Miller, C.: "The Legitimate vulnerability market: the secretive world of 0-day exploit sales," in *WEIS*, 2007.

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

[22] Ellis, R., and Stevens, Y.: "Bounty everything: Hackers and the making of the global bug marketplace," Available at SSRN 4009275, 2022.

[23] Gamero-Garrido, A., Savage, S., Levchenko, K., and Snoeren, A. C.: "Quantifying the pressure of legal risks on third-party vulnerability research," in *ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 1501-1513.

[24] Wunder, J., Kurtz, A., Eichenmüller, C., Gassmann, F., *et al.*: "Shedding light on CVSS scoring inconsistencies: A user-centric study on evaluating widespread security vulnerabilities," in *2024 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2024, pp. 1102-1121.

[25] Ayala, J., Ngo, S., and Garcia, J.: "A Deep Dive Into How Open-Source Project Maintainers Review and Resolve Bug Bounty Reports," in *2025 IEEE Symposium on Security and Privacy (SP)*, San Francisco, CA, USA, 2025, pp. 63-63, doi: 10.1109/SP61157.2025.00063.

[26] Piao, Y., Lolla, H., and Woods, D.: "The long shadow of the Computer Fraud and Abuse Act: Exploring user discussions on the legality of vulnerability research on Reddit," *Rossfest Festschrift*, 2025, pp. 111-119.

[27] Li, Y., and Zhao, L.: "Collaborating with bounty hunters: how to encourage white hat hackers' participation in vulnerability crowdsourcing programs through formal and relational governance," *Information & Management*, vol. 59, no. 4, 2022, article 103648.

[28] Leventhal, G.S.: "What should be done with equity theory? New approaches to the study of fairness in social relationships," in *Social Exchange: Advances in Theory and Research*, Springer US, Boston, MA, 1980, pp. 27–55.

[29] Ambrose, M.L., and Arnaud, A.: "Are procedural justice and distributive justice conceptually distinct?" in *Handbook of Organizational Justice*, Psychology Press, 2013, pp. 59–84.

[30] Adams, J. S.: "Inequity in social exchange," in *Advances in Experimental Social Psychology*, vol. 2, Academic Press, 1965, pp. 267-299.

[31] Fieseler, C., Bucher, E., and Hoffmann, C. P.: "Unfairness by design? The perceived fairness of digital labor on crowdworking platforms," *Journal of Business Ethics*, vol. 156, 2019, pp. 987-1005.

[32] Bies, R.J., and Moag, J.S.: "Interactional justice: Communication criteria of fairness," in R.J. Lewicki, B.H. Sheppard, and M.H. Bazerman (eds.), *Research on Negotiations in Organizations*, 1986.

[33] Colquitt, J.A.: "On the dimensionality of organizational justice: A construct validation of a measure," *Journal of Applied Psychology*, vol.86, no.3, 2001, pp. 386–400.

[34] Noordegraaf, J. E., and Weulen Kranenbarg, M.: "Why Do Young People Start and Continue with Ethical Hacking? A Qualitative Study on Individual and Social Aspects in the Lives of Ethical Hackers," *Criminology & Public Policy*, vol. 22, no. 4, 2023, pp. 803-824.

[35] Camp, L. J., and Wolfram, C.: "Pricing security: A market in vulnerabilities," in *Economics of Information Security*, Springer US, Boston, MA, 2004, pp. 17-34.

[36] Gandal, N., Zrahia, A., Markovich, S., and Riordan, M.: "The Simple Economics of an External Shock on a Crowdsourced 'Bug Bounty Platform'," CEPR Discussion Papers, No. 17443, 2022.

[37] Laszka, A., Zhao, M., and Grossklags, J.: "Banishing misaligned incentives for validating reports in bug-bounty platforms," in *Computer Security–ESORICS 2016: 21st European Symposium on Research in Computer Security, Proceedings, Part II*, vol. 21, Springer International Publishing, Heraklion, Greece, 2016.

[38] Kuehn, A., and Mueller, M.: "Analyzing bug bounty programs: An institutional perspective on the economics of software vulnerabilities," in *2014 TPRC Conference Paper*, 2014.

[39] Egelman, S., Herley, C., and Van Oorschot, P. C.: "Markets for zero-day exploits: Ethics and implications," in *New Security Paradigms Workshop*, 2013.

[40] Alexopoulos, N., Meneely, A., Arnouts, D., and Mühlhäuser, M.: "Who are vulnerability reporters? A large-scale empirical study on FLOSS," in *15th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, 2021, pp. 1-12.

[41] Dellago, M., Woods, D.W., and Simpson, A.C.: "Characterising 0-day exploit brokers," in *21st Workshop on the Economics of Information Security (WEIS)*, 2022.

[42] Hata, H., Guo, M., and Babar, M. A.: "Understanding the heterogeneity of contributors in bug bounty programs," in *2017 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, IEEE, 2017.

[43] Luna, D., Allodi, L., and Cremonini, M.: "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*, 2019.

[44] Maillart, T., Zhao, M., Grossklags, J., and Chuang, J.: "Given enough eyeballs, all bugs are shallow? Revisiting Eric Raymond with bug bounty programs," *Journal of Cybersecurity*, vol. 3, no. 2, 2017, pp. 81-90.

[45] Walshe, T., and Simpson, A. C.: "Coordinated vulnerability disclosure programme effectiveness: Issues and recommendations," *Computers & Security*, vol. 123, 2022, article 102936.

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

[47] Finifter, M., Akhawe, D., and Wagner, D.: "An empirical study of vulnerability rewards programs," in *22nd USENIX Security Symposium (USENIX Security 13)*, 2013, pp. 273-288.

[48] Aljedaani, W., Mkaouer, M. W., Eler, M. M., and Kessentini, M.: "Empirical Investigation of Accessibility Bug Reports in Mobile Platforms: A Chromium Case Study," in *CHI Conference on Human Factors in Computing Systems*, 2024, pp. 1-17.

[49] Atefi, S., Sivagnanam, A., Ayman, A., Grossklags, J., and Laszka, A.: "The benefits of vulnerability discovery and bug bounty programs: Case studies of Chromium and Firefox," in *ACM Web Conference*, 2023, pp. 2209-2219.

[50] Laszka, A., Zhao, M., Malbari, A., and Grossklags, J.: "The Rules of Engagement for Bug Bounty Programs," in *Financial Cryptography and Data Security: 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26–March 2, 2018, Revised Selected Papers*, vol. 22, pp. 138-159, Springer Berlin Heidelberg, 2018.

[51] Cavalcanti, Y. C., Neto, P. A. M. S., Lucrédio, D., Vale, T., *et al.*: "The bug report duplication problem: an exploratory study," *Software Quality Journal*, vol. 21, 2013, pp. 39-66.

[52] Jacobs, J., Romanosky, S., Suciu, O., Edwards, B., *et al.*: "Enhancing Vulnerability Prioritization: Data-driven Exploit Predictions with Community-Driven Insights," in *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, IEEE, 2023, pp. 194-206.

[53] Park, S., and Albert, K.: "A Researcher's Guide to Some Legal Risks of Security Research," a joint publication of the Cyberlaw Clinic at Harvard Law School and the Electronic Frontier Foundation, 2020.

[54] Etcovitch, D., and van der Merwe, T.: "Coming in from the Cold: A Safe Harbor from the CFAA and the DMCA § 1201 for Security Researchers," Berkman Klein Center Research Publication 2018-4, 2018.

[55] de Smale, S., van Dijk, R., Bouwman, X., van der Ham, J., *et al.*: "No One Drinks from the Firehose: How Organizations Filter and Prioritize Vulnerability Information," in *2023 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2023, pp. 1980-1996.

[56] Egelman, S., Herley, C., and Van Oorschot, P.C.: "Markets for zero-day exploits: Ethics and implications," in *Proceedings of the 2013 New Security Paradigms Workshop*, 2013, pp. 41–46.

[57] Perlroth, N.: "This is how they tell me the world ends: The cyberweapons arms race," Bloomsbury Publishing USA, 2021.

[58] Meakins, J.: "A zero-sum game: The zero-day market in 2018," *Journal of Cyber Policy*, vol.4, no.1, 2019, pp. 60–71, Taylor & Francis.

[59] Walshe, T., and Simpson, A.: "Towards a Greater Understanding of Coordinated Vulnerability Disclosure Policy Documents," *Digital Threats: Research and Practice*, vol. 4, no. 2, 2023, pp. 1-36.

[60] Stone-Gross, B., Kruegel, C., Almeroth, K., Moser, A., *et al.*: "Fire: Finding rogue networks," in *Annual Computer Security Applications Conference*, pp. 231–240, 2009.

[61] Wachs, J.: "Making Markets for Information Security: The Role of Online Platforms in Bug Bounty Programs," arXiv preprint arXiv:2204.06905, 2022.

[62] Hirschman, A. O.: "Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States," Harvard University Press, 1970.

[63] Maffie, M., and Gough, M. D.: "Bargaining Against the Machine: A Theory of Bargaining Power in the Gig Economy," in *Advances in Industrial and Labor Relations*, Emerald Publishing Limited, 2023, pp. 83-99.

[64] Munaiah, N., Camilo, F., Wigham, W., Meneely, A., *et al.*: "Do Bugs Foreshadow Vulnerabilities? An In-depth Study of the Chromium Project," *Empirical Software Engineering*, vol. 22, 2017, pp. 1305-1347.

[65] Bettenburg, N., Just, S., Schröter, A., Weiss, C., *et al.*: "What Makes a Good Bug Report?," in *16th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, 2008, pp. 308-318.

[66] Arneson, R.J.: "Equality and equal opportunity for welfare," *Philosophical Studies*, vol.56, no.1, 1989, pp. 77–93.

[67] Binns, R.: "On the apparent conflict between individual and group fairness," in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020, pp. 514–524.

[68] Tyler, T.R.: "Social justice," in Mikulincer, M., Shaver, P.R., Dovidio, J.F., *et al.* (eds.), *APA Handbook of Personality and Social Psychology; Vol. 2. Group Processes*, American Psychological Association, 2015, pp. 95–122.

[69] Fleisher, W.: "What's fair about individual fairness?" in *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 2021, pp. 480–490.

[70] Binns, R.: "Fairness in machine learning: Lessons from political philosophy," *Proceedings of Machine Learning Research*, 2017.

[71] Tyler, T., Degoey, P., and Smith, H.: "Understanding why the justice of group procedures matters: A test of the psychological dynamics of the group-value model," *Journal of Personality and Social Psychology*, vol.70, no.5, 1996, pp. 913–930.

[72] Colquitt, J.A., and Jackson, C.L.: "Justice in teams: The context sensitivity of justice rules across individual and team contexts," *Journal of Applied Social Psychology*, vol.36, no.4, 2006, pp. 868–899.

[73] Davies, S., and Roper, M.: "What's in a Bug Report?," in *8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, 2014, pp. 1-10.

[74] Mu, D., Cuevas, A., Yang, L., Hu, H., *et al.*: "Understanding the Reproducibility of Crowd-Reported Security Vulnerabilities," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 919-936.

[75] Rajbhandari, A., Zibran, M. F., and Eishita, F. Z.: "Security versus Performance Bugs: How Bugs Are Handled in the Chromium Project," in *2022 IEEE/ACIS 20th International Conference on Software Engineering Research, Management and Applications (SERA)*, IEEE, 2022, pp. 70-76.

[76] Guo, P. J., Zimmermann, T., Nagappan, N., and Murphy, B.: ""Not my bug!" and Other Reasons for Software Bug Report Reassignments," in *ACM 2011 Conference on Computer Supported Cooperative Work*, 2011, pp. 395-404.

[77] Franken, G., Van Goethem, T., Desmet, L., and Joosen, W.: "A Bug's Life: Analyzing the Lifecycle and Mitigation Process of Content Security Policy Bugs," in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 3673-3690.

[78] Elazari, A.: "Hacking the Law: Are Bug Bounties a True Safe Harbor?," in *Enigma 2018 (Enigma 2018)*, 2018.

[79] Haddix, J.: "The Darkest Side of Bug Bounty," presented at DEF CON 32, [Online Video], 2022. Available: https://www.classcentral.com/course/youtube-def-con-32-the-darkest-side-of-bug-bounty-jason-haddix-360301.

[80] Vlad, C.: "Don't be fooled by the mirage of bug bounty hunting," [LinkedIn Post], 2024. Available: https://www.linkedin.com/posts/cristivlad_def-con-32-the-darkest-side-of-bug-bounty-activity-7264938523558858753-PS-X/.

[81] "Microsoft Bug Bounty Program's (MSRC) response was poor: Initially, they misjudged and dismissed the issue entirely," Hacker News, 2021. Available: https://news.ycombinator.com/item?id=29472478.

[82] Le Coguic, G.: "Cons of Bug Bounty," 10degres, 2018. Available: https://10degres.net/cons-of-bug-bounty/.

[83] Gupta, V.: "Another Dark Reality of Bug Hunting," LinkedIn, 2024. Available: https://www.linkedin.com/pulse/another-dark-reality-bug-hunting-vijay-gupta--kyjhc/.

[84] Furnell, S.: "The Cybersecurity Workforce and Skills," *Computers & Security*, vol. 100, 2021, article 102080.

[85] Budd, J. W.: "Employment with a Human Face: Balancing Efficiency, Equity, and Voice," Cornell University Press, 2019.

[86] Colquitt, J.A., Greenberg, J., Zapata-Phelan, C.P.: "Organizational justice," in *The Oxford Handbook of Organizational Psychology*, vol.1, 2012, pp. 526–547.

[87] Beugré, C.D., and Baron, R.A.: "Perceptions of systemic justice: The effects of distributive, procedural, and interactional justice," *Journal of Applied Social Psychology*, vol.31, no.2, 2001, pp. 324–339.

[88] Greenberg, J.: "Organizational justice: Yesterday, today, and tomorrow," *Journal of Management*, vol.16, no.2, 1990, pp. 399–432.

[89] Colquitt, J.A., Greenberg, J., Zapata-Phelan, C.P.: "What is organizational justice? A historical overview," in *Handbook of Organizational Justice*, Psychology Press, 2013, pp. 3–56.

[90] Google: "Update to Policy Regarding Unactionable Reports and Duplicates," Chrome Vulnerability Reward Program Rules, 2023. Available: https://bughunters.google.com/about/rules/chrome-friends/5745167867576320/chrome-vulnerability-reward-program-rules#duplicate-reports.

[91] Al-Banna, M., Benatallah, B., Schlagwein, D., Bertino, E., *et al.*: "Friendly Hackers to the Rescue: How Organizations Perceive Crowdsourced Vulnerability Discovery," in *Pacific Asia Conference on Information Systems (PACIS'18)*, 2018, p. 230.

[92] HackerOne: "Response Target Metrics," HackerOne Documentation, 2023. Available: https://docs.hackerone.com/en/articles/8505145-response-target-metrics.

[93] Reddit User: "Which companies scammed you in bug bounty programs?" Reddit, discussion thread, 2023. Available: https://www.reddit.com/r/bugbounty/comments/1ice7z9/which_companies_scammed_you_in_bug_bounty_programs/.

[94] Deutsch, M.: "Equity, equality, and need: What determines which value will be used as the basis of distributive justice?" *Journal of Social Issues*, vol.31, no.3, 1975, pp. 137–149.

[95] Nowakowski, J.M., and Conlon, D.E.: "Organizational justice: Looking back, looking forward," *International Journal of Conflict Management*, vol.16, no.1, 2005, pp. 4–29.

[96] HackerOne: "Hacker Mediation," HackerOne Documentation, 2023. Available: https://docs.hackerone.com/en/articles/8466617-hacker-mediation.

[97] "Chromium Issue Tracker," The Chromium Projects. [Online]. Available: https://issues.chromium.org/issues.

[98] Shafigh, S., Benatallah, B., Rodríguez, C., and Al-Banna, M.: "Why Some Bug-Bounty Vulnerability Reports Are Invalid? Study of Bug-Bounty Reports and Developing an Out-of-Scope Taxonomy Model," in *15th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, 2021, pp. 1-6.

[99] Paul, R., Turzo, A.K., and Bosu, A.: "Why security defects go unnoticed during code reviews? A case-control study of the Chromium OS project," in *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, IEEE, 2021, pp. 1373–1385.

[100] Fan, Y., Xia, X., Lo, D., *et al.*: "Chaff from the wheat: Characterizing and determining valid bug reports," *IEEE Transactions on Software Engineering*, vol.46, no.5, 2018, pp. 495–525.

[101] Meneely, A., Rodriguez Tejeda, A.C., Spates, B., *et al.*: "An empirical investigation of socio-technical code review metrics and security vulnerabilities," in *Proceedings of the 6th International Workshop on Social Software Engineering*, 2014, pp. 37–44.

[102] Munaiah, N., Camilo, F., Wigham, W., *et al.*: "Do bugs foreshadow vulnerabilities? An in-depth study of the Chromium project," *Empirical Software Engineering*, vol.22, 2017, pp. 1305–1347.

[103] Aljedaani, W., Javed, Y., Alenezi, M., *et al.*: "LDA categorization of security bug reports in Chromium projects," in *Proceedings of the 2020 European Symposium on Software Engineering*, 2020, pp. 154–161.

A APPENDIX A

Table 4. Description and approaches of metrics

Approach	Metric	Description
Historical	Average First Response Time	The average duration between a report submission and the vendor's initial response.
	Average Triage Time	The average time from submission to the assignment of a reviewer.
	Average Fix Time	The average duration from bug submission to its resolution.
	Average Bounty Decision Time	The average time taken from report submission to the final bounty determination.
	Average Evaluation Time	The average time from the first assignment to reassignment or fixed.
	Average Dispute Resolution Time	The average time taken to resolve disputes or reevaluate cases.
	Average Payment Time	The average time taken to pay rewards.
	Reward-Severity Ratio	Whether the bounty matches the reward range for vulnerabilities of the declared severity.
	Dispute Resolution Rate	The ratio of disputes resolved amicably versus those unresolved or escalated.
	Previous Prosecution Record	The history of programs taking legal action against hunters.
Perception	Satisfaction Score	Hunters' perceptions of professionalism and respect in communication based on their feedback.
	Expectation Consistency	Survey assessing hunters' reward expectations compared to the actual rewards received..
	Hunter Comprehension	Questionnaires to assess hunters' understanding of policies and processes.
	Perceived Researcher Protection	Hunters' perceptions of the platforms' or safe harbors' protections.
Contractual	Clarity of Submission Guidelines	Whether submission guidelines clearly specify required information, and whether templates or examples are provided.
	Clarity of Scope Definition	Whether policies clearly outline in-scope and out-of-scope vulnerabilities with specific examples.
	Clarity of Legal Risk Disclosure	The existence of clear legal disclaimers and whether applicable laws or jurisdictions are specified.
	Policy Readability Test	The complexity and readability of policy texts.