PLATFORM STRATEGY RESEARCH SYMPOSIUM 2025
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0000-0000, EISSN 0000-0000

Platform Openness: Evidence from Bug Bounty Programs*

Aviram Zrahia
The Department of Public Policy, Tel Aviv University, aviramzrahia@mail.tau.ac.il

Sarit Markovich
Kellogg School of Management, Northwestern University, s-markovich@kellogg.northwestern.edu

Tobias Kretschmer
ISTO, LMU Munich and Nova SBE, Lisbon, t.kretschmer@lmu.de

Abstract. We study the role of openness on platform markets, using bug bounty programs as our empirical setting. We define openness as the organization's decision to run a public (open) or private (closed) program on a bug bounty platform, where private programs restrict access to selected researchers, and public programs allow unrestricted participation from anyone on the platform. While openness may expand participation, it can also reduce submission quality, introducing inefficiencies.

Leveraging a large-scale proprietary dataset from Bugcrowd, we analyze the performance differences between openness structures and the effects of transitioning from a closed to an open model. Our findings reveal that private programs outperform public ones, but their attractiveness and activity level decline over time. Transitioning to an open structure can revitalize these programs; however, this activity surge diminishes faster than other open programs. Despite this, its effect remains significant compared to the programs' earlier closed state, providing positive value to organizations.

Our results highlight the trade-offs in platform governance and provide insights for decision-makers designing their openness strategy. In the cybersecurity context, we offer actionable advice to organizations seeking to optimize vulnerability discovery strategies while balancing researcher engagement and the efficiency of their bug bounty program.

Key words: Openness Strategy, Network Effects, Platform Governance, Software Vulnerabilities, Bug Bounty Programs, Cybersecurity Economics.

1. Introduction

Two-sided markets facilitate interactions between two distinct user groups, such as buyers and sellers, software developers and users, or firms and workers, and typically benefit from cross-side network effects,

* Working Paper.

where the value to one group increases as participation from the other group grows. The literature on platform governance has extensively examined the decision of openness in two-sided markets (Kretschmer et al. 2022). In such markets, openness can generate positive feedback loops: reducing participation barriers leads to greater user adoption, which in turn attracts more participants from the other side, further reinforcing the platform's dominance (Gawer 2014). This effect has been widely studied in digital marketplaces, app ecosystems, and payment networks, where an open strategy often accelerates innovation, user engagement, and economic value creation (Boudreau 2010). However, the benefits of openness are less clear in markets without cross-side network effects, where additional participation does not inherently increase value for other user groups.

This paper explores the choice between open and closed structures using bug bounty programs hosted on a bug bounty platform as our empirical setting. Bug bounty platforms operate as two-sided markets, connecting organizations with ethical hackers to identify software vulnerabilities. While there are cross-side network effects on the platform level—more programs attract more researchers, which in turn attract more programs—individual bug bounty programs do not inherently benefit from cross-side network effects. Instead, increased participation may lead to greater competition among researchers, potentially discouraging high-quality contributors and thereby generating negative same-side network effects. This scenario raises a fundamental governance question: In the absence of cross-side network effects, should organizations adopt an open or closed platform structure? More generally, how should participants in a two-sided market market make their openness decision?

In markets with cross-side network effects, openness is often associated with increased innovation, greater market adoption, and enhanced value creation through complementor participation (Parker and Alstyne 2005, Gawer 2014). In contrast, in markets without such effects, openness may introduce diminishing returns by potentially reducing the quality and value of contributions without providing the benefit of increased participation.

To study these questions, we utilize a large proprietary panel dataset from Bugcrowd, a leading crowdsourced bug bounty platform.[1] Organizations active on the platform must choose whether to design their bug bounty program as a public program—a program that is open for all researchers—or as a private one, where participation is by invitation only. Public programs tend to attract a larger pool of researchers, leading to more contributions (submissions) and potentially a higher proportion of low-value reports. Private programs, by contrast, maintain a more selective researcher base, yielding fewer but potentially higher-quality contributions. This trade-off between public and private structures mirrors the quantity-quality trade-off associated with the open vs. closed decision in platform markets.

To support our analysis, we categorize submissions into three distinct outcomes: (i) unique vulnerability discoveries, (ii) duplicate submissions (previously reported vulnerabilities), and (iii) rejected submissions

[1] `https://www.bugcrowd.com/`

(non-security issues or invalid reports). These distinctions allow us to assess the relative quality and value of contributions across public and private programs.

Our statistical analysis highlights a trade-off between quantity and quality in public and private programs. Public programs facilitate broader researcher participation, yet the quality of submissions and researchers in private programs tend to be higher. Submission quality is assessed based on vulnerability severity (priority ranking) and outcome (unique, duplicate, or rejected). Researcher quality, in turn, is inferred from their proven expertise and historical performance—measured by their success in identifying high-priority vulnerabilities and receiving monetary rewards or acknowledgments. Additionally, our analysis confirms that, on average, private programs are more likely to yield unique discoveries and are less likely to receive a duplicate or rejected submission.

We further examine the implications of private-to-public openness transition by employing a differences-in-differences (DiD) framework to study the causal impact of programs switching from a closed to an open structure. Our findings reveal an interesting dynamic: while transitioning to a public program initially provides a boost in participation and contributions, this effect diminishes over time as interest in the program declines. Still, shifting to a public structure provides lasting effects after a program matures compared to programs that did not make that transition. Furthermore, programs that switch later in their lifecycle derive greater benefits, suggesting that timing is critical in maximizing the value of openness.

Unlike many digital platforms where openness is a platform-level decision, in bug bounty ecosystems, openness is determined by the organization at the individual program level. Our findings offer actionable insights for organizations and bug bounty platforms in designing an optimal openness strategy aligned with their objectives and constraints. Beyond bug bounty programs, our results have broader implications for platform governance and digital ecosystems. In markets without cross-side network effects, the choice of openness must be carefully calibrated to account for diminishing returns, competition effects, and cost considerations.

The remainder of this paper is structured as follows: Section 2 provides the theoretical background and related literature on bug bounty platforms. Section 3 details our dataset and methodology. Section 4 presents empirical findings, followed by a more profound statistical analysis in Section 5. Finally, Section 6 concludes with broader implications and future directions.

2. Background

2.1. Bug bounty programs

Bug bounty programs are structured arrangements between organizations and individual security researchers for exchanging vulnerabilities as products. Bug bounty platforms host multiple programs and facilitate this transaction. Governments and other policymakers promote them as a mainstream method for discovering and disclosing software vulnerabilities in various systems, products, and services, alongside other methods qualitatively evaluated by Zrahia (2024).

Organizations may provide monetary rewards for uniquely discovered vulnerabilities or acknowledge researchers for their findings through a monthly leaderboard or an all-time ranking. Paid programs are called Managed Bug Bounties (MBBs), while non-paying programs are known as Vulnerability Disclosure Programs (VDPs). We hereafter refer to this distinction as the *reward type*.

Submissions to a program are classified by the platform's or organization's professional teams into one of three classes. First, *Rejected* submissions are dismissed as invalid or beyond the program's scope. This category includes submissions that are not applicable, not reproducible, or out of scope.[2] Submissions that are not rejected fall under the category of *valid* submissions, which can be further divided into *unique* submissions—vulnerabilities that have not been previously identified and warrant a reward—and *duplicate* submissions, which refer to vulnerabilities that have already been discovered and reported by other researchers, and are typically not eligible for rewards. Duplicate and rejected submissions create inefficiencies, providing only limited value (highlighting individual researchers' capabilities) while consuming time and processing resources.

The described transaction can be viewed from a microeconomic perspective, similar to any other supply and demand market. In this view, researchers act as suppliers, while organizations represent the demand. The "product" consists of valid submissions, with pricing determined by the current market equilibrium.[3]

Figure B9 summarizes the interaction between all players and the submission workflow over Bugcrowd's platform (Zrahia et al. 2024). In line with this perspective, we hereafter describe key aspects related to the three market players: the platform, the organization, and the researchers.

The platform. Several key aspects characterize the governance of bug bounty platforms. First, these platforms act as a trustworthy intermediate entity, enabling a structured, safe, and legal discovery process while ensuring the interests of both parties are met (Miller 2007). In this intermediary role, the platforms reduce information asymmetries and transaction costs, enabling the creation of a market for security vulnerabilities (Zrahia 2024). Next, platforms create a framework for legal, regulatory, and privacy compliance by enforcing rules and terms of participation. Additionally, they support social welfare by facilitating the disclosure of vulnerabilities to the public post-discovery by offering a full or partial safe harbor policy—a practice that allows ethical researchers to submit their findings without fear of legal consequences. Next, they provide the technological infrastructure, submission workflow processes, and services, allowing efficient interaction between the organizations and the researchers. They also offer guidance, support, and training to both parties. For example, bounty pricing guidance to organizations[4] and security penetration

[2] The scope of a program delineates which vulnerabilities researchers are permitted to test, including the specific websites, applications, or devices that are considered in or out of scope. It also outlines the categories of vulnerabilities eligible for rewards and the testing methods allowed.

[3] Zrahia et al. (2024) offers a heuristic model describing bug bounty platforms in terms of supply and demand to study the impact of the COVID-19 exogenous shock on platform dynamics. It also provides more details regarding the platform, dataset, and submission workflow.

[4] See `https://www.bugcrowd.com/wp-content/uploads/2023/12/bugcrowd-whats-a-bug-worth.pdf`.

training to researchers.[5] Finally, some platforms also provide matching services, allowing organizations to find the best-fit researchers for their programs.[6]

The organization. A key attribute of bug bounty programs is their level of openness, which is determined by the organization's choice between a *public* and a *private* program. Public programs are open to all security researchers, allowing unrestricted participation, while private programs operate on an invitation-only basis, limiting access to researchers selected explicitly by the organization. This distinction means private programs have a closed supply-side structure, while public programs embody an open structure. Adopting an open or closed program structure is a governance decision an organization should make, influenced by its strategic goals and platform recommendations.

One of the considerations that might influence organizations when choosing between open and closed program structures is the processing costs. As elaborated by Zrahia et al. (2024), the number of valid submissions can be broken down into the total number of submissions multiplied by the accuracy of those submissions, defined as the ratio of valid submissions to total submissions. Higher accuracy values indicate that researchers are more thorough and precise in their submissions, which reduces the costs programs incur for processing these submissions. As shown later when discussing the quantity-quality tradeoff in section4.2, pre-selected researchers in private programs produce, on average, higher quality results than the entire researcher's cohort.

Another key characteristic that could change over time is the amount of payment per unique discovered vulnerability determined by the organization. This value is primarily correlated with the severity of the vulnerability discovered. However, several other factors also influence the payments. For example, the program's maturity plays a significant role—more mature programs tend to have fewer vulnerabilities, so they might appear less attractive to researchers. Therefore, organizations might offer higher rewards to attract key talent later in the program's lifecycle. Additionally, the extent to which the target was tested internally before the program was launched can affect the bounty amount, along with other considerations (Zrahia et al. 2024).

The researchers. We assume that a researcher's primary goal is to discover a vulnerability and produce a unique and valid submission that receives a monetary reward. The number of these submissions can be expressed as the product of the number of valid submissions and the probability of winning with a valid submission. This probability, influenced by the number of duplicate valid submissions, can be calculated using the ratio of paid submissions to valid submissions. A higher ratio indicates a greater likelihood of a researcher receiving payment for a valid submission, reflecting lower competition. Conversely, a lower ratio signifies intense competition due to the tournament structure of the bounty programs (Zrahia et al. 2024). In

[5] See: `https://www.bugcrowd.com/hackers/bugcrowd-university/`).

[6] In our setting, Bugcrowd uses a data-driven hacker selection and activation machine-learning (ML) technology to perform this matching. See: `https://www.bugcrowd.com/products/platform/crowdmatch/`.

our setting, only 18.5% of the submissions to MBB programs (those that offer payment by structure) were awarded with dollar payments. Similarly, of the 31,754 unique users who submitted to the platform since its inception, only 17.5% were awarded a monetary payment. The top 100 users account for nearly 23% of the submissions on the platform and earn 44% of the monetary rewards, while the top 10 users, a fraction of the research community, receive almost 14% of the platform's rewards (Table B13). Based on their significant contribution, organizations likely aim to attract the top user cohorts to their programs while competing with other programs for their work hours. Furthermore, although the top 10, 20, or 100 cohorts primarily focus on private programs, they still submit about 20% of their work to public programs and earn about 15% of their monetary rewards from public programs (Table B14). This highlights the potential benefits of opening a program to attract top researchers who may not have been invited when it was previously closed.

2.2. Theoretical background

The choice of the level of openness has been widely studied in the platform governance and strategy literature. Prior research has largely focused on markets characterized by network effects, where increased participation enhances platform value through positive feedback loops (Rochet and Tirole 2003, Parker and Alstyne 2005). A fundamental consideration in the literature on platform governance is the quantity-quality trade-off (Eisenmann et al. 2006, Zhu and Iansiti 2012)). Open platforms attract a larger user base, facilitating broader participation and innovation (Boudreau 2010). However, openness also reduces control over quality, leading to an influx of lower-value contributions (Franke and Von Hippel 2003). Closed platforms, by contrast, enforce stricter access controls, preserving quality at the expense of scale and diversity (Alexy et al. 2013).

This trade-off is particularly relevant in bug bounty programs, where organizations must balance attracting a large pool of security researchers (to increase vulnerability discovery) against maintaining submission quality (to reduce noise and operational costs). Public programs invite unrestricted participation, maximizing submission volume, while private programs curate participants, ensuring higher-quality discoveries at the cost of lower participation.

Zrahia et al. (2024) focus on the economics of bug bounty platforms as two-sided markets with cross-side network effects between organizations (demand-side) and security researchers (supply-side). Unlike traditional two-sided markets where complementors benefit from greater participation, bug bounty researchers do not directly gain from more participants. Instead, increased competition can reduce individual payoffs, discouraging engagement (Sridhar and Ng 2021, Arora et al. 2008, p. 1). The reason for this unique setting attribute lies in the tournament structure of bug bounty programs: only the first submission of a distinct vulnerability not already known to the organization qualifies the researcher for monetary reward or other forms of recognition. This distinctive dynamic differentiates bug bounty platforms from other digital ecosystems explored in platform literature.

While openness is generally advantageous in markets with strong network effects, its value in markets without such effects is more ambiguous (Kretschmer et al. 2022, Boudreau 2011). In the context of bug bounty programs, openness does not inherently increase value for existing participants, and excessive competition may even deter high-quality contributors. This dynamic challenges the assumption that openness is always beneficial.

Prior studies on bug bounty effectiveness have explored how program design impacts submission rates, researcher incentives, and vulnerability discovery (Miller 2023, Kuehn and Mueller 2014). Research suggests that private programs yield higher-quality submissions, attracting top-tier researchers who are less likely to submit duplicates or low-value findings (Callaway and Sant'Anna 2021). Public programs, in contrast, experience higher participation but lower efficiency, with firms bearing the cost of filtering redundant and invalid submissions (Maillart et al. 2017).

3. Data and methodology

Our empirical analysis is based on a comprehensive proprietary panel dataset of over 480,000 submissions acquired through a Data Transfer Agreement between Tel Aviv University and Bugcrowd. The dataset includes all vulnerability submission activity of all programs from 2012 through mid-2021. It contains anonymized (however detailed) program-, organization-, researcher-, and submission-level information, offering a unique depth and breadth of insights. Every submission to a program is dated and includes its outcome, priority, and the unique identification of the submitting researcher.[3]

We define an organization's openness decision as its choice between public (open) and private (closed) structures.[7] While most programs maintain the same structure throughout their existence, some transition between these types. We specifically consider 86 programs in our data that switched from private to public as transitioning from closed to open, allowing us to study the implications of increasing openness over time.

We structure our analysis into two stages. First, we examine in Section 4 the general performance characteristics of openness by comparing private programs to public ones. Next, we focus in Section 5 on the 86 programs that transitioned from private to public and evaluate four key performance metrics to measure the effect of openness.

3.1. Measures of program performance

To account for the quantity-quality trade-off that is typical to platform markets, we use four key metrics as proxies for program performance: (i) *submission count*, representing the total amount of research activity on a program, regardless of the submissions' outcomes; (ii) *number of unique vulnerabilities* discovered, reflecting the program's value to the organization; (iii) *number of new researchers*, indicating a program's attractiveness. Attracting new researchers is essential for generating new ideas based on the notion that

[7] While we categorize programs as either open or closed, in reality, openness is a continuous attribute, as private programs can manage and adjust the number of invited researchers over time.

"given enough eyeballs, all bugs are shallow."[8]; and (iv) *number of duplicate submissions*, which proxies for both program competitiveness and the operational overhead costs of managing the program.

3.2. Dataset preparation

We obtained all submission records from Bugcrowd's platform that are available in the existing dataset. Given our research focus on closed-to-open transformation, we removed records of irrelevant programs that changed their openness structure more than once or transitioned from public to private. As a result, our final dataset includes three mutually exclusive groups: (i) private programs that remained private, (ii) public programs that remained public, and (iii) private programs that switched to public and remained public. Programs in the latter cohort, referred to as "treated programs," switched their openness structure at different times (a scenario known as "staggered treatment"). Once a program has been treated, it remains treated permanently.

Next, we consolidated each of the four performance metrics discussed in Section 3.1 into a single weekly entry formatted as ISO 8601 week date (for example, 2019-W12). We aggregated the metrics weekly rather than monthly or quarterly to ensure we had enough observations to establish reliable estimations of pre-treatment and post-treatment outcome trends. Additionally, if a program switched from private to public in the middle of a week, the entire week is considered untreated to guarantee unbiased results. We list additional details on these performance measures and other key variables used throughout this paper in Table 1.

Finally, we completed missing weekly observations during each program's existence; if a program had no submissions during a specific week, we added that week to the timeline and assigned a zero value to all its program-level performance metrics. Although our dataset remains unbalanced (with a different number of observations per unit), it is complete, with no weekly gaps from the first to the last week of activity for each program unit.

3.3. Empirical strategy

By employing three distinct comparison strategies, we provide a robust understanding of the causal effect of the openness change. The first approach, Differences-in-differences (DiD), compares the change in outcomes over time for the treated group to the change in outcomes for the control group. DiD isolates the treatment effect by differencing out unobserved characteristics and time trends. Specifically, it leverages the pre-treatment period to establish a baseline difference between the groups. Then, it examines how this difference changes after the treatment (in the treated group). This approach is particularly valuable in addressing potential selection bias, as it compares within-unit changes over time.

[8] The idea, known as "Linus's law," suggests that collaborative development and testing of open-source software by many individuals can lead to less buggy (vulnerable) code (Raymond 1999).

Table 1 Variable definitions

Variable Name	Variable Definition	Expected Values
Switched	Indicator of whether a program switched from private to public	1=Switched 0=Otherwise
Treated period	Indicator of before/after switching	1=post-switch 0=pre- and never switched programs
Program openness	Program level indication of openness	0=Public 1=Private
Tenure	Week count since program started	0-max
Treated week	Age (in ISO 8601 week date) at which a program switched	0-max
Submissions count	Weekly aggregated count of submissions	0-max
Unique vulnerabilities	Weekly aggregated count of unique vulnerabilities	0-max
New researchers	Weekly aggregated count of distinct new researchers	0-max
Duplicates	Weekly aggregated count of duplicate submissions	0-max
Submission year	Calendar submission year to control for seasonality	2012-2021

The before-and-after comparison (within-treated) method estimates the average treatment effect on the treated (ATT) by comparing the average outcome of treated units during the post-treatment period to their average outcome in the pre-treatment period. To begin, we visualize the average outcome over time in relation to the treatment to illustrate its dynamic impact. Next, we use a regression model to estimate the effect while controlling for the program's duration and fixed effects.

Finally, the treated-untreated comparison (post-treatment) strategy compares the outcomes of switched programs in the post-treatment period with those of public programs that have remained open to all researchers. This approach uses a panel regression with time-varying treatment effects to identify the cross-sectional difference between treated and untreated units already in the treated state.[9]

4. Descriptive statistics

This section provides descriptive statistics highlighting key differences between private and public programs. Our data consists of 2,483 programs, with 2,377 consistently maintaining their openness structure (180 public and 2,197 private programs). The remaining 106 changed their level of openness at least once, either from private to public or vice versa. Among these, 86 switched only once from private to public and will be the focus of our analysis (see Table 2).

[9] We found no evidence indicating that units that transitioned to a public state differ systematically from those that have always maintained a public structure

Table 2 Count of programs by openness structure

	Public to private	Private to public	Public structure	Private structure
Switched once	4	86		
Switched twice	16	16		
Not switched			180	2197

[1] Of 2,483 programs, 16 have changed their openness structure twice and are therefore listed twice.

Both public and private programs show a nearly identical distribution of program types, with 52%-53% classified as MBBs and 47%-48% as VDPs. Yet, the majority of submissions are made to paid MBB programs (see Table B11).

The differences in activity distributions between private and public programs are clearly shown in Table 3. While 92% of all programs are private, the total number of submissions is nearly evenly split between private and public programs. Still, around 59% of unique submissions—those that provide value to organizations—are submitted to private programs. Moreover, nearly 80% of the total dollar rewards paid to researchers on the platform are associated with private programs. These statistics suggest a potential trade-off between the quantity and quality of submissions; the higher percentage of payments in private programs reflects the significance (quality) of the discovered vulnerabilities despite a nearly equivalent count of submissions in public programs. We further explore this aspect in Section 4.2.

Table 3 Distribution of program characteristics by openness (percentage)

Program openness	No. of programs[1]	Paid submissions[2]	Dollar rewards[2]
Public	7.7%	23.0%	20.7%
Private	92.4%	77.0%	79.3%
Total	100.0%	100.0%	100.0%
Program openness	**Total submissions[1]**	**Unique submissions[1]**	**High-priority unique submissions[1]**
Public	50.2%	41.3%	34.1%
Private	49.8%	58.8%	65.9%
Total	100.0%	100.0%	100.0%

[1] Based on 2,377 programs that have maintained their openness attribute throughout their existence.
[2] Based on 1,508 MBB programs that paid monetary rewards and have maintained their openness attribute throughout their existence.

4.1. Submission probabilities

While bug bounty programs can provide tremendous value to organizations, these come at a cost. Specifically, each submission bears processing costs associated with the effort needed to verify its validity, and determine if the vulnerability has already been identified by another researcher, and thus represents a duplicate submission. We, therefore, distinguish between three types of submissions: (i) uniquely discovered (valuable) vulnerabilities, (ii) duplicates of an already identified vulnerability, or (iii) invalid and therefore rejected submissions.

A logistic regression for each type shows that the odds of submission being unique are 1.8 times higher in private vs. public programs (Table B12). The odds of submission being duplicate or rejected are 25% and 33% lower in private programs, respectively. Figure 1 illustrates the predicted probabilities, clearly showing that private programs dominate all three variables; i.e., private programs are more likely to receive submissions of unique vulnerabilities and less likely to receive duplicate and rejected submissions, which are costly for the firm to process yet provide no real economic value.

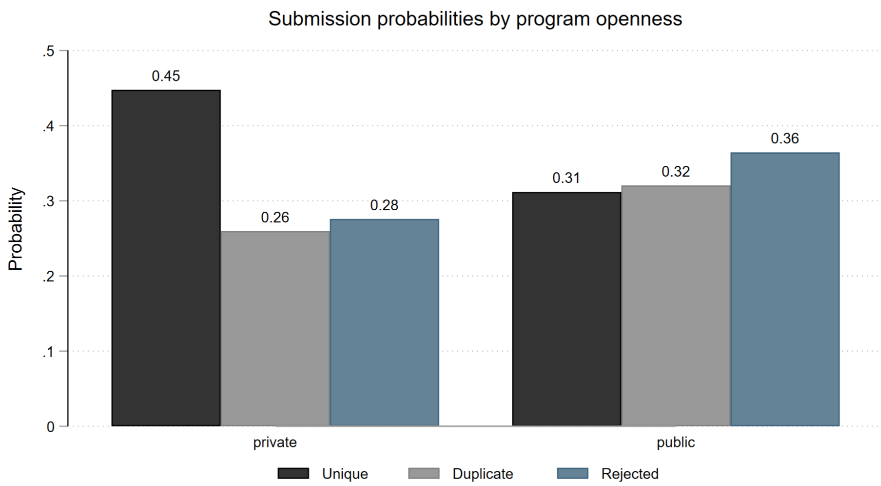

Figure 1 **Predicted submission probabilities (rounded) of private and public programs.**

An interesting aspect of organizations' decision on program openness is how the distribution of submissions and their associated probabilities change over time. Specifically, as a program matures, the likelihood of discovering unique vulnerabilities is expected to decline since researchers have likely already found the "low-hanging fruit." Similarly, interest of new researchers is also expected to wane over time as the chances of meaningful discoveries—and corresponding rewards—diminish.

Figure 2 presents the expected probabilities of the three submission types for public and private programs over time. Interestingly, while the likelihood of a unique submission for public programs decreases over time, it remains nearly steady in private programs. This may be due to public programs attracting more new researchers each week, resulting in more submissions, while private programs show a decline (Figure B12). This subsequently enhances competitiveness over each potential unique finding, leading to a lower likelihood. For both types of programs, the likelihood of receiving a duplicate submission decreases slightly over time, reflecting the declining number of submissions and a reduced probability of identifying a valid (correct) vulnerability. Conversely, the probability of receiving a rejected submission increases over time for both openness structures; however, this rate is significantly higher in public programs than in private ones. Since a rejected submission indicates an incorrect finding, we suspect that the higher quality of researchers participating in private programs mitigates this effect compared to public programs.

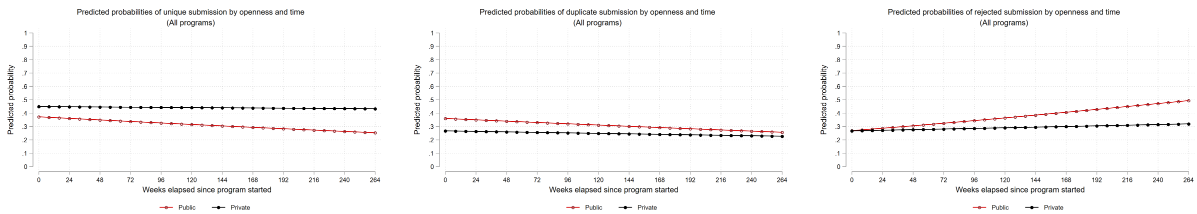

(a) Unique submissions (b) Duplicate submissions (c) Rejected submissions

Figure 2 Predicted probabilities of unique, duplicate, and rejected submissions by openness and time.

4.2. The quantity-quality trade-off

As typical in the literature, the open vs closed debate represents a quality vs quantity trade-off. Open platforms attract a larger user base by reducing entry barriers and fostering broad participation (Rochet and Tirole 2003). However, this openness often comes at the cost of quality control, as the influx of participants may lead to lower-value contributions or increased noise (Eisenmann et al. 2006). In contrast, closed platforms maintain stricter access controls, which can help preserve quality by ensuring that only vetted participants contribute, but at the expense of reduced scale and potential innovation (Zhu and Iansiti 2012).

This trade-off is particularly relevant in bug bounty programs, where public (open) programs attract more researchers but also more duplicates and low-quality submissions, whereas private (closed) programs receive fewer but higher-quality submissions. In our case, quality and quantity can be reflected in both the number and quality of submissions, as well as the number of researchers and their quality, as a researcher's professional competence impacts their ability to discover unique vulnerabilities that provide economic value to organizations.

The statistics above demonstrate that while public programs attract more submissions, their quality is lower. In terms of the number of researchers, we look at the number of new researchers a program attracts over time. As expected, private programs see a decline in new weekly researchers. Yet, the trend flips for public programs, likely reflecting platform expansion and the ease of access public programs offer.

To assess researchers' quality, we look at three different variables that combined serve as a strong proxy for skill level: (i) the number of unique vulnerabilities they discovered, (ii) the sum of priorities of the researchers' submissions,[10] and (iii) the total dollar payments they received. These measures combined allow us to consider the expertise needed to be the first to discover vulnerabilities, account for their severity (importance), and the researcher's overall past success on the platform.

As Table 4 shows, there is a striking difference in the quality of researchers. Specifically, the average number of unique submissions per user upon their first approach to the program is nine times higher in

[10] Critical (P1) submissions are assigned five quality points; high-priority submissions (P2) receive four points; moderate importance (P3) submissions get three points; less valuable P4 submissions earn two points; and the lowest priority submissions (P5) are assigned a single point. We acknowledge that this scoring system is ordinal, meaning that the intervals between priority levels do not necessarily reflect the actual differences in value.

private programs compared to public programs. Additionally, the average constructed quality score for users is ten times higher in private programs, and the average dollar earnings are over eighteen times greater than those in public programs. This difference confirms that private programs selectively invite researchers based on their skills and effectively maintain high-quality participation.

Table 4 Average user quality metrics by openness

Program openness	Unique vulnerabilities	User quality measure	Dollar earnings
Public	10.62	40.58	$1,480
Private	97.07	409.04	$27,908

[1] All measured quality variables reflect the average value of users in their first submission to any program.
[2] Based on 2,377 programs that have maintained their openness attribute throughout their existence.
[3] User quality measure is an ordinal scoring system of priorities as described in section 4.2.

Still, while skilled researchers may identify unique vulnerabilities earlier in the program's life cycle, thereby increasing economic value, mature programs might benefit from diverse new researchers. The decline in submissions for private programs over time supports this hypothesis, suggesting there may be advantages to transitioning from private to public programs after a certain period.

5. Empirical analysis: the effect of openness

The 86 programs that chose to open themselves and transitioned once from private to public offer a unique opportunity to further study the impact of openness by comparing the performance of programs that switched to those that remained private, as well as by analyzing performance before and after the change. Programs that changed their openness structure did so at different times. Figure 3 illustrates the timeline of transitioned programs; 27 switched within the first five months of their activity, whereas only eight programs transitioned more than two years into their existence.

We focus our analysis on three cohorts of programs: 180 public programs that remained public throughout their existence, 2,197 private programs that stayed private, and 86 private programs that transitioned to private at some point. Using four performance metrics, we compare the performance of these groups, making a distinction between pre- and post-transition periods for the switched programs.

Figure 4 presents the predicted values from a linear regression that captures the time trend within each cohort.[11] Private programs that transitioned to public exhibit pre-switch trends similar to those of private programs that did not make that change. The transition to public status results in an initial boost, followed by a steeper decline across all performance measures. In contrast to private programs, public programs experience a gradual increase over time, both in weekly submissions and in the number of new researchers (also visualized in Figure B12).

[11] Since these values are predictions, the figure does not perfectly reflect the data at every point. For example, no program will produce negative submissions post-transition despite the observed tenure trend in later weeks.

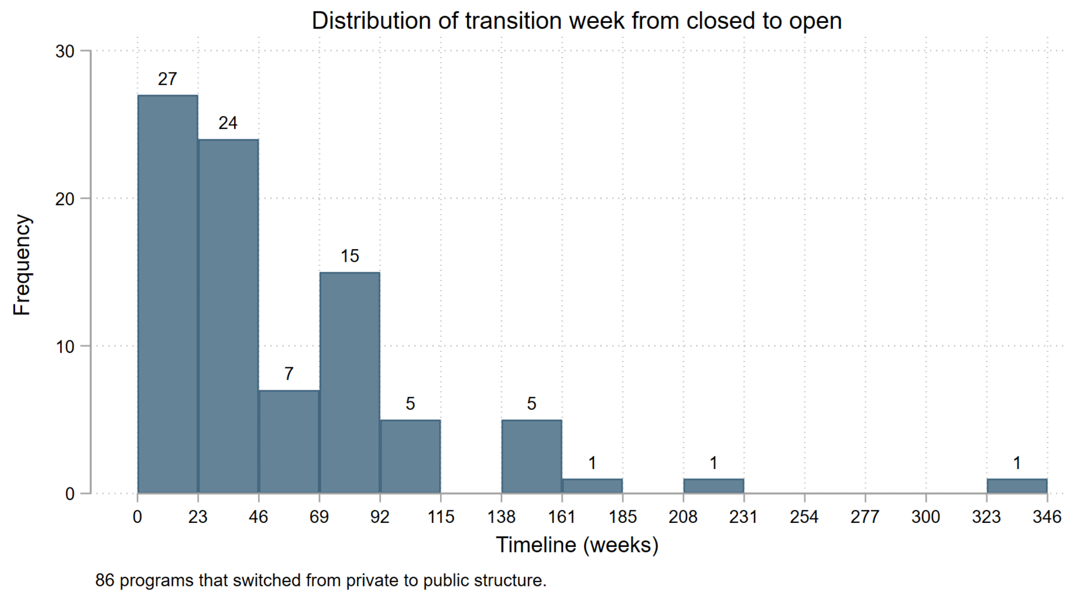

Figure 3 **Distribution of transition week from closed to open.**

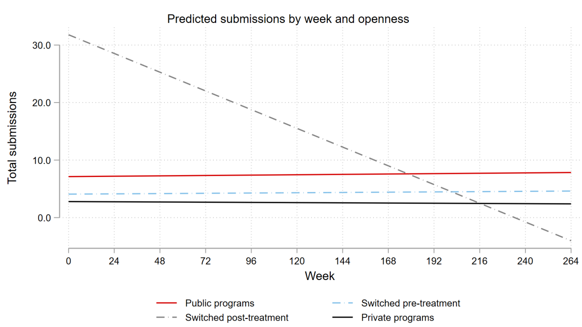

(a) Weekly submissions slope

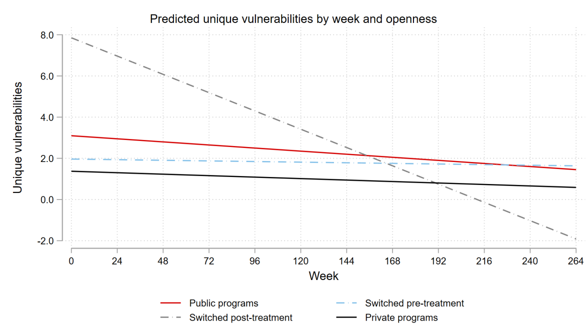

(b) Unique vulnerabilities slope

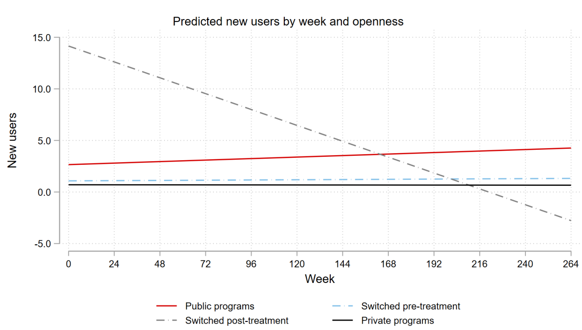

(c) New users slope

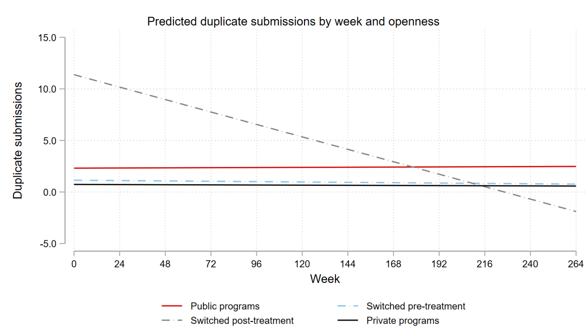

(d) Duplicate submissions slope

Figure 4 **Comparison of performance measures over time**

In the following three-stage analysis, we first examine private programs using a Differences-in-Differences (DiD) approach. This allows us to compare the performance of private programs that transitioned to public status with those that remained private. Next, we conduct an event study to analyze the performance of the transitioning programs both before and after the switch. Finally, we employ a panel data regression with time-varying treatment effects to compare public programs with those that transitioned to public status after the change.

5.1. Private programs: transitioning vs. staying private

Focusing only on private programs, we perform a Differences-in-Differences (DiD) analysis where the dummy variable *Switched* takes on the value of “1” once a program transitions from private to public structure and “0” otherwise. We control for *Submission year* to remove potential time-related confounding effects.

This analysis method addresses the challenges of treatment heterogeneity since the treatment’s causal effect might vary across different groups or over time (Gardner 2022). The interpretation of the results relies on the “parallel trends” assumption, stating that without treatment, the average outcomes for the treatment and control groups would follow parallel paths over time (Callaway and Sant’Anna 2021).[12]

Table 5 Differences in differences weekly measurements.

	(1)		(2)		(3)		(4)	
Switched	22.841***	(0.612)	5.915***	(0.246)	10.173***	(0.151)	7.435***	(0.320)
Tenure	-0.007*	(0.004)	-0.003*	(0.002)	-0.002**	(0.001)	-0.005**	(0.002)
Switched × Tenure	-0.031***	(0.005)	-0.011***	(0.002)	-0.017***	(0.001)	-0.005**	(0.002)
Submission year	0.124	(0.183)	-0.059	(0.073)	0.048	(0.045)	0.165*	(0.096)
Constant	-246.585	(368.813)	121.046	(148.261)	-96.224	(90.826)	-332.728*	(192.741)
Observations	73766		73766		73766		73766	
Adjusted R-squared	0.304		0.290		0.310		0.167	

[1] Regression (1): For total submissions. Regression (2): For unique vulnerabilities. Regression (3): For new distinct users. Regression (4): For duplicate submissions.
[2] Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

As Table 5 shows, the coefficient on *Switched* is positive and significant in all four regressions. Specifically, the regression results suggest that opening up by transitioning to a public structure leads, on average, to 23 additional weekly submissions, six identified vulnerabilities, ten new researchers joining the program, and seven duplicate submissions each week. That is, transitioning to a more open structure has a significant positive economic value to the firm, allowing it to tap into a larger pool of researchers. Moreover, the rising interest from new researchers and the increase in duplicate submissions indicate the program’s competitiveness has grown. Appendix A includes two additional DiD methods that provide robustness to our findings.

The statistics in Section 4 support the finding in Sridhar and Ng (2021) that “programs receive fewer valid reports as they grow older and bugs become harder to find.” To examine whether program aging might be

[12] To test this assumption, we conducted a regression analysis comparing pre-treated observations with private programs that remained closed, and we could not find any statistically significant differences.

the catalyst for transforming private programs into public programs and the impact of such a transformation, we include the variable *Tenure* in the regression, representing the number of weeks a program has been active.

As expected, a statistically significant negative tenure effect is observed across all performance proxies, indicating a decline in program interest over time. This results in fewer submissions, fewer new researchers, and, most importantly, fewer unique vulnerabilities found. Interestingly, interacting *Switched* with *Tenure* reveals that, with the exception of duplicate submissions, all performance proxies decline more sharply post-transition than pre-transition. That is, relative to programs that remained private, opening up gives a program a positive boost that fades rapidly. We note that the decline in performance over time may result in a plateau for programs that do not switch, leaving little room for further deterioration. In contrast, transitioning programs experience a surge in activity, allowing them to continue degrading over time.

Lastly, the submission year variable has a significant effect only on the number of duplicates—this reflects the already identified growth in the number of researchers over time and the resulting increase in competitiveness on the platform (Zrahia et al. 2024).

5.2. Before and after: performance changes in switched programs

Next, we focus only on the programs that transformed their openness and employ an event study analysis to compare outcomes before and after intervention within the same group.[13] We use the same four weekly performance measurements of total submissions, uniquely discovered vulnerabilities, new submitting users, and submission duplicates.

We start by visualizing the mean weekly values of our measurements over time in Figure 5. We employ a ten-week event window before and after the transition, which happens at t=0. As the figure shows, the transition increased the values of all four performance variables, although the positive effects diminish after 3-4 weeks, at which point the trend goes back to the pre-transition trend.

To explore these findings further, we ran a regression on all 86 switched programs across all weeks in our dataset. We employed fixed-effect and cluster standard errors at the program level to account for potential correlations in outcomes and errors within units.

As shown in Table 6, and consistent with the DiD analysis, switching to a public program provides a considerable performance boost. The tenure effect is also negative and statistically significant, suggesting a decline in all outcome metrics over time. Interestingly, the coefficient of the interaction between *Switched* and *Tenure* is statistically insignificant—possibly because weekly performance trends after the transition are similar to those observed before the switch. Indeed, as shown in Figure 5, once the initial post-switch boost fades, the change in all performance measures over time closely resembles that of the pre-transition period. Alternatively, given the small number of transitioning programs, we may not have enough statistical power to identify the effect.[14]

[13] An event study estimates dynamic treatment effects and provides a built-in graphical summary of results over time (Miller 2023).

[14] Including a non-linear relationship for *Tenure* in the regression also shows no effect.

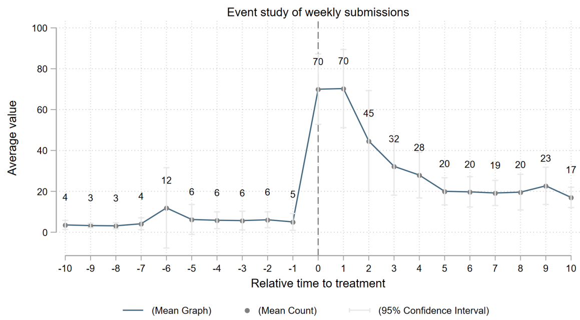

(a) Weekly submissions event study

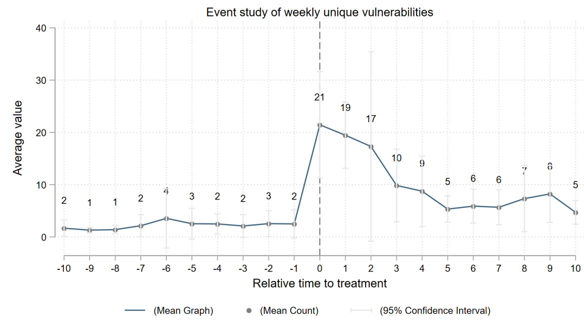

(b) Unique vulnerabilities event study

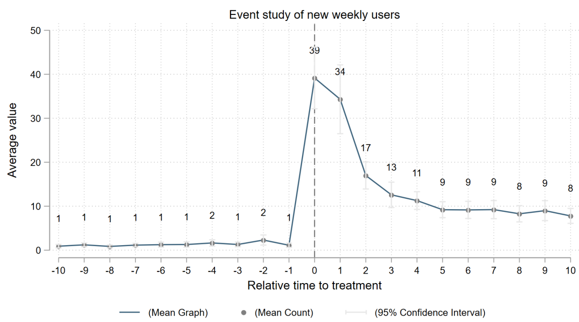

(c) New users event study

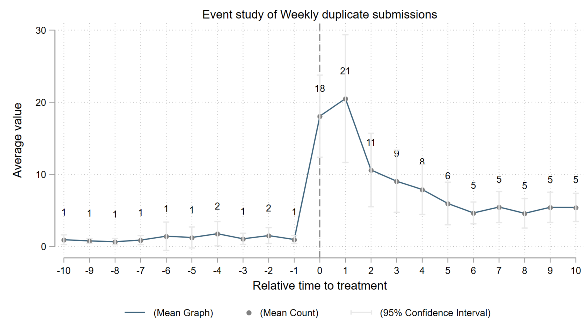

(d) Duplicate submissions event study

Figure 5 Event study weekly results

Table 6 Average weekly treatment effect on the treated (tenure view).

	(1)		(2)		(3)		(4)	
Switched	22.465***	(4.555)	5.797***	(1.990)	9.992***	(1.193)	7.285***	(1.801)
Tenure	-0.035**	(0.016)	-0.014*	(0.007)	-0.016**	(0.006)	-0.014*	(0.007)
Switched × Tenure	-0.007	(0.029)	-0.004	(0.009)	-0.005	(0.010)	0.005	(0.016)
Constant	2.934	(2.069)	1.298	(0.900)	0.955*	(0.481)	0.900	(0.813)
Observations	9971		9971		9971		9971	
Adjusted R-squared	0.057		0.024		0.141		0.021	

[1] Regression (1): For total submissions. Regression (2): For unique vulnerabilities. Regression (3): For new distinct users. Regression (4): For duplicate submissions.
[2] Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

5.3. Timing of transition

Given the variance in the timing at which organizations choose to transition, one may wonder whether there is an optimal program age that maximizes the value of switching. Thus, we define two new variables: the variable *Treated week*, which represents the age of the program at the time of its transition (measured as the number of weeks since inception), and *Treated period*, which serves as a dummy variable indicating whether the observation occurs pre- or post-treatment.

Table 7 shows the regression results where we continue to use fixed effects and cluster errors at the program level. The program's age at the time of transition is collinear with the fixed effects of the program; however, the positive and significant coefficient on its interaction with *Treated period* suggests that programs that transition later derive greater value from switching—exceeding the overall declining trend delineated by *Tenure*. Interestingly, the coefficient on the interaction is not significant for the number of duplicated submissions, suggesting that delaying the transition does not increase the costs associated with handling additional duplicate submissions. While this finding may lead organizations to consider delaying the transition of their program later in its lifecycle, such a delay comes with the cost of postponing the additional activity benefits that come with opening up sooner.

Table 7 Average weekly treatment effect on the treated (treated week timeline)

	(1)		(2)		(3)		(4)	
Treated period	14.068***	(3.644)	2.965***	(1.072)	5.997***	(1.138)	4.663**	(1.916)
Treated week	0.000	(.)	0.000	(.)	0.000	(.)	0.000	(.)
Treated period × Treated week	0.134***	(0.049)	0.044***	(0.011)	0.061***	(0.021)	0.050	(0.030)
Tenure	-0.055***	(0.020)	-0.021***	(0.007)	-0.027***	(0.008)	-0.017**	(0.008)
Constant	5.269***	(1.898)	2.097***	(0.623)	2.085***	(0.521)	1.574*	(0.874)
Observations	9971		9971		9971		9971	
Adjusted R-squared	0.064		0.030		0.159		0.024	

[1] Regression (1): For total submissions. Regression (2): For unique vulnerabilities. Regression (3): For new distinct users. Regression (4): For duplicate submissions.
[2] Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

5.4. How opened private programs compare to public ones

We conclude the analysis by comparing the post-switch performance of transitioned programs with that of public programs of the same age that began and remained public. To accomplish this, we focus solely on post-treatment observations of the switched programs. We conduct a regression that includes the *Switched* dummy variable, which takes on the value of "1" for transitioning programs and "0" otherwise, along with the *Tenure* continuous variable that represents the age of the program in weeks. The interaction of these two variables captures differences in performance between the two cohorts over time.

Table 8 demonstrates perfect collinearity between the *Switched* dummy variable and the fixed effects of the programs, as expected. However, the key finding is that, across all four performance measures, programs that transitioned from private to public age more rapidly post-treatment than those that began as public and maintained that structure. The predicted model of weekly submissions and new users in Figure B12 suggests that private programs mature faster than public ones. This pattern is further reinforced by the fact that private programs generally operate for a shorter duration, with submissions and unique discovered vulnerabilities to these programs occurring earlier in their lifecycle (Figure B10 and Figure B11 respectively).

Consequently, the notable difference in performance decline over time between public and transitioned programs (post-treatment) can be attributed to two main factors. First, these private programs experienced a steeper decline in performance before transitioning. Second, there was a performance boost after they opened up, which enabled the more significant decline observed after the transition.

Table 8 Weekly measurements by tenure: Treated vs. public programs.

	(1)		(2)		(3)		(4)	
Switched	0.000	(.)	0.000	(.)	0.000	(.)	0.000	(.)
Tenure	0.003	(0.004)	-0.006***	(0.002)	0.006***	(0.002)	0.001	(0.002)
Switched × Tenure	-0.138***	(0.034)	-0.031***	(0.008)	-0.070***	(0.014)	-0.051***	(0.015)
Constant	11.666***	(0.736)	3.972***	(0.209)	4.772***	(0.326)	3.981***	(0.315)
Observations	28729		28729		28729		28729	
Adjusted R-squared	0.008		0.010		0.020		0.003	

[1] Comparing post-treatment observations of 86 treated units, to 180 public programs with the same tenure week.
[2] Regression (1): For total submissions. Regression (2): For unique vulnerabilities. Regression (3): For new distinct users. Regression (4): For duplicate submissions.
[3] Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

6. Discussion

Our findings provide valuable insights into firms' strategic decisions with respect to openness and its corresponding impact on user engagement, quality, and overall program performance. The results underscore the trade-offs between private and public programs and highlight key differences in the operation and performance of both structures, as well as the dynamic nature of program transitions.

The quantity-quality trade-off. A central theme in the literature on openness is the well-documented trade-off between quantity and quality. Public programs, by design, attract a broader pool of researchers, resulting in a higher volume of submissions that continues even after the program has matured. However, the quality of these submissions is generally lower, as reflected in the higher proportion of duplicate and rejected reports. In contrast, private programs maintain a more selective researcher base, leading to fewer but higher-quality submissions. This aligns with prior platform governance literature suggesting that openness fosters participation at the cost of quality control (Eisenmann et al. 2006, Rochet and Tirole 2003).

Furthermore, private programs attract a considerably higher proportion of successful researchers (compared to public programs). These programs tend to operate for shorter periods than public programs, with vulnerability submissions occurring earlier in their lifecycle. This activity pattern aligns with the assumption that experienced researchers prioritize engagement when the likelihood of making meaningful discoveries and receiving rewards is higher. Our finding suggests a trade-off: while closed programs benefit from a concentration of experienced participants, open programs offer broader accessibility, potentially fostering researcher diversity and program longevity. This trade-off raises an important question for organizations about prioritizing researcher quality or maximizing participation diversity.

Performance over time. Our results indicate that private programs, though initially more effective in attracting skilled researchers and valuable submissions, experience a decline in activity over time. This decline is likely due to the exhaustion of "low-hanging fruit" easy-to-discover vulnerabilities, leading to diminishing incentives for continued participation. On the other hand, public programs show a more sustained level of engagement, with a steady influx of new researchers and submissions. However, the observed increase in the probability of rejected submissions over time (Figure 2c) suggests that public programs may become increasingly inefficient as they mature, requiring organizations to devote more resources to filtering and processing low-value reports.

We observe a statistically significant negative tenure effect within private programs, indicating that program performance deteriorates over time as researcher engagement declines. These dynamics may motivate organizations to switch from a private to a public structure once the program matures, hoping a public structure can inject renewed activity. Transitioned programs experience an initial surge across all performance measures. The increased visibility and accessibility of public programs attract a larger researcher base, boosting other performance measures, including the number of unique vulnerabilities discovered. However, following the initial boost, performance metrics declined more rapidly over time compared to public programs that were always open. Still, our robust DiD analysis suggests a meaningful growth in all performance measures post-treatment, compared to non-switched private programs, suggesting that opening up revitalizes engagement and provides additional value to the organization.

Our results further suggest that the timing of transitions influences program effectiveness. Later-stage treatments yield a higher effect on three of the four key metrics (except for duplicate submissions). However, delaying the transition to an open structure may cause stagnation in activity, emphasizing the cost of time. Notably, private programs that did not transition to public structure appear to have reached a plateau, leaving little room for further decline.

Figure 6 illustrates a qualitative timeline view of organizational openness strategic decisions.[15] This dynamic underscores the need for strategic decision-making in designing the lifecycle of bug bounty programs. Organizations should be aware that closed programs perform better than their open counterparts; however, this superior performance diminishes more swiftly over time. Consequently, we argue that transitioning from a closed to an open structure once the activity level degrades may enhance performance.

Openness in markets without cross-side network effects. Unlike traditional two-sided platform markets, where cross-side network effects increase value through greater participation, bug bounty programs do not exhibit these effects. In fact, they may even experience negative same-side network effects. The overall success of a bug bounty program largely depends on the expertise of researchers; however, excessive

[15] The diagram intentionally omits a possible organizational decision to switch from an open to a closed structure. Three of the four programs that made this transition changed in the first week of their operation, possibly indicating a configuration error. Our research focuses solely on the motivations for transforming from a seemingly higher-performing closed structure to an open structure.

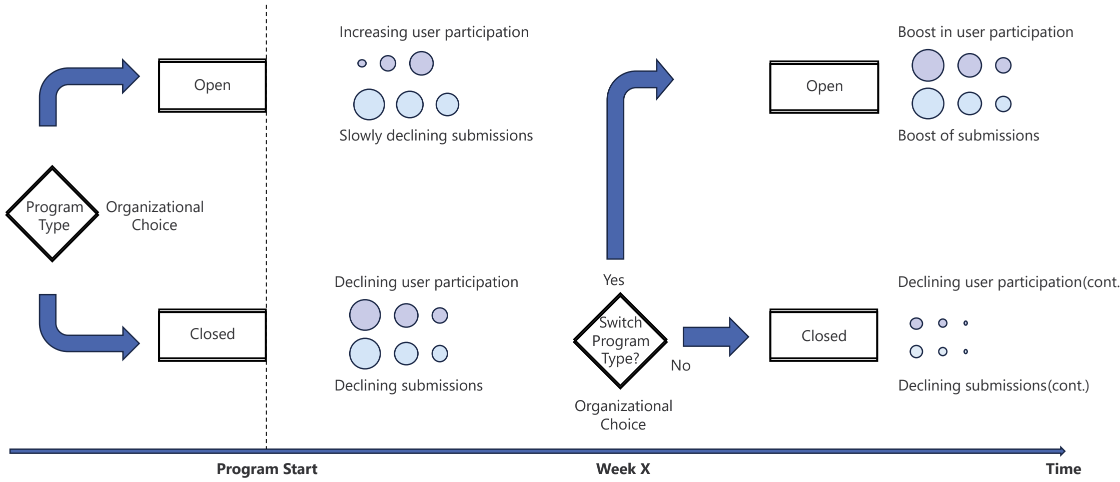

Figure 6 **A qualitative timeline view of organizational openness strategic decision.**

competition within an open program can discourage participation by diminishing the chances of individual success. Therefore, in markets without cross-side network effects, the quantity-quality trade-off plays out differently: although openness enhances activity, as demonstrated by our performance measures, it comes with a cost. Organizations must, therefore, balance the cost of processing the increase in lower-quality submissions against the potential benefits of attracting a larger, more diverse research base that will generate new findings.

Our findings suggest that a hybrid approach—starting with a closed program before transitioning to an open structure—may yield the highest returns. This strategy minimizes the costs associated with low-quality submissions during the program's early high-activity phase while ensuring the necessary boost in activity by opening up once the initial engagement declines. This transition leverages the benefits of both models, combining the early-stage efficiency of private programs with the expanded reach and diversity of public ones.

Managerial implications. Our findings offer several key takeaways for decision-makers. First, organizations should recognize that although private programs yield higher-quality submissions, they may eventually face researcher fatigue and declining engagement. The decision to transition to a public program should be strategically timed to rejuvenate participation without overwhelming resources with low-value submissions. Moreover, the top researchers submit around 20% of their work to public programs and receive approximately 15% of their monetary rewards from these programs (Table B14), highlighting the potential advantages of opening a program to attract top researchers who might not have been invited when it was previously closed.

Second, platform operators must consider the operational costs associated with different program structures. Public programs require robust triage and filtering mechanisms to efficiently manage the influx of lower-quality submissions. Platform investments in automation, dynamic and tiered reward structures, and researcher reputation systems may mitigate the downside of open participation while still leveraging its benefits.

Third, firms should take a dynamic approach to openness rather than treating it as a one-time decision. Programs can begin as private to build a strong foundation of high-quality research and later transition to public access when engagement wanes. Additionally, hybrid models—such as phased openings where only top-performing researchers gradually gain initial access to public programs—could provide a practical middle ground between exclusivity and scale.

Finally, in markets without strong cross-side network effects, choosing between open and closed participation should be guided primarily by cost-benefit analysis rather than expectations based on positive network effects. Managers should carefully assess whether the increase in activity justifies the additional processing costs and whether alternative incentives, such as targeted recruitment of skilled researchers or promotional offers, may be more effective in sustaining engagement over time.

Acknowledgments

The authors gratefully acknowledge Bugcrowd for sharing their data anonymously to enable our analysis.

References

Alexy O, George G, Salter AJ (2013) Cui Bono? The Selective Revealing of Knowledge and Its Implications for Innovative Activity. *Academy of Management Review* 38(2):270–291, ISSN 03637425, URL `http://dx.doi.org/10.5465/AMR.2011.0193`.

Arora A, Telang R, Hao X (2008) Optimal Policy for Software Vulnerability Disclosure. *Management Science* 54(4):642–656, ISSN 00251909, URL `http://dx.doi.org/10.1287/MNSC.1070.0771`.

Borusyak K, Jaravel X, Spiess J (2024) Revisiting Event-Study Designs: Robust and Efficient Estimation. *The Review of Economic Studies* 91(6):3253–3285, ISSN 0034-6527, URL `http://dx.doi.org/10.1093/RESTUD/RDAE007`.

Boudreau K (2010) Open Platform Strategies and Innovation: Granting Access vs. Devolving Control. *Management Science* 56(10):1849–1872, ISSN 00251909, URL `http://dx.doi.org/10.1287/MNSC.1100.1215`.

Boudreau KJ (2011) Let a Thousand Flowers Bloom? An Early Look at Large Numbers of Software App Developers and Patterns of Innovation. *Organization Science* 23(5):1409–1427, ISSN 10477039, URL `http://dx.doi.org/10.1287/ORSC.1110.0678`.

Butts K, Gardner J (2022) did2s: Two-Stage Difference-in-Differences. *The R Journal* 14(3):162–173, ISSN 20734859, URL `http://dx.doi.org/10.32614/RJ-2022-048/`.

Callaway B, Sant'Anna PH (2021) Difference-in-Differences with multiple time periods. *Journal of Econometrics* 225(2):200–230, ISSN 0304-4076, URL `http://dx.doi.org/10.1016/J.JECONOM.2020.12.001`.

Eisenmann T, Parker G, Alstyne MWV (2006) Strategies for Two-Sided Markets. *Harvard Business Review* 84(10):92–101, ISSN 00178012.

Franke N, Von Hippel E (2003) Satisfying heterogeneous user needs via innovation toolkits: the case of Apache security software. *Research Policy* 32(7):1199–1215, ISSN 0048-7333, URL http://dx.doi.org/10.1016/S0048-7333(03)00049-0.

Gardner J (2022) Two-stage differences in differences, URL https://arxiv.org/abs/2207.05943v1.

Gawer A (2014) Bridging differing perspectives on technological platforms: Toward an integrative framework. *Research Policy* 43(7):1239–1249, ISSN 0048-7333, URL http://dx.doi.org/10.1016/J.RESPOL.2014.03.006.

Kretschmer T, Leiponen A, Schilling M, Vasudeva G (2022) Platform ecosystems as meta-organizations: Implications for platform strategies. *Strategic Management Journal* 43(3):405–424, ISSN 1097-0266, URL http://dx.doi.org/10.1002/SMJ.3250.

Kuehn A, Mueller M (2014) Analyzing Bug Bounty Programs: An Institutional Perspective on the Economics of Software Vulnerabilities. *2014 TPRC Conference Paper*, URL http://dx.doi.org/10.2139/SSRN.2418812.

Maillart T, Zhao M, Grossklags J, Chuang J (2017) Given enough eyeballs, all bugs are shallow? Revisiting Eric Raymond with bug bounty programs. *Journal of Cybersecurity* 3(2):81–90, ISSN 20572093, URL http://dx.doi.org/10.1093/cybsec/tyx008.

Miller C (2007) The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. *Independent Security Evaluators*, 1–10 (Citeseer), URL http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.139.5718.

Miller DL (2023) An Introductory Guide to Event Study Models. *Journal of Economic Perspectives* 37(2):203–30, ISSN 0895-3309, URL http://dx.doi.org/10.1257/JEP.37.2.203.

Parker GG, Alstyne MWV (2005) Two-Sided Network Effects: A Theory of Information Product Design. *Management Science* 51(10):1494–1504, ISSN 0025-1909, URL http://dx.doi.org/10.1287/MNSC.1050.0400.

Raymond E (1999) The cathedral and the bazaar. *Knowledge, Technology & Policy* 12(3):23–49, ISSN 1874-6314, URL http://dx.doi.org/10.1007/s12130-999-1026-0.

Rochet JC, Tirole J (2003) Platform competition in two-sided markets. *Journal of the European Economic Association* 1(4), ISSN 15424766, URL http://dx.doi.org/10.1162/154247603322493212.

Sridhar K, Ng M (2021) Hacking for good: Leveraging HackerOne data to develop an economic model of Bug Bounties. *Journal of Cybersecurity* 7(1), ISSN 2057-2085, URL http://dx.doi.org/10.1093/CYBSEC/TYAB007.

Zhu F, Iansiti M (2012) Entry into platform-based markets. *Strategic Management Journal* 33(1), ISSN 01432095, URL http://dx.doi.org/10.1002/smj.941.

Zrahia A (2024) Navigating vulnerability markets and bug bounty programs: A public policy perspective. *Internet Policy Review* 13(1), ISSN 21976775, URL `http://dx.doi.org/10.14763/2024.1.1740`.

Zrahia A, Gandal N, Markovich S, Riordan M (2024) The simple economics of an external shock to a bug bounty platform. *Journal of Cybersecurity* 10(1):6, ISSN 2057-2085, URL `http://dx.doi.org/10.1093/CYBSEC/TYAE006`.

Appendices

A. Differences-in-differences robustness

We have chosen two additional DiD methods that would suit our setting and provide robustness to our findings: the two-stage differences-in-differences method developed by (Gardner 2022) and the DiD imputation of (Borusyak et al. 2024). Both methods base an initial estimation of group and time effects from the untreated observations. These estimates predict what would have happened to the treated units had they not received it. Furthermore, they both support our unbalanced panel dataset, which might be challenging to some recently introduced code implementations for DiD models with staggered treatment adoption and heterogeneous causal effects.

A.1. Two-stage DiD

We employ a two-stage least squares (2SLS) Difference-in-Differences (DiD) estimator using the Stata did2s implementation developed by Butts and Gardner (2022). This approach first isolates fixed effects of program and time using control programs and the not-yet-treated observations of the treatment group. In the second stage, we estimate the average treatment effects by comparing treated and untreated outcomes after removing the first-stage effects.

The regression models include the four weekly metrics discussed earlier as the dependent variable. In the first stage, we control for the organization's fixed effects, the year of observation, program tenure (measured in weeks), and reward type (classified as MBB or VDP). In the second stage, we regress the residual outcome on the *Switched* dummy, which equals "1" for any week following the openness change and "0" otherwise, to estimate the treatment effects. Standard errors are clustered at the program level to account for potential correlation within units.

Our findings, presented in Table A9, indicate that private programs that transitioned their openness structure have, on average, an increase of 16 weekly submissions. Additionally, these programs uncover three more unique vulnerabilities each week, attract seven more distinct new users weekly, and experience an increase of six duplicate submissions per week. All results have been rounded and compared to untreated private programs.

Table A9 Differences in differences weekly measurements two-stage DID method.

	(1)		(2)		(3)		(4)	
Switched	15.597***	(2.692)	3.271***	(0.628)	6.673***	(0.655)	5.762***	(1.483)
Observations	74301		74301		74301		74301	

[1] Regression (1): For total submissions. Regression (2): For unique vulnerabilities. Regression (3): For new distinct users. Regression (4): For duplicate submissions.
[2] Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

We further illustrated the results in Figure A7, estimating the examined metrics ten weeks before and after the treatment. The effect is the largest immediately after the switch to public structure and dampens over time.

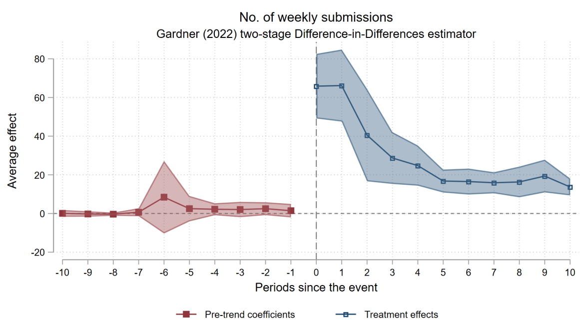

(a) Total submissions effect

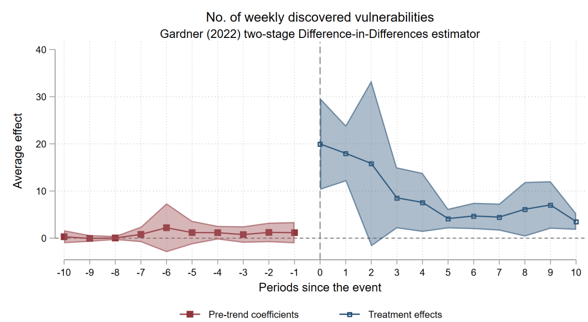

(b) Unique vulnerabilities effect

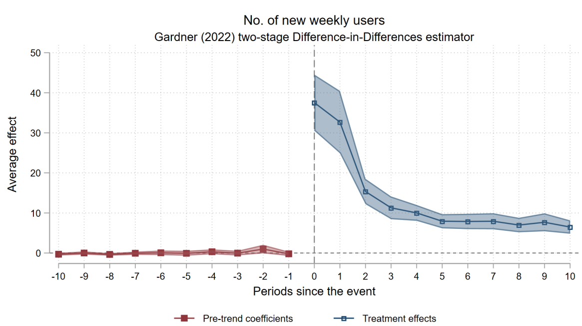

(c) New users effect

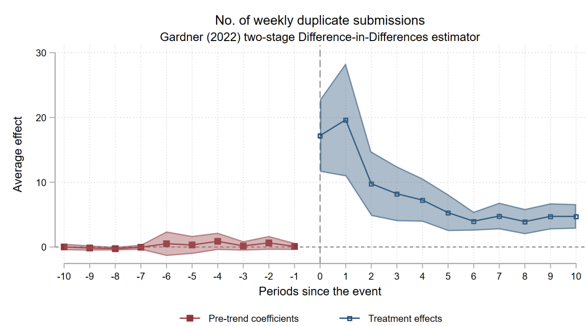

(d) Duplicate submissions effect

Figure A7 Weekly results two-stage differences-in-differences effect

A.2. DiD imputation

The DiD imputation approach utilizes a two-way fixed-effects ordinary least squares regression, implemented in three stages. First, unit and period fixed effects are estimated using only untreated observations. These estimates are then employed to impute the untreated potential outcomes for treated observations. Finally, the imputed treatment effects are aggregated to estimate the average of the heterogeneous treatment effects.

The analysis uses the did_imputation Stata command developed by Borusyak et al. (2024). The weekly metrics discussed earlier are the dependent variables used in four models. Additionally, we incorporate fixed effects for the organization administering the program, the year of observation, and the reward type of program (classified as either MBB or VDP). Finally, to account for factors that may influence the outcome, we added a *Tenure* control variable representing the number of weeks the program has been active for each submission.

Our findings, summarized in Table A10, show (after rounding) the same weekly increase in all outcomes following treatment as presented in Section A.1. Figure A8 visually illustrates these results by estimating the examined metrics ten weeks before and after the treatment.

Table A10 Differences in differences weekly measurements (DID imputation method).

	(1)		(2)		(3)		(4)	
Switched	15.967***	(1.621)	3.443***	(0.290)	6.833***	(0.329)	5.834***	(0.851)
Tenure	-0.003	(0.004)	-0.002	(0.002)	-0.001	(0.001)	0.000	(0.001)
Observations	74298		74298		74298		74298	

[1] Regression (1): For total submissions. Regression (2): For unique vulnerabilities. Regression (3): For new distinct users. Regression (4): For duplicate submissions.
[2] Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

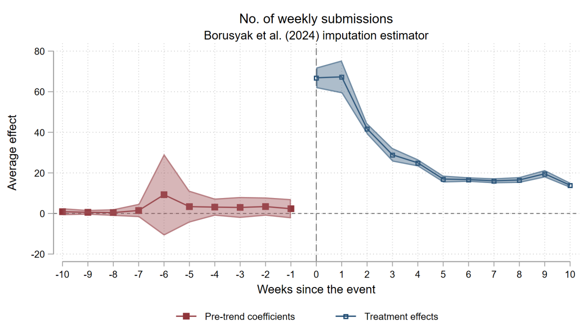

(a) Weekly submissions effect

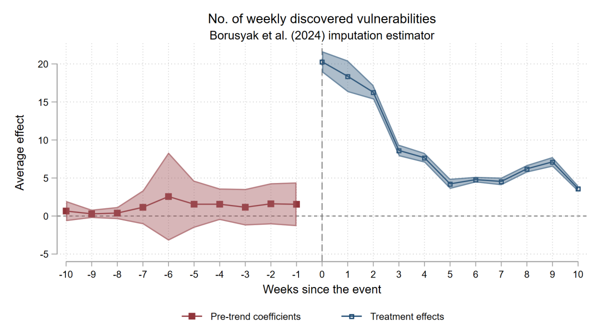

(b) Unique vulnerabilities effect

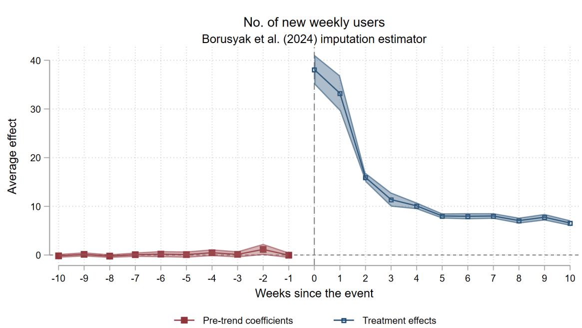

(c) New users effect

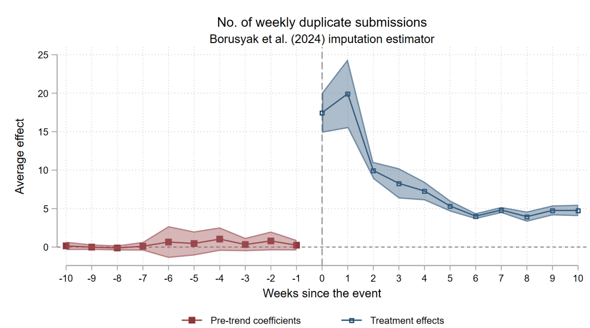

(d) Duplicate submissions effect

Figure A8 Weekly results differences-in-differences imputation effect

B. Supporting information

Table B11 Distribution of program openness by type

Program openness	Count of programs		Count of submissions		Count of unique vulnerabilities	
	VDP	MBB	VDP	MBB	VDP	MBB
Public	51.7%	48.3%	22.1%	77.9%	23.0%	77.0%
Private	52.7%	47.3%	32.2%	67.8%	32.7%	67.3%
Total	52.6%	47.4%	27.1%	72.9%	28.7%	71.3%

[1] Based on 2,377 programs that have maintained their openness attribute throughout their existence.

Table B12 The probability of submission results by program openness.

	(1)		(2)		(3)	
Private	1.791***	(0.013)	0.744***	(0.006)	0.665***	(0.005)
Observations	349000		349000		349000	
Pseudo R-squared	0.015		0.004		0.007	

[1] Regression (1): For unique submissions. Regression (2): For duplicate submissions. Regression (3): For rejected submissions.
[2] Based on 2,377 programs that have maintained their openness attribute throughout their existence.
Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

Table B13 Distribution of top user submissions and monetary rewards (percentage)[1].

User cohort	% of total monetary rewards[2]			% of total submissions[2]		
	Public	Private	Total	Public	Private	Total
Top 10	11.3%	14.6%	13.9%	1.5%	6.4%	3.9%
Top 20	14.6%	22.2%	20.6%	2.0%	9.9%	6.0%
Top 100	31.3%	47.4%	44.1%	5.7%	22.7%	14.2%

[1] Based on submissions of the 2,377 programs that have maintained their openness attribute throughout their existence. Users' rank was determined based on their aggregated USD monetary rewards.
[2] Percentage measures relative to the total activity of all users in each column.

Table B14 Distribution of top user submissions and monetary rewards by openness (percentage)[1].

User cohort	% of total monetary rewards by cohort[2]			% of submissions by cohort[2]		
	Public	Private	Total	Public	Private	Total
Top 10	16.9%	83.1%	100.0%	18.5%	81.5%	100.0%
Top 20	14.7%	85.3%	100.0%	17.1%	82.9%	100.0%
Top 100	14.7%	85.3%	100.0%	20.2%	79.8%	100.0%
All users	20.7%	79.3%	100.0%	50.1%	49.9%	100.0%

[1] Based on submissions of the 2,377 programs that have maintained their openness attribute throughout their existence. Users' rank was determined based on their aggregated USD monetary rewards.
[2] Percentage measures relative to the activity of each user cohort.

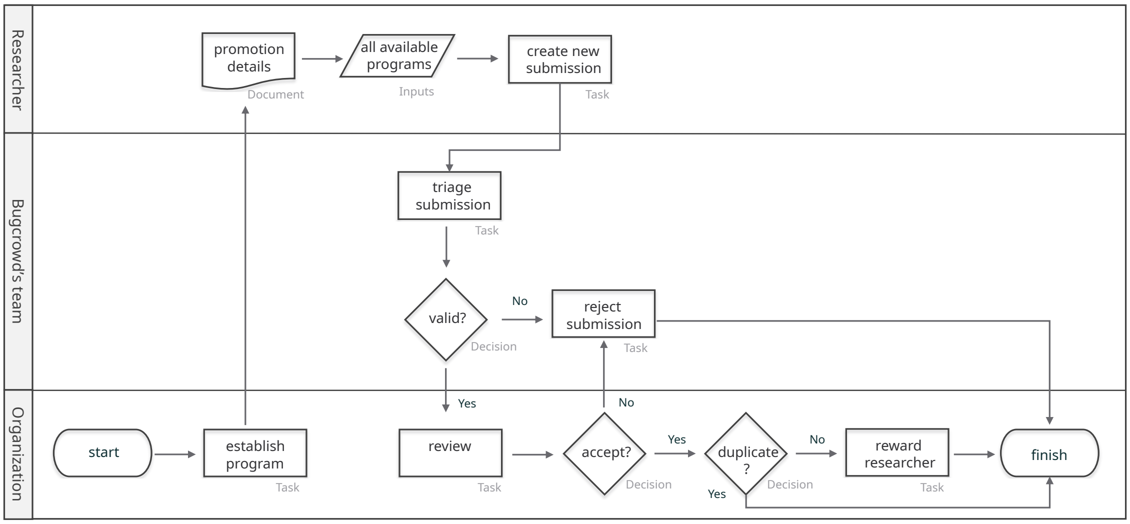

Figure B9 Bugcrowd's platform submission workflow (Zrahia et al. 2024).

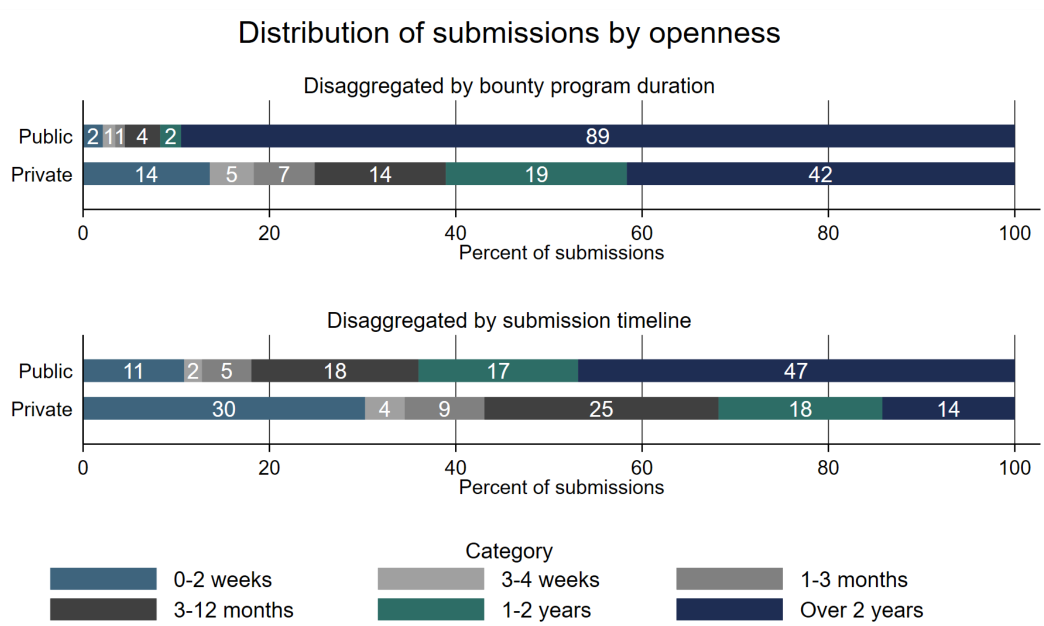

Figure B10 Total submissions by openness across program and submission timeline durations.

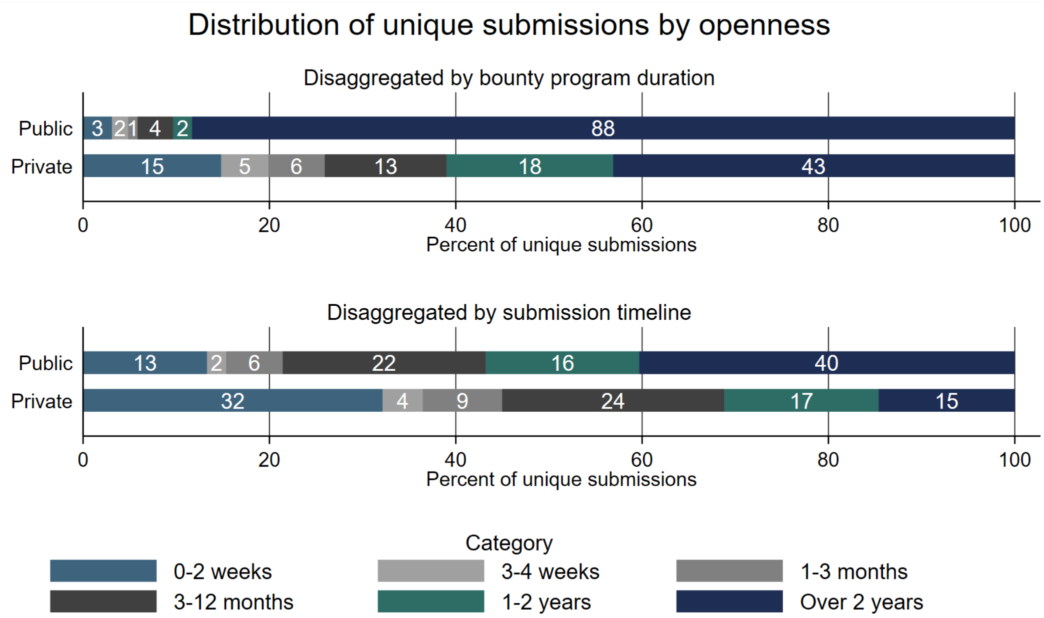

Figure B11 **Unique submissions by openness across program and submission timeline durations.**

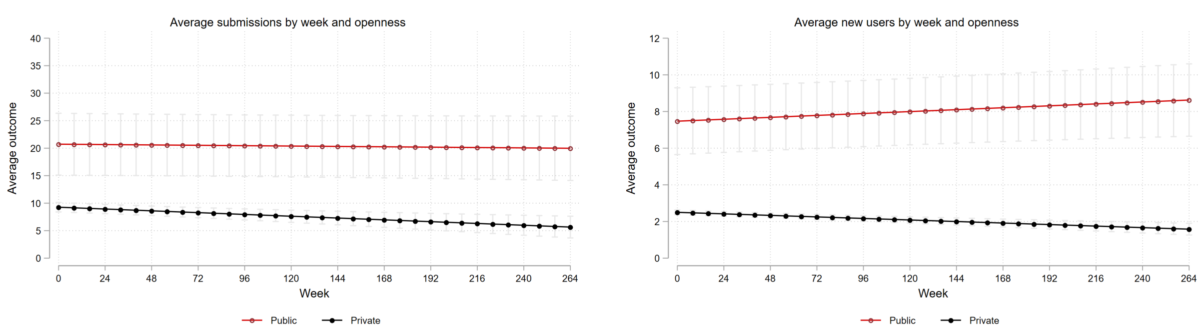

(a) Weekly submissions over time

(b) Weekly new distinct users over time

Figure B12 **Average predicted count of Weekly submissions and new users over time by openness**