KNOWLEDGE WORKER STRATEGY:
Effort Allocation Dynamics in Cybersecurity Crowdsourcing Platforms

Dana Etgar Itzhaki
Tel Aviv University

Neil Gandal
Tel Aviv University

Michael Riordan
Columbia University

May 2025

Abstract

This paper examines how knowledge workers—specifically, ethical hackers—allocate effort within cybersecurity crowdsourcing platforms. Using empirical data from Bugcrowd, a leading vulnerability disclosure platform, we explore how effort diversification across tasks influences outcomes under competitive, reward-based incentives. We develop a theoretical model grounded in knowledge economy principles to predict when researchers choose to specialize versus diversify. The analysis reveals that high-skill researchers benefit more from diversification, while others achieve better outcomes through focused effort. Our findings contribute to the understanding of digital labor markets, strategic decision-making, and productivity dynamics in knowledge-intensive ecosystems.

1. Introduction

> *"The new education must teach the individual how to classify and reclassify information... Tomorrow's illiterate will not be the man who can't read ; he will be the man who has not learned how to learn."* — **Alvin Toffler**

The world is changing rapidly. In modern digital economies, the value of work increasingly depends not on physical labor, but on the ability to process, apply, and adapt knowledge. This shift has generated intense academic interest in knowledge workers—individuals whose productivity hinges on intellectual expertise, creativity, and responsiveness to complex information environments. At the same time, digital platforms have transformed how these workers engage with markets, clients, and each other.

Traditional economic theory assumes rational expectations: individuals form beliefs about the future based on available information and use these beliefs to guide decision-making (Weizsäcker, 2010). However, in fast-changing digital environments, where uncertainty, competition, and rapid innovation dominate, this assumption becomes more fragile. Workers must make strategic choices without full information, relying instead on heuristics, learning, and dynamic adaptation. Understanding how they allocate resources in such conditions is central to analyzing productivity and strategic behavior in the modern economy.

etgaritzhki@mail.tau.ac.il
gandal@tauex.tau.ac.il
mhr21@columbia.edu

In this paper, we examine a specific but revealing case: ethical hackers who participate in cybersecurity crowdsourcing platforms, such as Bugcrowd. These researchers, also known as white-hat hackers, embody many of the traits central to knowledge economy debates. They operate autonomously, adapt to high uncertainty, and compete based on expertise. Crucially, their work is structured around digital tournaments, where only a subset of participants are rewarded based on performance. This environment provides a unique opportunity to study strategic decision-making under competitive pressure.

By analyzing ethical hackers as knowledge workers, we aim to connect insights from cybersecurity economics and the broader literature on knowledge labor. We focus on a central question: should a researcher allocate all effort to one task, or diversify across multiple tasks? How does this decision depend on skill and reward structure? Our theoretical model formalizes the tradeoff, and our empirical analysis uses real-world data to test its predictions.

We find that effort diversification depends on both the researcher's quality and the expected prize: high-quality researchers are more likely to benefit from diversification. This work contributes not only to understanding bug bounty dynamics but also to larger questions about productivity, incentives, and competition in knowledge-based labor markets.

The remainder of the paper is structured as follows. Section 2 reviews the literature on vulnerability markets and the knowledge economy, with particular attention to knowledge worker productivity and labor dynamics. Section 3 introduces the Bugcrowd platform and describes the empirical dataset. Section 4 presents a theoretical model of effort allocation and diversification, including its assumptions and predictions. Section 5 analyzes real-world data in light of the model, highlighting patterns in researcher behavior and success rates. Section 6 concludes by summarizing key findings and their implications for digital labor markets and cybersecurity strategy.

2. Related Literature

2.1 The Market for Vulnerabilities

A market for vulnerabilities is a structured or informal exchange where security flaws in software, hardware, or networks are identified, priced, and transferred between buyers and sellers. These transactions are often based on the flaws' potential exploitability, impact, and legal or ethical considerations. In the past decade, several studies have analyzed data from crowdsourcing platforms to better understand the dynamics of this market.

Zrahia et al. (2024) utilized the same dataset as this study to examine the impact of an exogenous external shock (COVID-19) on Bugcrowd. They found that there was an immediate and very large effect on the supply side (researchers) yet a much smaller demand effect (number of programs). The equilibrium outcome was a large increase in duplicate valid submissions, resulting in a lower probability of winning a monetary reward, and a corresponding decrease in the expected reward for a valid submission. Sridhar and Ng (2021) analyzed a data set from HackerOne platform. Their findings show that security researchers are motivated more by non-monetary factors (with a price elasticity of supply between 0.1 and 0.2), and factors like company revenue and brand profile don't significantly influence the number of valid vulnerability reports received, while program age negatively impacts report volume due to the increasing difficulty of finding new bugs.

Zhao et al. (2015) studied publicly available data of two representative web vulnerability discovery ecosystems (Wooyun and HackerOne) and showed that white hat communities in both ecosystems continuously grow, and monetary incentives have a significantly positive cor-

relation with the number of vulnerabilities reported. Maillart et al. (2017) have analyzed a data set of public bounty programs and found researchers tend to switch to newly launched bounty programs at the expense of existing ones. Subramanian and Malladi (2020) studied 41 public bounty programs and examined issues involved with their implementation. Algarni and Malaiya (2014) used an open vulnerability database to study the career, motivation, and methods of the most successful researchers. They concluded that a major percentage of vulnerabilities are discovered by individuals external to firms, and that financial reward is a major motivation, especially to researchers in Eastern Europe.

2.2 Knowledge Economy and Knowledge Workers

While the literature on vulnerability markets provides important insights into researcher behavior, it often treats participants as rational agents abstracted from their broader roles as knowledge workers. To address this, we draw on foundational work in the knowledge economy to better understand how white-hat hackers navigate digital labor platforms.

The knowledge economy (KE) is characterized by the centrality of information, expertise, and innovation in economic production. Houghton and Sheehan (2000) describe it as a system where intangible assets, rather than physical capital, drive productivity and value creation. Drucker (1999) famously argued that knowledge workers, whose output depends on intellectual skill rather than manual labor, would become the most valuable assets of organizations in the 21st century.

In cybersecurity platforms, ethical hackers clearly embody this role: they operate independently, apply specialized expertise, and compete based on insight rather than routine execution. Brynjolfsson et al. (2014) extend this discussion by describing a power-law economy, where a small number of highly skilled individuals capture a disproportionate share of value, much like top performers on tournament-based platforms such as Bugcrowd.

Powell and Snellman (2004) highlight the tension between worker autonomy and managerial control in knowledge-intensive systems. This tension is evident in crowdsourcing platforms, where researchers exercise discretion over task choice but remain constrained by platform structures and incentive mechanisms. Binney (2001) provides a multidimensional view of knowledge management, including collaboration, innovation, and sharing—activities that are also present in hacker communities.

Recent advances in AI and automation further shape this environment. For example, Wiles et al. (2024) show that generative AI not only automates routine tasks but also raises the strategic and cognitive demands placed on knowledge workers. As their tools evolve, ethical hackers must continuously adapt to remain effective and competitive in this rapidly changing landscape.

These perspectives frame ethical hackers not merely as users of digital labor platforms but as autonomous, strategic actors making effort allocation decisions under uncertainty. While a few studies have employed formal economic frameworks to analyze researcher behavior (Sridhar and Ng, 2021; Zrahia et al., 2024), this area remains underexplored. This paper aims to help fill that gap.

3. Bug Bounty Platform and Data

"It takes A crowd to defend A crowd" [1]

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

Many organizations lack the resources and diversified skills to find hidden vulnerabilities before attackers do. Unfortunately, using reactive tools alone leads to noisy, low-impact results that miss emerging risks. Even sophisticated companies can misjudge the creativity, patience, and diverse skills of today's attackers. Crowdsourcing emerged to address the skills gap and the imbalance between attackers and defenders by incentivizing ethical hackers to report critical bugs.

Bug bounty programs are a structured and legal way for security researchers to be rewarded for finding software vulnerabilities. The programs enable organizations to get in touch with cybersecurity experts ("white hat" hackers) whose knowledge complements that of the organizations' own development and testing teams. From the security researchers' side, these programs offer an opportunity to be rewarded legally for the vulnerabilities they find.

Products and services that bring together different groups of users are often referred to by economists as "two-sided markets" or "two-sided networks", (Rochet and Tirole (2006)). A two-sided network involves two different user groups that interact through an intermediary or platform. The value for one group of users typically depends on the size and engagement of the other group. A key feature of two-sided networks is the cross-side network effect. This means that the value of the network to one group of users increases as the number of users in the other group grows.

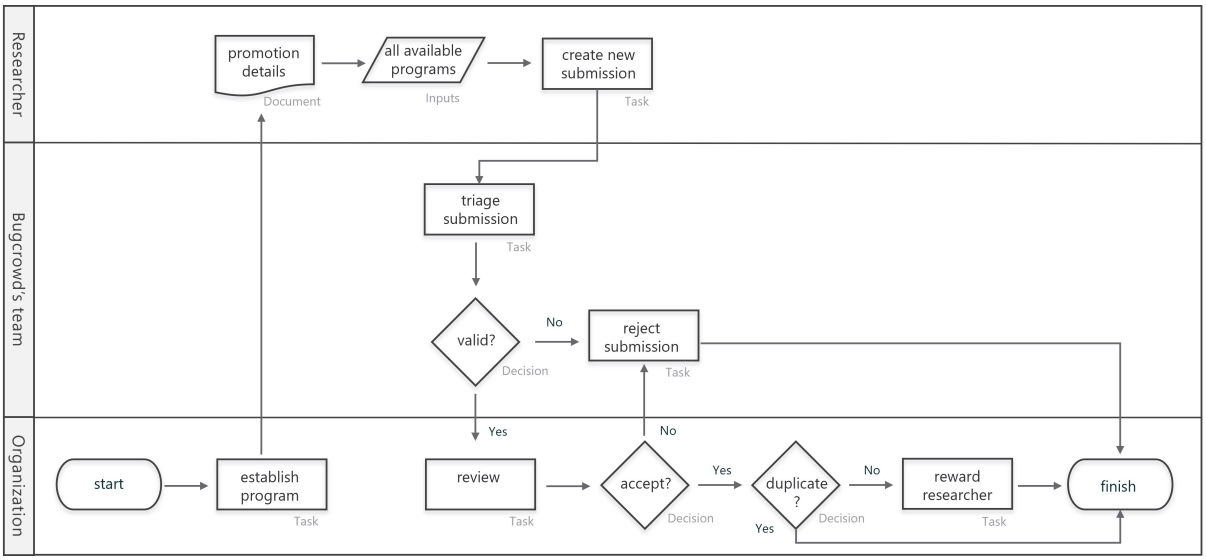

Figure 1: Submission workflow in Bugcrowd platform

Bugcrowd is one of the two top leading platform for crowd sourcing. Figure 1 (Zrahia et al. (2024)) details the workflow for submissions over Bugcrowd's platform. Prior to starting a program, the organization defines its objectives and goals, including the exact list of software programs to be tested (web applications, APIs, mobile versions, etc.). The next step is shaping the researcher engagement plan, and specifically the program's duration (continuous or ad-hoc), researchers' access (public or private), the payment range per vulnerability (by priority), and more. Submissions are categorized according to a priority scale of P1 to P5 where P1 are critical vulnerabilities and P5 are informational weaknesses which may not even be fixed. The platform provides a well-defined Vulnerability Rating Taxonomy (VRT) for researchers to determine the priority of their submission. Once the program is launched, organizations have their teams ready to process the incoming submissions, after they have been verified, triaged (prioritized) and screened for duplicates and relevancy by the platform's team. Valid

vulnerabilities are then integrated into the existing Software Development Lifecycle (SDLC) tools to be fixed, and related reward payouts are processed accordingly.

Bugcrowd offers two types of programs; Managed Bug Bounty programs (MBBs) give a monetary reward to the first researcher to submit a unique valid vulnerability as well as points.[2] Bug bounty platforms thus creates a tournament-like arrangement between researchers in the platform.

The data set, taken from Bugcrowd platform, spans from 2012 and includes all bug submission activity through May 2021. During that time, Bugcrowd's platform has hosted more than 2,400 programs offered by more than 1,000 organizations, and attracted more than 30,000 active researchers who made at least one submission to a program. We have detailed data on almost 500,000 submissions made during that time. The data on submissions specifies: Researcher that made the submission; Date and time of the submission (by seconds); The program and whether it is private or public; Whether the submission was valid, and if so if it was paid or a duplicate; The amount paid (in US dollars) and the amount of points awarded.

Table 1: Activity in the Platform

Year	Submissions	Active Users	Active Bounties	Total Rewards ($)	Success Rate
2013	962	229	9	20,625	10
2014	13,745	1,174	39	171,117	6
2015	13,656	1,737	93	575,422	16
2016	22,185	3,046	161	2,878,811	22
2017	29,002	3,560	247	5,089,244	26
2018	38,050	4,983	369	8,669,754	27
2019	69,515	7,258	512	11,979,593	20
2020	107,005	11,215	699	11,475,884	14
2021	29,783	5,898	439	4,067,639	16

The data set also contains information on the organizations side: firm size, country of origin, and when it first joined the platform. Many firms run simultaneously more than one program and for each we have its status, start/end dates, and whether it is open to everyone or only to selected researchers. While only the first researcher to discover a valid vulnerability is awarded a monetary payment, the data set also records duplicate valid submissions.

Table 1 provides a broad view of the overall activity on the platform. The trend shows an increase in activity, with more users, submissions, and bounties over the years. However, it is notable that the total rewards decreased between 2020 and 2021, as Zrahia et al. (2024) showed in their work.

One important limitation of our dataset is that it includes only data from the Bugcrowd platform. While Bugcrowd is a leading player in the vulnerability disclosure ecosystem, other platforms such as HackerOne or Synack may differ in terms of access policies, workflows, and researcher composition. As such, our findings may not fully generalize across all bug bounty ecosystems.

Nonetheless, this paper aims to contribute to a broader understanding of labor patterns within digital knowledge work. Bug bounty platforms offer a valuable empirical window into

[2]The second type of program, Vulnerability Disclosure Programs (VDPs) rewards hackers with points, but no monetary awards. We focus on MBB programs in this research.

how autonomous knowledge workers, operating in competitive, skill-based markets, strategically allocate effort. The insights we derive from this case study may help illuminate common dynamics in other sectors of knowledge-intensive labor.

4. Theoretical Framework: Returns to Diversification in Bug Bounty Work

This section presents a two-tiered theoretical framework to analyze strategic effort allocation by ethical hackers on bug bounty platforms. Our goal is to explore the economic rationale behind diversification and specialization. We begin with a generalized model that captures the decision environment abstractly, and then introduce a tractable toy model to derive analytical insights and connect to empirical observations.

4.1 Generalized Framework

We consider a population of heterogeneous agents (ethical hackers), each with a unified finite effort budget E, which can be distributed across multiple bounty programs. The goal of each agent is to maximize expected utility, defined as the sum of expected rewards across programs minus the cost of total effort.

Let agent j allocate efforts $e_{1j}, e_{2j}, \ldots, e_{Nj}$ to N programs, subject to $\sum_{i=1}^{N} e_i \leq E, \forall_j$. The success probability function for task i is $s_i(e_{ij})$, an increasing function that may include regions of both increasing and decreasing marginal returns.

Agent quality is captured by a parameter λ_j, where lower λ indicates higher ability. The cost function is $C_j = \lambda_j h\left(\sum_{i=1}^{N} e_{ij}\right)$, where $h(\cdot)$ is increasing and convex.

The agent j optimization problem is:

$$\max_{\{e_i\}} \sum_{i=1}^{N} R_i s_i(e_i) - \lambda_j h\left(\sum_{i=1}^{N} e_i\right)$$

To study returns to diversification, we compare different feasible effort allocations. Rather than posing a binary choice between complete specialization and full diversification, we analyze the continuum of partial allocations to multiple tasks. For example, we define the return to a marginal increase in diversification from an allocation $\mathbf{e} = (e_1, \ldots, e_k, 0, \ldots, 0)$ by evaluating the utility gain from reallocating a small portion of effort to an additional task $k+1$, holding total effort constant. This approach captures the nuanced trade-offs agents face in realistic platform environments.

5. An Illustrative Example of Effort Diversification

We now present a simplified example of the general model with two bounty programs. Each researcher allocates a total effort budget $E = 1$, deciding whether to specialize in one program or diversify effort across two. Let R be the expected reward from a successful submission (assumed equal across programs), and let the probability of success be a concave increasing function of effort:

$$s(e) = \frac{e^2}{2} - \frac{e^3}{3}$$

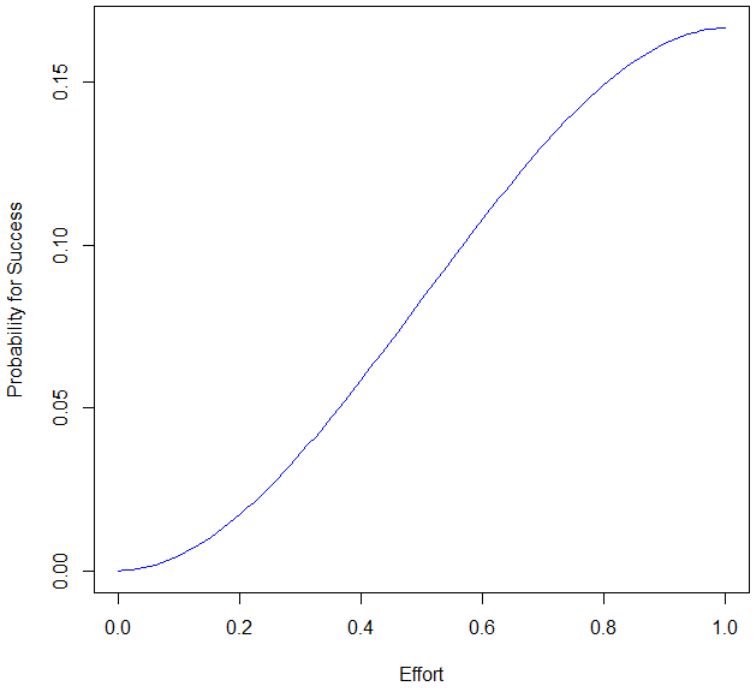

Figure 2: Success probability function.

The cost of exerting effort e is:

$$c(e) = \frac{\lambda e^3}{6}$$

where $\lambda > 0$ captures the researcher's inefficiency (higher λ implies greater marginal cost of effort). We use a cubic cost function to reflect increasing difficulty in effort scaling, particularly relevant in high-skill knowledge work such as vulnerability discovery, where sustained cognitive effort results in sharply rising fatigue or failure rates. This functional form captures the intuition that while initial effort may be relatively inexpensive, maintaining high levels of mental engagement becomes disproportionately more taxing—both cognitively and strategically.

The researcher's utility is:

$$U = R \cdot s(e) - c(e)$$

5.1 Specialization Strategy

The researcher allocates all effort to a single program ($e = 1$):

$$U_1 = R\left(\frac{1}{2} - \frac{1}{3}\right) - \frac{\lambda}{6} = \frac{R}{6} - \frac{\lambda}{6}$$

5.2 Diversification Strategy

Effort is split evenly across two programs ($e = 0.5$ per program):

$$U_2 = 2 \cdot \left[R \left(\frac{0.25}{2} - \frac{0.125}{3} \right) \right] - \frac{\lambda}{6} = \frac{R}{6} - \frac{\lambda}{6}$$

At this point, both strategies yield equal utility. However, to explore optimal effort choices more generally, we derive utility-maximizing effort allocations for each strategy.

5.2.1 Optimized Specialization

$$U_1(e) = R \left(\frac{e^2}{2} - \frac{e^3}{3} \right) - \frac{\lambda e^3}{6}$$

Maximizing with respect to e, we find:

$$\frac{dU_1}{de} = R(e - e^2) - \frac{\lambda e^2}{2} = 0 \quad \Rightarrow \quad e_1 = \frac{2R}{2R + \lambda}$$

Substituting back:

$$U_1 = \frac{2R^3}{3(2R + \lambda)^2}$$

5.2.2 Optimized Diversification

$$U_2(e) = 2R \cdot s \left(\frac{e}{2} \right) - \frac{\lambda e^3}{6} = R \left(\frac{e^2}{4} - \frac{e^3}{12} \right) - \frac{\lambda e^3}{6}$$

Taking the derivative:

$$\frac{dU_2}{de} = \frac{R}{2} e - \left(\frac{R}{4} + \frac{\lambda}{2} \right) e^2 = 0 \quad \Rightarrow \quad e_2 = \frac{2R}{R + 2\lambda}$$

Substitute to get:

$$U_2 = \frac{R^3}{3(R + 2\lambda)^2}$$

5.3 Comparative Statics

The return to diversification is:

$$\Delta U = U_2 - U_1 = \frac{R^3(2R^2 - 4R\lambda - 7\lambda^2)}{3(2R + \lambda)^2(R + 2\lambda)^2}$$

This expression reveals that diversification dominates when:

$$R > \frac{1}{2}(2\lambda + 3\sqrt{2\lambda})$$

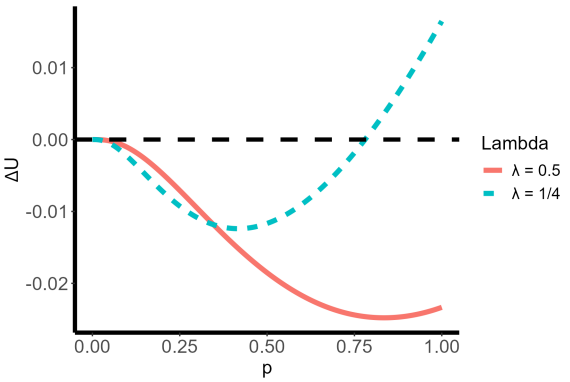

Figure 3: Return to diversification

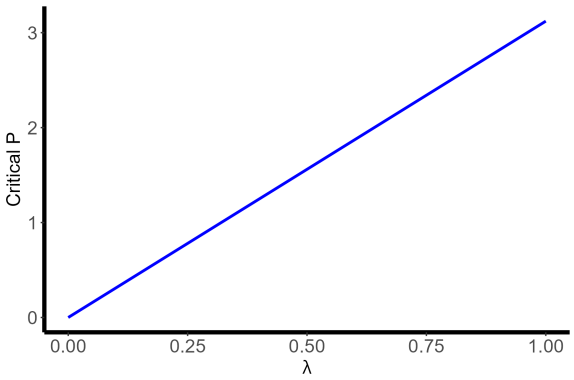

Figure 4: Critical value of R

Interpretation. This example illustrates how the returns to diversification depend on both the expected prize R and the researcher's efficiency parameter λ. Researchers with higher expected success rates or lower marginal costs of effort (i.e., higher-quality researchers) benefit more from spreading their effort across multiple tasks. This is because they are more likely to convert partial effort into successful outcomes. In contrast, lower-quality researchers face steep marginal costs and may fail to generate rewards when effort is diluted. As such, diversification emerges as a strategic advantage primarily for the elite group—those with the skill, experience, or insight to navigate multiple projects effectively and allocate effort where returns are highest.

5.4 Numerical Illustration

To illustrate the model's implications, we present a comparison of optimal effort allocation under different parameter values. Table 2 reports the values of total effort for a single project (e_1) and for two projects (e_2) under varying levels of the researcher's effort cost parameter λ and expected prize R. Since e_2 represents total effort across two projects, the effort per project in the diversified case is given by $\frac{e_2}{2}$.

Expected Prize (R)	**Effort Cost** (λ)	**Single Project Effort** (e_1)	**Two Projects Effort** (e_2)
2.0	0.1	0.98	0.95
2.0	0.25	0.94	0.89
2.0	0.5	0.89	0.80
1.0	0.1	0.91	0.83
1.0	0.25	0.80	0.71
1.0	0.5	0.67	0.57
0.5	0.1	0.80	0.67
0.5	0.25	0.67	0.50
0.5	0.5	0.50	0.33

Table 2: Optimal effort allocation under varying levels of researcher inefficiency (λ) and expected prize (R).

The relationship between effort cost (λ), expected prize (R), and optimal effort allocation (e_1, e_2) reveals how researchers strategically manage their participation. As λ increases, both e_1 and e_2 decline, indicating that higher marginal effort costs discourage intensive engagement. However, when λ is low, the gap between e_1 and e_2 is minimal—suggesting that high-quality researchers (low λ) can sustain meaningful effort across multiple projects.

The level of the prize R also plays a key role: higher rewards incentivize greater effort in both strategies, and the utility gap between specialization and diversification narrows. Conversely, at low prize levels, overall effort declines, and the relative drop in e_2 is steeper, reflecting reduced viability of diversification under weak incentives. Taken together, the numerical results support the notion that diversification is most attractive to skilled researchers operating under favorable conditions—namely, low effort costs and high potential returns.

6. Empirical Results and Discussion

This section examines the relationship between effort, success, and diversification using real-world data from the Bugcrowd platform, providing empirical support for the theoretical model.

Our dataset includes 31,754 researchers. A striking feature is the skewed participation distribution: 14,212 submitted only once, and a small fraction made thousands of submissions, with the most prolific submitting over 6,000 times. This heterogeneity aligns with the model's assumption of variation in researcher quality.

Figure 5 illustrates an important result. This plot explores the relationship between effort and success, as shown in Figure 3, using real-world data. The data is divided into two-week intervals. Effort is defined as the total number of submissions made by a single researcher within each two-week period; only submissions for paying bounties are counted. Each dot in the plot represents the total number of submissions made by a researcher on the platform during a two-week interval. Success is measured by the number of submissions that received a monetary reward. A few classifications for effort were tested, all of which calculated the number of submissions within a specific period of time. Each method produced similar results.

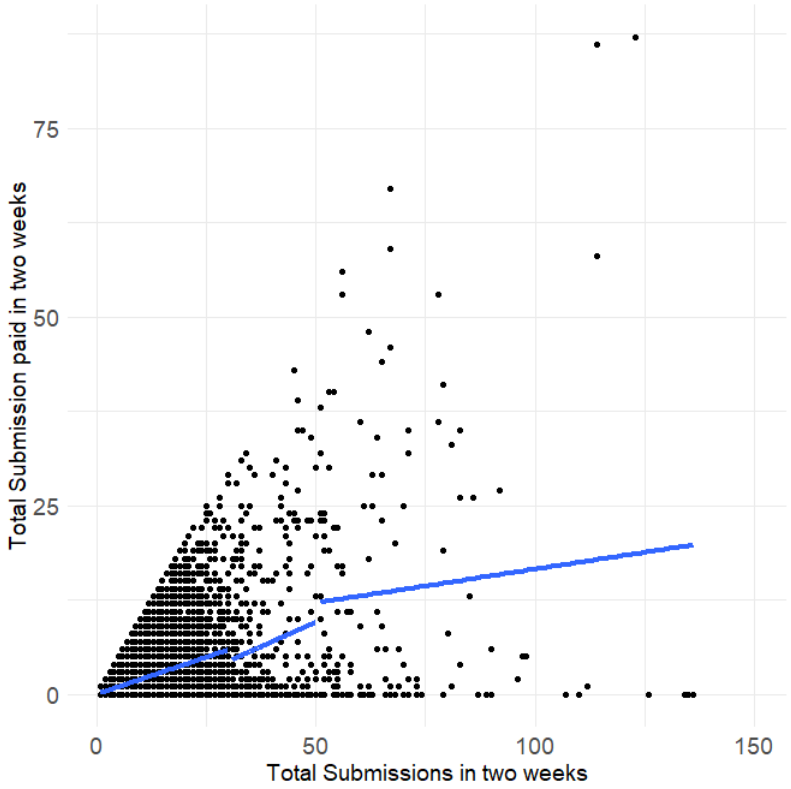

Figure 5: Success relative to effort in real-world data

As predicted by the theoretical model, the empirical analysis shows that success increases

as effort increases. This aligns with the assumption that researchers allocate their effort strategically to maximize their expected reward. Moreover, the curve highlights distinct regions with increasing and decreasing marginal returns to effort. In the increasing marginal returns phase, additional effort yields disproportionately higher success rates. This aligns with the model's prediction that, at lower effort levels, researchers face a learning curve or knowledge accumulation effect, where each additional unit of effort contributes significantly to performance. This could be due to factors such as increased familiarity with the platform, improved problem-solving strategies, or network effects where accumulated expertise leads to more efficient submissions.

Conversely, in the decreasing marginal returns phase, the model predicts that at higher levels of effort, additional work leads to diminishing gains in success. This could be due to cognitive fatigue, saturation of available vulnerabilities, or increased competition, which makes each additional submission less likely to yield a reward. The model incorporates this by allowing effort to have a nonlinear effect on expected outcomes, where beyond a certain threshold, the probability of success per additional unit of effort declines.

We then explore diversification: do researchers perform better when focusing on a single bounty or distributing effort across several? Figure 6 shows that while submission diversity increases with overall activity, average rewards do not follow a monotonic trend.

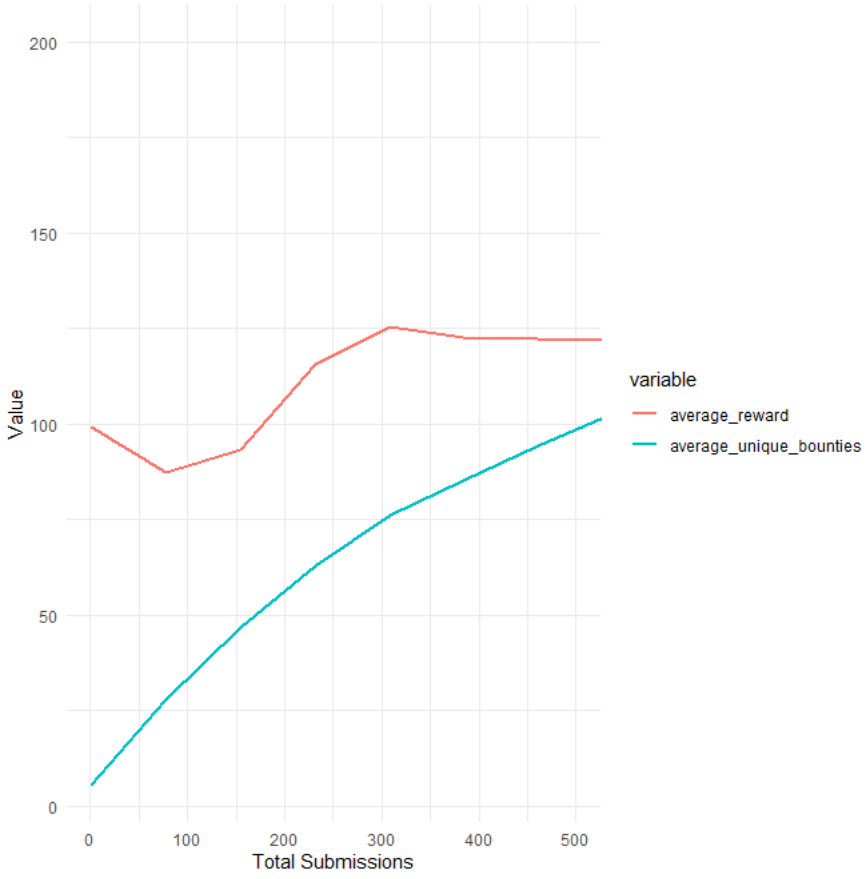

Figure 6: Comparison of average number of bounties and reward levels

Among researchers who submitted twice, those focusing on one bounty earned more than those diversifying across two. However, in a targeted analysis of a one-month high-traffic period in 2019 (chosen for its elevated activity and pre-COVID consistency), we observe a striking reversal among the top 10% of researchers ranked by success rate: those who diversified earned

significantly more. Specifically, they averaged $3,200 compared to $794 for those focusing on a single bounty. This supports the theoretical prediction that diversification is more profitable for highly skilled individuals.

Table 3 summarizes these findings:

	Didn't diversify	**Diversified**
Total population	$203	$128
Top 10%	$794	$3,200

Table 3: Reward comparison by diversification strategy

We also observe a highly skewed distribution of platform activity. Figure 7 shows that about 20% of researchers account for 85% of submissions, a classic long-tail distribution. This suggests a Pareto-like distribution of quality, with high performers driving most of the output.

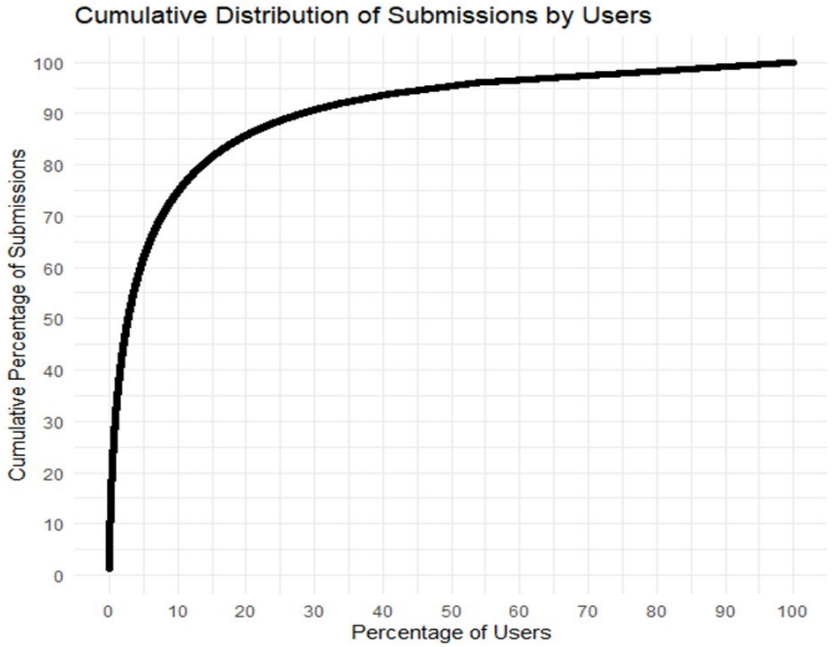

Figure 7: CDF of researcher submissions

These findings reinforce the idea that effort and success are shaped by strategic behavior in a competitive ecosystem. Skill heterogeneity and reward structures create strong incentives for high performers to diversify, while lower-performing researchers tend to specialize due to limited capacity or higher marginal costs. This highlights the role of economic decision-making even in technical and knowledge-based domains like cybersecurity.

7. Conclusion

This paper contributes to the understanding of cybersecurity labor markets through the lens of economic theory. By modeling ethical hackers as knowledge workers operating in a competitive platform economy, we frame effort allocation as a strategic decision governed by incentives, costs, and skill.

The theoretical model demonstrates how researchers optimize effort under uncertainty, and the empirical results from Bugcrowd data validate these predictions. We find that effort positively correlates with success but exhibits diminishing marginal returns. Diversification is not universally beneficial—only top-performing researchers appear to profit from it.

Our findings emphasize the relevance of economic models in analyzing behavior within cybersecurity platforms. While most research in this field comes from technical or managerial perspectives, our approach foregrounds incentives, optimization, and efficiency—concepts that are central to labor economics and platform studies.

Looking forward, this work lays the foundation for extensions incorporating dynamic strategies, learning, and cross-platform interactions. More advanced statistical modeling could further explore the determinants of success and uncover latent researcher traits. As digital labor markets evolve, particularly in high-skill sectors like cybersecurity, economic frameworks will be increasingly essential for designing effective platforms and policies.

References

Algarni, A. M. and Malaiya, Y. K. (2014). Software vulnerability markets: Discoverers and buyers. *International Journal of Computer and Information Engineering*, 8(3):480–490.

Binney, D. (2001). The knowledge management spectrum–understanding the km landscape. *Journal of knowledge management*, 5(1):33–42.

Brynjolfsson, E., McAfee, A., and Spence, M. (2014). New world order: labor, capital, and ideas in the power law economy. *Foreign Affairs*, 93(4):44–53.

Drucker, P. F. (1999). Knowledge-worker productivity: The biggest challenge. *California management review*, 41(2):79–94.

Houghton, J. and Sheehan, P. (2000). *A primer on the knowledge economy*. Victoria University.

Maillart, T., Zhao, M., Grossklags, J., and Chuang, J. (2017). Given enough eyeballs, all bugs are shallow? revisiting eric raymond with bug bounty programs. *Journal of Cybersecurity*, 3(2):81–90.

Powell, W. W. and Snellman, K. (2004). The knowledge economy. *Annu. Rev. Sociol.*, 30(1):199–220.

Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667.

Sridhar, K. and Ng, M. (2021). Hacking for good: Leveraging hackerone data to develop an economic model of bug bounties. *Journal of Cybersecurity*, 7(1):tyab007.

Subramanian, H. C. and Malladi, S. (2020). Bug bounty marketplaces and enabling responsible vulnerability disclosure: An empirical analysis. *Journal of Database Management (JDM)*, 31(1):38–63.

Weizsäcker, G. (2010). Do we follow others when we should? a simple test of rational expectations. *American Economic Review*, 100(5):2340–2360.

Wiles, E., Krayer, L., Abbadi, M., Awasthi, U., Kennedy, R., Mishkin, P., Sack, D., and Candelon, F. (2024). Genai as an exoskeleton: Experimental evidence on knowledge workers using genai on new skills. *Available at SSRN 4944588*.

Zhao, M., Grossklags, J., and Liu, P. (2015). An empirical study of web vulnerability discovery ecosystems. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 1105–1117.

Zrahia, A., Gandal, N., Markovich, S., and Riordan, M. H. (2024). The simple economics of an external shock to a bug bounty platform. *Journal of Cybersecurity*, 10(1):10–1093.