The Simple Economics of an External Shock to a Bug Bounty Platform

Aviram Zrahia[1,*], Neil Gandal[2], Sarit Markovich[3], and Michael Riordan[4]

[1]The Department of Public Policy, Tel Aviv University
[2]Berglas School of Economics, Tel Aviv University and University of Tulsa
[3]Kellogg School of Management, Northwestern University
[4]The Department of Economics, Columbia University

[*]Corresponding author: aviramzrahia@mail.tau.ac.il

Abstract

We first provide background on the "nuts and bolts" of a bug bounty platform: a two-sided marketplace that connects firms and individual security researchers ("ethical" hackers) to facilitate the discovery of software vulnerabilities. Researchers get acknowledged for valid submissions, but only the first submission of a distinct vulnerability is rewarded money in this tournament-like setting. We then empirically examine the effect of an exogenous external shock (COVID-19) on Bugcrowd, one of the leading platforms. The shock presumably reduced the opportunity set for many security researchers who might have lost their jobs or been placed on a leave of absence. We show that the exogenous shock led to a huge rightward shift in the supply curve and increased the number of submissions and new researchers on the platform. During the COVID period, there was a significant growth in duplicate (already known) valid submissions, leading to a lower probability of winning a monetary reward. The supply increase resulted in a significant decline in the equilibrium price of valid submissions, mostly due to this duplicate submission supply-side effect. The results suggest that had there been a larger increase in the number of firms and bug bounty programs on the platform, many more unique software vulnerabilities could have been discovered.

Key words: Bug Bounty Platforms, Software Vulnerabilities, Exogenous Shock, COVID-19.

Introduction

While cybersecurity attacks occur for many reasons and in different ways, exploiting software vulnerabilities (also known as bugs) is one of the primary attack vectors. Bug bounty programs and platforms utilize crowd-sourcing to find these bugs, with the notion that "Given enough eyeballs, all bugs are shallow" [1, p. 29].

Bug bounty programs are a structured and legal way for security researchers to be rewarded for finding software vulnerabilities. These programs enable organizations to connect with ethical hackers (hereafter, researchers) whose cybersecurity expertise and knowledge complement that of the organizations' development and testing teams. Such programs allow researchers to be rewarded legally for the vulnerabilities they find.

Bug bounty platforms are two-sided markets that bring bug bounty programs and researchers together. The programs are structured as tournaments where companies pay monetary rewards only for unique vulnerabilities found. Top researchers might get invited to private programs where only selected researchers can participate, thereby increasing the probability of being the first to find and report a vulnerability.

The popularity of bug bounty programs and platforms is constantly increasing, potentially building the foundation for an overall higher demand. Governmental agencies have begun to use bug bounty programs. In 2021, the US-based Cybersecurity & Infrastructure Security Agency (CISA) launched a federal civilian enterprise-wide crowdsourced vulnerability disclosure policy platform.[1] Additionally, cyber insurance firms have begun to recognize the benefits of firms participating in a bug bounty program.[2]

[1] `https://www.cisa.gov/blog/2021/07/29/cisa-announces-new-vulnerability-disclosure-policy-vdp-platform`.

[2] Marsh, a large global insurance firm, identifies as part of its "cyber catalyst" program security products that reduce cyber risks. Participation in the HackerOne bug bounty platform is considered a "certified" product that can lower cyber insurance

Bug bounty programs and platforms are part of a more general "gig economy" trend where enterprises supplement labor, and workers supplement income with gig work. Gig work platforms provide enterprises with skilled and flexible labor access and workers with opportunities to compete in global job markets. Furthermore, they facilitate "bridge employment" (i.e., temporary employment between career jobs) and provide income opportunities in downtimes when the market does not accommodate full-time employment. Indeed, studying the ride-sharing market, Koustas [2] finds that, on average, driving for gig platforms replaced 73% of lost income from a primary job. Moreover, taking advantage of gig work platforms during bad times helps overcome periods of income volatility. Similarly, Collins et al. [3] find that workers typically start new platform work in times of a personal income crisis.

Employing a unique data-set provided by Bugcrowd,[3] we document the effect of an exogenous shock – the COVID-19 pandemic – on the market for vulnerabilities within the bug bounty platform. The data covers the 2017-21 period, and we focus each year on the three months from March to May. We examine the impact of the COVID-19 shock represented by the 2020 period on the demand for vulnerabilities by participating organizations and the supply of vulnerabilities from active researchers.

Since its launch, Bugcrowd's platform has hosted more than 2,400 programs offered by more than 1,000 organizations while attracting more than 30,000 researchers who made at least one submission to a program. We have data on both valid and invalid submissions; valid submissions are correctly identified vulnerabilities within the defined scope of a program. While only the first researcher to discover a valid vulnerability is eligible for a monetary incentive, our data also includes details on duplicate valid submissions. In these instances, a researcher accurately identified a vulnerability but was not the first to report it and, therefore, was not compensated. We can calculate the average payment for valid submissions by considering both paid and duplicate submissions. This variable turns out to be a key in comprehending the impact of the COVID-19 shock.

We develop a heuristic supply-and-demand model following the notion that information security can be explained more clearly using the language of microeconomics [4]. We define valid submissions as the relevant product and the average monetary reward for valid submissions as the price. Our analysis shows that the COVID-19 shock substantially shifted the supply curve by significantly increasing the number of active researchers and submissions and shifted the demand curve more moderately by slightly increasing the growth in programs with submission activity. These shifts, combined, significantly increased the number of valid submissions primarily due to the substantial growth in the number of duplicates, reflecting the much more significant supply curve shift. Consequently, there was a large decrease in the average equilibrium price for valid submissions because valid duplicates do not receive a monetary reward. Our regression analysis suggests that the COVID-19 main effect on the paid rewards was mostly for high-priority (critical) valid submissions, and it was significant both statistically and economically: our results show a $640 average reduction in payments for critical vulnerabilities during the shock period.

The COVID-19 shock "threw" the market for vulnerabilities out of a previously more-or-less stable equilibrium while impacting all ecosystem players. It presumably reduced the outside opportunity set for researchers who lost their jobs or were placed on a leave of absence during that period. These researchers likely had more time on hand to look for bug bounty program vulnerabilities. On the demand side, there was a smaller increase in new programs relative to previous periods, even though many organizations adapted to the pandemic by allowing their employees to work from home, potentially enabling "black hat" hackers to take advantage of the increased security vulnerabilities of the less effective home security systems and the newly deployed remote access solutions.[4] This is likely because of a relatively long start-up time for new bug bounty programs.

First, we examine the implications for the researcher community. The ratio of paid to valid submissions, representing the probability of being paid, decreased over time, reflecting a slightly increased competitive trend. However, this ratio fell dramatically during the COVID period and bounced back in the following 2021 period. The reduction in the average equilibrium price for a valid submission due to the COVID-19 shock presumably dampened the incentives of individual researchers to search for vulnerabilities.[5] Specifically, there was less than a one-in-four chance of being paid for a valid submission in 2020, compared to an average of one-in-three in 2019 and 2021. The dramatic reduction in the paid-to-valid submissions ratio in 2020 is solely due to the supply side, and it accounts for most of the equilibrium price decline for valid submissions.

We also examine the implications for the general public. The discovery of vulnerabilities in products and services (followed by patching) has a positive externality on the direct consumers and the public. Consequently, the following counterfactual argument elaborated in the Policy and Platform Implications section can be made: Had the demand response met the supply in the COVID period, the total number of distinct submissions could have been 64% greater than the actual count. This counterfactual implies there might have been a missed opportunity to examine more software and find more unique vulnerabilities during the COVID-19 shock period.

Furthermore, the COVID-19 shock provides a unique opportunity to address key public policy issues associated with crowd-sourcing and the gig economy. An often-mentioned benefit of the gig economy is that the response to an external shock on the supply side should be almost instantaneous. Here, we show that this was indeed the case, and we quantify the effects of the increased supply of new researchers and submissions, who likely joined when the value of outside options was lowered. However, for the gig economy to act as a genuine means to mitigate large income drops during bad times, one needs to encourage a significant response from the demand side.

prices (`https://www.marsh.com/us/services/cyber-risk/products/cyber-catalyst.html`).

[3] `https://bugcrowd.com/`.

[4] Consequently, many organizations experienced a significant increase in the number and severity of security incidents as the attack surface expanded. The change in attack patterns following COVID-19 has been documented extensively by market analysts, cyber-security vendors, and governmental agencies. See also Lallie et al. [5].

[5] Indeed, the total number of submissions fell dramatically in 2021 relative to 2020.

Finally, we look at the implications of the COVID-19 shock on the organizations and the platform. Though the vast increase in the number of submissions did not result in a significant growth in the number of unique vulnerabilities discovered, it did increase the costs associated with submission processing. These transaction costs may be split between the platform and the organization as determined by their specific contract and approved procedures.[6] Regardless of the party who bore this cost, it increased significantly during COVID-19 because many more duplicate submissions were processed, and their value to the organization (already aware of these discovered vulnerabilities) is zero.[7]

The paper proceeds as follows: In the Background section, we elaborate more on the real-world dynamics of vulnerabilities and the literature associated with bug bounty programs and platforms. The Bugcrowd Platform section includes details on engagement rules and submission workflow. The Data section describes our data-set, and the Heuristic Model section suggests a supply-and-demand model that motivates our analysis. The Effect of COVID-19 Exogenous Shock section presents our main results related to the platform's supply side, demand side, and equilibrium properties. Finally, the Policy and Platform Implications section briefly discusses the findings and their meaning.

Background

The life cycle of a vulnerability (or "bug") starts with its creation during coding. Assuming that adversaries do not find the vulnerability first, it will likely become known to the vendor either by internal testing or due to responsible disclosure done by a researcher, also known as a white-hat or ethical hacker.[8] Once discovered and verified, a patch that eliminates the vulnerability will be offered to all users of the affected product. This process is similar whether the vulnerability is found in a software product or an online service. For a product, the vendor will most likely release a technical security notification to its customers (pre-scheduled or emergency) detailing the importance and associated risks of the patched vulnerability and the affected software versions. The vulnerability will also be listed in publicly available feeds such as CVE and NVD.[9]

There are markets for vulnerabilities as a product, both from the adversarial and defense perspectives. In this paper, we focus only on the legal defensive market, sometimes referred to as the "white market," rather than on the "black market" for exploits.[10]

Bug Bounty Platforms as Market Intermediaries

As Malladi and Subramanian [9] note, there are three categories of security crowd-sourcing markets for vulnerabilities.

- The first category is institutional bug bounty programs hosted directly by software vendors who set their own policies and compensation plans. They solicit external researchers to find bugs in their products for monetary and non-monetary incentives. While this is a feasible option for large firms, it typically is not cost-effective for most firms.
- The second category is via private intermediaries that purchase vulnerabilities from researchers to sell them further downstream.
- The third category, which is the focus of this paper, is bug bounty platforms (intermediaries) that connect organizations and security researchers via a "two-sided" network or platform.

Economists often refer to products and services that bring together different groups of users as "two-sided markets" or "two-sided networks" [10]. These markets take many forms. Some common examples of such two sides brought together by a platform owner are buyers and sellers (Amazon), media consumers and advertisers (Facebook), application developers and device makers (Apple iOS), or users and application developers (Google Cloud Platform).[11]

Platform-based markets are typically characterized by indirect (cross-side) network effects [12, 13, 14], where each side's perceived value of the platform increases with the number of users on the other side.[12] In general, the platform provides the infrastructure and rules of engagement to attract both sides of the market. Two-sided platforms create value and improve economic efficiency [10, 16], possibly by reducing the transaction costs faced by distinct groups of participants.[13]

Bug bounty platforms are two-sided markets connecting organizations that want to crowd-source software security with researchers. Ideally, a platform hosts many programs for many organizations and has many high-quality active researchers. The researcher who is the first to find and report a novel vulnerability receives a payment (bounty). Bug bounty platforms thus create a tournament-like arrangement and establish the rules of engagement and submission procedures. However, the firm is the one that determines the program structure, testing scope, and the range of rewards to be paid to researchers.

In two-sided bug bounty platforms, individual researchers are sellers, the organizations initiating the bounty programs are buyers, and the discovered vulnerabilities are products. The demand comes from firms interested in protecting their software products or services against exploits used by adversaries. The

[6] Our data-set does not include details on the agreements, prices, or costs associated with the platform-organization joint activity.

[7] In line with Herley [6], these costs and other externalities should not be neglected.

[8] If a firm discovers its own vulnerabilities, there is no market for "white hat hackers." See Choi et al. [7] for a theoretical model that addresses disclosure in this setting.

[9] CVE® is a list of publicly known cyber-security vulnerabilities maintained by the MITRE Corporation. It feeds NIST's US National Vulnerability Database (NVD), which adds more context.

[10] If an adversary discovers the vulnerability before the firm, they might produce a zero-day (0-day) exploit, best defined as an "exploit without a patch." There is a "black market" for zero-day vulnerabilities, as described by Ablon and Bogart [8].

[11] Arce [11] considers cloud services to be a two-sided market.

[12] New platforms are often confronted with the problem that both sides will only join the platform when they expect sufficient numbers of the other group to join. This initial problem of getting all sides of the market on board is called the chicken-and-egg problem [15].

[13] This is in line with Munger [17], who argues that reducing transaction costs is the driving force of the sharing economy.

market's supply side consists of researchers eager to get paid for their expertise.

The magnitude of the paid bounties is at the company's discretion and depends primarily on the severity of the vulnerability found. Payments, however, are also affected by factors such as the program's maturity – the more mature the program, the harder it is to find new vulnerabilities, and thus the reward is higher – how well was the target tested internally before the program was launched, and more. In addition to monetary payments, the researchers are rewarded with reputation points that determine their relative rank within the platform and may help them receive invitations to work in private bounty programs, where participation is by invitation only.

Literature on Bug Bounty Programs and Platforms

Empirical work on bounty programs has examined vulnerability trends, responses by hackers, and reward structures of participating organizations. Zhao et al. [18] 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. Furthermore, monetary incentives have a significant positive correlation with the number of vulnerabilities reported. Maillart et al. [19] have analyzed a data-set of public bounty programs and found that researchers tend to switch to newly launched bounty programs at the expense of existing ones. Malladi and Subramanian [9] studied 41 public bounty programs and examined issues involved with their implementation. Algarni and Malaiya [20] used an open vulnerability database to explore the most successful researchers' careers, motivations, and methods. They concluded that individuals external to firms discovered a significant percentage of vulnerabilities and that financial reward is a primary motivation for participation, especially for Eastern European researchers. Sridhar and Ng [21] utilize a proprietary data-set to empirically show that bounty platforms are cost-effective for companies to improve their security posture and that hackers are relatively price insensitive.

The Bugcrowd Platform

Rules of Engagement

The rules of engagement between researchers and organizations on a bug bounty platform are structured to benefit both sides: they encourage researchers to practice responsible disclosure of high-value vulnerabilities and ensure the timely response and payment of organizations once a valid and unique bug is submitted. Submissions are either accepted if they are valid or rejected if they are invalid. Valid submissions can be unique or duplicate, i.e., submissions that report vulnerabilities other researchers have already registered.

The researched Bugcrowd platform offers two types of programs. Managed Bug Bounty programs (MBBs) provide a monetary reward and points to the first researcher who submits a unique, valid vulnerability. Researchers who later find the same vulnerability receive only points for their valid "duplicate" submission.[14] On the other hand, Vulnerability Disclosure Programs (VDPs) only offer points to researchers without any monetary rewards. The points researchers earn are reported in a monthly and all-time leaderboard, reflecting their ranking.[15] The "recognition" researchers receive for finding meaningful (high-priority) vulnerabilities and their demonstrated skills and interests may increase their likelihood of receiving invitations to work on private MBB programs.

Organizations pre-select the program type, considering the differences between VDPs and MBBs. VDPs are viewed as a cost-predictable, baseline security structure that encourages anyone to report any vulnerability, thus meeting compliance requirements. In contrast, MBBs incentivize discovering high-priority vulnerabilities and may be used to test more specific assets.[16]

We accounted for both MBB and VDP programs when analyzing the COVID-19 effect on supply and demand. However, we only accounted for MBB programs that reward researchers with dollar payments when referring to equilibrium price effects and their related regressions. Our results are qualitatively unchanged if we separately analyze the point rewards of VDP programs, as shown in Appendix B.

The rest of this section will detail the submission workflow and the bounty pricing dynamics of the researched platform.

Submission Workflow

The workflow for submissions over Bugcrowd's platform is outlined in Figure 1. Before starting a program, the organization must make two strategic decisions. The first decision is whether it will be an MBB or VDP. The second decision is whether it will be private or public. The program's goals and scope are then defined, including the specific software components to be tested, such as web applications, APIs, or mobile versions. Next, the organization develops a researcher engagement plan and determines the program's duration (continuous or ad-hoc). The organization also establishes the payment range for a unique submission in advance based on its priority. This information is shared with researchers through the program's portal on the platform. Thus, they know the potential rewards they could receive for a unique discovery in a particular priority and program. This actual payment represents the economic value of a submission to the organization.

Submissions are categorized according to a priority scale of P1 to P5, with P1 being critical vulnerabilities and P5 being informational weaknesses that may not even be fixed. The platform provides a well-defined Vulnerability Rating Taxonomy (VRT) for researchers to determine the priority of their submission.[17]

Once the program is launched, organizations should process incoming submissions after they have been verified, triaged

[14] In rare cases, organizations may reward researchers who submit valid duplicate entries to an MBB program to acknowledge the researcher's effort and motivate their future work. We have included these payments in our analysis, but the results are qualitatively unchanged if we exclude these "runner-up" payments.

[15] `https://bugcrowd.com/leaderboard` (lower rank is better).

[16] `https://www.bugcrowd.com/blog/vulnerability-disclosure-program-or-managed-bug-bounty-how-to-determine-which-program-is-best-for-you/`.

[17] A resource outlining Bugcrowd's baseline priority rating for common vulnerabilities: `https://bugcrowd.com/vulnerability-rating-taxonomy/`.

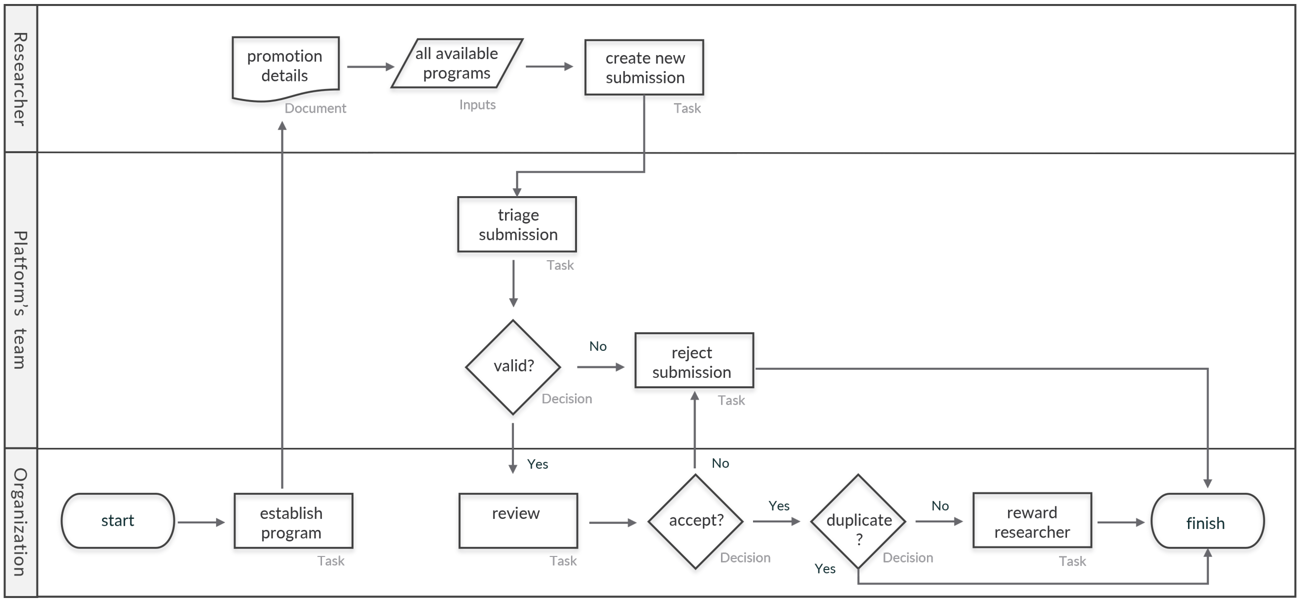

Figure 1: Bugcrowd's platform submission workflow.

(prioritized), and screened for duplicates and relevancy by the platform's team.[18] Unique, valid vulnerabilities are then integrated into the existing Software Development Life-cycle (SDLC) tools to be fixed, and related reward payouts are processed accordingly.

Data

The paper employs a unique data-set obtained through a Data Transfer Agreement (DTA) between Tel Aviv University and Bugcrowd. The data spans the entire period the company has existed, i.e., from 2012, and includes all bug submission activities through May 2021. It contains information on the programs, organizations, researchers, and bug submissions.

The data on programs covers public programs, which are open to all researchers, as well as private programs, where only invited researchers are allowed to participate and submit vulnerabilities. There are trade-offs between the two types of programs: There are likely more "eyeballs" looking for vulnerabilities in public programs, while the limited competition in private programs may lead to higher-quality researchers and more research effort. As far as we know, this paper is among the first to empirically analyze a bug bounty data-set that includes researcher attributes and detailed submissions from private programs. Table C1 shows changes in active and new private and public programs across entire calendar year periods. Table C2 shows the yearly distribution of active and new MBB and VDP programs.

The organizational data includes the firm size, country of origin, and when it first joined the platform. Many firms run more than one program simultaneously, and for each, we have its status, start/end dates, and whether it is public or private.

The data on researchers include characteristics such as country of origin, date of first submission, and relative rank (partially reflecting past success).

The data on submissions specifies, among other things, the following:

- The identification of the submitter;
- Date and time of the submission;
- The program type (MBB or VDP) and whether it is private or public;
- Whether the submission was valid, and if so, whether it was unique or a duplicate;
- The amount paid (in US dollars);
- The number of points rewarded;

To focus on the effect of the exogenous COVID-19 shock, we examine five time periods from 2017-21.[19] Each period includes three full months of activity from March 1 to May 31 of the respective year. We argue that the COVID-19 shock was the strongest during the 2020 three-month period, as indicated by multiple governmental resources. The United States Department of Transportation statistics show vehicle miles traveled were the lowest during this period.[20] Also, the United Nations statistics division report shows global manufacturing was the lowest during the first few months of the pandemic.[21] According to the US Bureau for Labor Statistics, the initial pandemic-related shock was primarily apparent as an increase in separations of employees from their payroll for any reason. It rose from 5.7 million in February 2020 to a 16.3 million historical high in March 2020, and 11.5 million (second highest) in April

[18] The process is described in `https://docs.bugcrowd.com/customers/getting-started/with-bugcrowd/`.

[19] The platform was much smaller before 2017; hence, we started our examination that year.

[20] `https://www.bts.gov/covid-19/daily-vehicle-travel`.

[21] `https://unstats.un.org/unsd/ccsa/`.

2020.[22] Figure 2 best illustrates the shock showing three-month moving averages of flows into and out of employment from 2017 through 2021. Hence, we define March 1 to May 31 of 2020 as the "COVID period" and include the same period in all other years to exclude seasonality effects.

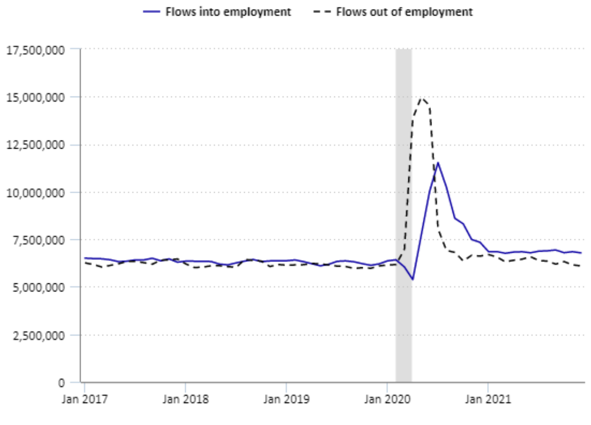

Shaded area represents a recession as determined by the National Bureau of Economic Research.
Source: U.S. Bureau of Labor Statistics.

Figure 2: Flows into and out of employment from 2017-21, three-month moving average, seasonality adjusted (US Bureau of Labor Statistics).

Heuristic Model

Using the data-set provided by Bugcrowd, we examine the effect of an exogenous external shock (COVID-19) on the Bugcrowd platform. We use a heuristic demand, supply, and equilibrium price model to motivate our analysis. First, we define the product and its price while referring to the "tournament" structure of program rewards.

We define the product as a valid vulnerability submission by a researcher to a participating program. While looking for a vulnerability, researchers are imperfectly informed about the activities of other researchers who may be searching for the same vulnerability. As a result, there may be duplicate valid submissions, only one of which earns a payment. Therefore, the product's price is the expected payment, which we define empirically as the average payment for a valid submission. The underlying simplifying assumption is that ex-ante, every valid submission has an equal chance of succeeding.

With these definitions in hand, consider the heuristic model in Figure 3.[23] The horizontal axis measures the total number of valid submissions over some period, and the vertical axis measures the average payment for a valid submission. The upward-sloping supply curve recognizes that if the price is higher, then researchers will devote more effort to searching for vulnerabilities, and more researchers will participate in more programs. The downward-sloping demand curve similarly recognizes that, under more favorable terms, companies might expand the scope of programs to cover more vulnerabilities and submit more programs to the platform. The equilibrium price brings these two sets of incentives into balance.

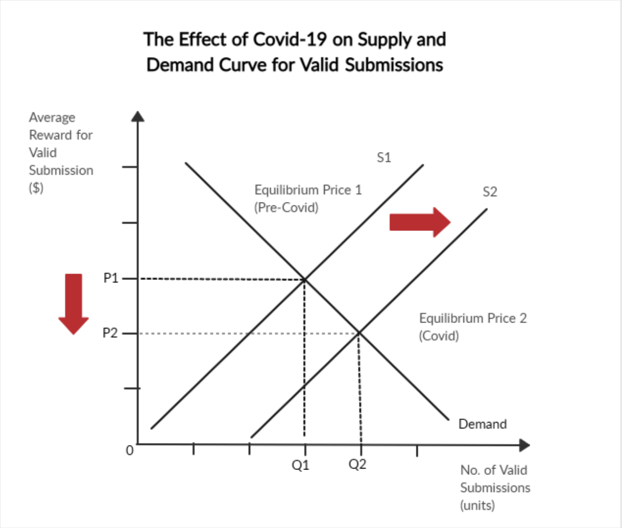
The Effect of Covid-19 on Supply and Demand Curve for Valid Submissions
Average Reward for Valid Submission ($)
S1
S2
Equilibrium Price 1 (Pre-Covid)
P1
P2
Equilibrium Price 2 (Covid)
Demand
0
Q1
Q2
No. of Valid Submissions (units)

Figure 3: Supply and demand curves for valid submissions on a bug bounty platform.

Submissions are either valid or not; valid submissions can be either unique or duplicate. A feature of our framework is that the number of valid submissions (quantity) can be decomposed into the total number of submissions multiplied by the accuracy of submissions, defined as the ratio of valid submissions to total submissions. Bugcrowd follows accuracy, as higher values indicate that the researchers are more vigilant and concise with their submissions, lowering the programs' costs associated with processing submissions. The number of discovered vulnerabilities, i.e., unique and valid submissions that are paid a monetary reward, in turn, equals the number of valid submissions times the probability of winning with a valid submission. This probability depends on the number of duplicate valid submissions and is calculated as the paid-to-valid submissions ratio.[24]

We examine how the COVID-19 shock impacted these and other relevant variables derived from the Bugcrowd data-set. Based on the heuristic model, our main conclusion is that the COVID-19 shock shifted the supply curve to the right, had a relatively minor impact on demand, and resulted in a significant decline in the average price of a valid submission, as illustrated in Figure 3.

[22] `https://www.bls.gov/blog/2022/labor-market-dynamics-during-the-covid-19-pandemic.htm`.

[23] Figure 3 also shows a rightward shift of the supply curve associated with the COVID-19 shock, as discussed later.

[24] Our primary focus is on the equilibrium effects of a shift in the supply curve, for which valid submissions are a natural quantity variable. On the other hand, *unique* and valid submissions are a more natural quantity for the demand curve because organizations have no value for duplicates. Hence, an appropriate transformation, accounting for the probability of winning, is required to re-express demand in terms of the number of valid submissions.

Table 1. Summary of propositions: The effect of COVID-19 exogenous shock across the heuristic model components.

Research Questions	Related Findings
Supply Side	
What was the COVID-19 effect on the supply of submissions and active researchers?	Proposition 1: There was a significant increase in the number of submissions and active researchers during the COVID-19 shock. Both values dropped back post-COVID.
Demand Side	
How did COVID-19 affect the demand generated by private and public active programs?	Proposition 2: The number of active private programs grew significantly during the COVID period, while public programs showed no evidence of a COVID-19 effect.
Equilibrium Price	
How did COVID-19 affect the probability of earning a monetary reward (researcher incentives)?	Proposition 3: The probability of earning a reward given a valid submission declined during the COVID period and bounced back post-COVID.
What was the COVID-19 effect on the average rewards for paid submissions?	Proposition 4: The COVID-19 effect on rewards of paid submissions was insignificant, indicating a lack of meaningful changes in the organizations' compensation strategy.
What was the COVID-19 effect on expected rewards for valid submissions?	Proposition 5: There was an economically meaningful negative effect on rewards of high-priority valid submissions during the COVID period.
What was the relative supply-side vs. demand-side effect on the expected rewards for valid submissions?	Proposition 6: The decline in the expected rewards during the COVID period is mostly associated with supply-side effects.
Productivity View	
What was the COVID-19 effect on the productivity of uniquely discovered vulnerabilities?	Proposition 7: The massive increase in submissions did not result in a significant growth in the number of unique vulnerabilities discovered.

The Effect of COVID-19 Exogenous Shock

In this section, we empirically explore the effect of the COVID-19 shock on the supply and demand sides of the market and then discuss overall equilibrium changes. Before we begin, it is important to note that there is very little (if any) market power on either side of the market. The top 100 researchers in 2020 earned only 24% of total rewards, and the top 10 programs accounted for only 33% of annual 2020 payments (Table C3). We analyzed the three-month data (March 1 - May 31) for 2017-21 to overcome any seasonality effects. Table 1 summarizes our research questions and related findings.

The Supply Side

We start with the most prominent effect of the COVID-19 shock. Figure 4 shows year-over-year (YoY) changes in the number of researchers and submissions during the three-month periods for 2017-21. While the supply side of the platform experienced steady growth from 2017 to 2019, there was a significant increase in supply-side activity during the 2020 COVID period. Specifically, total submissions increased from 21,157 in 2019 to 53,098 in 2020, representing a 151% growth. This growth rate during the COVID period is much larger than in previous periods. Moreover, in the subsequent 2021 period, there was a significant drop in the number of submissions to 30,955 submissions. In terms of the number of active researchers, there was a 72% growth in the number of active researchers between 2019 and 2020, compared with a 43% growth between 2017-18 and a 53% growth between 2018-19.

Proposition 1. *There was a significant increase in the number of submissions and active researchers during the COVID-19 shock. Both values dropped back post-COVID.*

Table C4 clearly shows that this vast increase in submissions was primarily driven by researchers who joined the platform in 2020; they made 20,118 submissions (38% of the total). By comparison, in 2019 and 2021, there were only between 6,000 and 7,000 submissions made by new researchers.

The number of valid submissions surged 155% during the COVID period, from 13,083 to 33,304. An increase that completely dwarfs the increases between 2017-18 (38%) and 2018-19 (51%). Furthermore, in 2021, the number dropped to 16,460 valid submissions (a 51% decline). Interestingly, this huge increase in valid submissions was primarily driven by researchers who joined during the COVID period and made 13,163 valid submissions. In comparison, new researchers in 2019 and 2021 made 2,952 and 2,939 valid submissions during their first year on the platform.[25]

Geographically, the most significant increase in submissions during the COVID period came from researchers in India and

[25] The average number of submissions for an active researcher increased from 7.18 in 2019 to 10.45 in the 2020 COVID period (a 45% growth), further emphasizing the supply curve shift.

Turkey. In particular, submissions from India soared from 9,335 in 2019 to 31,673 in 2020, an increase of more than 200%. Submissions from Turkey skyrocketed from just 472 in 2019 to 7,724 in 2020. A significant decline followed this surge in the subsequent period of 2021.

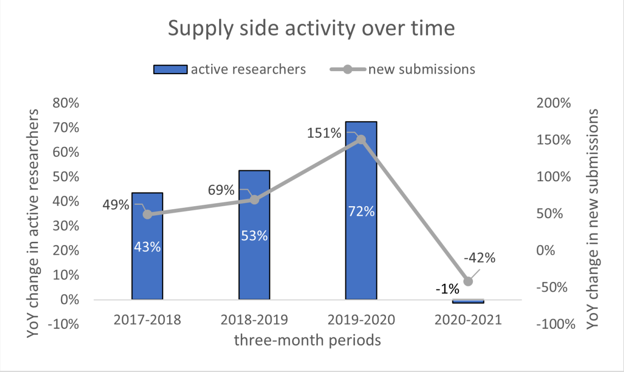

Figure 4: YoY changes in active researchers and the number of submissions, three-month periods for 2017-21.

The analysis above suggests that the COVID-19 shock attracted many more researchers to participate in bug bounty programs, likely because there was a decrease in the option of a full-time job outside. If this is at least a partial driver of the effect we observe, one would expect the activity level of the researchers who joined during the COVID period to plummet significantly in 2021, when many companies were hiring again. We, therefore, look at the number of submissions of the COVID period cohort in the post-COVID era and compare it to the number of submissions of other cohorts who joined the platform before the pandemic.

Figure 5 displays the YoY changes in the activity level of researchers on the platform (the number of submissions) based on the year they joined. The graph highlights the differences in activity levels during COVID and post-COVID (i.e., the number of submissions in 2021 compared to 2020) as well as between the first and second years of researchers' activity.

Researchers who joined the platform during the COVID period made significantly fewer submissions in 2021. Specifically, they made 89% fewer submissions in this post-COVID period. This decline was exceptionally large relative to other cohorts in their second year on the platform. For comparison, there was a 53%, 43%, and 62% drop for the 2017, 2018, and 2019 cohorts between their first and second years.

The Demand Side

Next, we examine the demand side of the platform. Table C5 presents the changes in active and new programs for various program types across the three-month periods between 2017-21. During the COVID period, there has been a rise in the number of new programs and active programs (those that received at least one submission) compared to previous years. Interestingly, the increase is not symmetric across public and private programs.

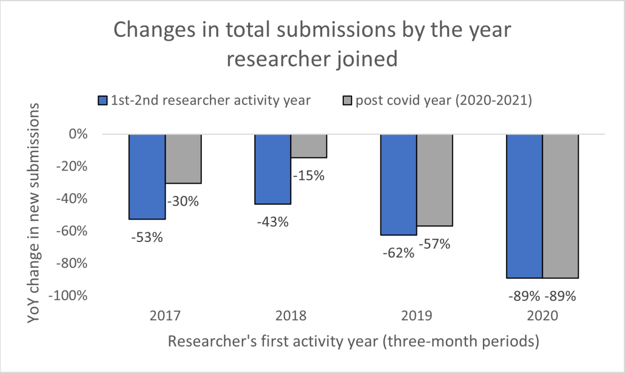

Figure 5: YoY changes in submissions for researchers who joined the platform during the three-month periods of 2017-20. Changes were measured between the COVID and post-COVID periods and the first and second years of joining.

Proposition 2. *The number of active private programs grew significantly during the COVID period, while public programs showed no evidence of a COVID-19 effect.*

Table C5 shows an overall upward trend in the total number of active private programs, averaging 63% a year before the COVID period. Between the 2019 and 2020 periods, however, there was a 91% increase in active private programs, followed by a moderate rise of 31% in the following year. In contrast, public programs had a more moderate overall upward trend but showed no evidence of a COVID effect in 2020. Given that most new researchers did not initially have access to private programs, the supply increase greatly exceeded the demand.

Another possible avenue for demand-side effects is a change in rewards programs pay for discovered vulnerabilities. The Bugcrowd platform recommends ranges for payments rewarding different priority submissions. Furthermore, it encourages established programs, for which new vulnerabilities may be more challenging to discover, pay at the higher end of the range,[26] and add bonuses outside the range for particularly significant vulnerabilities.[27] However, programs may decide what exact reward to pay within each range and activate some judgment about classifying the priority of submissions. Given such discretion, the average reward per paid submission might be viewed as a demand-side variable, possibly shifting during the COVID period. We defer a discussion of this possibility to the next section.

Equilibrium Outcome

The discussion above suggests that the COVID-19 shock shifted the supply curve to the right – dramatically increasing the number of researchers on the platform. On the demand side, there was a moderate increase in active programs. The overall effect in equilibrium outcomes may manifest itself as a decrease in the probability of winning a reward, a reduction of actual rewards

[26] https://www.bugcrowd.com/resources/guide/bugcrowds-defensive-vulnerability-pricing-model/.

[27] https://docs.bugcrowd.com/customers/submission-management/rewarding/.

paid by programs, or a combination of both. We will discuss each in turn and then look at the combined effect referring to MBB programs that compensate researchers with cash rewards.

Probability of Earning a Monetary Reward

Given the tournament structure of the market, an increase in the number of valid submissions implies an increase in competition for being the first to find a vulnerability and thus may decrease the probability of winning a reward. Consequently, this may negatively affect the researcher's motivation to dedicate research time and submit their findings.

Indeed, Figure 6 shows that the COVID-19 shock dramatically affected the probability of winning a reward, represented by the ratio of paid submissions to valid submissions. This ratio ranged between 37-39 percent in 2017 and 2018 and 33-30 percent in 2019 and 2021, reflecting a slightly increased competition among researchers relative to previous years. However, the ratio fell significantly to 23% in 2020. It seemed the shock essentially "threw" the platform out of the equilibrium it had maintained for most of its existence.

Proposition 3. *The probability of earning a reward given a valid submission declined during the COVID period and bounced back post-COVID.*

This increased competition is observed for both public and private programs but was especially dramatic in the case of public programs, where the probability of winning a reward given a valid submission was 18% in 2019 and 2021 but fell to only 11% in 2020.[28]

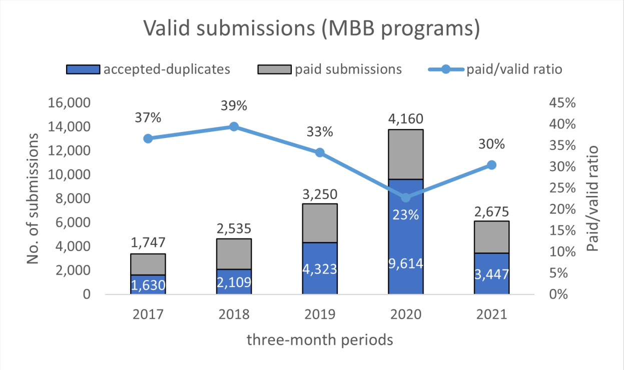

Figure 6: Valid submissions during the three-month periods of 2017-21. The paid-to-valid ratio indicates the probability of earning a reward given a valid submission.

The COVID Effect on Submission Rewards

We further examine the increase in competition during COVID-19 by looking at the actual submission payments. The first regression in Table 2 captures the effect and significance of the COVID-19 shock on actual rewards by examining the interaction of the COVID variable with submission priority for paid submissions. The COVID dummy variable equals 1 for 2020 and 0 for all other three-month periods. The P1-P5 priority definitions are by Bugcrowd and reflect the risk associated with the discovered vulnerability.[29] The coefficient of the COVID interaction with submission priority is statistically insignificant for all submission priorities.

Proposition 4. *The COVID effect on rewards of paid submissions was insignificant, indicating a lack of meaningful changes in the organizations' compensation strategy.*

Given that the probability of earning a reward for a valid submission declined during the COVID period (Proposition 3), we now look at the expected rewards for valid submissions. The second regression in Table 2 measures the effect and significance of the COVID-19 shock on those rewards by interacting the variable of interest, COVID, with submission priority. As expected, compared to the lowest priority P5 baseline, payments increase with the submission priority. The COVID main effect for P1 valid submissions is a statistically significant average reduction of $640. The effect diminishes as submission priority declines: $216 for P2, $36 for P3, and statistically insignificant for P4 and P5.[30]

Proposition 5. *There was an economically meaningful negative effect on rewards of high-priority valid submissions during the COVID period.*

Our results suggest that the COVID effect on expected rewards of valid submissions is present mainly in high-priority submissions, emphasizing the missed opportunity to benefit the organizations and the public with essential vulnerability fixes. Furthermore, it highlights the need for the platform to control the supply-demand equilibrium and maintain incentives for the leading researchers to participate in the future.

As detailed in Appendix A, the results we report are robust if we replace the selected COVID dummy with the continuous government pandemic "Stringency Index" offered by Hale et al. [22]. Furthermore, our findings are robust under different periods, accounting for all the submissions between January 2017 and May 2021 (the end of our data-set).

The Main Driving Factor

In general, the statistically significant and meaningful decline in expected rewards shown in The COVID Effect on Submission Rewards subsection can be driven by (i) an increase in the number of valid submissions (a supply-side effect), (ii) a decrease in monetary rewards for paid submissions (both demand and supply side effects), or a combination of both. More concretely,

[28] In the case of private programs, the probability of earning a reward given a valid submission fell from 53% in 2019 to 37% in 2020.

[29] `https://www.bugcrowd.com/glossary/vulnerability-priority/`.

[30] No priority was listed for a small number of rewarded submissions. Unlike our descriptive statistics in The Main Driving Factor section, the regressions do not include submissions with a missing priority field. Regardless, all our results remain qualitatively unchanged.

Table 2. COVID-19 effects regression table across the three-month periods. The COVID-19 shock is defined as the year 2020 period (a binary variable).

	(1)		(2)	
Priority=1	2628.066***	(139.257)	1669.594***	(22.737)
Priority=2	1255.319***	(135.640)	823.738***	(16.938)
Priority=3	285.728**	(133.953)	212.481***	(13.054)
Priority=4	5.875	(134.472)	48.069***	(12.580)
COVID period	-58.784	(214.659)	-0.869	(15.247)
Priority=1 × COVID period	-212.821	(230.430)	-640.963***	(38.111)
Priority=2 × COVID period	-148.791	(221.447)	-216.703***	(28.695)
Priority=3 × COVID period	95.445	(218.237)	-36.693*	(21.632)
Priority=4 × COVID period	69.775	(218.888)	-18.904	(19.596)
Constant	154.925	(132.325)	2.133	(9.532)
Observations	13302		46031	
Adjusted R-squared	0.256		0.175	

Regression (1): COVID-19 effect on actual rewards (paid submissions) interacted with P1-P5 submission priority.

Regression (2): COVID-19 effect on expected rewards (valid submissions) interacted with P1-P5 submission priority.

Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

the average reward for valid submissions can be decomposed as follows:

Equation 1. *(average reward for valid submissions) = (average reward for paid submissions) * (paid submissions/valid submissions)*

A combination of supply and demand factors most likely drives changes in the first term. Changes in the second term (the ratio) arguably are due exclusively to the supply side, i.e., the number of unique vulnerabilities discovered and the number of valid duplicates.[31] To populate the equation, we empirically measure below the average reward for paid and valid submissions and use the paid-to-valid submission ratio already shown in Figure 6. Since rewards vary substantially across the importance (priority) of vulnerabilities, we delineated our data by vulnerability priority, grouping P1 and P2 vulnerabilities into a "higher-priority" category and P3-P5 vulnerabilities into a "lower-priority" category.

As Figure 7 shows, the average reward for valid submissions ranged between $256 and $340 during the 2017-19 period but fell dramatically to just $166 in 2020 (a decline of 46% relative to 2019). The average reward then rose to $246 in 2021. The trend for higher-priority and lower-priority vulnerabilities was similar but more prominent for high-priority vulnerabilities (see the second regression of Table 2).

We present the average monetary reward for paid submissions in Figure 8. As the figure shows, the average reward decreased from $989 in 2019 to $744 in 2020 (a decline of 21%). However, as shown in the first regression of Table 2, this decline was not statistically significant.

During the COVID period, the paid-to-valid submissions ratio fell significantly from 33% to 23% (a decline of 32%), while the monetary rewards for paid submissions decreased by only 21%). Therefore, we conclude that the decline in the expected reward between 2019 and 2020 associated solely with supply-side effects was 52% larger than the remaining residual decline, attributed to a combination of supply and demand factors.[32]

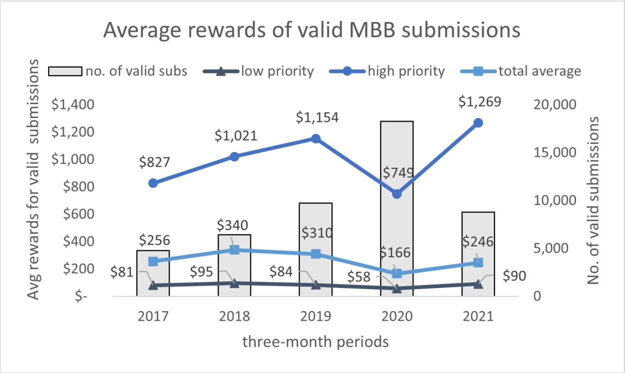

Figure 7: Expected average rewards of MBB submissions by priority during the three-month periods of 2017-21 in Bugcrowd's bug bounty platform.

Proposition 6. *The decline in the expected rewards during the COVID period is mostly associated with supply-side effects.*

The demand and supply responses differ across higher and lower-priority vulnerabilities. Specifically, while there was a meaningful 23% drop in the average actual payment for higher-priority paid submissions, rewards for lower-priority submissions slightly increased (Figure 8). This strongly suggests that a shift in the supply curve almost completely drove the decrease in expected reward for valid lower-priority submissions. For higher-priority vulnerabilities, the equilibrium outcome is likely a result of combined supply and demand responses.

[31] We cannot think of demand-side factors that would cause a *decline* in the paid-to-valid ratio.

[32] In Appendix B, we examine the VDP program type separately and show that our results are qualitatively unchanged.

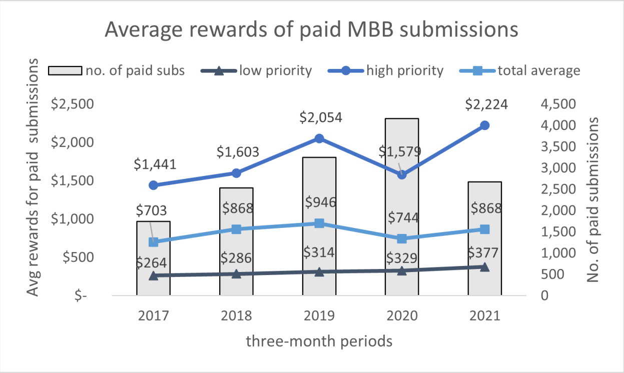

Figure 8: Actual average rewards of MBB submissions by priority during the three-month periods of 2017-21 in Bugcrowd's bug bounty platform.

In addition, as Table C6 shows, the demand and supply responses differ across public and private programs. For private programs, the average actual payment dropped 26% between 2019-20, while for public programs, it dropped only 10%, indicating the effect associated partially with the demand side for public programs is significantly smaller. This aligns with our demand-side findings, which show that active private programs exhibited a larger growth than public programs during the COVID period. Furthermore, the paid-to-valid submission ratio for public programs declined by 40% compared to a 29% decline for private programs. Given that most new researchers who joined the platform during the COVID period initially only had access to public programs, the increase in supply for public programs greatly exceeded the rise in demand.

Unique Vulnerabilities Discovered

While the number of valid submissions increased tremendously during the COVID period, there was no significant increase in the number of paid submissions between 2019 and 2020 (Figure 6).[33] Since the number of paid submissions in MBB programs is essentially the number of unique vulnerabilities discovered, it reflects the increase in the number of unique vulnerabilities discovered. The percentage increase in the number of paid vulnerabilities was essentially the same between 2018 and 2019 as the percentage increase between 2019 and 2020 (28%). This was primarily because, during the COVID period, more researchers were competing over the same vulnerabilities, especially in public programs. Although the number of valid submissions to public programs skyrocketed from 5,558 in 2019 to 10,226 in 2020, the number of unique vulnerabilities discovered (and hence paid for) was virtually unchanged (1,021 in 2019 vs. 1,136 in 2020.)

Proposition 7. *The massive increase in the number of submissions did not result in a significant growth in the number of unique vulnerabilities discovered.*

Policy and Platform Implications

The exogenous shock we examined was significant and led to a huge response on the supply side and a much smaller one on the demand side. As a consequence, the equilibrium price fell dramatically during the COVID period. One would expect other bug bounty platforms to experience similar shifts in demand, supply, and equilibrium price. Moreover, we believe the results do not depend on the specific rules of engagement and tournament structure that bug bounty programs exhibit. The COVID-19 shock size is dramatic and does not occur often. Still, there are examples of significant shocks that dramatically changed the equilibrium of an existing ecosystem. For instance, consider the entry of ride-sharing services (like Uber and Lyft) into the previously tightly regulated market, primarily limited to a fixed number of taxis. Similar to our case, there was a huge increase in supply, which led to a fall in the equilibrium price of ride services [23, 24]. The change in the ride-sharing platform was permanent – and led to a new equilibrium that persisted over time. The exogenous shock in the platform we studied was temporary, and while some of the effects remained once the COVID pandemic was over, gradually, over time, supply fell, demand increased, and expected prices rose.

Regarding the market for vulnerabilities, past research has suggested that the grey market for sharing exploits and vulnerabilities is more lucrative than the black market, and both are distinctly more lucrative than the white market [25]. The significant supply response we identified from the COVID-19 shock implies that more bug bounty programs and larger platforms might affect this dynamic. In particular, increased demand for researchers and vulnerability submissions would likely drive more transactions in the white market rather than the black or grey markets.

The vast number of valid submissions in the 2020 COVID period suggests that had there been a more significant increase in demand, many more unique vulnerabilities would have been likely discovered. If the demand response to the shock in both MBB and VDP program types had been similar to the supply response, the number of unique vulnerabilities identified by researchers in 2020 could have been 64% higher than the actual number.[34] This "counterfactual" illustrated in Figure 9 seems reasonable since all software contains vulnerabilities, and "there will always be more vulnerabilities in a given piece of software" [26, p. 17]. This assumption also implies that some level of duplication is desirable. However, the relatively low ratio of unique to valid submissions during the COVID period might discourage researchers from searching hard for vulnerabilities.

From a gig economy perspective, platforms like Uber, UpWork, and Bugcrowd offer tremendous opportunities for employees during turbulent times. Indeed, our analysis suggests that many researchers had reached out to bug bounty platforms during the COVID period when the market experienced a massive decline in job opportunities. The flexibility and low barriers to entry of gig work platforms allowed skilled researchers to supplement earnings by finding vulnerabilities.

From a policy perspective, the limited demand response in terms of an increase in the number of programs may be seen

[33] The number rose from 2,535 in 2018 to 3,250 in 2019 and 4,160 in 2020, before falling back to 2,675 in 2021.

[34] As shown in Table C7, given a similar response, the ratio of unique-to-valid submissions would have been 57% (the average for 2019 and 2021) rather than falling to 35%.

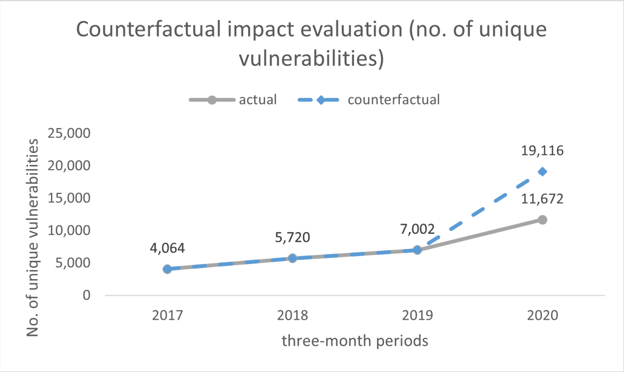

Figure 9: Counterfactual impact evaluation for the number of unique vulnerabilities discovered in the 2020 COVID period.

as a missed opportunity. A policy encouraging the quick launch of new programs would likely have benefited researchers and organizations. Discovering and patching vulnerabilities in products and services also creates a positive external benefit (beyond the platform) because a compromised asset in one firm, due to a vulnerability, might cause a supply chain attack that affects other organizations and individuals elsewhere. As Moore [27, p. 107] writes, "Insecurity creates negative externalities." See also Varian [28] and Jean Camp and Wolfram [29], who were among the first to make this point.

Policy-makers could sponsor, encourage, or mandate bug bounty programs in governmental agencies. The platform could incentivize firms to introduce new programs and offer the means and underlying technology to do it quickly. In addition, organizations could widen the scope of existing programs to include more product/service areas, thus increasing the possible number of vulnerabilities to discover.

The platform could also help with information asymmetry caused by researchers not knowing where their peers are active and which programs are overwhelmed with submissions due to the tournament structure of bounty programs. If the platform signals researchers how busy a program is, it might help rebalance supply with demand. The signal could indicate, for example, that a program has a low paid-to-valid ratio, which means a high percentage of duplicates. Finally, researchers may cooperate to reduce duplicates and work in teams, utilizing collaboration options on the platform.[35] Taken together, these policies might even encourage some black-hat researchers to join the program and work legally.

Acknowledgments

We are very grateful to the Area Editor and two anonymous referees whose many comments and suggestions significantly improved the paper. We thank Casey Ellis, the founder of Bugcrowd, and his team for sharing their raw data and adding context to our paper. We also thank seminar participants at the EIEF, Reichman University, Tel-Aviv University, and the 2022 IIOC for their helpful comments and suggestions.

Funding

We thank "The Foerder Institute For Economic Research" and "The Sapir Center for Development" at Tel Aviv University for research funding.

References

1. Eric Raymond. The cathedral and the bazaar. *Knowledge, Technology & Policy*, 12(3):23–49, 1999. ISSN 1874-6314. doi: 10.1007/s12130-999-1026-0. URL `https://doi.org/10.1007/s12130-999-1026-0`.
2. Dmitri K Koustas. Consumption Insurance and Multiple Jobs: Evidence from Rideshare Drivers. 2018. URL `https://uchicago.app.box.com/v/DKoustas-RideSmoothing-WP`.
3. B Collins, A Garin, E Jackson, D Koustas, and ... Is gig work replacing traditional employment? Evidence from two decades of tax returns. 2019. URL `https://www.irs.gov/pub/irs-soi/19rpgigworkreplacingtraditionalemployment.pdf`.
4. Ross Anderson. Why information security is hard - An economic perspective. *Proceedings - Annual Computer Security Applications Conference, ACSAC*, 2001-January:358–365, 2001. ISSN 10639527. doi: 10.1109/ACSAC.2001.991552.
5. Harjinder Singh Lallie, Lynsay A. Shepherd, Jason R.C. Nurse, Arnau Erola, Gregory Epiphaniou, Carsten Maple, and Xavier Bellekens. Cyber security in the age of COVID-19: A timeline and analysis of cyber-crime and cyber-attacks during the pandemic. *Computers and Security*, 105:102248, 6 2021. ISSN 01674048. doi: 10.1016/j.cose.2021.102248.
6. Cormac Herley. So long, and no thanks for the externalities: The rational rejection of security advice by users. *Proceedings New Security Paradigms Workshop*, pages 133–144, 2009. doi: 10.1145/1719030.1719050. URL `https://dl.acm.org/doi/10.1145/1719030.1719050`.
7. Jay Pil Choi, Chaim Fershtman, and Neil Gandal. Network security: Vulnerabilities and disclosure policy. *Journal of Industrial Economics*, 58(4):868–894, 2010. ISSN 00221821. doi: 10.1111/j.1467-6451.2010.00435.x.
8. Lillian Ablon and Andy Bogart. Zero Days, Thousands of Nights: The Life and Times of Zero-Day Vulnerabilities and Their Exploits. Technical report, RAND Corporation, 2017. URL `https://www.rand.org/pubs/research_reports/RR1751.html`.
9. Suresh S. Malladi and Hemang C. Subramanian. Bug Bounty Programs for Cybersecurity: Practices, Issues, and Recommendations. *IEEE Software*, 37(1):31–39, 2020. ISSN 19374194. doi: 10.1109/MS.2018.2880508.
10. Jean Charles Rochet and Jean Tirole. Two-sided markets: A progress report. In *RAND Journal of Economics*, volume 37, 2006. doi: 10.1111/j.1756-2171.2006.tb00036.x.
11. Daniel G. Arce. Cybersecurity and platform competition in the cloud. *Computers & Security*, 93:101774, 6 2020. ISSN 0167-4048. doi: 10.1016/J.COSE.2020.101774. URL `https://linkinghub.elsevier.com/retrieve/pii/S0167404820300584`.
12. Feng Zhu and Marco Iansiti. Entry into platform-based markets. *Strategic Management Journal*, 33(1), 2012. ISSN

[35] See, for example, the researchers' collaboration options offered by Bugcrowd: `https://docs.bugcrowd.com/researchers/reporting-managing-submissions/reporting-a-bug/submission-collaboration/`.

01432095. doi: 10.1002/smj.941.

13. Jean Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 2003. ISSN 15424766. doi: 10.1162/154247603322493212.
14. Geoffrey G. Parker and Marshall W. Van Alstyne. Two-Sided Network Effects: A Theory of Information Product Design. *https://doi.org/10.1287/mnsc.1050.0400*, 51(10): 1494–1504, 10 2005. ISSN 0025-1909. doi: 10.1287/MNSC.1050.0400. URL `https://pubsonline.informs.org/doi/abs/10.1287/mnsc.1050.0400`.
15. Bernard Caillaud and Bruno Jullien. Chicken & Egg: Competition among Intermediation Service Providers. *The RAND Journal of Economics*, 34(2), 2003. ISSN 07416261. doi: 10.2307/1593720.
16. Paul Belleflamme and Martin Peitz. Platform competition: Who benefits from multihoming? *International Journal of Industrial Organization*, 64, 2019. ISSN 01677187. doi: 10.1016/j.ijindorg.2018.03.014.
17. Michael Munger. Coase and the sharing economy. In *Forever contemporary: the economics of Ronald Coase*, pages 187–208. Institute for Economic Affairs London, 2015.
18. Mingyi Zhao, Jens Grossklags, and Peng Liu. An empirical study of web vulnerability discovery ecosystems. In *Proceedings of the ACM Conference on Computer and Communications Security*, volume 2015-Octob, pages 1105–1117. Association for Computing Machinery, 10 2015. ISBN 9781450338325. doi: 10.1145/2810103.2813704.
19. Thomas Maillart, Mingyi Zhao, Jens Grossklags, and John Chuang. Given enough eyeballs, all bugs are shallow? Revisiting Eric Raymond with bug bounty programs. *Journal of Cybersecurity*, 3(2):81–90, 2017. ISSN 20572093. doi: 10.1093/cybsec/tyx008. URL `https://doi.org/10.1093/cybsec/tyx008`.
20. Abdullah M Algarni and Yashwant K Malaiya. Most Successful Vulnerability Discoverers: Motivation and Methods. *Proceedings of the International Conference on Security and Management (SAM)*, page 1, 2014. URL `https://search.proquest.com/docview/1524243342`.
21. Kiran Sridhar and Ming Ng. Hacking for good: Leveraging HackerOne data to develop an economic model of Bug Bounties. *Journal of Cybersecurity*, 7(1), 12 2021. ISSN 2057-2085. doi: 10.1093/CYBSEC/TYAB007. URL `https://academic.oup.com/cybersecurity/article/7/1/tyab007/6168453`.
22. Thomas Hale, Noam Angrist, Rafael Goldszmidt, Beatriz Kira, Anna Petherick, Toby Phillips, Samuel Webster, Emily Cameron-Blake, Laura Hallas, Saptarshi Majumdar, and Helen Tatlow. A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour 2021 5:4*, 5(4):529–538, 3 2021. ISSN 2397-3374. doi: 10.1038/s41562-021-01079-8. URL `https://www.nature.com/articles/s41562-021-01079-8`.
23. Jeremy Horpedahl. Ideology Über Alles? Economics bloggers on Uber, Lyft, and other transportation network companies. *Econ Journal Watch*, 12(3), 2015.
24. Farshad Kooti, Nemanja Djuric, Mihajlo Grbovic, Vladan Radosavljevic, Luca Maria Aiello, and Kristina Lerman. Analyzing uber's ride-sharing economy. *26th International World Wide Web Conference 2017, WWW 2017 Companion*, pages 574–582, 2017. doi: 10.1145/3041021.3054194. URL `https://dl.acm.org/doi/10.1145/3041021.3054194`.
25. Lillian Ablon and Martin Libicki. Hackers' Bazaar: The Markets for Cybercrime Tools and Stolen Data. *Defense Counsel Journal*, 82(2):143–152, 2015. ISSN 0895-0016. doi: 10.12690/0161-8202-82.2.143.
26. Jonathan M Spring. An Analysis of How Many Undiscovered Vulnerabilities Remain in Information Systems. Technical report, CARNEGIE-MELLON UNIV PITTSBURGH PA, 2022. URL `https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=875310`.
27. Tyler Moore. The economics of cybersecurity: Principles and policy options. *International Journal of Critical Infrastructure Protection*, 3(3-4):103–117, 2010. ISSN 18745482. doi: 10.1016/j.ijcip.2010.10.002. URL `https://www.sciencedirect.com/science/article/pii/S1874548210000429`.
28. Hal Varian. System Reliability and Free Riding. In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, pages 1–15. Springer US, Boston, MA, 2004. ISBN 978-1-4020-8090-6. doi: 10.1007/1-4020-8090-5{_}1. URL `https://doi.org/10.1007/1-4020-8090-5_1`.
29. L Jean Camp and Catherine Wolfram. Pricing security: A market in vulnerabilities. In *Economics of information security*, pages 17–34. Springer, 2004.

Appendices

COVID Effect Robustness Checks

We test for robustness by replacing the COVID binary variable with the government pandemic "Stringency Index" by Hale et al. [22]. This index, listed per day and by country as part of The Oxford COVID-19 Government Response Tracker (OxCGRT), "records the strictness of 'lockdown style' policies that primarily restrict people's behaviour."[36] The index is an integer value between 0 and 100, with 0 indicating the pre-COVID era and 100 the strictest possible governmental policies. We set the Stringency Index based on the submission date and the researcher's country of origin.

In the first regression of Table A1, we interact the index with priority across the five periods from 2017-21. The priority coefficient increases consistency for higher priority (P5 lowest priority being the baseline), reflecting the platform's official payment structure favoring higher priority submissions. The Stringency Index's main effect for P1 valid submissions is an average reduction of $4.6 in payments for each index point increase. The effect for P2 submission declines by $1 on average for each index point increase. For the lower P3-P5 submissions, the effect of the Stringency Index is insignificant.

Figure A1 illustrates this effect based on the average index value across the 2017-19 periods (pre-COVID era where the index was zero), the COVID 2020 period (with an average index of 77.15), and the 2021 period (with an average index of 67.51). Hence, we interpret the COVID average effect for P1 submissions as a decline of $354 in dollar payments for the 2020 period and a $310 decrease for the 2021 period. For P2 submissions,

[36] `https://www.bsg.ox.ac.uk/research/covid-19-government-response-tracker`.

Table A1. COVID-19 effects regression table. COVID-19 is defined by the OxCGRT Stringency Index (a continuous variable).

	(1)		(2)	
Priority=1	1595.091***	(24.959)	1615.731***	(11.104)
Priority=2	787.695***	(19.356)	645.992***	(8.705)
Priority=3	206.623***	(15.436)	207.423***	(6.924)
Priority=4	47.560***	(14.711)	44.231***	(6.527)
Stringency Index	-0.020	(0.186)	-0.030	(0.101)
Priority=1 × Stringency Index	-4.569***	(0.475)	-2.179***	(0.267)
Priority=2 × Stringency Index	-1.097***	(0.347)	-0.406**	(0.189)
Priority=3 × Stringency Index	-0.190	(0.261)	-0.080	(0.143)
Priority=4 × Stringency Index	-0.179	(0.238)	-0.142	(0.128)
Constant	2.709	(11.521)	3.820	(5.164)
Observations	45774		152999	
Adjusted R-squared	0.169		0.204	

Regression (1): Stringency Index interacted with P1-P5 submission priority across the three-month periods.

Regression (2): Stringency Index interacted with P1-P5 priority across all submissions between January 2017 through May 2021 (the end of our data-set).

Standard errors in parentheses. * p <0.10, ** p <0.05, *** p <0.01.

the decline due to the COVID effect is $86 for 2020 and $75 for the 2021 period.

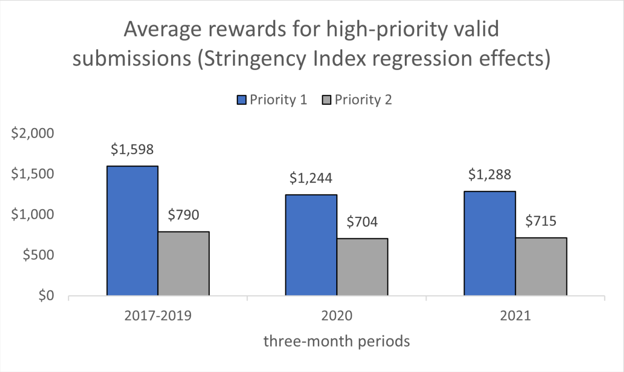

Figure A1: Table A1 regression (1) effects: average rewards for valid submissions, based on the average Stringency Index for the three-month periods (pre-COVID 2017-19, COVID 2020, and post-COVID 2021.

For further robustness, we look at all valid submissions during the entire period between January 2017 and May 2021 rather than just the three-month submissions. The results, detailed in Table A1, are qualitatively unchanged. As before, we interact the Stringency Index with submission priority to explore whether the COVID effect differs. The index's main effect for P1 valid submissions is an average reduction of $2.2 in payments for each index point increase. The average index value across all of 2020 is 64.47, including the pre-COVID months of that year. The average index value for the 2021 submissions is higher (67.31) partially because our data-set ends in June 2021 and does not reflect all the pandemic wave cycles. Hence, we interpret the COVID average effect as a decline of $142 and $149 in dollar payments for the P1 valid submissions of 2020 and 2021, respectively. The effect is lower for P2 priority submissions (a decline of $28 and $29 in payments for 2020 and 2021, respectively) and insignificant for P3-P5 priorities.

Note that the COVID effect is higher in regression (2) of Table 2 when measured by the COVID dummy across the five three-month periods, compared to Table A1 regression (1), which uses the Stringency Index for the same periods. We argue that while the index reflects the governmental policies and responses, it does not capture the private market response nor the labor market dynamics reflected in Figure 2. Furthermore, the COVID effect in Table A1 is lower when measured across all submissions in regression (2) compared to the three-month periods in regression (1). This difference illustrates the dynamics of the 2020 shock period compared to the ongoing COVID situation: Though the average Stringency Index is nearly identical, the effect is higher during the three-month shock period. Subject to their business sector and industry, many companies may have responded to the initial shock by adopting technologies such as remote access and video conferencing. We may assume many security researchers who generate the supply on the bug bounty platform work for those tech-savvy companies. Therefore, limiting the analysis to the first three months of the pandemic will better reflect the shock in this context.

Separating Monetary and Point Rewards

Below, we complete the analysis of the two different program types: (i) Managed Bug Bounty Programs (MBBs) as analyzed in the Equilibrium Outcome section, and (ii) Vulnerability Disclosure Programs (VDPs) analyzed below.

The results for VDPs are particularly interesting as point rewards bear no direct monetary costs to organizations. The number of points rewarded for a unique discovered vulnerability is predefined per priority by the platform. VDPs also reward valid duplicates but with fewer points.[37] However, Bugcrowd removed the point rewards for low-priority duplicate

[37] `https://docs.bugcrowd.com/researchers/receiving-rewards/getting-rewarded/`.

submissions between the COVID 2020 and post-COVID 2021 periods to prevent “point harvesting,” which yields no value to programs.[38]

Figure B1 shows that the average points rewarded for unique VDP submissions are nearly identical between 2019 and 2021, reflecting the fixed-point reward system and lack of demand-side effect. However, average points for valid submissions fell from eight points during the 2017-19 periods to just four points in 2020 before bouncing back to 13 points in 2021.

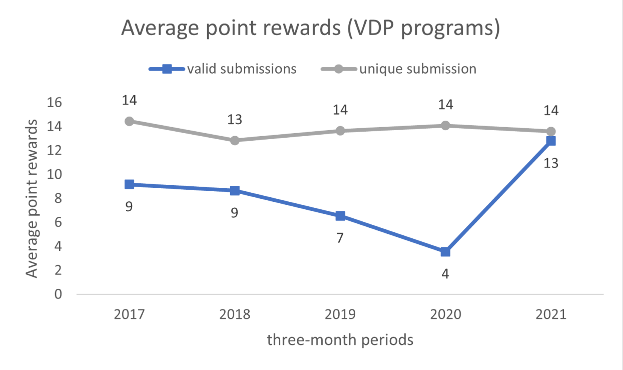

Figure B1: Average rewards for unique and valid submissions, for VDPs only, during the three-month periods of 2017-21 in Bugcrowd’s bug bounty platform.

The 50% points decline between 2019 and 2020 could be attributed to two supply-side factors. The first and more prominent is the 58% drop in the ratio of unique-to-valid VDP submissions (see Figure B2).[39] The second is a slightly higher percentage of low-priority valid submissions in the COVID period compared to 2019. The average point increase in 2021 was presumably affected by the increasing unique-to-valid submission ratio that bounced back to 62%, but also by the change in the points reward policy, which eliminated the rewards for low-priority duplicate submissions.

We summarize by noting that, given that organizations could not alter the number of points rewarded in VDPs, the equilibrium price change during the COVID period is attributed solely to supply-side factors.

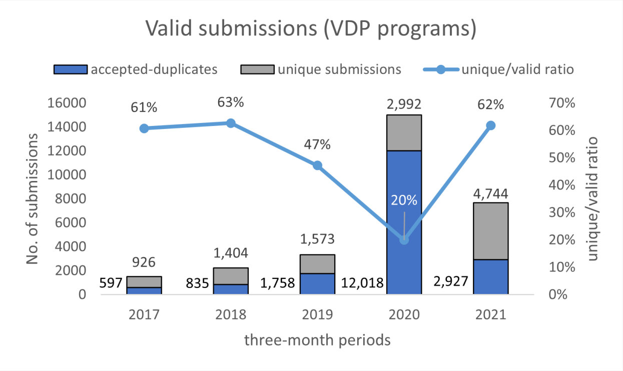

Figure B2: Valid submissions for VDP programs only (unique and duplicate) during the three-month periods of 2017-21.

[38] For more information on this policy change, see `https://www.bugcrowd.com/blog/update-to-bugcrowd-points-system/`. After the end date of our data-set, Bugcrowd completely disabled points on VDPs, as noted in `https://www.bugcrowd.com/blog/how-bugcrowd-sees-vulnerability-disclosure-programs-and-points/`.

[39] The unique-to-valid ratio indicates the probability of earning more points due to a unique vulnerability discovered.

Supporting Tables

Table C1. Program perspective: changes in active and new program types (private and public) across entire calendar year periods (all programs).

Program type	No. of active programs[1]				No. of new programs[2]			
	2017	2018	2019	2020	2017	2018	2019	2020
Public	105	132	145	151	35	33	33	25
Private	252	357	525	1,155	189	243	336	821
Total	357	489	670	1,306	224	276	369	846

[1] Programs with one or more submissions during the period.

[2] Programs that their first submission occurred during the period.

Table C2. Program perspective: changes in active and new program types (MBB and VDP) across entire calendar year periods (all programs).

Program type	No. of active programs[1]				No. of new programs[2]			
	2017	2018	2019	2020	2017	2018	2019	2020
VDP	110	120	158	607	80	74	93	490
MBB	247	369	512	699	144	202	276	356
Total	357	489	670	1306	224	276	369	846

[1] Programs with one or more submissions during the period.

[2] Programs that their first submission occurred during the period.

Table C3. Market perspective: Percentage of total rewards for top 100 researchers and top 10 programs across entire calendar year periods.

Researcher cohort	Percentage of total rewards[1]				Program cohort	Percentage of total rewards[2]			
	2017	2018	2019	2020		2017	2018	2019	2020
Top 100 researchers	43%	34%	23%	24%	Top 10 programs	42%	39%	42%	33%
Other researchers	57%	66%	77%	76%	Other programs	58%	61%	58%	67%

[1] A yearly view on total rewards share of top 100 researchers (supply side).

[2] A yearly view on total rewards share of top 10 programs (demand side).

Table C4. Researcher experience: submissions and rewards by the year a researcher joined, across the three-month periods (all programs).

Year researcher joined[1]	No. of submissions					Percentage of cohort submissions from yearly total				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
2017 researchers	1,630	770	695	614	427	19%	6%	3%	1%	1%
2018 researchers		3,182	1,806	1,270	1,085		25%	9%	2%	4%
2019 researchers			6,142	2,308	995			29%	4%	3%
2020 researchers				20,118	2,177				38%	7%
2021 researchers					6,832					22%
Year researcher joined	**No. of valid submissions[2]**					**Valid submissions / total submissions ratio[3]**				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
2017 researchers	917	500	476	389	281	56%	65%	68%	63%	66%
2018 researchers		1,857	1,118	851	687		35%	47%	54%	63%
2019 researchers			2,952	1,415	595			23%	26%	60%
2020 researchers				13,163	1,092				65%	50%
2021 researchers					2,939					43%

[1] Only researchers who joined during the three-month periods are listed in this view.

[2] Valid submissions are the sum of paid and accepted-duplicate submissions.

[3] The valid-to-total submissions ratio reflects the accuracy level of submissions.

Table C5. Program perspective: changes in active and new programs across the three-month periods (all programs).

Public programs	No. of active public programs					Percentage from total active public programs				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
Active from past period[1]	23	27	24	31	35	28%	25%	20%	24%	24%
New between periods[2]	53	73	85	95	107	64%	68%	72%	73%	72%
New this period[3]	7	7	9	5	6	8%	7%	8%	4%	4%
Total public active	83	107	118	131	148	100%	100%	100%	100%	100%
Private programs	**No. of active private programs**					**Percentage from total active private programs**				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
Active from past period	11	18	33	60	91	10%	11%	12%	12%	13%
New between periods	56	84	152	271	412	51%	50%	56%	52%	61%
New this period	43	66	86	186	174	39%	39%	32%	36%	26%
Total private active	110	168	271	517	677	100%	100%	100%	100%	100%

[1] Programs that were active in the previous three-month period, with one or more submissions during the current period.

[2] Programs that started between the previous three-month and current periods.

[3] New programs with first submission during the period.

Table C6. Program perspective: Changes in public-private programs across the three-month periods (MBB programs only).

Program Type	Average rewards for paid submissions					Paid/valid submission ratio[1]				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
Public	$ 698	$ 845	$ 777	$ 701	$ 790	26%	23%	18%	11%	18%
Private	$ 706	$ 879	$ 1,023	$ 761	$ 901	51%	61%	53%	37%	43%
Total	$ 703	$ 868	$ 946	$ 744	$ 868	37%	39%	33%	23%	30%

[1] The ratio of paid submissions to valid submissions represents the probability of being paid, given the submission is correct.

Table C7. Program perspective: Changes in accepted-duplicate, unique, and valid submissions across the three-month periods (all programs).

	Valid submissions					Year-Over-Year growth (percentage change)			
	2017	2018	2019	2020	2021	2017-2018	2018-2019	2019-2020	2020-2021
Accepted-duplicate submissions	2,227	2,944	6,081	21,632	6,374	32%	107%	256%	-71%
Unique submissions	4,064	5,720	7,002	11,672	10,086	41%	22%	67%	-14%
Valid submissions	6,291	8,664	13,083	33,304	16,460	38%	51%	155%	-51%
Unique/valid ratio	65%	66%	54%	35%	61%	2%	-19%	-35%	75%