

(DRAFT)

The Economics of and the Effect of an External Shock on a Crowdsourced “Bug Bounty Platform”

Aviram Zrahia¹, Neil Gandal², Sarit Markovich³, and Michael Riordan⁴

¹Tel Aviv University, aviramzrahia@mail.tau.ac.il

²Tel Aviv University, gandal@tauex.tau.ac.il

³Northwestern University, s-markovich@kellogg.northwestern.edu

⁴Columbia University, mhr21@columbia.edu

Abstract

In this paper, 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 find and be rewarded for discovering software vulnerabilities. We then empirically examine the effect of an exogenous external shock (Covid-19) on Bugcrowd, one of the two largest “two-sided” bug bounty platforms. The shock reduced the opportunity set for many security researchers who either lost their jobs or were placed on a leave of absence. We show that the exogenous shock led to a huge rightward (downward) shift in the supply curve and to an increase both in the number of submissions and new researchers on the platform. The results suggest that had there been a larger increase in number of firms with bug bounty programs on the platform, many more unique software vulnerabilities would have been discovered. We quantify the benefits to the platform from the exogenous shock which enables us to shed light on the benefits associated with the gig economy.

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

1. Introduction

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.

“Two-sided” bug bounty platforms connect organizations that want to crowdsource part of their software security with ethical security researchers (hereafter, researchers). Such platforms provide organizations with access to a wide group of talented researchers with a wide range of knowledge and skills, thereby increasing the probability of finding vulnerabilities. Companies pay only for unique vulnerabilities found. Top researchers enjoy the opportunity to be invited to private programs where only selected researchers can participate, thereby increasing the probability of being the first to find and report

a vulnerability. Bugcrowd and HackerOne are the leading bug bounty platforms, and both began operations in 2012.

Bug bounty programs and bug bounty platforms are part of a more general trend to a “gig economy” - where enterprises supplement labor, and workers supplement income, with gig work. From the enterprises perspective, these platforms provide access to skilled and flexible labor. Gig work platforms creates opportunities for workers to access and compete in global job markets. Furthermore, they facilitate “bridge employment” (i.e., temporary employment between career jobs) and provide income opportunities in down times when the market does not accommodate full-time employment. Indeed, studying the ridesharing market, Koustas [2018] finds that, on average, driving for gig platforms replaced 73% of lost income from a main job. Moreover, taking advantage of gig work platforms during bad times helps overcome periods of income volatility. Similarly, Collins et al. [2019] find that workers typically start new platform work in times of a personal income crisis.

Employing a unique data set provided by Bugcrowd, we document the effect of an exogenous shock—the Covid-19 pandemic—on the market for vulnerabilities within the bug bounty platform. The data cover years 2017-2021, and for each year we focus on the three-month period from March to May. The 2020 period corresponds to the Covid shock. We examine the impact of the Covid shock on both the demand for vulnerabilities by participating organizations, and the supply of vulnerabilities from active researchers.

Since it was launched, Bugcrowd’s platform has hosted more than 2,400 programs offered by more than 1000 organizations, and attracted more than 30,000 active researchers who made at least one submission to a program. The data set records payments for valid submissions, which are for vulnerabilities within the defined scope of a program. While only the first researcher to discover a valid vulnerability is awarded a monetary payment, the data set also records duplicate valid submissions. A duplicate valid submission means that the researcher correctly identified a valid vulnerability, but was not first, and therefore did not receive a monetary award. Accounting for both paid submissions and duplicate valid submissions enables us to compute the average payment for valid submissions, which turns out to be a key variable for understanding the effects of the Covid shock.

Our analysis reveals that the Covid shock impacted both the supply and demand for valid vulnerabilities, but impacted supply much more dramatically. On the supply side, the shock greatly increased the number of researchers participating on the Bugcrowd platform. This makes sense if the shock reduced the outside opportunity set for researchers who either lost their jobs or were placed on a leave of absence during that period, because these researchers had more time on hand to look for vulnerabilities in bug bounty programs. On the demand side, there was a small increase in new programs relative to previous periods. This makes sense if organizations adapted to the pandemic by allowing their employees to work from home, enabling “black hat” hackers potentially to take advantage of the increased security vulnerabilities of the less effective home security systems and the newly deployed remote access solutions.¹

We interpret the Covid shock with a heuristic supply-and-demand model, by interpreting valid submissions as the relevant product and the average monetary award for valid submissions as the price. The Covid shock substantially shifted the supply curve by greatly increasing the number of active researchers, and shifted the demand curve more moderately by slightly increasing the growth in active programs. These shifts combined to increase greatly the number of valid submissions. Interestingly, this increase in quantity was mostly due to a huge increase in the number of duplicate valid submissions, reflecting the much more significant supply curve shift. Consequently, there was a large decrease in the average equilibrium price for a valid submission, because valid duplicates do not receive a monetary reward.

That is, the Covid shock “threw” the market for vulnerabilities out of a previously more-or-less stable equilibrium. The

reduction in equilibrium price due to the Covid shock presumably dampened the incentives of individual researchers to search for vulnerabilities. There was a one-in-six chance of being paid for a valid submission in 2020, compared to a slightly larger than a one-in-three chance in 2019 and 2021. Indeed the total number of submissions fell dramatically in 2021, relative to 2020. The ratio of paid submissions to total valid submissions was in the 48-49 percent range for 2017-2018 and in the 36-37 percent range in 2019 and 2021. The latter numbers reflect slightly increased competition trend among researchers, but not a dramatic change. If the demand response would have increased in 2020 so as to keep the ratio of paid submissions to total valid submissions in the 36-37 percent range, as in 2019 and 2021, rather than falling to 16%, the total number of paid submissions would have been more than double the actual number in 2020.

The Covid shock provides an opportunity to address key public policy issues associated with crowdsourcing and the “gig” economy. We show what happened in the “white” market for vulnerabilities, when the value of outside options was lowered. An often mentioned benefit of the gig economy (freelance work as opposed to permanent jobs) is that the response from an external shock should be almost instantaneous on the supply side. Here we show that this was the case, and we quantify effects from the increased supply of new researchers and submissions.

The paper proceeds as follows: In section 2 we elaborate more on the real-world dynamics of vulnerabilities, and the literature associated with bug bounty programs and platforms. Section 3 includes details on Bugcrowd’s platform submission workflow. We then discuss the data in section 4, and a heuristic supply-and-demand model in section 5. Section 6 looks at key empirical properties of the platform. In section 7, we examine the effect of the shock on the supply side, demand side, and equilibrium properties of the platform. Finally, sections 8 and 9 briefly discuss the findings and their meaning.

2. 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.² 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 on-line service. For a product, the vendor will most likely release a technical security notification to its customers (either 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.³

¹ 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 have been documented extensively by market analysts, cyber-security vendors and governmental agencies. See Lallie et al. [2021].

² If the firm finds some of its own vulnerabilities there is no “white hat hacker” market. See Choi et al. [2010] for a theoretical model that addresses disclosure in this setting.

³ CVE® is a list of entries of publicly known cyber-security vulnerabilities maintained by the MITRE Corporation. It feeds NIST’s U.S. National Vulnerability Database (NVD) which adds more context.

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.⁴

2.1. Bug Bounty Platforms as Market Intermediaries

As Malladi and Subramanian [2020] note, there are three categories of security crowd-sourcing markets for vulnerabilities.

- The first category is institutional bug bounty programs which are hosted directly by software vendors who set their own policies and compensation plan. 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. Here intermediaries connect organizations and security researchers via a “two-sided” network or platform.

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]. These platforms take many forms. In general, the platform provides the infrastructure and rules of engagement in order to attract both sides of the market. Many of today’s most valuable firms, including Apple, Amazon and Google, are platforms or two-sided markets. Two-sided platforms create value and improve economic efficiency [Rochet and Tirole, 2006, Belleflamme and Peitz, 2019]. Some common examples of such two sides brought together by a platform owner are buyers and sellers (Amazon), media consumers and advertisers (Facebook), or application developers and device makers (Apple iOS). Two-sided markets can generate value by reducing the transaction costs faced by distinct groups of participants. Platform based markets are typically characterised by indirect (cross-side) network effects [Zhu and Iansiti, 2012, Rochet and Tirole, 2003], as each side’s perceived value of the platform increases with the number of users on the other side.⁵

Bug bounty platforms are two-sided markets as they connect organizations that want to crowdsource their software security with researchers. Ideally, a platform hosts many programs for multiple organizations and has many high quality researchers. The researcher who first finds and reports a novel vulnerability receives a payment (bounty). Bug bounty platforms thus create a tournament-like arrangement. The program structure, scope,

and rewards are often determined by the firm, but the rules of engagement and procedures are established by the platform.

In the 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 (who wrote the code or own the on-line service) and are interested in protecting against exploits by adversaries. The supply side of the market 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 maturity of the program—the more mature the program, the harder it is to find new vulnerabilities and thus the higher the reward, 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 which determine their relative rank within the platform and may enable them to receive invitations to work in private bounty programs—programs where participation is by invitation only.

2.2. 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. [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 correlation 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. Malladi and Subramanian [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.

None of these studies had access to private bug bounty programs, which made up to 88% of the bounty programs activity on Bugcrowd’s researched platform during the full calendar years we examine. See Table B1.

3. The Bugcrowd Platform

3.1. Rules of Engagement

The rules of engagement between a hacker and an organization 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 bug has been submitted. Submissions are either valid (correct) or invalid. Valid submissions can then be either paid or duplicate, i.e., submissions that report vulnerabilities that have already been reported by other researchers.

Bugcrowd offers two types of programs. Managed Bug Bounty programs (MBBs) give a monetary reward to the first

⁴ If an adversary discovers the vulnerability before the firm, they might produce a zero-day (0-day) exploit, which is best defined as an “exploit without a patch”. There is a “black market” for zero day vulnerabilities as described by Ablon and Bogart [2017].

⁵ 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 referred to as the chicken-and-egg problem by Caillaud and Jullien [2003].

researcher to submit a unique valid vulnerability as well as points. Later researchers who find the same vulnerability in an MBB program receive only points (but no monetary award) for their valid “duplicate” submission.⁶ Points increase the likelihood that a researcher will receive an invitation to work on a private program.⁷ The second type of program, Vulnerability Disclosure Programs (VDPs) reward hackers with points, but no monetary awards. The “recognition” they get for finding a meaningful (high priority) unique vulnerability, increases the likelihood that they will receive invitations to work on private MBB programs.

In the main analysis below, we included MBB programs and VDPs. Our results are qualitatively unchanged if we exclude VDP programs, as we show in Appendix A. In that appendix, we examine the two program types separately and show that the results are qualitatively unchanged.

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

3.2. Submission Workflow

Figure 1 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.⁸ 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.⁹ Valid vulnerabilities are then integrated into the existing Software Development Lifecycle (SDLC) tools to be fixed, and related reward payouts are processed accordingly.

⁶ In some rare cases organisations do give a monetary reward for a valid duplicate submission to an MBB program, in order 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.

⁷ The points are accumulated per researcher and reported in monthly and all-time leadership boards: <https://bugcrowd.com/leaderboard> (lower rank is better).

⁸ A resource outlining Bugcrowd’s baseline priority rating, including certain edge cases, for common vulnerabilities: <https://bugcrowd.com/vulnerability-rating-taxonomy/>.

⁹ In some cases a vulnerability can be marked duplicate during the triage stage and eventually be rejected (based on <https://docs.bugcrowd.com/customers/getting-started/with-bugcrowd/>).

4. Data

The paper employs a unique data set obtained from Bugcrowd.¹⁰ The data spans the full period for which the company has been in existence, i.e., from 2012 and includes all bug submission activity through May 2021. Since we are interested in the effect of the Covid-19 shock, we will primarily focus on data from 2017-2021. We have detailed data on submissions, researchers, and programs, for public programs—programs open for all researchers—as well as private programs, where only invited researchers are allowed to participate and submit vulnerabilities. This distinction is especially important given that more than 90% of the new programs during the 2017-2021 period were private. Our data set was obtained through a Data Transfer Agreement (DTA) between Tel Aviv University and Bugcrowd. As far as we know, data which includes researcher attributes and detailed submissions from private programs, was not available to scholars up until now.

The data set contains information on the demand side (organizations/programs), the supply side (researchers), and the product (bug submissions). The organizational data includes: 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. 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:

- Researcher that made the submission;
- Date and time of the submission;
- 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);
- The amount of points awarded;

In order to focus on the effect of the exogenous Covid-19 shock, we examine five three-month time periods from 2017-2021. Each time period includes three full months of activity from March 1 to May 31 of the respective year. We chose this period of the year since the pandemic was declared in early March 2020 and many countries enforced lock-downs during this time. Thus the Covid shock was strongest during this period. We define the March 1 to May 31 2020 period to be the “Covid period,” and in order to exclude seasonality effects, we include the same period in all other years as well. In what follows, unless otherwise noted, the analysis is for the three full months from March 1 to May 31.

5. 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 model of demand, supply, and equilibrium price to motivate our analysis. To begin, we define the product and its price with reference to the “tournament” structure of program rewards. We define the product to be a valid vulnerability submission by a researcher to a participating

¹⁰ <https://bugcrowd.com/>.

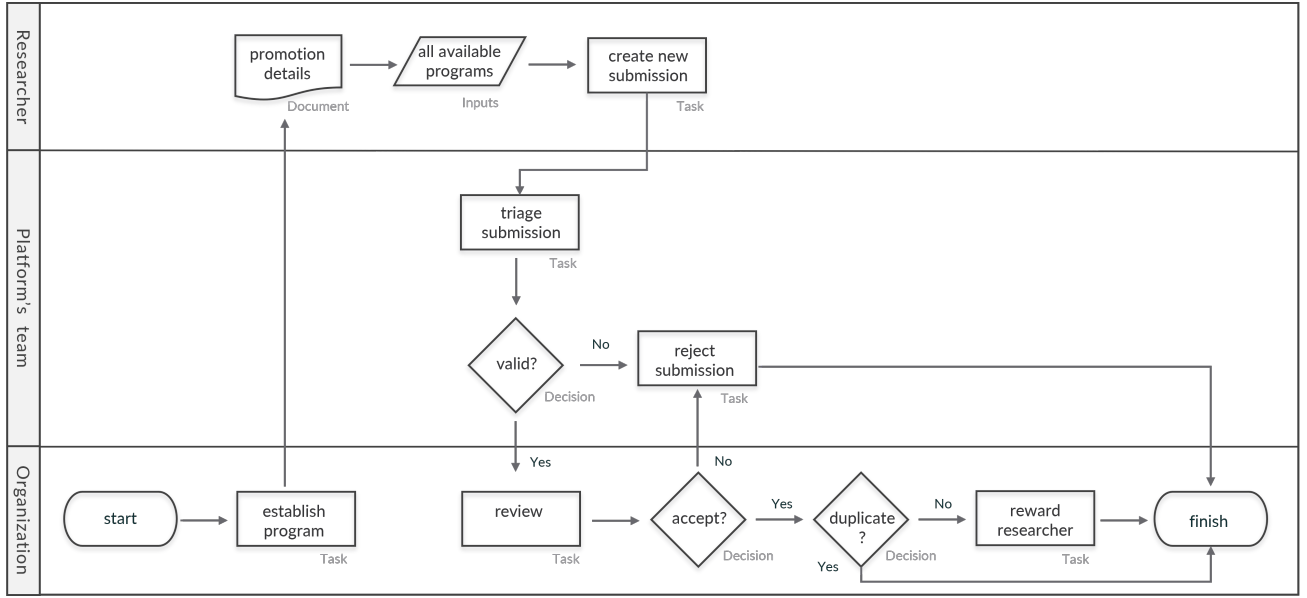


Figure 1: Bugcrowd's platform submission workflow.

program. Researchers, while looking for a vulnerability, 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. The price of the product, therefore, is the expected payment, which empirically we define 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 2.¹¹ The horizontal axis measures the total number of valid submissions over some time period, and the vertical axis measures the average payment for a valid submission. The upward sloping supply curve recognizes that, if 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 might submit more programs to the platform. The equilibrium price brings these two sets of incentives into balance.

Submissions are either valid or not, and valid submissions can be paid 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, which in turn lowers the programs' costs associated with processing submission. The number of discovered vulnerabilities, i.e., unique valid submissions that are paid a monetary award, in turn, equals the number of valid submissions times the probability of winning with a valid submission.




Figure 2: Supply and demand curves for valid submissions on a bug bounty platform.

This probability depends on the number of duplicate valid submissions, and it is calculated as the paid-to-valid submissions ratio.

We examine how the Covid shock impacted these and other relevant variables derived from the Bugcrowd data set. Our main conclusion, based on the heuristic model, is that the Covid 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 for a valid submission, as illustrated in Figure 2.

¹¹ Figure 2 also shows a rightward shift of the supply curve associated with the Covid shock. See discussion below.

6. Key Properties of the Platform

Before we examine the results associated with the Covid shock, we present some descriptive statistics on the platform based on the full-year data for 2017-2020.

As Table B2 shows, most of the demand (the bug bounty programs) comes from organizations from the United States (U.S.), while a plurality of active researchers—defined as researchers with at least one submission—are from India. U.S. organization represented 68% of active programs in 2020, and similar values in other time periods as well. From the supply side we see a steady increase over the years in the percentage of active researchers and submissions coming from Asian countries. In 2020, Asian countries represented 58% of active researchers and 70% of submissions, with India being the most dominant country in this region (46% of all active researchers, and 60% of all submissions in 2020).

These statistics are consistent with findings in “The Online Labour Index”, an economic indicator that provides an online gig economy equivalent to conventional labor market statistics.¹² According to their data, roughly half of the gig economy labor demand on selected digital platforms originates from the United States [Kässi and Lehdonvirta, 2018], one third of all online freelancers are from India, and 15% from Bangladesh¹³

In terms of programs, organizations are shifting their activity away from public programs to private programs (see Table B1). Specifically, in 2017, public programs accounted for nearly 29% of all active programs, but represented only 12% in 2021. The share of new private programs launched grew from 84% in 2017 to a staggering 97% in 2020.¹⁴

Bug bounties are considered a great way for researchers to gain experience and build reputation. This suggests that researchers may not continue to work on a bug platform for long and, hence, the platform would experience high attrition rate. To examine this, we categorized researchers by the year of their first submission period and looked at their attrition rate between that period and subsequent years. Indeed, attrition rates are very high, especially between the first and second years. For researchers who joined during 2017-2020, the attrition rate between the first and second years is between 71 and 76 percent. Attrition rates between the second and third years then decrease to 30-35 percent. While the high attrition rates may be explained in part by an increase in researchers outside opportunity as they build experience and reputation, it may also be a result of researchers being discouraged by low returns for their efforts.

On the other hand, those who stay get better over time. For all cohorts of researchers (2017, 2018, 2019, and 2020), there was a continuous rise in the paid-to-valid ratio each year. See Table B3.


¹² The Online Labour Index (OLI) is derived from the iLabour research project at the Oxford Internet Institute: <https://ilabour.oii.ox.ac.uk/online-labour-index/>.

¹³ OLI 2020 update: <https://ilabour.oii.ox.ac.uk/online-labour-index2020/>.


¹⁴ From the firm standpoint, there is a trade-off between public and private programs. There are more eyeballs with public programs, but also there is greater competition, which can be discouraging for researchers. Apparently the latter effect is more important.

While rewards on the platform are skewed, such that top researchers earn a disproportional share of total monetary rewards, this share has been declining suggesting that no researchers have market power. As Figure 3a shows, in 2017, the top 100 researchers (at that point in time) earned 43% of the total rewards. The top 100 researchers in 2020 (at that point in time) earned only 23% of the total rewards. The top 500 researchers exhibit a similar trend.

On the demand side, Figure 3b shows the share of total monetary rewards paid by the top 10 programs. These programs accounted for 42% of the total annual payments in 2017 and only 33% of annual payments in 2020. Taken together, these values and trends suggest very little (if any) market power on either side of the market.



(a) Top 100 Researchers



(b) Top 10 Programs

Figure 3: A yearly view on total rewards share of top 100 researchers (supply) and top 10 programs (demand) in Bugcrowd’s platform.

7. The Effect of Covid-19 Exogenous Shock

Next, we explore the effect of the Covid shock on the supply and demand sides of the market and then discuss overall equilibrium changes. We use for the analysis the 3-month period data (March 1 - May 31) for 2017-2021.

7.1. The Supply Side

We start with the most prominent effects of the Covid shock:

Observation 1 *The Covid shock resulted in a significant increase in the number of researchers, and an especially large increase in the number of total submissions.*

Figure 4 shows year-over-year (YoY) changes in the number of researchers and submissions during the three-months periods for 2017-2021. While the platform experienced steady growth on the supply side from 2017-2019, supply side activity increased significantly during the Covid 2020 period.

Specifically, the number of total submissions increased from 21,157 in 2019 to 53,098 in 2020—a 151% growth. This growth rate during the Covid period is much larger than the growth observed between 2017-2018 (49%), and 2018-2019 (69%). Moreover, in the subsequent 2021 period, there is a large drop in the number of submission to 30,955 submissions. Table B4 clearly shows that this huge increase in the number of submissions was primarily driven by researchers who joined the platform in 2020; they made 20,118 submissions (38% of total submissions). By comparison, in 2019 and 2021, there were only between 6,000 and 7,000 submission made by new researchers.

In terms of the number of active researchers, there was a 72% growth in the number of active researchers between 2019 and 2020 as compared with a 43% growth between 2017-2018 and a 53% growth between 2018-2019.

The number of valid submissions surged 172% during the Covid period, from 9,525 to 25,864. An increase that completely dwarfs the increases between 2017-2018 (35%), and 2018-2019 (64%). Furthermore, in 2021, the number dropped to 10,195 valid submissions, virtually the same level as in 2019. Interestingly, this huge increase in the number of valid submissions was primarily driven by researchers who joined during the Covid period, who made 10,737 valid submissions. In comparison, new researchers in 2019 and 2021 made 1,979 and 1,166 valid submissions during their first year on the platform, respectively. See Table B4.

Geographically, the largest increase in submissions, during the Covid period, came from researchers from India and Turkey. In particular, submissions from India soared from 9,335 in 2019 to 31,673 in 2020, an increase of more than 22,000. Submissions from Turkey skyrocketed from just 472 in 2019 to 7,724 in 2020. This surge was then followed by a significant decline in the subsequent period of 2021.




Figure 4: YoY changes in active researchers and number of submissions, 3-months periods for 2017-2021.

7.1.1. Post-Covid

The analysis above suggests that the Covid shock attracted many more researchers to participate in bug bounty programs, likely because there was a decrease in the outside option of a full-time job. If this is indeed, 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.

Observation 2

Researchers who joined the platform during the Covid period made many fewer submissions in 2021. This decline was exceptionally large relative to other cohorts in their second year on the platform.

Figure 5 shows YoY changes in researchers' activity level (i.e., number of submissions) based on the year the researcher joined the platform. The figure shows changes in activity levels during Covid and post-Covid (i.e., number of submission in 2021 relative to 2020) as well as between the first and second year of researchers' activity.

As the figure shows, the 2020 researchers decreased their activity levels dramatically in 2021 (post-Covid). Specifically, researchers who joined the platform during the Covid period made 89% fewer submissions in the post-Covid 2021 period. For comparison, there was a 53%, 43%, and 62% drop respectively for the 2017, 2018, and 2019 cohorts between their first and second years..




Figure 5: YoY changes in submissions for researchers who joined Bugcrowd during the 3-month periods of 2017-2020. Changes were measured between the Covid and post-Covid 3-months periods, and the first and second year of joining.

7.2. The Demand Side

Given the large response on the supply side, next, we examine the demand side of the platform. Table B5 shows the changes in active and new programs across the three-month periods between 2017-2021 for the different program types. As expected, relative to previous years, there is an increase in the number of

active programs as well as in the number of new programs on the platform. Interestingly, the increase is not symmetric across public and private programs.

Observation 3 *Active private programs exhibited a larger growth during the Covid period, as compared to previous year. Public programs had a more moderate overall upward trend, yet showed no evidence of a Covid effect in 2020.*

As Table B5 shows, there is an overall upward trend in the total number of active private programs averaging 57% a year before the Covid period. Between the 2019 and 2020 periods, however, there was a 91% increase in the number of active private programs, followed by a correction in the next year. Correspondingly, the percentage share of *new* private programs rose from 32% in 2019 to 36% in 2020, and then dropped to 26% in 2021. 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 increase in supply greatly exceeded the increase in demand.

Another possible avenue for demand-side effects is in the rewards paid by programs for discovered vulnerabilities. The Bugcrowd platform establishes ranges for payments rewarding different priority submissions, but the programs have discretion about what exact reward to pay within each range, as well some discretion about classifying the priority of submissions. Furthermore, Bugcrowd encourages established programs, for which new vulnerabilities may be more difficult to discover, to pay at the higher end of the range,¹⁵ and also to pay bonuses outside the range for particularly significant vulnerabilities.¹⁶ Given such discretion, the average reward per paid submission might be viewed as demand-side variable, possibly shifting during the Covid period. We defer a discussion of this possibility to the next section.

7.3. Equilibrium Outcome

The discussion above suggests that the Covid shock shifted the supply curve to the right—dramatically increasing the number of researchers on the platform. On the demand side there was a more moderate increase in active programs and a potential response in pricing. Below, we look at equilibrium changes over time in (i) the ratio of paid submissions divided by total valid submissions, (ii) the awards (both expected and actual rewards for vulnerabilities over time,) and (iii) the number of unique vulnerabilities found.

7.3.1. Probability of Earning a Monetary Reward

Observation 4 *The probability of earning a monetary reward, given a valid submission, declined during Covid and bounced back post Covid.*

Figure 6 shows that the Covid shock seemed to essentially “throw” the platform out of the equilibrium it had maintained for most of its existence.¹⁷ The ratio of paid submissions to total valid submissions was in the 48-49 percent range for 2017-2018 and in the 36-37 percent range in 2019 and 2021.¹⁸ In 2020, the ratio of paid submissions to total valid submissions fell dramatically to 16%.

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 was in the 20-22 percent range in 2019 and 2021 but was only 7 percent in 2020!¹⁹




Figure 6: Valid submissions during the 3-month periods of 2017-2021. Paid-to-valid ratio indicates the probability of earning a reward given the submission is valid.

7.3.2. Expected and Actual Rewards

The increase in competition during Covid can be also seen by looking at expected rewards, measured empirically by the average reward for valid submissions.

Observation 5 *Expected rewards dropped tremendously, during the Covid period, and bounced back in 2021.*

As Figure 7 shows, the average reward for valid submissions ranged between \$341 - \$404 during the 2017-2019 period, but fell dramatically to just \$122 in 2020 (a decline of 66% relative to 2019). The average reward then rose to \$319 in 2021. The trend for higher-priority and lower-priority vulnerabilities was similar.

In general, a decline in expected rewards can be driven by (i) an increase in the number of valid submissions (a supply side effect), and/or (ii) a decrease in monetary rewards for paid submissions (both demand and supply side effects), or a combination of both. More concretely, the average reward for valid

¹⁵ <https://www.bugcrowd.com/resources/guide/bugcrowds-defensive-vulnerability-pricing-model/>.

¹⁶ <https://docs.bugcrowd.com/customers/submission-management/rewarding/>.

¹⁷ This refers to the period beginning in 2017. The platform was much smaller before that year.

¹⁸ The 2019 and 2021 numbers reflect slightly increased competition among researchers relative to 2017 and 2018, but not a dramatic change.

¹⁹ In the case of private programs, the probability of earning a reward given a valid submission, fell from 57% in 2019 to 37% in 2020.




Figure 7: Expected rewards by submission priority during the 3-month periods of 2017-2021 in Bugcrowd's Bug Bounty platform.

submissions can be decomposed as follows: $(\text{average reward for paid submissions}) * (\text{paid submissions/valid submissions})$. Changes in the first term are most likely driven by a combination of supply and demand factors. Changes in the right 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.²⁰

We present the average monetary reward for paid submissions in Figure 8. As the figure shows, in 2020, monetary rewards for paid submissions decreased from an average of \$989 in 2019 to \$746 in 2020 (a decline of 25% relative to 2019). During the same time, the paid-to-valid submissions ratio fell dramatically from 36% to 16% (a decline of 55%). We conclude that the decline in the expected award between 2019 and 2020 associated solely with supply-side effects (a decline of 55% in the paid-to-valid ratio) was twice as large as the remaining residual decline (25% decrease in actual paid rewards) which was due to a combination of supply and demand factors.²¹




Figure 8: Actual average rewards by submission priority.

²⁰ We cannot think of demand-side factors that would cause a decline in the paid-to-valid ratio.

²¹ We examine in Appendix A the two program types (MBB and VDP) separately and show that these results are qualitatively unchanged.

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.²² As Figure 8 shows, the demand and supply responses differ across higher and lower priority vulnerabilities. Specifically, while there was a meaningful 33% drop in the average actual payment for higher-priority paid submissions, rewards for lower-priority submissions slightly increased. This strongly suggests that the decrease in expected reward for valid lower-priority submissions was almost completely driven by a shift of the supply curve. For higher-priority vulnerabilities, the equilibrium outcome is likely a result of a combined supply and demand responses (with the supply response being the dominant driver, as discussed earlier).

In addition, the demand and supply responses differ across public and private programs. For private programs, the average actual payment dropped 31% between 2019-2020, while for public programs it dropped only 6%, indicating the effect associated with demand side for public programs is significantly smaller. This is in-line with our demand side observation, which shows that private programs exhibited a larger growth during the Covid period. Furthermore, The paid-to-valid submission ratio for public programs declined by 68% compared to a 36% decline for private programs. Given that most new researchers, who joined the platform during Covid period, initially only had access to public programs, the increase in supply for public programs greatly exceeded the increase in demand.

7.3.3. Unique Vulnerabilities Discovered

Observation 6 The percentage increase in the actual number of paid awards between 2019 and 2020 was very similar to the percentage increase in the number of paid submissions between 2018 and 2019.

As Figure 6 shows, while the number of valid submissions increased tremendously during Covid, there was not a large increase in the number of paid submissions between 2019 and 2020. Since the number of paid submissions is essentially the number of unique vulnerabilities discovered, this means that there was not a large increase in the number of unique vulnerabilities discovered. The number rose from 2,857 in 2018 to 3,444 in 2019 and 4,232 in 2020, before falling back to 3,821 in 2021. The percentage increase in the number of unique vulnerabilities discovered was essentially the same between 2018 and 2019 (21%) as the percentage increase between 2019 and 2020 (23%).

This was primarily because during the Covid period, there were more researchers competing over the same vulnerabilities, especially in public programs. Despite the fact that the number of valid submissions to public programs skyrocketed from 5,417 in 2019 to 17,429 in 2020, the number of unique vulnerabilities discovered (and hence paid for) was virtually unchanged (1,109 in 2019 vs. 1,145 in 2020.)

²² The P1-P5 definitions are by Bugcrowd. For a small number of rewarded submissions, there was no priority listed. Therefore, we mapped them into high-priority or low-priority categories where possible, based on their value compared to the average amount rewarded per priority category, for the same program during the same calendar year.

7.4. Summary

The incredibly large number of valid submissions in 2020 suggest that had there been a larger increase in number of firms with bug bounty programs on the platform, many more unique vulnerabilities would have been likely discovered.²³ If the demand response to the shock had been similar to the supply response such that the ratio of paid to valid submissions remained at 36-37 percent as in 2019 and 2021, rather than the 16% during 2020, the number of unique vulnerabilities identified by researchers in 2020 could have been more than double the actual number in 2020. This “counterfactual” seems reasonable since all software contains vulnerabilities, and that there are always more vulnerabilities to find. Many researchers indeed believe that “with the complexity of current hardware and software systems arising from billions of transistors and millions of lines of code, there are effectively an infinite number of unknown vulnerabilities.”²⁴ This also implies that some level of duplication is desirable. However a 16% ratio of paid to valid submissions during the 2020 (Covid) seems too low to encourage researchers to search hard for vulnerabilities.

8. Policy Implications

Our results bear important policy implications for the market for vulnerabilities, as well as, more generally, for the gig economy.

In terms of 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 [Ablon and Libicki, 2015]. The large supply response we identified from the Covid shock suggests that more bug bounty programs and larger bug bounty platforms could change this dynamic. In particular, an increase in demand for researchers and vulnerability submissions would likely drive more transactions to take place in the white market rather than the black or grey markets. Governmental agencies have begun to use bug bounty programs. The Cybersecurity & Infrastructure Security Agency (CISA) in the US, has announced in 2021, a vulnerability disclosure policy platform. The platform (provided by Bugcrowd and EnDyna), allows agencies to list systems in scope for their vulnerability disclosure policies, so security researchers may try to find vulnerabilities in agency websites and submit reports for analysis.²⁵

More generally, gig work platforms like Uber, UpWork, and Bugcrowd, offer a tremendous opportunity for employees during turbulent times [Koustas, 2018, Collins et al., 2019]. Indeed, our analysis suggests that many researchers have reached out to bug bounty platforms during the Covid period when the market experienced a huge decline in job opportunities and thus many researchers’ outside option was reduced. The flexibility and low barriers to entry of gig work platforms provided

skilled researchers with an opportunity to supplement earning by finding vulnerabilities. Our analysis suggests that many of the new researchers who took advantage of this opportunity during the crisis, decreased their activity level once market conditions improved.

From a policy perspective, the relatively small demand response in terms of an increase in the number of programs might be seen as a missed opportunity. Specifically, given the large number of researchers working on finding vulnerabilities, more of the effort could have been directed toward finding new unique vulnerabilities as opposed to finding duplicate vulnerabilities. A policy that encourages the introduction of new programs may have benefited both researchers and organizations and could have resulted in increased productivity.

9. Conclusion

The paper examines the effect of the Covid-19 pandemic’s exogenous shock on activity on the Bugcrowd bug bounty platform. We find 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.

To the best of our knowledge, this is the first study to analyze a large, detailed data set of bounty activity on a bug bounty platform which includes data on private programs, as well as duplicate submissions. Hence, it allows for a more comprehensive analysis, given that private programs represent more than 90% of newly offered programs in recent year.

10. Acknowledgments

We thank Casey Ellis, the founder and CTO of Bugcrowd, and his phenomenal team for sharing with us 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 helpful comments and suggestions.

References

- L. Ablon and A. 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.
 - L. Ablon and M. 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.
 - A. M. Algarni and Y. 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>.
 - P. Belleflamme and M. Peitz. Platform competition: Who benefits from multihoming? *International Journal of Industrial Organization*, 64, 2019. ISSN 01677187. doi: 10.1016/j.ijindorg.2018.03.014.
- ²³ This statement holds for private programs as well, since the paid to valid ratio fell by 36% from 2019 to 2020 for private programs.
- ²⁴ Quote from <https://www.sigarch.org/lets-keep-it-to-ourselves-dont-disclose-vulnerabilities/>.
- ²⁵ <https://www.cisa.gov/blog/2021/07/29/cisa-announces-new-vulnerability-disclosure-policy-vdp-platform>.

- B. Caillaud and B. Jullien. Chicken & Egg: Competition among Intermediation Service Providers. *The RAND Journal of Economics*, 34(2), 2003. ISSN 07416261. doi: 10.2307/1593720.
- J. P. Choi, C. Fershtman, and N. 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.
- B. Collins, A. Garin, E. Jackson, D. Koustas, and ... Is gig work replacing traditional employment? Evidence from two decades of tax returns. *Unpublished paper, IRS ...*, 2019.
- O. Kässi and V. Lehdonvirta. Online labour index: Measuring the online gig economy for policy and research. *Technological Forecasting and Social Change*, 137:241–248, 12 2018. ISSN 00401625. doi: 10.1016/j.techfore.2018.07.056.
- D. K. Koustas. Consumption Insurance and Multiple Jobs: Evidence from Rideshare Drivers. *Working Paper*, 2018.
- H. S. Lallie, L. A. Shepherd, J. R. Nurse, A. Erola, G. Epiphaniou, C. Maple, and X. Bellekens. Cyber security in the age of COVID-19: A timeline and analysis of cyber-crime and cyber-attacks during the pandemic. *Computers & Security*, 105:102248, 6 2021. ISSN 0167-4048. doi: 10.1016/J.COSE.2021.102248.
- T. Maillart, M. Zhao, J. Grossklags, and J. 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>.
- S. S. Malladi and H. 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.
- J. C. Rochet and J. Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 2003. ISSN 15424766. doi: 10.1162/154247603322493212.
- J. C. Rochet and J. Tirole. Two-sided markets: A progress report. In *RAND Journal of Economics*, volume 37, 2006. doi: 10.1111/j.1756-2171.2006.tb00036.x.
- M. Zhao, J. Grossklags, and P. 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.
- F. Zhu and M. Iansiti. Entry into platform-based markets. *Strategic Management Journal*, 33(1), 2012. ISSN 01432095. doi: 10.1002/smj.941.

Appendices

A. Separating Monetary and Point Rewards

Below we separately analyze the two different program types: (i) Managed Bug Bounty Programs (MBBs), and (ii) Vulnerability Disclosure Programs (VDPs). The reward “currency” in this analysis is US dollars for MBBs and points for VDPs. As we show, the results for VDPs are especially interesting, as point rewards bear no monetary costs.

We start with MBBs. As Figure 9 shows, our results remain qualitatively unchanged when we limit our data to submissions to MBBs only: average rewards for valid submissions fall dramatically in 2020 (a decline of 45%) and bounce back in 2021. All in all, the decrease that was due solely to the supply side response (30% decrease in paid-to-valid ratio), was 41% larger than the remaining residual decline (21%) in the average reward in paid submissions. See Figure 10 for paid-to-valid submissions analysis.




Figure 9: Average rewards for paid and valid submissions, for MBBs only, during the 3-month periods of 2017–2021 in Bugcrowd’s Bug Bounty platform.




Figure 10: Valid submissions for MBB programs only (paid and accepted-duplicates), during the 3-month periods of 2017–2021. Paid-to-valid ratio indicates the probability of earning a monetary reward.

Average rewards for higher- and lower-priority vulnerabilities follow a pattern identical to the one we discussed in our main analysis. In particular, when we limit our data to only MBBs, the decline in average rewards for lower-priority valid submissions is still solely due to the supply side response.

Moving to VDPs, these programs award researchers with points which are used for ranking and can be also valuable with respect to invitations to private programs. Consequently, researchers in general and new researchers in particular find it beneficial to put effort into submissions to these programs.

Given that VDPs entail no actual monetary reward, and the number of points awarded for a unique discovered vulnerability ("paid submissions") is predefined and based on the submission's priority, one would expect the pandemic to have no effect on average rewards. Indeed, as Figure 11 shows, average points rewarded for paid VDP submissions remained practically the same during all periods. Still, average points for valid submissions fell from 6-8 points, during the 2017-2019 periods, to just 2 points in 2020, before bouncing back to 7 points in 2021.²⁶ This means that the decline in 2020 in expected points awards in all VDPs was solely due to the increase in the supply side. The large increase in submissions to VDPs during the Covid period resulted in a 70% drop in the ratio of paid-to-valid VDP submissions, and as a result in drop of 70% in the average reward for valid submissions (see Figure 12).




Figure 11: Average rewards for paid and valid submissions, for VDPs only, during the 3-month periods of 2017-2021 in Bugcrowd's Bug Bounty platform.




Figure 12: Valid submissions for VDP programs only (points paid and accepted-duplicates), during the 3-month periods of 2017-2021. Paid-to-valid ratio indicates the probability of earning points as a result of a unique vulnerability discovered.

²⁶ The trend for higher- and lower-priority vulnerabilities was similar.

B. Supporting Tables

Table B1. Program perspective: changes in active and new programs across full calendar year periods.

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

¹Programs with one or more submissions during the period.

²Programs which their first submission occurred during the period.

Table B2. Top countries view: geographical diversity of active researchers, active programs, and submissions, across full calendar year periods.

Top countries	United States ¹				India ²			
	2017	2018	2019	2020	2017	2018	2019	2020
Percentage of active programs	71%	75%	71%	68%				
Percentage of active researchers					30%	33%	38%	46%
Percentage of total submissions					40%	41%	49%	60%

¹Demand side top country.

²Supply side top country.

Table B3. Researcher experience: changes in submission success rate and paid to valid submission ratio, across full calendar year periods.

Year researcher joined	Valid / total submission ratio ¹				Paid / valid submissions ratio ²			
	2017	2018	2019	2020	2017	2018	2019	2020
2017 researchers	30%	40%	53%	48%	29%	42%	43%	48%
2018 researchers		32%	45%	46%		31%	34%	39%
2019 researchers			43%	43%			15%	20%
2020 researchers				40%				7%

¹Valid to total submission ratio, reflects the accuracy of submissions for the researchers' cohort.

²Paid to valid submission ratio reflects the chances to win the tournament given the submission is correct.

Table B4. Researcher experience: submissions and rewards by year researcher joined, across the three-months periods.

Year researcher joined ¹	No. of submissions					Percentage of 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	428	285	375	313	193	26%	37%	54%	51%	45%
2018 researchers		1,008	762	605	521		32%	42%	48%	48%
2019 researchers			1,979	1,104	396			32%	48%	40%
2020 researchers				10,737	769				53%	35%
2021 researchers					1,166					17%

¹Only researchers who joined during the three-months periods are listed in this view

²Valid submissions are the sum of paid submissions and accepted-duplicates.

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

Table B5. Program perspective: changes in active and new programs across the three-months periods.

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	81	33	60	91	10%	48%	12%	12%	13%
New between periods	56	21	152	271	412	51%	13%	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%

¹Programs which were active in the previous 3-month period, with one or more submissions during the period.

²Programs which started between the previous 3-month period to the period.

³New programs with first submission during the period.