(DRAFT - not for citation or distribution)

The Effect of an External Shock (Covid-19) on a Crowdsourced “Bug Bounty Platform”

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

[1]Tel Aviv University, aviramzrahia@mail.tau.ac.il
[2]Tel Aviv University, gandal@tauex.tau.ac.il
[3]Northwestern University, s-markovich@kellogg.northwestern.edu
[4]Columbia University, mhr21@columbia.edu

Abstract

In this paper, we 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 new researchers on the platform and of the quality of the security researchers. 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 to trade vulnerabilities as products between firms and individual security researchers. The program enables organizations to get in touch with cyber security 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 their software security with ethical hackers (or security researchers). Bugcrowd and HackerOne are the leading Bug Bounty Platforms. HackerOne was founded in 2012, while Bugcrowd began operations in 2013.

The Bugcrowd platform hosts more than 2,400 bug bounty programs, offered by more than 1,000 organisations and government agencies. More than 30,000 hackers/researchers have made at least one submission on the Bugcrowd platform. In these Bug Bounty platforms, the first researcher to identify a particular vulnerability in the software receives a monetary reward.

Importantly, in the data set provided by Bugcrowd that we employ in this paper, correct duplicate submissions are also recorded. A correct duplicate means that the researcher correctly identified a particular vulnerability, but since he/she was not the first researcher to identify that vulnerability, there is no monetary award.

1.1. Defining the Product

Using the data set provided by Bugcrowd, we examine the effect of an exogenous external shock (Covid-19) on the Bugcrowd platform. In order to decompose the 2020 Covid effect into a supply curve shift, a demand curve shift, and an equilibrium price response, we must define the product and address the “tournament” structure of the program.

The product is a (discovered) vulnerability submitted by a researcher. This interpretation, however is complicated by the possibility of duplicate correct submissions, only one of which earns a payment.

The horizontal access of the demand and supply space for this product is the number of correct submissions, and the vertical axis is the expected average payment for a correct submission. The quantity thus includes duplicate submissions

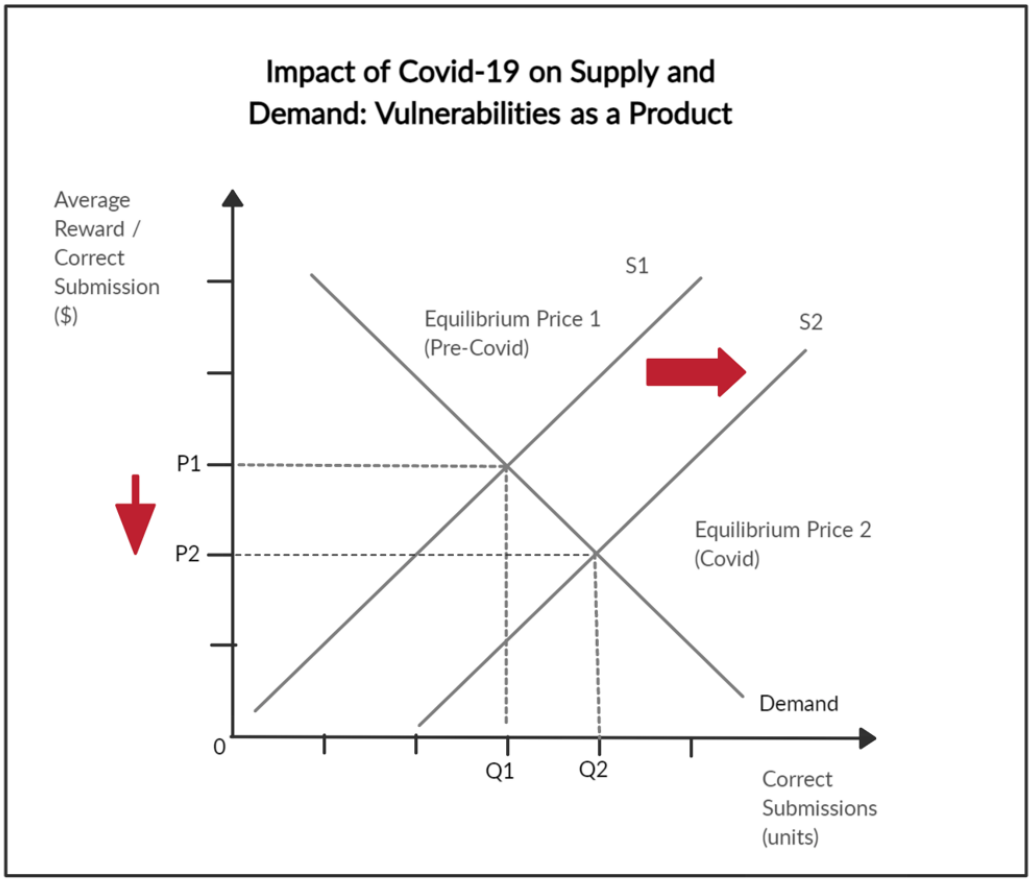
Impact of Covid-19 on Supply and Demand: Vulnerabilities as a Product
Average Reward / Correct Submission ($)
S1
S2
Equilibrium Price 1 (Pre-Covid)
Equilibrium Price 2 (Covid)
P1
P2
Demand
0
Q1
Q2
Correct Submissions (units)

Figure 1: Supply and demand curves for product vulnerabilities over a bug bounty platform.

and the expected price accounts for the fact that duplicate submission are unpaid. See Figure 1.[1]

The underlying idea behind this framework is that every correct submission has an equal chance of succeeding, so the relevant price is really the expected payment.

A feature of this labelling is that the number of correct submissions (quantity) is decomposed into the total number of submissions multiplied by the quality of submissions. The number of discovered vulnerabilities (i.e. successful submissions that result in a monetary award) in turn equals to the number of correct submissions times the probability of winning with a correct submission (which depends on the number of duplicates). Alternatively, the number of discovered vulnerabilities is equal to the number of total submissions times the success rate, where the success rate is the number of paid submissions divided by the total number of submissions.

1.2. Our Analysis and Key Results

In order to focus on the effect of the exogenous Covid shock, we analyze 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 lockdowns during this time. Thus the Covid shock was strongest during this period. Hence, we define the March 1 to May 31 2020 period to be the “Covid” period.

On the supply side, the shock reduced the opportunity set for many security researchers who either lost their jobs or were placed on a leave of absence during the Covid shock. We show that exogenous Covid shock led to a significant rightward shift of the supply curve.

In particular, there was a huge increase in the number of correct duplicates, which rose by 255 percent from 6,081 in 2019 to 21,632 in 2020, before falling to 6,374 in 2021. This increase was primarily driven by new researchers to the platform during the Covid period. They submitted 11,126 correct duplicates in 2020.

As a consequence of this supply side shift, the ratio of paid submissions to total correct submissions (paid submissions + accepted duplicate submissions) fell dramatically from 36 percent in 2019 to just 16 percent in 2020. It rebounded to 37 percent in 2021.

The new researchers to the Bugcrowd platform were of high quality: The percentages of their vulnerability submissions that were correct was much higher than those from other cohorts. Additionally, the ratio of their correct duplicates to total submissions was higher than other cohorts of researchers.

[1] Figure 1 also shows a rightward shift of the supply curve associated with the Covid shock. See discussion below.

On the demand side, organisations had to adapt to the pandemic and in many cases permitted their employees to work from home during Covid-19. This created an opportunity for “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. As a consequence, many companies experienced a significant growth in number and severity of security incidents as the attack surface expanded.[2]

Although in theory, this shock thus provided incentives for additional firms to join bug bounty platforms and have their software examined for vulnerabilities, we show that the shock led to just a small rightward shift of the demand curve.[3]

In terms, of equilibrium changes, the large rightward shift of the supply curve as a result of the shock led to a dramatic fall in the expected price (award) per correct submission. The expected average award per correct submission ranged from $341 - $404 during the 2017-2019 period, but fell dramatically to just $122 in 2020. The expected average award rose to $319 in 2021.

If the demand response had increased in 2020 so that the ratio of paid submissions to total correct submissions had remained in the 36-37 percent range as in 2019 and 2021, rather than falling to 16 percent, the total number of identified vulnerabilities would have been more than double the actual number in 2020.

The shock provides an opportunity to address key public policy issues associated with crowdsourcing and the “gig” economy. We show what happened when the value of outside options were lowered in the “white” market for vulnerabilities. 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. Here we show that this was the case, and we quantify effects from the increased supply of high quality researchers.

The paper proceeds as follows: In section 2 we elaborate more on the real-world dynamics of vulnerabilities and bug bounty platforms. Section 3 includes details on Bugcrowd’s platform submission workflow. We discuss the data in section 4. Section 5 looks at key empirical properties of the platform. In section 6, we examine the effect of the shock on the supply side, demand side, and equilibrium properties of the platform. Finally, section 7 provides brief conclusions.

2. Background

2.1. Vulnerabilities as a Product (VaaP)

The vulnerability (or “bug”) life cycle 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 security researcher, also known as a white-hat or ethical hacker. In such a case a patch that eliminates the vulnerability will be offered to all researchers of the affected 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.[4]

There are markets for vulnerabilities as a product (VaaP), both from the adversarial and defense perspectives. In this paper we focus only on the legal defensive market sometimes referred to as ”white market”,[5] rather than on the ”black market” for exploits.[6]

2.2. 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 (hackers) to find bugs in their products for a monetary and non-monetary incentive. While this is a feasible option for the largest firms, this is typically 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 groups of researchers 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

[2] The change in attack patterns following Covid-19 have been documented extensively by market analysts, cyber-security vendors and governmental agencies. See also Lallie et al. [2021].

[3] This slower response is possibly due in part to the time it takes for a firm to join the bug bounty platform and to get a bug bounty program up and running.

[4] 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.

[5] Of course, it is possible that the software firm finds some of its own vulnerabilities ex-post, i.e., after the code is written. In such a case there is no ‘white hat hacker” market. See Choi et al. [2010] for a theoretical model that addresses the setting in which the firm finds its own vulnerabilities ex-post.

[6] 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].

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 researchers on the other side.[7]

Bug Bounty Platforms are two-sided markets and they connect organizations that want to crowd-source their software security with ethical hackers (or security researchers). Ideally, a platform hosts many programs for multiple organizations and has many high quality researchers. The researcher who first finds a novel vulnerability report receives a payment (bounty).

Bug Bounty Platforms 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) and are interested in protecting against exploits by adversaries. Finding vulnerabilities in this market is done via crowd-sourcing. The supply side of the market consists of researchers ("white hat hackers") who are eager to get paid for their expertise.

A security researcher in this context is a skilled computer expert/hacker who uses his/her technical knowledge to identify vulnerabilities for rewards. The magnitude of the paid bounties depends on the severity of the vulnerability found, and in addition to monetary payments, the researchers are rewarded with reputation points which determine their relative rank and may enable them to receive invitations to work in private bounty programs.

2.3. Literature on Bug Bounty Programs

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 percent of the bounty programs activity on Bugcrowd's researched platform during the full calendar years we examine. See Table A1.

3. The Bugcrowd Platform

Bugcrowd's Bug Bounty program rewards hackers for valid accepted bug submissions with a monetary reward. Only the first researcher who finds a vulnerability qualifies for a monetary reward (in addition to points). Later researchers who find the same bug receive points (but no monetary award) for their correct "duplicate" submission.[8] The rest of this section will detail the submission workflow and the bounty pricing dynamics of the researched Crowdcontrol^{TM} platform.

3.1. Submission Workflow

The rules of engagement between a hacker and an organization on Bugcrowd's platform are structured to benefit both sides: on one hand they encourage researchers to practice responsible disclosure of high value vulnerabilities, and on the other hand ensure the timely response and payment of organizations once a valid bug has been submitted.

Figure 2 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?), selective researchers' access (public or private?), the payment range per vulnerability (by priority) and more. Once the program is launched, organizations have their teams ready to process the incoming submissions, after they have been triaged (prioritized) and screened for duplicates and relevancy by the platform's team. Valid vulnerabilities are then integrated into the existing Software Development Lifecycle (SDLC) tools to be fixed, and related reward payouts are processed accordingly.

4. Data

The paper employs a unique data-set obtained from Bugcrowd.[9] The data spans the full period for which the company has been in existence, i.e., from 2013 and includes all bug submission activity through May 2021. Since we are interested in the effect of the Covid-19 shock, we primarily will focus on data from 2017-2021.

The data covers not only public programs open to any researcher, but also private bounty programs which account for

[7] New platforms are often confronted with the problem that both sides will only enter the platform market 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 the literature on platforms [Caillaud and Jullien, 2003].

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

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

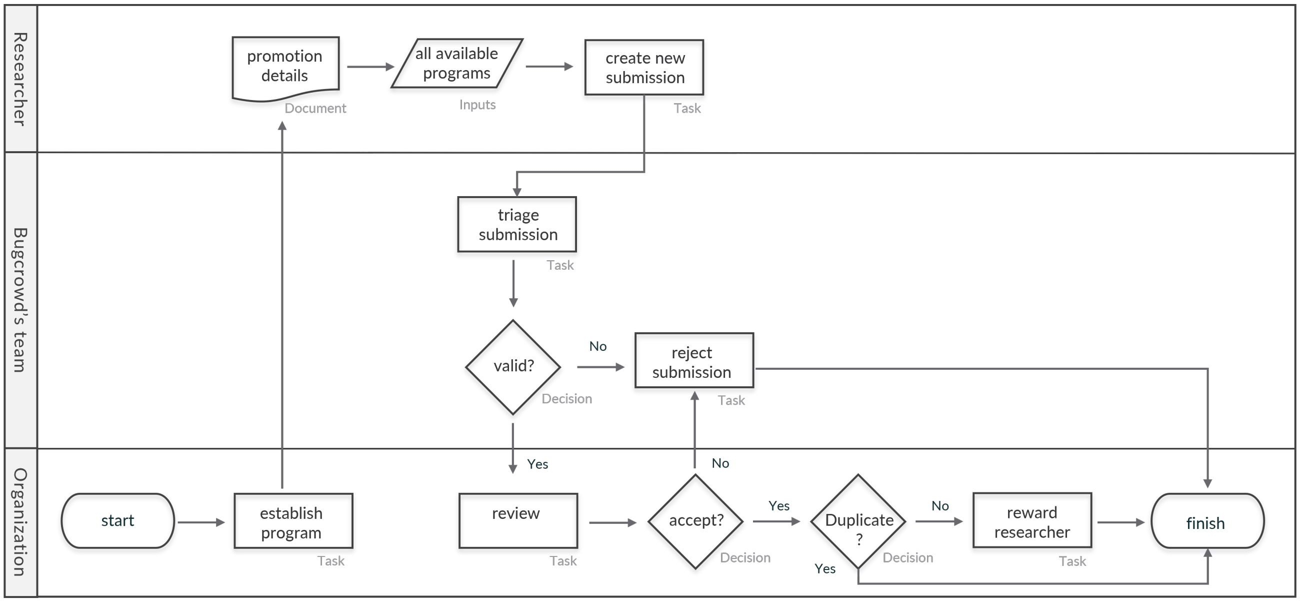

Figure 2: Bugcrowd's platform submission workflow. In some cases a vulnerability can be marked duplicate during the triage stage and eventually be rejected (based on `https://docs.bugcrowd.com/customers/getting-started/with-bugcrowd/`)

the majority of the new vulnerability programs.[10] The data includes complete information for every public and private bug bounty program during this period since the platform has been in existence. Only selected successful researchers are invited to participate in private programs, while public programs are available to every registered researcher.

Up until now, detailed data from such a platform at the level of the submission by each researcher was not available to scholars. Previous scholars have had some access to public bounty programs. Given the increasing importance of private bug bounty programs, it is important to analyze data from these programs as well. Our data was obtained through a Data Transfer Agreement (DTA) between Tel Aviv University and Bugcrowd.

The data set contains information on the demand side (organization / program), the supply side (researchers), and the product (bug submissions). The organizational data includes the firm size, country of origin, and when it first joined the platform. Many firms run simultaneously more than one project (bug bounty program) and for each we have its status, start/end dates, and whether it is open to everyone or only to selected researchers (a public or a private program). Data on the researcher includes characteristics such as country of origin, day of first submission, relative rank (reflecting successes), and whether can participate in private programs. The data on submissions specifies the actual bug submission and its outcome (correct, not correct, duplicate). It details the submission timeline and processing dates, bug severity, and the program type (public or private), the monetary award (if any), and the points (which impact the researcher's rank).

The unique identifiers for each submission, researcher, organization, and program allow the formation of a panel data set at the level of a submission. For each submission, we thus have the following data:

- We know which researcher made the submission.
- When the submission was made (day/time).
- To which bug bounty program the submission was made.
- Whether the submission was correct, i.e., did it find a valid vulnerability within the program's scope.
- The amount in USD paid for each vulnerability.
- Whether the submission was an accepted "duplicate", i.e., it found a valid vulnerability, but someone else had submitted it previously.
- The amount of points awarded. Bugcrowd has an explicit point system.
- The country in which the security researcher is located.
- The location (country) of the organization which initiated each program.

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 lockdowns during this time. Thus the Covid shock was strongest during this period. Hence, we define the March 1 to May 31 2020 period to be the "Covid" period.

In order to exclude seasonality effects, we include the same period in all other years as well. Thus, the analysis in the paper is for the three full months from March 1 to May 31 unless noted otherwise.

[10] More than 90 percent of the new programs during the 2017-2021 period are private programs.

5. Key Properties of the Platform

Before we examine the results associated with the Covid shock, we discuss several key empirical properties of the platform based on the full year data 2017-2020.

Observation 1 *Organizations from the United States generate most of the demand (the bug bounty programs), while most of the active researchers are from India.*

Over the years, the percentage of active researchers and submissions from Asian countries grew steadily. In 2020 it reached 58 percent of the active researchers and 70 percent of the submissions, with India as the most dominant country in this region (46% of all active researchers, and 60% of all submissions in 2020).

From the demand perspective, organizations based in the United States operate the vast majority of programs, with 68 percent of active programs in 2020, and similar values in other time periods as well. See Table A2.

These properties are consistent with the findings from “The Online Labour Index”, an economic indicator that provides an online gig economy equivalent to conventional labor market statistics.[11] 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 come from India, and 15 percent come from Bangladesh[12]

Observation 2 *Organisations are moving their activity away from public programs to private programs.*

The evolution of programs (the demand side) over time detailed in Table A1 shows a clear trend towards private programs. In 2017 public programs accounted for nearly 29 percent of all active programs, while in 2021, such programs accounted for 12 percent only. The share of new private programs launched over the years grew from 84 percent in 2017 to 97 percent in 2020.

Observation 3 *Although the number of active researchers is growing steadily, the researcher attrition rate (no submissions) for the platform is very high.*

We categorized researchers by the year of their first submission period and analyzed their attrition rate between that period and subsequent years. For researchers who joined in 2017-2020, the attrition rate between the first and second years is between 71 and 76 percent. In addition, the attrition rate between the second and third years is 30-35 percent. See Table A3.

Observation 4 *For researchers who remain on the platform, the submission success rate, as measured by paid submissions divided by total submissions, increases over time.*

In Table A4 we categorized researchers by the year of their first submission period. We then calculated their success rate during that year and subsequent years. The table shows that for all cohorts of researchers (2017, 2018, 2019, and 2020), there was a continuous rise in the submission success rate each year.[13]

Observation 5 *There is very little (if any) market power on either side of the market.*

On the supply side, Table A5 shows that the percent of submissions and total rewards earned by the top 100 researchers according to Bugcrowd’s point-based rank methodology, has been declining over time. In 2017, the top 100 researchers (at that point in time) earned 43 percent of the total rewards. In 2020, the top 100 researchers (at that point in time) earned only 24 percent of the total rewards. Thus the top 100 researchers (which is already a large number) earn a very small portion of the rewards on the platform and this portion has been declining. Figure 3 shows the total reward share of the top 100 researchers.

On the demand side, Figure 3 shows the reward share paid by the top 10 programs. These programs accounted on average for 43 percent of the yearly payments in 2017 and 34 percent of the total payments in 2018. In 2019 and 2020, the top 10 programs accounted for just 23-24 percent of the rewards in 2020. Thus, most of the programs pay small awards relative to the total awards earned on the platform and this percentage has been falling over time.

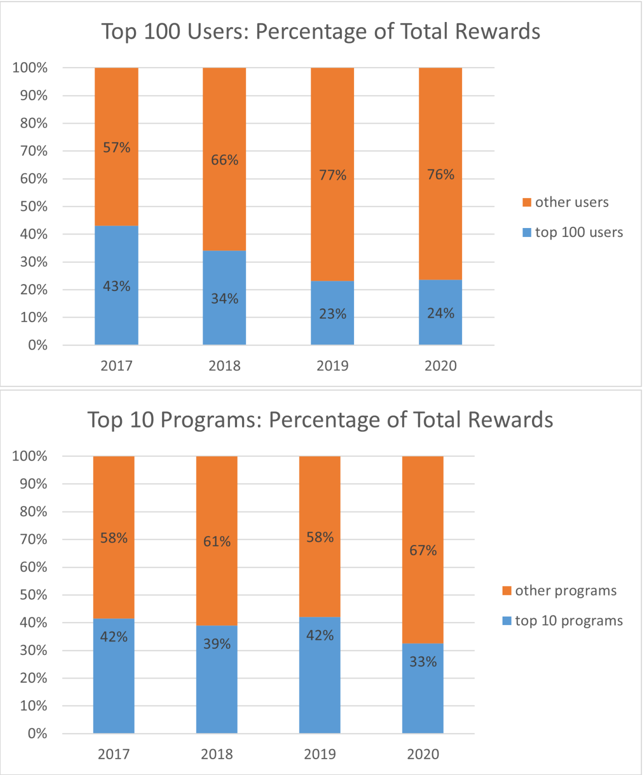

Figure 3: A yearly view on total rewards share of top 100 researchers (supply) and top 10 programs (demand) in Bugcrowd’s platform.

[11] 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/`.

[12] OLI 2020 update: `https://ilabour.oii.ox.ac.uk/onlinelabourindex2020/`.

[13] There is one exception: for 2017 researchers, the submission success rate declined slightly after 2019.

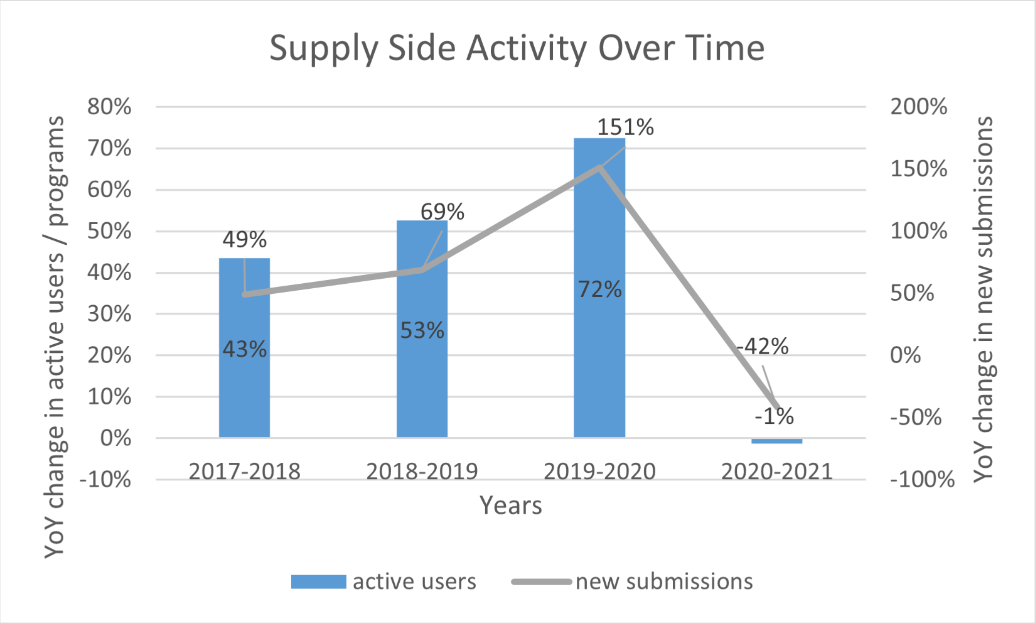

Figure 4: Year Over Year (YoY) changes in active researchers and new submissions, between the 3-month periods of 2017-2021 in Bugcrowd's platform.

6. Effect of Covid-19 Exogenous Shock

In this section, we explore the effects of the shock on the supply and demand sides of the market and examine the equilibrium changes in the platform as well. This analysis is conducted using the 3-month period data (March 1 - May 31) for 2017-2021.

6.1. The Supply Side

Observation 6 *The Covid-19 shock led to a significant increase in the number of researchers, and an especially large increase in the number of total submissions.*

Figure 4 shows the Year-Over-Year (YoY) change in the number of researchers and submissions during the three-month periods of 2017-2021. While the platform experienced steady growth on the supply side from 2017-2019, supply side activity increased significantly during Covid 2020 period:

- In terms of numbers of submissions, there were 53,098 submissions in the relevant three month period of 2020, versus 21,157 in the same period of 2019. This 151 percent growth during the Covid period is much larger than the changes between 2017-2018 (49 percent), and 2018-2019 (69 percent). In the subsequent 2021 period year, this number dropped to 30,945 submissions.
- This huge increase in the number of submissions was primarily driven by 2020 researchers, who made 20,118 submissions (38 percent of total submissions). By comparison, in 2019 and 2021, there were only between 6,000 and 7,000 submission made by new researchers.
- When comparing 2019 and 2020, there was a 72 percent growth in the number of researchers who made at least one submission (denoted active researchers), versus a 43 percent growth between 2017-2018 and a 53 percent growth between 2018-2019.

While the large increase in the number of researchers and total submissions is important, the key metric of interest on the supply side is the quality of new researchers and the number of correct submissions.[14]

Observation 7 *The number of correct submissions grew dramatically during the Covid shock in 2020. Further, there was a large increase in the percentage of "accepted duplicates", relative to the total number of submissions. These increases were primarily due to the new researchers who joined the platform in 2020.*

- As shown in Figure 5, there were 25,864 correct submissions in the relevant three month period of 2020, versus 9,525 in the same period of 2019. This 172 percent growth during the Covid period dwarfs the changes between 2017-2018 (35 percent), and 2018-2019 (64 percent). In the 2021 period, this number dropped to 10,195 correct submissions, virtually the same number as in 2019.
- This huge increase in the number of correct submissions was primarily driven by 2020 researchers, who made 11,441 correct submissions. In comparison, new researchers in 2019 made 2,249 correct submissions in that year. In 2021, new researchers made 1,492 correct submissions. See Table A4.
- The Covid shock led to steep increase in the percentage of "accepted duplicates", relative to the total number of submissions.The percent of accepted duplicate submissions out

[14] Recall that Bugcrowd keeps track of "accepted duplicates", which are (I) submissions that correctly identify vulnerabilities but (II) were not the first submission to identify that vulnerability. That is, another researcher has made a correct submission previously regarding that vulnerability, and hence received the monetary award for that particular contribution.

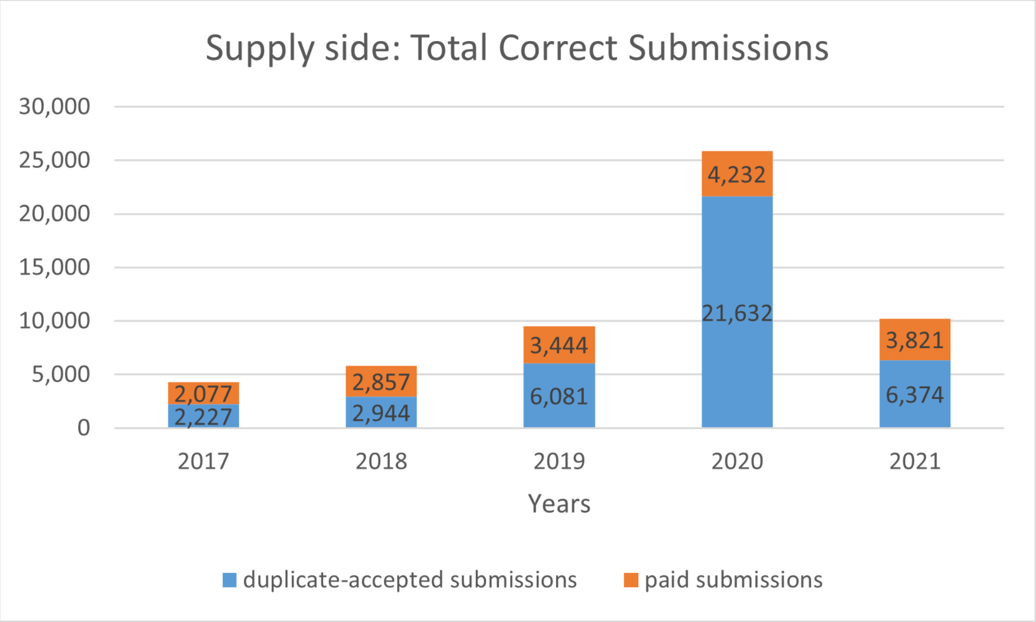

Figure 5: Total correct submissions (paid and correct duplicates), during the 3-month periods of 2017-2021.

of total submissions grew from an average of slightly above 26% during the 2017-2019 period, to 41 percent in 2020. See Table A6.

- For new researchers in 2020, the ratio of accepted duplicates to total submissions was 55 percent, well above the same ratio in that year for other cohorts. The corresponding percentages in 2020 for other cohorts were as follows: 25% correct duplicates for 2017 researchers; 33% duplicates for 2018 researchers; and 44% duplicates for 2019 researchers. See Table A4.

6.1.1. Public vs. Private Bug Bounty Programs

We now examine public programs and private programs separately, since new researchers typically cannot make submissions to private programs in their first few months on the platform (recall that participating in private programs requires an "invitation.")

Observation 8 *There was a very large increase in the percent of correct duplicates submitted to public programs. There was no such increase for private bug bounty programs.*

As shown in Table A6, during the Covid period, accepted duplicates in public bounty programs totaled 16,284 out of 35,554 total submissions. Thus the percent of correct duplicates was 46 percent in 2020. By comparison, from 2017 to 2019, the percent of correct duplicates in public programs ranged from 22-31 percent. In 2021, the percent of correct duplicates in public programs fell to 19 percent.

By comparison, in the case of private bug bounty programs, from 2017 to 2019, the percent of correct duplicates ranged from 22-32 percent. During the Covid period, this percentage was 30 percent. In 2021, the percent of correct duplicates in private programs was 22 percent.

Observation 9 *There was a huge increase in the number of submissions from India and Turkey during the Covid period, followed by a significant decline in the subsequent period of 2021.*

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.

6.1.2. Intensity of Work

Did the "intensity of work" change during the Covid-19 period? We define intense work to be three or more submissions of the same researcher to a particular bounty program during the same day. We created an "intensity" dummy variable that equaled one if the submission was part of "intense activity". Otherwise, it was equal to zero. We then calculated the intensity rate as the percent of submissions that are part of an intense activity effort divided by the total submissions.

Observation 10 *The work intensity of researchers in public programs increased during the Covid period.*

The ratio of intense submissions out of total submissions grew from 31 percent in pre-Covid 2019 to 37 percent in Covid 2020 period, and dropped back to 28 percent in the following post-Covid 2021 period. These changes are primarily in public programs. For private programs, the intensity rate was the same for 2019, 2020, and 2021. See Table A7.

Since most new researchers were not able to compete in private programs, the increase in intense activity in public programs in 2020 reflects the behavior of the new researchers, many of whom had fewer outside opportunities during that period. These researchers not only submitted more, but also worked more intensely during Covid-19, implying they had more time at hand to focus on the bug bounty task.

6.1.3. The Aftermath on the Supply Side

Observation 11

Those who joined the platform in 2020 made many fewer submissions in 2021, and this decline was exceptional relative to other cohorts in their second year.

In the post Covid period of 2021, 2020 researchers submitted significantly less than any other researcher group in their second year on the platform. Figure 6 shows the Year Over Year (YoY) changes in new submissions by the year the researcher joined. Researchers who joined the platform during Covid 2020 made 89 percent fewer submission in the following post-Covid 2021 period. The submission decline for other cohorts in their second year relative to their first year dropped only by 38 percent (2017 researchers), 43 percent (2018 researchers) and 62 percent (2019 researchers). Thus many of the new researchers in the Covid-19 period left the platform as other opportunities returned.

6.2. The Demand Side

We now briefly examine the demand side of the platform.

Observation 12 *Relative to previous years, there was just slightly more growth in demand (in terms of the number of bug bounty programs) during the Covid period than during previous periods.*

Table A8 shows the changes in active and new programs across the three-month periods between 2017-2021. There was a 67 percent increased in the number of bug bounty programs between 2019-2020, which was above the 42 percent growth between 2017-2018 and a 41 percent growth between 2018-2019. But the increase in programs in 2021 was just 27 percent. Thus, overall the yearly growth between 2019 and 2021 in the number of bug bounty programs is slightly below 46 percent, virtually unchanged from the 2017-2019 period.

As opposed to the decline in active researchers and submissions in the post-Covid 2021 period, the increase in the number of active programs is likely due in part to the long setup time required for a firm to plan, budget, and execute a program. Joining the platform and submitting a bug report is an instant and cost-less effort for a new researcher, while it takes much longer for a program to launch. Organizations may have started their internal process for introducing a new program during the three-month period in 2020, but the program did not begin during that period.

6.3. Equilibrium Changes in the Platform

In this section we show the changes in the average price for vulnerabilities (average rewards) and the intense competitiveness during Covid 2020 period.

Observation 13 *The huge supply shock immediately following Covid led to a steep decline in the expected average payment for correct submissions.*

The ratio of paid submissions to total correct submissions fell dramatically during the Covid shock. This percentage was 48 in 2017, 49 in 2018 and 36 in 2019. In 2020, the ratio of paid submissions out of total correct submissions was just 16 percent. It rose to 37 percent in 2021. See Figure 7.

The large rightward shift of the supply curve as a result of the shock thus led to a dramatic fall in the expected price (award) per correct submission. The expected average award per correct submission ranged from $341 - $404 during the 2017-2019 period, but fell dramatically to just $122 in 2020. The expected average award rose to $319 in 2021. See Figure 8.

Observation 14 *There was little difference in the dramatic decline in the expected average award per correct submission between high priority and low priority vulnerabilities in 2020.*

In Bugcrowd's platform researcher submissions are categorized by a priority scale of P1 to P5 (critical to non-exploitable weaknesses accordingly), based on a well-defined Vulnerability Rating Taxonomy (VRT).[15]

We then delineated our analysis by the importance of the vulnerability. We grouped P1 and P2 vulnerabilities into the high priority category and P3-P5 vulnerabilities into the low priority category.[16]

In the pre-shock periods of 2017-2019, expected average rewards per correct submissions ranged from $121-$124 for low

[15] A resource outlining Bugcrowd's baseline priority rating, including certain edge cases, for common vulnerabilities: `https://bugcrowd.com/vulnerability-rating-taxonomy/`.

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

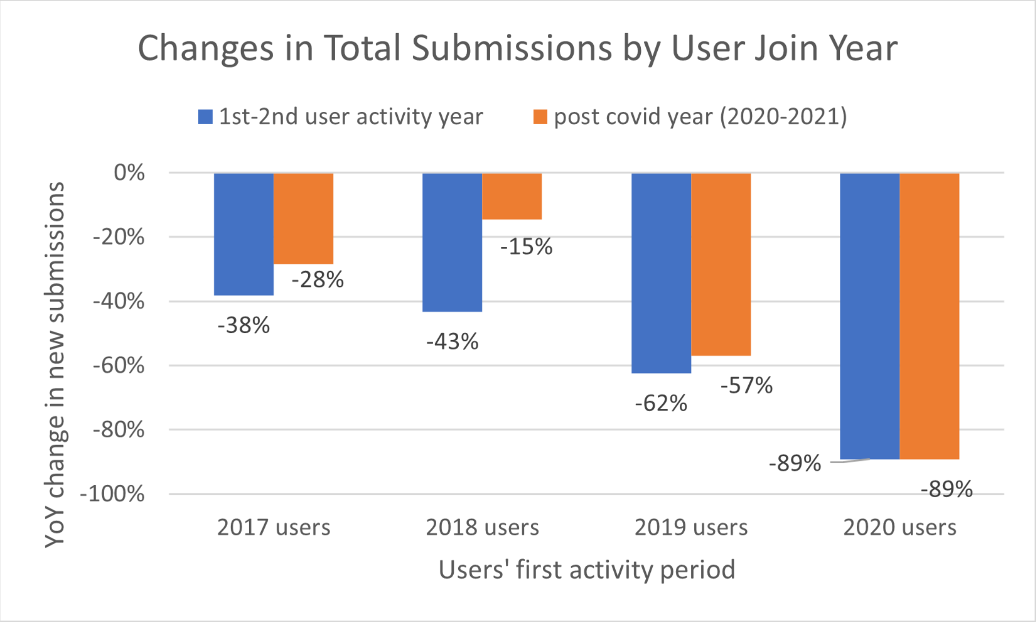

Figure 6: Year Over Year (YoY) changes in new submissions, looking at researchers who joined Bugcrowd's Bug Bounty platform during the 3-month periods of 2017-2020. Changes were measured between the first year to the second year of each researcher group activity, and between the Covid 2020 and post Covid 2021 3-months periods.

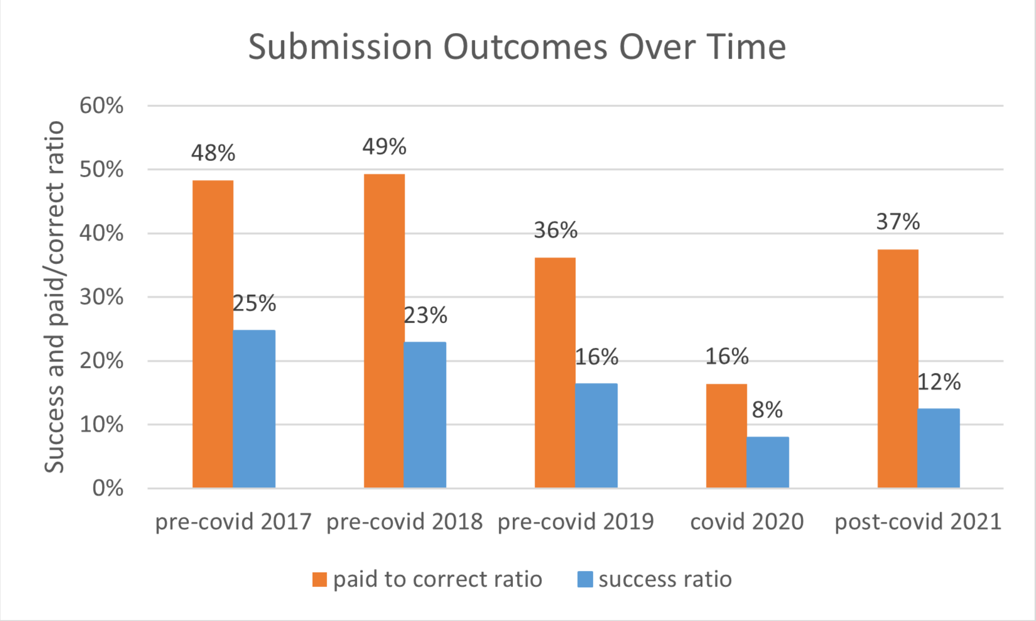

Figure 7: Success rate and paid-to-correct ratio, during the 3-month periods of 2017-2021 in Bugcrowd's Bug Bounty platform.

priority vulnerabilities. During the Covid shock, this expected payment decline to $42, before rising in 2021 to $121.

There was a similar decline in the expected average reward for high priority vulnerabilities. In the case of high priority submissions, the expected average reward per correct submissions was $1520 in 2019, but fell to just $629 in 2020, before increasing to $1,255 in 2021.

Observation 15 *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.*

Overall, the number of paid submissions rose by approximately 23 percent from 3,444 in 2019 to 4,232 in 2020. The percentage increase in the number of paid submissions between

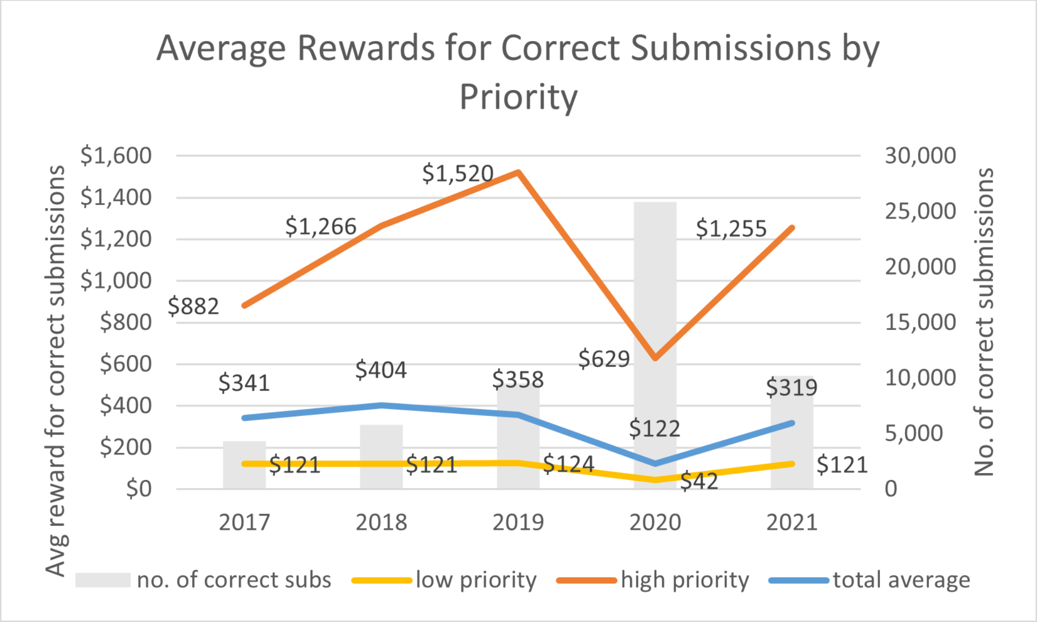

Figure 8: Total correct submissions, and average rewards by submission priority during the 3-month periods of 2017-2021 in Bugcrowd's Bug Bounty platform.

2019 and 2020 is thus very similar to the percentage increase (21 percent) in the number of paid submissions between 2018 and 2019.

6.4. Policy Issues

Given the number and correct submission activity of the high quality researchers who joined the platform in 2020, there was a missed opportunity to examine more organizational software.

If the demand response to the shock had been similar to the supply response such that ratio of paid to correct submissions of 36-37 percent in 2019 and 2021, rather than the 16 percent during 2020, the number of unique vulnerabilities identified by researchers in 2020 would have been more than double the actual number in 2020.

This "counterfactual" seems reasonable since (i) the quality of the new researchers in 2020 was very high and (ii) all software contains vulnerabilities.[17]

7. Discussion and Conclusion

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.

In this paper, we provided background on and examined the effect of the Covid-19 pandemic's exogenous shock on the Bugcrowd bug bounty platform. We found that there was an immediate and very large effect on the supply side. We quantified what happens when outside options are reduced and/or hackers find themselves with more time on their hands. In terms of the gig economy, the Covid shock enabled many skilled researchers (primarily from India and Turkey) an opportunity to try to earn rewards from finding vulnerabilities, and many took advantage of this opportunity. On the demand side, the relatively small response to the exogenous shock appears to have been a missed opportunity, given the large number of high quality researchers working on finding vulnerabilities.

Past research has suggested 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 (and larger) bug bounty platforms could change this dynamic. In such a case, more transactions would take place in the white market rather than the black or grey markets.

8. Acknowledgments

We thank Casey Ellis, the CTO and founder of Bugcrowd, and his phenomenal team for sharing with us their data-set and adding context to our findings.

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`.

[17] This assumes that there are always more vulnerabilities to find. Many researchers 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." Quote from `https://www.sigarch.org/lets-keep-it-to-ourselves-dont-disclose-vulnerabilities/`.

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.

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.

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.

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. Supporting tables

Table A1. Program perspective: changes in active and new programs across full calendar year periods.

program type	No. of active programs[1]				Percentage of active programs			
	2017	2018	2019	2020	2017	2018	2019	2020
public	105	132	145	151	29%	27%	22%	12%
private	252	357	525	1,155	71%	73%	78%	88%
total	357	489	670	1,306	100%	100%	100%	100%
program type	**No. of new programs[2]**				**Percentage of new programs**			
	2017	2018	2019	2020	2017	2018	2019	2020
public	35	33	33	25	16%	12%	9%	3%
private	189	243	336	821	84%	88%	91%	97%
total	224	276	369	846	100%	100%	100%	100%

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

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

Table A2. Top countries view: geographical diversity of active researchers, active programs, and submissions, across full calendar year period.

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%

[1] Demand side top country.

[2] Supply side top country.

Table A3. Researcher perspective: changes in active researchers and attrition rate across full calendar year periods.

Year researcher joined	Active researchers[1]				Researcher attrition rate		
	2017	2018	2019	2020	2017-2018	2018-2019	2019-2020
2012 researchers	3	4	3	2	33%	-25%	-33%
2013 researchers	112	91	77	80	-19%	-15%	4%
2014 researchers	193	162	131	124	-16%	-19%	-5%
2015 researchers	257	213	178	133	-17%	-16%	-25%
2016 researchers	609	396	338	291	-35%	-15%	-14%
2017 researchers	2,650	648	454	386	-76%	-30%	-15%
2018 researchers		3,969	1,069	721		-73%	-33%
2019 researchers			5,973	1,753			-71%
2020 researchers				9,553			
total	3,824	5,483	8,223	13,043	43%	50%	59%

[1] Researchers who submitted at least once during the full calendar year period.

Table A4. Researcher experience: submissions and rewards by year researcher joined, across the three-months periods.

researcher join year	No. of submissions					Percentage of submissions from yearly total				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
2017 researchers	2,988	1,848	1,747	1,671	1,196	36%	15%	8%	3%	4%
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 paid submissions[1]					Submission success rate[2]				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
2017 researchers	546	553	601	545	317	18%	30%	34%	33%	27%
2018 researchers		283	271	236	276		9%	15%	19%	25%
2019 researchers			329	213	197			5%	9%	20%
2020 researchers				315	334				2%	15%
2021 researchers					187					3%
Year researcher joined	No. of accepted-duplicate submissions[3]					Accepted-duplicates / total submissions ratio				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
2017 researchers	679	446	536	419	325	23%	24%	31%	25%	27%
2018 researchers		801	584	417	292		25%	32%	33%	27%
2019 researchers			1,920	1,013	240			31%	44%	24%
2020 researchers				11,126	563				55%	26%
2021 researchers					1,305					19%
Year researcher joined	No. of correct submissions[4]					Correct submissions / total submissions ratio				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
2017 researchers	1,225	999	1,137	964	642	41%	54%	65%	58%	54%
2018 researchers		1,084	855	653	568		34%	47%	51%	52%
2019 researchers			2,249	1,226	437			37%	53%	44%
2020 researchers				11,441	897				57%	41%
2021 researchers					1,492					22%

[1] Paid submissions are those which were both accepted and rewarded.

[2] Submission success rate is defined as the number of submissions that received monetary awards divided by the total submissions.

[3] Accepted-duplicates are correctly identified vulnerabilities which were already submitted by another researcher.

[3] Correct submissions are the sum of paid submissions and accepted-duplicates.

Table A5. Rank analysis: total submissions and rewards by researcher rank, across the full calendar year.

top 100 and top 500[1]	percentage of total submissions				percentage of total rewards			
	2017	2018	2019	2020	2017	2018	2019	2020
top 100	28%	19%	8%	6%	43%	34%	23%	24%
top 500	48%	36%	22%	17%	63%	53%	43%	48%
all ranks[2]	percentage of total submissions				percentage of total rewards			
	2017	2018	2019	2020	2017	2018	2019	2020
top 1000	54%	43%	28%	23%	71%	59%	50%	57%
rank > 1000	15%	20%	20%	28%	16%	19%	25%	26%
un-ranked (new)	31%	37%	52%	49%	14%	22%	25%	16%

[1] Top 100 and top 500 share of total submissions and total rewards (rank categories overlap).

[2] Total submissions and total rewards percentage share of three rank categories: ranks 1-1000, ranks > 1000, no rank (new researchers).

Table A6. Program perspective: changes in submissions per program type across the three-month periods.

program type	No. of submissions					Percentage of submissions				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
public	4,701	7,573	13,893	35,554	15,525	56%	60%	66%	67%	50%
private	3,711	4,954	7,264	17,544	15,430	44%	40%	34%	33%	50%
total	8,412	12,527	21,157	53,098	30,955	100%	100%	100%	100%	100%
program type	No. of accepted-duplicate submissions[1]					Ratio of accepted-duplicate / total submissions				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
public	1,046	1,862	4,308	16,284	3,015	22%	25%	31%	46%	19%
private	1,181	1,082	1,773	5,348	3,359	32%	22%	24%	30%	22%
total	2,227	2,944	6,081	21,632	6,374	26%	24%	29%	41%	21%
program type	No. of paid submissions[2]					Average rewards per paid submission				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
public	709	823	1,109	1,145	850	$705	$847	$745	$702	$795
private	1,368	2,034	2,335	3,087	2,971	$708	$809	$1,105	$762	$867
total	2,077	2,857	3,444	4,232	3,821	$707	$820	$989	$746	$851

[1] Duplicate and accepted submissions are technically correct submissions which were already submitted by another researcher, and therefore were not rewarded in cash.

[2] Paid submissions are those which were both accepted and rewarded. Accepted submissions which were already submitted by another researcher are considered duplicate and therefore are not rewarded.

Table A7. Researcher perspective: changes in researchers' intense activity per program type, across the three-month periods.

program type	No. of total submissions					YoY change of total submissions			
	2017	2018	2019	2020	2021	2017-2018	2018-2019	2019-2020	2020-2021
public	4,701	7,573	13,893	35,554	15,525	61%	83%	156%	-56%
private	3,711	4,954	7,264	17,544	15,430	33%	47%	142%	-12%
total	8,412	12,527	21,157	53,098	30,955	49%	69%	151%	-42%
program type	No. of intense activity submissions[1]					YoY change of intense submissions			
	2017	2018	2019	2020	2021	2017-2018	2018-2019	2019-2020	2020-2021
public	1,200	1,940	3,841	12,889	3,064	62%	98%	236%	-76%
private	1,718	2,238	2,715	6,518	5,702	30%	21%	140%	-13%
total	2,918	4,178	6,556	19,407	8,766	43%	57%	196%	-55%
program type	Intensity rate[2]					YoY change of intensity rate			
	2017	2018	2019	2020	2021	2017-2018	2018-2019	2019-2020	2020-2021
public	26%	26%	28%	36%	20%	0%	8%	31%	-46%
private	46%	45%	37%	37%	37%	-2%	-17%	-1%	-1%
total	35%	33%	31%	37%	28%	-4%	-7%	18%	-23%

[1] Intense activity is defined as a focused effort of the same researcher submitting three or more vulnerabilities to the same program on the same day.

[2] Intensity rate is defined as the percent of submissions that are part of an intense activity effort divided by the total submissions.

Table A8. Program perspective: changes in active and new programs across the three-months periods.

program type	No. of active programs[1]					percentage of active programs[2]				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
public	83	107	118	131	148	43%	39%	30%	20%	18%
private	110	168	271	517	677	57%	61%	70%	80%	82%
total	193	275	389	648	825	100%	100%	100%	100%	100%
program type	No. of new programs[2]					perentage of new programs				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
public	7	7	9	5	6	14%	10%	9%	3%	3%
private	43	66	86	186	174	86%	90%	91%	97%	97%
total	50	73	95	191	180	100%	100%	100%	100%	100%

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

[2] Programs with first submission during the period.

Table A9. Submission perspective: changes in paid submissions per submission priority, across the three-month periods.

priority[1]	No. of submissions					percentage of submissions				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
low priority	5,263	7,833	13,696	37,125	20,067	63%	63%	65%	70%	65%
high priority	1,719	2,281	3,172	8,459	4,752	20%	18%	15%	16%	15%
no priority	1,430	2,413	4,289	7,514	6,136	17%	19%	20%	14%	20%
total	8,412	12,527	21,157	53,098	30,955	100%	100%	100%	100%	100%
priority[1]	No. of duplicate accepted submissions[2]					Success rate[3]				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
low priority	1,654	2,415	5,229	19,299	5,481	25%	23%	16%	8%	14%
high priority	519	474	713	2,182	823	41%	40%	31%	15%	20%
no priority	54	55	139	151	70	5%	6%	5%	1%	0%
total	2,227	2,944	6,081	21,632	6,374	25%	23%	16%	8%	12%
priority[1]	No. of paid submissions					Competitiveness level[4]				
	2017	2018	2019	2020	2021	2017	2018	2019	2020	2021
low priority	1,312	1,808	2,235	2,885	2,853	56%	57%	70%	87%	66%
high priority	699	911	999	1,244	961	43%	34%	42%	64%	46%
no priority	66	138	210	103	7	45%	28%	40%	59%	91%
total	2,077	2,857	3,444	4,232	3,821	52%	51%	64%	84%	63%

[1] A small number of rewarded submissions with no priority assigned to them by Bugcrowd were mapped into high-priority or low-priority categories when possible, based on their value compared to the average rewarded amount per priority on the same program during the same calendar year.

[2] Duplicate and accepted submissions are technically correct submissions which were already submitted by another researcher, and therefore were not rewarded in cash.

[3] Submission success rate is defined as the number of submissions that received monetary awards divided by the total submissions.

[4] Defined as "1-probability of receiving a monetary award when submitting a valid vulnerability".