

Understanding Incentivized Mobile App Installs on Google Play Store

Shehroze Farooqi
University of Iowa/ICSI

Álvaro Feal
IMDEA Networks Institute /
Universidad Carlos III de Madrid

Tobias Lauinger
NYU

Damon McCoy
NYU

Zubair Shafiq
University of California, Davis

Narseo Vallina-Rodriguez
IMDEA Networks Institute/ICSI

ABSTRACT

“Incentivized” advertising platforms allow mobile app developers to acquire new users by directly paying users to install and engage with mobile apps (e.g., create an account, make in-app purchases). Incentivized installs are banned by the Apple App Store and discouraged by the Google Play Store because they can manipulate app store metrics (e.g., install counts, appearance in top charts). Yet, many organizations still offer incentivized install services for Android apps. In this paper, we present the first study to understand the ecosystem of incentivized mobile app install campaigns in Android and its broader ramifications through a series of measurements. We identify incentivized install campaigns that require users to install an app and perform in-app tasks targeting manipulation of a wide variety of user engagement metrics (e.g., daily active users, user session lengths) and revenue. Our results suggest that these artificially inflated metrics can be effective in improving app store metrics as well as helping mobile app developers to attract funding from venture capitalists. Our study also indicates lax enforcement of the Google Play Store’s existing policies to prevent these behaviors. It further motivates the need for stricter policing of incentivized install campaigns. Our proposed measurements can also be leveraged by the Google Play Store to identify potential policy violations.

CCS CONCEPTS

• **Security and privacy** → *Domain-specific security and privacy architectures; Economics of security and privacy.*

KEYWORDS

Mobile apps, Incentivized Installs, Spam Installs, Reputation Manipulation

ACM Reference Format:

Shehroze Farooqi, Álvaro Feal, Tobias Lauinger, Damon McCoy, Zubair Shafiq, and Narseo Vallina-Rodriguez. 2020. Understanding Incentivized Mobile App Installs on Google Play Store. In *ACM Internet Measurement Conference (IMC '20)*, October 27–29, 2020, Virtual Event, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3419394.3423662>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '20, October 27–29, 2020, Virtual Event, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8138-3/20/10...\$15.00

<https://doi.org/10.1145/3419394.3423662>

Conference (IMC '20), October 27–29, 2020, Virtual Event, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3419394.3423662>

1 INTRODUCTION

Popular app stores such as the Google Play Store and the Apple App Store list over 2 million mobile apps [29, 38]. The proliferation of mobile apps has increased the competition among app developers to improve the “visibility” of their apps in app stores’ searches and top charts to acquire new users. Creating and retaining a solid user base is critical for maximizing the revenue of mobile apps, mostly through in-app advertising or in-app purchases [4, 6].¹ Furthermore, popular mobile apps with strong user engagement metrics (e.g., install counts, daily active users) and revenue will also be in a better position to attract funding from venture capitalists (VCs) [5, 35, 62].

Developers spent more than 14 billion dollars in 2019 on various types of advertising campaigns for promoting mobile app installs [56]. Oftentimes, this is done through traditional advertising-based models [16]. Another type of advertising is “incentivized” mobile app install campaigns, where users install an advertised mobile app and perform in-app tasks in exchange for monetary (e.g., gift card, PayPal balance) or non-monetary (e.g., in-app points or virtual currency) rewards [67]. Mobile app developers pay incentivized install platforms (or IIPs) to advertise an incentivized install offer, which is distributed to end users through a network of “affiliate” apps. These incentivized install campaigns are orders of magnitude cheaper in comparison to regular, non-incentivized mobile app install campaigns — on average, a mobile app install through incentivized advertising costs around \$0.06 (as we show later in the paper), as opposed to \$1.22 for non-incentivized installations [9].

The incentivized install ecosystem is highly sophisticated and opaque. Incentivized installs have a poor reputation since users (potentially “crowd workers”) are likely performing actions for a monetary reward rather than based on a genuine interest in an app. Furthermore, the ability to allocate specific in-app tasks to users (e.g., creating an account, completing a level of a game, performing in-app purchases) through incentivized install campaigns enables efficient manipulation of targeted user engagement metrics (e.g., number of active users, user session lengths) and revenues. These inflated user engagement metrics can help apps manipulate app store metrics (install counts, visibility), which might lead to potential violations of app store policies [2, 11]. In fact, app developers

¹Google Play Store and Apple App Store had an estimated revenue of 83 billion dollars in 2019 [54].

might abuse these schemes to deceptively monetize through in-app advertising or obtain VC funding based on inflated user engagement or revenue metrics. It is noteworthy that incentivized installs are banned by the Apple App Store and discouraged by the Google Play Store. Note that the Google Play Store allows their use as long as they are not used to manipulate app store metrics (e.g., install counts, visibility) [43]. As a result, IIPs often do not support the Apple App Store, but potentially generate millions of USD in revenue on the Google Play Store [7, 14].

It can be challenging for app stores to detect incentivized installs since these installs and user actions resemble that of authentic organic users [55]. Prior work [23, 47–49, 65, 66] has studied manipulation of app store metrics on the Google Play Store, focusing primarily on detecting installs driven by automation or by low-cost labor recruited outside of the app store. However, to the best of our knowledge, there has been no systematic investigation of incentivized installs. As a result, we lack a good understanding of the incentivized install ecosystem and the effectiveness of these campaigns to manipulate user engagement metrics. More importantly, it is unclear whether or not existing app store defenses are able to detect the manipulation of app store metrics (e.g., install counts, visibility).

To fill this gap, this paper studies the ecosystem of incentivized installs, its prevalence, and its potential impact. We narrow our focus on studying the ecosystem of incentivized installs offering monetary rewards on the Google Play Store. We measure the effectiveness of these campaigns and shed light on app stores’ existing defenses [42, 43] to detect incentivized installs. First, we perform measurements by purchasing incentivized installs from multiple IIPs for our purpose-built honey app that we publish on the Google Play Store. This allows us to understand the behavior and effectiveness of incentivized installs through the lens of mobile app developers and to evaluate existing enforcement by Google Play Store. Then, we perform longitudinal measurements to monitor incentivized install campaigns of mobile apps on IIPs over a period of three months, at scale. This complementary perspective allows us to characterize incentivized install campaigns of advertised mobile apps, and to understand the potential benefits of running these campaigns.

The main contributions of our study are the following:

- (1) Our measurements suggest that most of the users are crowd workers installing apps advertised on IIPs to earn money. The results suggest a lack of enforcement from the Google Play Store to detect these practices since our purchased installs increased the total install count of our honey app on the Google Play Store from 0 to over 1,000.
- (2) Through longitudinal measurements, we identify campaigns of 922 apps that advertised a variety of incentivized install offers. We find that more than half of these offers require users to perform in-app tasks targeting manipulation of a wide variety of user engagement metrics (e.g., install counts, increase user session lengths) and revenues. These behaviors seem to contradict Google Play Store’s policies.
- (3) Our results show that there is a correlation between apps advertised on IIPs and the improvement of app store metrics on the Google Play Store. In comparison to a baseline dataset of

regular Android apps, apps advertised on IIPs show up to 8 \times increase in install counts and appear in top charts up to 2 \times more frequently. Our results also demonstrate that there is a correlation between advertising apps on IIPs and developers of advertised apps raising funding: apps advertised on IIPs tend to raise funding up to 2 \times more frequently.

- (4) Finally, we find that artificial in-app activity generated through incentivized install campaigns provides opportunities to developers of mobile apps to artificially increase gross revenue through arbitrage and advertising. For example, more than 60% of apps requiring users to perform in-app tasks integrate 5 or more advertising libraries.

Our results show that IIPs can have negative effects on the app store and beyond (e.g., deceptively raise funding and erode consumer trust) due to their ability to manipulate user engagement metrics and revenue. In fact, some IIPs even openly advertise services that seem to violate Google Play Store’s policies. Our study motivates the need for stricter policies against IIPs, similar to those implemented in the Apple App Store [22, 53]. Our measurement methodology to monitor IIPs can also be leveraged by the Google Play Store to identify potential policy violations.

2 INCENTIVIZED INSTALLS

In this section, we explain the ecosystem behind incentivized install campaigns and its dynamics. The ecosystem of incentivized mobile app installs involves developers, incentivized install platforms (IIPs), affiliate apps, users, and third-party mediators. Figure 1 describes the flow of the offer and money in an incentivized install campaign, highlighting the role of different stakeholders.

2.1 Stakeholders

App developer. A mobile app developer (or a third-party like marketing companies hired to improve the usage metrics of an app) runs an incentivized install campaign by paying an IIP (e.g., fyber.com, ayetstudios.com, offertoro.com) to advertise an offer. The offer contains app details (name and app store URL), a payout, and a description of actions (e.g., “Install and Register,” “Install and Reach level 10”) that a user must perform to complete the offer. This offer is then distributed through affiliate apps, where users browse offers and select an offer to work on. Once a user completes an offer and the offer completion is certified by a third-party mediator, the payment is disbursed to all the stakeholders.

Incentivized Install Platforms (IIPs). IIPs provide an easy and organized way for app developers to manage and run their incentivized advertising campaigns. An IIP aggregates offers from developers in the form of an *offer wall*. These offers are then advertised to a network of affiliate apps that integrate the offer wall of the IIP. IIPs allow developers to scale their incentivized advertising campaigns by providing access to users from different affiliate apps and third-party mediation services. From the point of view of developers, there is a spectrum of IIPs depending on their review process, required monetary commitments, and level of compliance with app store policies.

On one end, we find **vetted IIPs** (e.g., fyber.com, offertoro.com, hangmyads.com) that have a stringent review process to vet developers. In most cases, they require developers to provide extensive

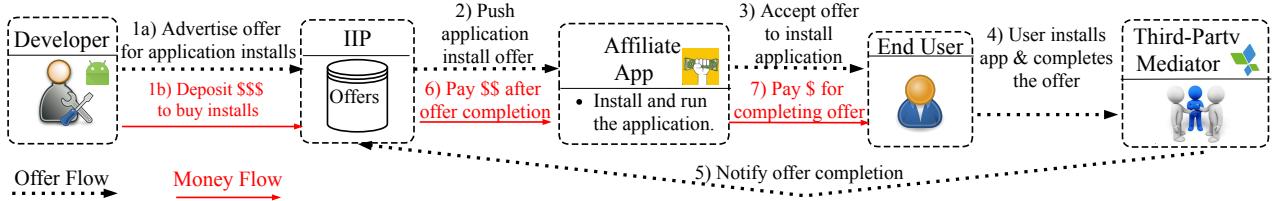


Figure 1: Workflow of an incentivized install campaign. An app developer runs an incentivized install campaign by paying an IIP to advertise an offer which is then distributed through an affiliate app where users browse offers and select an offer to work on. Once a user completes an offer (and completion has been certified by a third-party mediator), the payment is disbursed to all the stakeholders.



Figure 2: Screenshot of RankApp showing how they publicly advertise improvement of app's rank on Google Play Store as a service.

documentation (e.g., valid TAX id, bank account) and make significant upfront monetary commitments (sometimes as high as thousands of dollars). As a result, we expect vetted IIPs to be used predominantly by more established developers who can fulfill these requirements. Moreover, vetted IIPs tend to stay compliant with app store policies.

On the other end, we find **unvetted IIPs** (e.g., rankapp.org, mopeak.com, cpimobile.com) that usually do not have a review process to vet developers. The barrier to entry is comparatively low since no documentation is required and a developer can pay as little as 20 dollars to start a campaign. Moreover, unvetted IIPs appear to put less emphasis on compliance with app store policies. For example, RankApp publicly advertises that their incentivized install campaigns can help an app improve its rank in Google Play Store (see Figure 2). This may be in violation of Google Play Store's policy against manipulation of app store metrics [11].

We are able to easily identify several operational IIPs by querying Google search with relevant keywords (e.g., “incentivized installs” and “mobile offer walls”) and by investigating relevant discussion threads on Reddit and Quora. By reviewing the websites of seven IIPs (that we later study in Section 4) and also attempting to register as a developer with these IIPs, we are able to classify them as

IIP	Type	Home URL
Fyber	Vetted	fyber.com
OfferToro	Vetted	offertoro.com
AdscendMedia	Vetted	adscendmedia.com
HangMyAds	Vetted	hangmyads.com
AdGem	Vetted	adgem.com
ayeT-Studios	Unvetted	ayetstudios.com
RankApp	Unvetted	rankapp.org

Table 1: Characterization of different IIPs identified in our study by reviewing their websites and attempting to register with them as a developer.

“vetted” and “unvetted” as shown in Table 1 following the previous definitions.

Affiliate Apps. IIPs typically rely on multiple affiliate apps for disseminating advertised offers at scale. An affiliate app signs up with one or more IIPs to integrate their offer walls through a purpose-specific SDK. This enables the users of the affiliate app to access the offers listed in different offer walls. After a user completes an offer listed in the offer wall, the IIP keeps a fraction of the payout and releases the remaining payout to the affiliate app which, in turn, keeps a fraction of the payout and releases the remaining payout to the user. Affiliate apps offer either non-monetary awards (e.g., advertising these offers as a means to advance in the game) or monetary incentives (e.g., gift cards) to their users for completing these offers. We primarily focus on affiliate apps that offer monetary incentives in our study. As we discuss next, disbursement of the payout depends on certification by a third-party mediator.

Third-party Mediator. A third-party mediator (or attribution service) [68] is an entity trusted both by the developer and the IIP.² The advertised app integrates a purpose-specific SDK of a third-party mediator (e.g., appsflyer.com, kochava.com, adjust.com) to track offer completion. Many of these services also offer analytics and anti-fraud products. This mediator charges the developer either a fixed amount or a per-user rate. For instance, appsflyer.com charges 0.03 USD/user.

2.2 Offer Types

We divide offers in two categories based on the set of actions required from a user to complete an offer.

²Some IIPs also offer attribution services to track offer completion. However, developers may opt for third-party mediators to reduce the risk of fraud.

- *No activity* offers require users to only install and open the advertised app (e.g., “Install and Launch”). This type of incentivized installs can help an app developer to manipulate app store metrics (e.g., install counts, visibility). However, *no activity* offers are unlikely to manipulate other user engagement metrics (e.g., average user session length) because a user who only installs the app and opens it once (as required by the task) does not contribute to an increase in session length (discussed in Section 3).
- *Activity* offers require users to perform additional in-app tasks in the advertised app after installation. These tasks allow developers to manipulate additional user engagement metrics and require a non-trivial amount of time and effort from users. A developer can create *activity* install offers to inflate the number of registered users (e.g., “Install and Register”), user session lengths (e.g., “Install and Reach Level 10”), and revenue (e.g., “Install and make a \$4.99 in-app purchase”). These actions result in artificially generated user engagement that could lead to monetization opportunities (discussed in Section 4). For instance, requiring users to spend a longer amount of time in the app allows monetization through in-app advertising and increases the metrics related to user session length. Likewise, inflated user engagement metrics (e.g., active users) might help developers raise funding from VCs.

3 MEASUREMENTS OF INCENTIVIZED INSTALL PLATFORMS

In this section, we study IIPs by purchasing incentivized installs from them. These measurements not only allow us to get the first-hand experience of IIPs as a mobile developer but also help us measure the quality of installs (e.g., user engagement and retention). Finally, these measurements also shed light on Google Play’s existing enforcement to detect these incentivized installs. Next, we discuss our purpose-built honey app designed to purchase incentivized installs from IIPs.

3.1 Honey Mobile App

We customize an open-source “voice memos saving” Android app and publish it on the Google Play Store to serve as our honey app. We add instrumentation to upload metadata to our server. Specifically, our honey app collects information about user in-app activity (e.g., clicks on voice memo record button) and device information (e.g., list of other installed apps, device build,³ WiFi SSIDs, the /24 block of the public IPv4 address, and signals to identify whether the device is rooted⁴). This information is uploaded to our server whenever the user opens our honey app or clicks the voice memo record button. This information helps us measure user engagement and characteristics of users who install our honey app. We complement our information with the analytics provided by Google Play Store’s developer console.

Ethics. We received approval from our Institutional Review Board (IRB) before conducting experiments using our honey app. In our experiment, we try to adhere to the proposed guidelines in the Menlo report [17]. However, we note that obtaining direct consent

from users would not have been feasible since it can influence user actions and findings. Nevertheless, we provide a clear privacy policy to inform users of the nature and purpose of our data collection. We do not collect device information that might uniquely identify users (e.g., we do not collect the IMEI/IMSI of the mobile device, and we drop the last octet of the IPv4 address). While we do collect the wireless network’s SSID, which can be considered personal data, we only store a hashed value. This is enough for our purpose of identifying whether several phones are connected to the same network. We acknowledge that while the collection of installed packages in user devices can have privacy risks, it allows us to identify affiliate apps that integrate IIPs (as we discuss later in Section 4) and further understand users’ objectives. We also take security measures to protect the collected data, and communication with our server happens over encrypted channels. We only use the collected data for the purpose of our research. We also received approval from our institution to pay for IIPs and share billing information (e.g., tax ID) when hiring the services of vetted IIPs.

3.2 Measuring the Install Campaigns

We arbitrarily pick one vetted (Fyber) and two unvetted (ayeT-Studios and RankApp) IIPs (from Table 1) and purchase 500 *no activity* installs for our honey app. Our incentivized install campaigns across these three IIPs are spread over time such that no two campaigns deliver installs at the same time. We use analytics provided by Google Play Store’s developer console [12] to measure the delivery of installs by each IIP. Note that we use Google Play Store’s developer console to verify that we do not receive any organic installs (i.e., installs from app store search or top charts) during our incentivized install campaigns. Thus, we can confidently attribute installs received during each campaign to its respective IIP.

We combine the information collected by our honey app with analytics from Google Play Store’s developer console to analyze user acquisition and user engagement:

- **User acquisition:** Our honey app received a total of 1,679 installs, including 626 from Fyber, 550 from ayeT-Studios, and 503 from RankApp based on Google Play Store’s developer console analytics. The completion time of the campaigns varies between IIPs. While Fyber and ayeT-Studios are able to deliver all the installs within two hours of the campaign launch, RankApp takes more than 24 hours to deliver all the installs. The number of installs for Fyber and ayeT-Studios recorded by our server match Google Play’s developer console analytics. However, our server is missing telemetry from 45% (226 out of 503) of the installs purchased from RankApp. This indicates that a large fraction of RankApp users likely did not open the app since our server will not receive telemetry data from an install unless the app is opened.
- **User engagement:** We analyze user engagement beyond the completion of our advertised offer that only required users to install and open the app to earn the reward. Specifically, we measure how many users click on the voice memo record button (i.e., the only functionality in the app). Overall, we observe extremely low engagement. In general, less than half of the users clicked on the recording button. When looking at user engagement across

³We look for strings (e.g., generic, genymotion) to detect emulators.

⁴We use the open source library RootBeer [10].

IIPs, we observe lower engagement in RankApp users: 6% of RankApp users clicked on the record button vs. 44% of Fyber and ayeT-Studios users. Nevertheless, user engagement quickly fades over time. One day after installation, only a handful of users (4 Fyber, 2 RankApp, 1 ayeT-Studios) clicked the record button.

Our analysis shows that most of the users who install our honey app do the bare minimum effort to complete the offer. This reaffirms our initial hypothesis that most users are likely not genuinely interested in the apps advertised in the offer walls of IIPs. Instead, the majority seem to install the honey app to receive the incentive (payout).

Incentivized Users. Users who install our honey app could be bots who are installing advertised apps from the offer walls in an automated way (similar to click fraud [30]). We find that some of our installs show signs of automation. Our honey app is installed 4 times on an emulator (2 from RankApp and 2 from Fyber) and 7 of the devices that install our honey app (2 from Fyber, 4 from ayeT-Studios, 1 from RankApp) connect from ASNs of popular cloud services (e.g., Digital Ocean) when eyeball ASNs would be expected. We also find cases where users seem to deploy device farms [34] with presumably rooted devices to scale their operations for earning money. For example, we record 20 installs from different devices behind the same /24 block. 18 out of these 20 installs are from rooted phones that also share the same WiFi SSID.

We also analyze 17,454 apps that are installed on the devices of potentially real users who installed our honey app to identify their objectives. The most popular apps among the users of all three IIPs include several affiliate apps that integrate offer walls from IIPs. The most popular affiliate app among the users of RankApp is eu.gcashapp (37% users), ayeT-Studios is com.ayet.cashpirate (20% users), and Fyber is proxima.makemoney.android (9% users). Our manual analysis shows that the names of many apps contain keywords such as “money”, “reward”, or “cash”. We find that a large fraction of RankApp (98%), ayeT-Studios (72%), and Fyber (42%) users have installed at least one affiliate app that contains such keywords. This indicates that most of the users are likely semi-professional crowd workers who seek to earn money through these schemes. As we discuss later, we instrument a set of these affiliate apps to identify other mobile apps that run incentivized install campaigns on IIPs.

Note that we do not purchase installs from more IIPs due to our limited budget. Thus, our results may lack completeness. Nevertheless, our selection contains at least one IIP from the vetted and unvetted categories to reflect the dynamics of both types of IIPs in our results.

Takeaway. Our measurements of installs provided by three IIPs suggest that many of the users were likely not genuinely interested in our app and were rather crowd workers installing apps advertised on IIPs to earn money. These installs did increase the public install count of our honey app from 0 to 1,000, which indicates that they are effective for manipulating an app’s Google Play Store install count metric. This demonstrates a lack of enforcement by Google Play Store to detect incentivized installs.

4 MEASURING INCENTIVIZED INSTALL CAMPAIGNS IN THE WILD

In this section, we study incentivized install campaigns of mobile apps in the wild.

- (1) We develop a monitoring infrastructure to find mobile apps running incentivized install campaigns on IIPs (Section 4.1).
- (2) We characterize incentivized install campaigns across different IIPs in terms of: (1) how much users are being paid to complete different types of incentivized offers, and (2) which mobile apps are being advertised (Section 4.2).
- (3) We aim to understand how different types of campaigns could benefit developers in terms of improving (1) user engagement and app store visibility, (2) monetization, and (3) likelihood to raise funding (Section 4.3).

4.1 IIP Monitoring Infrastructure

To study IIPs in the wild, we cannot rely on our honey app because it does not allow us to monitor incentivized install offers for other mobile apps. To this end, we can either directly sign up with IIPs using a purpose-built honey affiliate app or instrument real-world affiliate apps. We pick the latter option as it is more scalable: it allows us to monitor many different IIPs integrated by various affiliate apps. Next, we explain our approach to identify and instrument real-world affiliate apps.

Instrumentation of Affiliate Apps. We instrument affiliate apps to extract offer walls of IIPs that are integrated inside these apps. Figure 3 shows the two main components of our instrumentation:

- (1) We implement a UI fuzzer based on Appium [15] to automate UI interactions with an affiliate app on an Android phone. This allows us to milk the affiliate apps without any human in the loop. We note that each affiliate app may have multiple offer walls that are organized in separate tabs. Our UI fuzzer sequentially opens all of the tabs to load the offer walls and then it scrolls through the offer wall to make sure that all the offers are loaded.
- (2) We intercept the network traffic generated by offer walls on the Android phone by configuring a mitmproxy [40] proxy server.⁵ Specifically, we parse the HTTP responses that are intercepted by the mitmproxy as a result of the stimuli generated by our UI fuzzer when loading an offer wall. These responses typically include offer details in JSON format containing offer description, payout, and the advertised app’s Google Play Store profile.

From the list of affiliate apps identified in Section 3, we select 8 affiliate apps that are available on Google Play Store and are fairly popular. Using the above methodology, we instrument these 8 affiliate apps as listed in Table 2. These affiliate apps have millions of installs, with as many as 10M+ installs for the most popular app. Note that we do not automate more affiliate apps since it entails significant engineering effort. We acknowledge that our instrumentation of 8 affiliate apps to monitor IIPs may lack completeness. Nevertheless, we are able to identify advertised apps from the offer

⁵Note that all offer walls use TLS encryption in their traffic. We decrypt this traffic by installing a self-signed certificate on the Android phone since none of the offer walls uses certificate pinning.

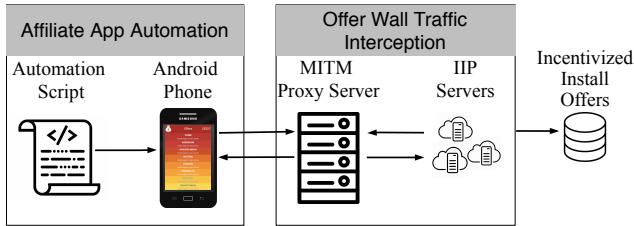


Figure 3: Infrastructure to monitor apps advertised on IIPs.

App Package Name	Installs	Fyber	AdscendMedia	AdGem	OfferToro	HangMyAds	RankApp	ayeT-Studios
com.mobvantage.	10M+	✓	✗	✓	✗	✓	✗	✓
CashForApps		✓	✓	✗	✗	✗	✗	✗
proxima.makemoney	5M+	✓	✓	✗	✗	✗	✗	✗
.android								
proxima.moneyapp	1M+	✓	✗	✗	✗	✗	✗	✗
.android								
com.biggash.app	1M+	✗	✓	✗	✓	✗	✗	✗
com.ayet.cashpirate	1M+	✓	✗	✗	✗	✗	✗	✓
eu.makemoney	1M+	✗	✓	✗	✗	✗	✓	✗
com.growthrich.	1M+	✗	✓	✗	✗	✗	✓	✗
makemoney								
make.money.easy	100K+	✓	✓	✗	✗	✗	✗	✓

Table 2: List of monitored affiliate apps and the offer walls of IIPs integrated inside them. We instrument these eight affiliate apps to find Google Play Store apps advertised on vetted and unvetted IIPs. RankApp and ayeT-Studios are unvetted IIPs while the remaining services are vetted IIPs.

walls of 7 IIPs integrated inside these 8 affiliate apps. Most of the instrumented affiliate apps integrate multiple IIPs, with the most popular one integrating offer walls from 4 different IIPs. We note that all of the 8 affiliate apps integrate at least one offer wall from vetted IIPs (Fyber, HangMyAds, AdscendMedia, AdGem, OfferToro), but most (5 out of 8) of them also integrate at least one offer wall from unvetted IIPs (ayeT-Studios, RankApp). We further increase our coverage of apps advertised on the monitored IIPs by running our milkers from the following eight countries: USA, UK, Spain, Israel, Canada, Germany, India, and Russia using datacenter VPN proxies offered by luminati.io.

Dataset. We use our instrumentation of affiliate apps to monitor IIPs over a period of three months from March-June 2019. We identify a total of 2,126 offers from 922 unique advertised apps that are advertised on 7 IIPs, both vetted and unvetted services. We note that offer payouts use different point systems across different affiliate apps. We normalize offer payouts of different affiliate apps by converting their points to equivalent dollar amounts.⁶

Overall, we identify a total of 1,128 unique offer descriptions in these 2,126 offers. We manually label offer descriptions into two

⁶By analyzing affiliate apps, we convert these reward points to an equivalent offer payout in USD that can be redeemed through gift cards (e.g., PayPal, Amazon) inside the affiliate app.

Offer Type	% of offers (N = 2,126)	Average payout
No activity	47%	\$0.06
Activity	53%	\$0.52
Activity (Usage)	37%	\$0.50
Activity (Registration)	11%	\$0.34
Activity (Purchase)	5%	\$2.98

Table 3: Prevalence of different types of incentivized install offers and their average payouts.

offer types (*no activity* and *activity*) according to the classification described earlier in Section 2. Our manual analysis of these offer descriptions further reveals that some *activity* offers are targeting specific user engagement metrics. Hence, we further divide *activity* offers into three subcategories: (1) *Registration* if the offer requires users to register an account, (2) *Purchase* if the offer requires users to make in-app purchase, and (3) *Usage* if the offer requires users to perform any other action.

4.2 Characterizing Incentivized Install Offers

We analyze offers and advertised mobile apps across different vetted and unvetted IIPs.

Offer Types and Payouts. Table 3 reports the presence of different types of incentivized install offers and their average payouts. We observe that *activity* offers are slightly more popular than *no activity* offers. 53% are *activity* offers while 47% are *no activity* offers. We observe that the average payout of *activity* offers is 9× higher than that of *no activity* offers. We further observe differences among the payouts of the subcategories of *activity* offers. The average payout of *purchase* offers is 9× and 6× higher than that of *registration* and *usage* offers, respectively. This is not surprising because users have to put in more effort and even spend money (to make in-app purchases) for completing *activity* offers. We conclude that manipulating advanced engagement metrics through *activity* offers is generally much more expensive than manipulating simple install counts through *no activity* offers.

Table 4 shows that there are differences in offer payouts across vetted and unvetted IIPs. The lower payouts in unvetted IIPs are because developers mostly advertise *no activity* offers on unvetted IIPs. Specifically, one unvetted IIP (RankApp) has no *activity* offers and an average payout of \$0.02, and the other unvetted IIP (ayeT-Studios) has only 29% *activity* offers and an average payout of \$0.05. In contrast, vetted IIPs advertise much more *activity* offers with higher payouts. These findings show that incentivized install campaigns on vetted IIPs are more interested in manipulating advanced engagement metrics.

Advertised Apps and Developers. Next, we compare advertised mobile apps and their developers across vetted and unvetted IIPs. Table 4 lists some basic statistics of advertised mobile apps and their developers by crawling their Google Play Store’s profile to collect metadata such as their install counts (Google reports installs in bins of a lower-bound “minimum” number of installs), release date, genre, and developer details (e.g., company name, websites, mailing address, developer ID). We note that mobile apps from a wide variety of app genres advertised in several IIPs are published on Google

IIP Name	IIP Type	Median Offer Payout	Offer Type		Number of Apps	Number of Developers	Number of Countries	Number of Genres	Median Install Counts	Median Age of Apps (Days)
			No activity	Activity						
RankApp	Unvetted	\$0.02	100.0%	0%	152	114	39	20	100	33
ayeT-Studios	Unvetted	\$0.05	71%	29%	392	351	44	51	1K	70
Fyber	Vetted	\$0.19	24%	76%	378	319	40	36	1M	777
AdscendMedia	Vetted	\$0.12	9%	91%	104	79	27	21	500K	722
AdGem	Vetted	\$1.71	16%	84%	28	27	15	8	500K	854
HangMyAds	Vetted	\$0.40	23%	77%	27	27	17	9	1M	699
OfferToro	Vetted	\$0.09	52%	48%	140	131	34	19	500K	557

Table 4: Summary of Google Play Store’s metadata and offers’ metadata of apps identified on vetted and unvetted IIPs. Mobile apps from a wide variety of app genres advertised on several IIPs are published on Google Play Store by hundreds of developers based in different countries. We also note that less-established (newer, less popular) apps are more likely to run *no activity* incentivized install campaigns and use unvetted IIPs rather than vetted IIPs presumably because of their higher barrier of entry.

Play Store by hundreds of developers based in different countries.⁷ As an example of unvetted IIPs, 392 apps of 51 genres advertised on ayeT-Studios are published on Google Play Store by 351 developers based in 44 countries. As an example of vetted IIPs, 378 apps of 36 genres advertised on Fyber are published on Google Play Store by 319 developers based in 40 countries. It is noteworthy that some of the popular apps are from mainstream developers. For example, we identify “Apple Music” from Apple and “LinkedIn: Job Search & Business News” from LinkedIn on vetted IIPs. Similarly, we identify “TikTok - Make Your Day” from TikTok and “Fiverr - Freelance Services” from Fiverr on unvetted IIPs. It is unclear whether these developers are purchasing incentivized installs themselves or if they are potentially deceived by third parties (e.g., marketing companies) or IIPs, which run incentivized installs to fulfill regular mobile app install advertising campaigns [56] (as we discuss later in Section 5).

Table 4 shows that there are significant differences in the age of the apps⁸ and the popularity of apps advertised on vetted and unvetted IIPs. For example, the median install count of apps advertised on both unvetted IIPs is 100 and 1,000, respectively, while it ranges between 500,000 and 1,000,000 for vetted IIPs. As another example, the median age of apps advertised on both unvetted IIPs is 33 and 70 days, respectively, while it ranges between 557 and 854 days for vetted IIPs. Overall, we conclude that less-established (e.g., newer or less popular) apps are more likely to use unvetted IIPs rather than vetted IIPs presumably because of their higher barrier of entry. Next, we show how well-established apps differently benefit from running *activity* campaigns on vetted IIPs as compared to less-established apps that generally run *no activity* campaigns on unvetted IIPs.

4.3 Impact of Incentivized Install Campaigns

In this section, we study how developers benefit from different types of incentivized campaigns on vetted and unvetted IIPs. First, we measure the impact of incentivized installs on app store metrics. Second, we analyze how some developers monetize different types of *activity* offers in their incentivized campaigns. Finally, we study

⁷Note that we uniquely identify a developer using the developer’s ID and identify the developer’s country by parsing their mailing addresses from their app’s Google Play Store profile. Unique apps are identified by their package name.

⁸We measure the age of an app by computing the difference between the start of the advertised app’s campaign date and the date in which the app was released on the Google Play Store.

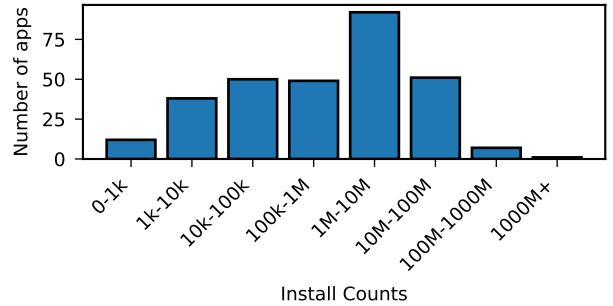


Figure 4: Install counts of the baseline apps.

whether incentivized installs might help developers raise funding. Note that some confounding factors (e.g., non-incentivized installs) may have an effect on the advertised apps to improve app store metrics or raise funding. To mitigate the impact of such factors, we compare the set of apps running incentivized install campaigns with a set of 300 baseline apps published on Google Play. Our baseline apps are sampled from the set of apps commonly used by Lumen users [51], a crowd-sourced privacy tool. Figure 4 shows that our baseline contains a comprehensive set of apps of all types of popularity ranging from less than 1K installs to over 1000M installs. Note that our sample of baseline apps has no overlap with the apps advertised on IIPs that are identified in our study. We also use the chi-squared test of independence [39] to determine whether there is a statistically significant difference in improved app store metrics and fund raising between apps observed in incentivized install campaigns, and those in the baseline. We acknowledge that any correlation we might find in this section does not necessarily imply causation because there could be other confounding factors beyond our purview (e.g., non-incentivized install campaigns).

4.3.1 Impact on App Store Metrics. We analyze whether incentivized install campaigns from vetted and unvetted IIPs may be able to manipulate app store metrics. To this end, we crawl Google Play Store profiles of apps to collect their install counts. We also crawl Google Play Store “top charts” that list trending apps (e.g., top gaming apps, top grossing apps). We periodically collect this data every other day from March 2019 to June 2019. This allows us to

App Set	No Increase	Increase
Baseline ($N = 300$)	294 (98%)	6 (2%)
Vetted ($N = 492$)	431 (88%)	61 (12%)
Unvetted ($N = 538$)	450 (84%)	88 (16%)

Table 5: Comparing apps that increased install counts from vetted and unvetted IIPs with baseline apps. Through the chi-squared test of independence, we find that there is indeed a correlation between apps advertised on both vetted and unvetted IIPs and an increase in the install counts of advertised apps.

measure how IIP campaigns on vetted and unvetted IIPs help apps increase their install counts and visibility on Google Play Store.

Increase in Install Counts. For this analysis, we compare the increase in install counts of apps advertised on vetted and unvetted IIPs with our set of baseline apps. Specifically, we check whether or not an app’s install count increases by the end of the incentivized install campaign as compared to the start of the campaign. For a fair comparison, we monitor install counts for baseline apps over the duration of 25 days – which is the average incentivized install campaign duration. Table 5 compares the number of apps from vetted IIPs, unvetted IIPs, and baseline that increased their install counts. To determine whether incentivized install campaigns significantly impact the increase in install counts, we use the chi-squared test of independence [3, 39]. We conduct two separate tests, “vetted vs. baseline” and “unvetted vs. baseline” to compare vetted and unvetted IIPs with the baseline. Our null and alternate hypotheses are:

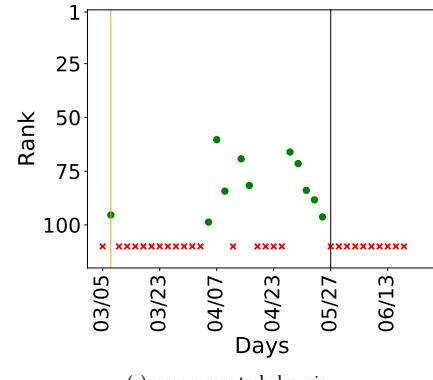
- H_0 : The proportion of apps whose install count increases is independent of incentivized install campaigns.
- H_a : The proportion of apps whose install count increases is associated with incentivized install campaigns.

For vetted vs. baseline, $\chi^2 = 26.0$ and $p = 3.378e^{-7}$. For unvetted vs baseline, $\chi^2 = 39.9$ and $p = 0.000$. Thus, we are able to reject the null hypothesis for both tests at the significance level of 0.05. We conclude that there is a correlation between apps advertised on vetted and unvetted IIPs and an increase in the install counts of advertised apps. Comparing vetted and unvetted IIPs, we note that the likelihood of an increase in install counts (with respect to the baseline) for vetted IIPs is 6× while it is 8× for unvetted IIPs. These findings seem counter-intuitive since apps using unvetted IIPs, despite cheaper offer payouts as shown in Table 4, are benefiting more than vetted IIPs. Nevertheless, as we discuss next, apps that use vetted services benefit in different ways, other than simply increasing their install counts.

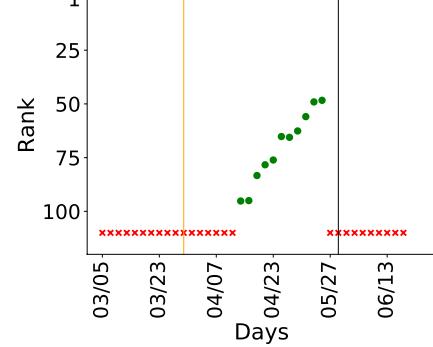
Appearance in Top Charts. For this analysis, we compare top chart appearances – *i.e.*, their app store visibility – of apps advertised on vetted and unvetted IIPs with our set of baseline apps. We check whether or not an app appears in the top charts by the end of their install campaign (*i.e.*, the last day we observe them in our crawl). To minimize bias, we exclude advertised apps that already appeared in top charts before the start of their campaign (*i.e.*, the first day we observe them in our crawl) and baseline apps that appeared in top charts at the start of our crawls. After this filtering, we are left with 320 apps from vetted IIPs, 484 apps from unvetted

App Set	Not Present	Present
Baseline ($N = 261$)	253 (96.9%)	8 (3.1%)
Vetted ($N = 320$)	296 (92.5%)	24 (7.5%)
Unvetted ($N = 484$)	472 (97.5%)	12 (2.5%)

Table 6: Comparing the appearance of advertised apps from vetted and unvetted IIPs with baseline apps in top charts. Through the chi-squared test of independence, we find that there is a correlation between apps advertised on vetted IIPs and the appearance of an app in the top chart.



(a) com.mmm.trebelmusic



(b) com.camelgames.wof

Figure 5: Case studies of the appearance of advertised apps in top charts during their campaign. The yellow line indicates the start of the campaign and the black line indicates the end of the campaign. Each red cross indicates the days when an app is not observed in a particular chart and each green dot indicates the days when an app has appeared. The y axis shows the percentile rank of the app.

IIPs, and 261 baseline apps. For a fair comparison, we again monitor visibility for the baseline apps over the duration of 25 days, which is the average incentivized install campaign duration. Table 6 reports the number of these apps from vetted IIPs, unvetted IIPs, and baseline apps that appeared in top charts. To determine whether incentivized install campaigns significantly impact app store visibility, we use the chi-squared test of independence. Our hypotheses are:

- H_0 : The proportion of apps that appear in top charts is independent of incentivized install campaigns.
- H_a : The proportion of apps that appear in top charts is associated with incentivized install campaigns.

For vetted vs. baseline, $\chi^2 = 5.43$ and $p = 0.02$. For unvetted vs baseline, $\chi^2 = 0.22$ and $p = 0.64$. Thus, we are able to only reject the null hypothesis for the vetted vs. baseline test at the significance level of 0.05. We conclude that there is a correlation between apps advertised on vetted IIPs and the appearance of an app in top charts. Mobile apps advertised on vetted IIPs appear in top charts up to 2× more frequently in comparison to baseline apps as shown in Table 6. However, we are unable to conclude whether advertising apps on unvetted IIPs affects the chances of appearance in top charts. Vetted IIPs not only help apps increase their install counts but can also help them appear in top charts. Note that even though developers generally pay more for *activity* offers on vetted IIPs, it appears worthwhile because Google Play Store places apps in top charts based on user engagement metrics [44], which cannot be inflated with *no activity* offers on unvetted IIPs.

To further illustrate this point, we analyze apps that use different types of *activity* offers as listed in Table 3. Figure 5 shows the temporal evolution of the rank of two example apps in the top charts. In Figure 5(a), the app “TREBEL - Free Music Downloads & Offline Play” starts appearing in the top-games chart after the start of its incentivized install campaign using *registration* and *usage* offers. Specifically, the description of these offers is “Install and register” and “Install, register, and download a song.” This suggests that manipulation of more targeted user engagement metrics through *registration* and *usage* offers can help apps appear in top charts. In Figure 5(b), the app “World on Fire” appears in the top-grossing chart a few days after the start of its incentivized install campaign using *purchase* offers. Specifically, the description of one of these *purchase* offers is “Install & Make any purchase.” This suggests that these *purchase* offers manipulated revenue through fake purchases, which in turn helped “World on Fire” appear in the “Top Grossing Apps” chart. We conclude that inflating user engagement metrics (e.g., time spent in the app) and revenue through *activity* offers may help apps appear in relevant top charts. As we discuss next, these *activity* offers also seem to provide monetization opportunities to developers.

Summary. Our results show differences in the effectiveness of running incentivized install campaigns on vetted and unvetted IIPs. While unvetted IIPs seem to be more effective in targeting a primitive engagement metric (*i.e.*, install counts), vetted IIPs can be more effective in manipulating advanced user engagement metrics and app store visibility. Our results also show that Google Play Store’s detection and filtering systems [42, 43] likely fail to effectively enforce their policy against app store manipulation by incentivized installs [11].

4.3.2 Monetization. We discuss how *activity* offers provide different opportunities for monetization through advertisement and arbitrage.

Advertising. In-app advertising is a popular way to monetize mobile apps [50]. Intuitively, there is a direct relationship between user engagement and revenue earned from advertisements. Higher user engagement provides developers more opportunities to show

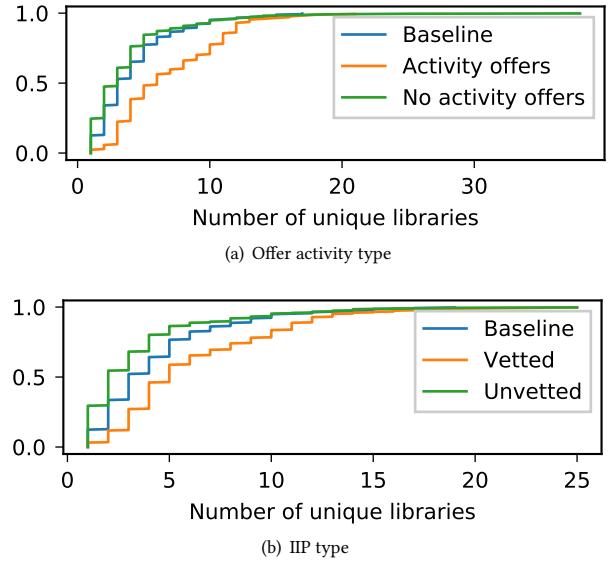


Figure 6: Distribution of unique ad libraries across apps.

ads, something known as “active eyeballs” in Ad Tech jargon [45]. Many *activity* offers require users to engage with the app for a long time period by asking them to complete a series of difficult tasks such as “Install & Reach level 10”. This user engagement can be directly monetized by displaying in-app ads to the user. As a proxy for measuring advertising revenues, we analyze whether apps using *activity* offers have more capability to monetize through advertising libraries. To this end, we analyze the number of advertisement libraries embedded in apps. We download APKs of baseline and advertised apps to perform static analysis⁹ using LibRadar [37].

Our results show that apps using *activity* offers have more ad libraries than apps using *no activity* offers. Figure 6(a) shows that 60% and 25% of apps with *activity* and *no activity* offers, respectively, have 5 or more ad libraries. We find well-known advertising vendors such as Google AdMob, AppLovin, and ChartBoost. We also find advertisers that serve the role of IIP (*e.g.*, Fyber). A similar pattern emerges when comparing vetted, unvetted, and baseline apps. Figure 6(b) shows that 55%, 20%, and 35% of vetted, unvetted, and baseline apps, respectively, have 5 or more ad libraries. This is expected because *activity* offers are more common on vetted IIPs in comparison to unvetted (as reported in Table 3).

Our results indicate two different strategies of incentivized install campaigns. On one hand, we find that mobile apps advertised on vetted IIPs have more advertisement libraries, giving these apps more opportunities to generate revenue since their incentivized install campaigns are more likely to advertise *activity* offers. On the

⁹We note that static analysis may miss some advertising libraries due to code obfuscation and dynamic code loading. Furthermore, we acknowledge that our analysis offers an upper bound on in-app advertising, as the mere presence of an advertising library does not guarantee that its code will be executed at run-time.

other hand, incentivized install campaigns from unvetted IIPs advertise *no activity* offers that do not require additional engagement from users and include fewer advertising libraries.

Arbitrage. We observe that developers can craft *activity* offers in a way that allows them to monetize through arbitrage. In arbitrage, a developer advertises an *activity* offer that requires users to complete additional tasks within the app such as filling surveys, watching videos, or even installing other apps (similar to affiliate apps). In return, the developer earns money through commissions after these offers within the app are completed. One example of arbitrage is the following offer that advertises an app “Cash Time - Earn Money - Android”. This *activity* offer pays \$0.67 to users for reaching 850 points in the app by completing additional tasks (e.g., watching videos, completing surveys, shopping for deals) in the app. Unfortunately, we cannot estimate the revenue earned through arbitrage to understand profitability. However, even if arbitrage is not profitable, these offers would likely enable an app to increase revenue-based metrics, such as revenue growth.

We manually analyze the descriptions of all *activity* offers and find that 3.9% of apps (36 out of 922) use arbitrage-based *activity* offers. We observe a higher percentage of such offers in vetted IIPs as compared to unvetted IIPs. Specifically, 7% of apps (35 out of 492) from vetted IIPs while only 2% of apps (10 out of 538) from unvetted IIPs use arbitrage-based *activity* offers. This is again expected because *activity* offers are more common on vetted IIPs in comparison to unvetted (as reported in Table 3).

Summary. Overall, our results show that incentivized install campaigns that advertise *activity* offers have opportunities to monetize through advertising and arbitrage. It is unclear whether these monetization strategies are sufficient to directly recuperate the cost of their incentivized install campaigns. Next, we explore other indirect avenues of monetization such as raising funding from venture capitalists.

4.3.3 Investor Funding. There is an inherent tension between growth and profitability [13] of mobile app startups. Many app developers are more interested in strong user growth as well as bolstering user engagement metrics because, in the long term, they can help developers raise funding from investors such as venture capitalists (VCs) [5, 62]. In this section, we study whether any of the 922 apps in our dataset are able to raise funding after launching incentivized install campaigns.

To this end, we use the Crunchbase [24] database that provides us with access to the list of companies that have raised funding. We download a snapshot of the Crunchbase database in October 2019 (a few months after our initial data collection). By searching for developer information from Google Play Store, we match 23% of 922 apps to their developers in the Crunchbase database. Crunchbase provides funding information such as the name of the funding organization, type of organization (e.g., angel investor, VC investor), amount of funding, and type of funding (e.g., seed round, series A, series B). We are able to match 39% of apps (192 out of 492) from vetted IIPs, 15% of apps (79 out of 538) from unvetted IIPs, and 27% of baseline apps (82 out of 300) in the Crunchbase database. We surmise that mobile apps advertised on vetted IIPs have a higher percentage of matched developers over unvetted IIPs since vetted

IIPs are used by more established developers (as discussed in Section 4.1). In fact, our manual analysis shows that mobile apps of the unmatched developers from unvetted IIPs often do not provide useful information in their Google Play Store profile (e.g., link to their website). Note that the lack of completeness of Crunchbase database may also result in the exclusion of some unmatched developers. To mitigate the impact of this limitation, we only compare those advertised apps and baseline apps that are matched in the Crunchbase database to perform further analysis.

We analyze whether advertising apps on vetted and unvetted IIPs has any effect on the funding raised by them. To this end, we compare the funding raised by developers whose mobile apps are advertised on vetted and unvetted IIPs with baseline apps. Specifically, we check whether or not a mobile app raises funding after running the incentivized install campaign(s). For comparison, we extract the same information for our baseline apps. Table 7 reports the number of these mobile apps from vetted IIPs, unvetted IIPs, and from the baseline whose developers are able to raise funding. To determine whether incentivized install campaigns significantly impact the funding raised by apps, we again use the chi-squared test of independence. We conduct two separate tests, “vetted vs. baseline,” and “unvetted vs. baseline,” to compare vetted and unvetted IIPs. Our hypotheses are:

- H_0 : The proportion of apps that raise funding is independent of incentivized install campaign.
- H_a : The proportion of apps that raise funding is associated with incentivized install campaign.

For vetted vs. baseline, $\chi^2 = 4.7$ and $p = 0.03$. For unvetted vs. baseline, $\chi^2 = 2.8$ and $p = 0.10$. Thus, we are able to only reject the null hypothesis for the vetted vs. baseline at the significance level of 0.05. We conclude that there is a correlation between advertising apps on vetted IIPs and developers of advertised apps raising funding. However, we are unable to conclude whether advertising apps on unvetted IIPs affects the chances of raising funding.

We further compare the offer types and payout amounts of incentivized install campaigns for the apps that raise funding versus those that do not. Table 8 presents a breakdown of offer types and payouts of 30 apps from vetted IIPs that raised funding after their campaigns. We observe that these apps are roughly equally likely to advertise *activity* and *no activity* offers. However, the incentivized install campaigns of the apps that raised funding have much higher payouts. For example, the average payout of *no activity* offers advertised by apps that raised funding is twice the average payout of all offers reported in Table 3. We surmise that the developers interested in raising funding need to aggressively acquire new users, and thus are willing to pay more.

We also look at two case studies of apps that raise VC funding after incentivized install campaigns. First, the *Password Manager* app by Dashlane raised multiple rounds of funding after launching incentivized install campaigns. The app was advertised on Fyber from March 12-27 and then May 7-14 using *activity* offers, which require users to install the app, create an account, and save at least two passwords in return for a reward around \$0.26-\$0.66. Coincidentally, Dashlane raised funding twice, each time about two weeks after each campaign: \$30 million on April 12 [46], and \$110 million on May 30 [64]. Second, *Droom: Used & New Car, Bike*,

App Set	Funding Raised	No Funding Raised
Baseline ($N = 82$)	5 (6.1%)	77 (93.9%)
Vetted ($N = 192$)	30 (15.6%)	162 (84.4%)
Unvetted ($N = 79$)	11 (13.9%)	68 (86.1%)

Table 7: Developers of mobile apps raising funding after campaigns using vetted and unvetted IIPs compared with baseline apps. Through the chi-squared test of independence, we find that there is a correlation between apps advertised on vetted IIPs and successfully raising funding after a campaign.

Offer Type	Percentage of Apps ($N = 30$)	Average Payout
No activity	67%	\$0.12
Activity	63%	\$0.92

Table 8: Breakdown of offer types and payouts of apps advertised on vetted IIPs that raised funding after their campaign.

Insurance, Loan & RTO, an e-commerce app that sells used and new vehicles, raised series F funding after appearing on IIPs. The app was advertised on Fyber and AdscendMedia from March 14–June 8 using *no activity* and *activity* offers, which require users to install the app and register a new account in return for a reward around \$0.02–\$0.06. Droom also raised funding about two weeks after its incentivized install campaign: Droom raised \$10 million on 20th June [59]. These two case studies indicate that incentivized installs could have misled investors who funded these apps (or the companies behind them). In the same vein, incentivized install campaigns of the remaining apps from vetted IIPs could also have misled investors since they also raised funding after the start of their incentivized install campaigns.

Next, we study developers matched with the Crunchbase database that are public companies [21]. Public companies generally report their user growth and engagement metrics on a quarterly basis. These metrics often play an important role in stock price targets from analysts and eventual stock price movements [28]. Thus, manipulation of these metrics through IIP campaigns can potentially boost the stock price of a publicly traded company. From the matching of developers from the Crunchbase database, we find that developers of 28 advertised mobile apps from vetted and unvetted IIPs are publicly traded companies. For example, Redfin, a publicly traded real estate brokerage company, reported a 27% increase in their mobile traffic during the second quarter of 2019 in comparison to the second quarter of 2018 [52]. Note that the incentivized install campaign of Redfin’s mobile app (“Redfin Real Estate”) advertised a *no activity* offer and a *usage* offer during the second quarter of 2019 as observed in our data collection. As another example, the quarterly report of IGG, a publicly traded mobile game company, praised its app “Lords Mobile” for staying in the top 5 of the top charts in 54 countries during the second quarter of 2019 [1]. We observed a total of 18 incentivized *usage* offers for this app over the duration of 3 months; these incentivized install campaigns could potentially have helped the app maintain its top ranks.

These public companies could be purchasing incentivized installs to artificially inflate their quarterly reports. Nevertheless, it is also possible that these companies are being defrauded by IIPs, as we discuss next in Section 5.

Summary. Overall, our results show a correlation between advertising apps on IIPs and developers of advertised apps raising funding from investors. This suggests that inflation of user engagement metrics through incentivized install campaigns may have misled investors.

5 DISCUSSION

Our analysis characterizes for the first time (and at scale) an unexplored part of the mobile ecosystem, and of alternative mobile app monetization strategies. We reveal a correlation between the apps participating in incentivized install campaigns and (1) their increase in install counts, (2) their appearance on Google Play Store’s top charts, and (3) in some cases, a correlation with successful investor rounds. The ability of such campaigns to boost install counts (and active users) can make apps appear more attractive than they actually are by manipulating Google Play Store’s rankings. This practice, prohibited on the Apple App Store, could be considered as potentially deceptive to Android users and investors alike, as it could give an unfair advantage to apps that are willing to pay users for installs over those that rely on organic growth strategies.

Yet, this does not imply that apps appearing on incentivized install services are willingly engaging in such activities. It is plausible that some app developers are themselves being deceived by third-party marketing organizations offering mechanisms to attract new users or improve their brand value. For example, marketers might attempt to arbitrage advertising campaigns by charging for more expensive non-incentivized installs, yet delivering lower-quality incentivized installs [8]. In fact, during our measurements, we observed campaigns involving apps published by well-known brands such as Amazon, Apple, and LinkedIn. As we discuss next, we disclose our findings to app developers.

5.1 Responsible Disclosure

We contacted the developers of popular apps advertised on vetted and unvetted IIPs for responsible disclosure. We contacted only 136 popular apps, each with 5M+ installs, out of 922 total apps since this was a manual process using the contact email address from their Google Play Store profile. We informed the developers that their apps were advertised as part of an incentivized install campaign. At the time of writing, we have received responses from three developers, all of whom were unaware of their apps participating in such campaigns. They also indicated that they are being defrauded. For example, one developer indicated that they contracted multiple external marketing organizations to acquire non-incentivized installs. It is possible that one of these marketing organizations advertised the developer’s app on IIPs during the time of our data collection.

We also disclosed our findings to Google since our results indicate that many of the advertised apps and IIPs are potentially violating Google Play Store’s policy. Other than the receipt of acknowledgement, we have so far not received any other feedback from Google either corroborating or refuting our findings.

5.2 Google Play Store’s Policy and Enforcement

Google Play Store’s policy [11] states that “*Developers must not attempt to manipulate the placement of any apps in Google Play.*” The presence and behavior of incentivized install platforms seem to directly contradict this requirement. In fact, some incentivized install platforms like RankApp (see Figure 2) are operational on Google Play Store despite openly advertising services that seem to run afoul of these policies.

Our successful purchase of incentivized installs in Section 3 suggested that Google Play Store’s detection and filtering systems [42, 43] might fail to effectively enforce their policy against manipulation of app store metrics. In order to measure potential enforcement, we revisit our collected data and look for a *decrease* in the install counts of advertised apps after the start of their campaign. Such a decrease would be an indicator that Google Play Store has identified and removed incentivized installs.¹⁰ Overall, we observed no decrease in the install counts of baseline apps or the apps advertised on vetted IIPs. For unvetted IIPs, we do occasionally observe a decrease in install counts during our data collection, but it affects only 2% of the advertised apps. One example was the app “Phonebook - Contacts manager” advertised on ayeT-Studios, whose install count decreased from 1,000 to 500.

Based on these observations, it seems that the effectiveness of enforcement on the Google Play Store is rather limited. We advocate moving towards a prohibitive stance regarding incentivized installs, similar to the Apple App Store [22], to impede such manipulation attempts while increasing the trustworthiness of top charts and install counts on the Google Play Store. Furthermore, our measurement methodology can also help the Google Play Store in improving its existing defenses by detecting apps that benefit from incentivized campaigns. Specifically, our proposed measurements can provide a ground truth of apps to help train machine learning models in detecting the lockstep behavior of users who perform similar in-app activities to complete the offer [18, 20, 33]. To foster follow-up research, we have also publicly shared our crawled data on <https://www.github.com/shehrozef/IncentInstalls>.

5.3 Limitations

Due to the limitations of our methodology, we cannot currently explore other aspects of the IIP ecosystem:

- We lack Google Play Store data of apps’ profiles and top charts outside of our crawl dates. Having historical and precise metadata for every app – including their presence in top charts – would have allowed us to better characterize the use and impact of IIPs on app store metrics.
- We lack complete visibility of the whole incentivized install ecosystem. We purchased installs for our honey app from three IIPs. We did not purchase installs from more IIPs due to budget limitations. In our in-the-wild measurement, our instrumentation of 8 affiliate apps allowed us to monitor 7 IIPs. We did not instrument more affiliate apps since it entails significant engineering effort.

¹⁰We acknowledge that if an app already has a large number of installs (e.g., 1M), we may not be able to observe the impact of install filtering.

- Comparing results of purchasing incentivized installs of our honey app with non-incentivized installs would help us better understand and distinguish the install quality and Google Play Store’s existing enforcement. We did not purchase non-incentivized installs due to budget limitations.
- We cannot estimate the amount of money spent by an incentivized install campaign of a mobile app because we lack data on the number of offers completed by users. We are only able to estimate the payout to complete an advertised offer that is observable in the offer wall’s traffic data.
- Our work focuses on studying IIP activity through affiliate apps. While other methods to distribute offers exist (such as micro service platforms), they fall out of the scope of this paper and require developing different methodologies that could be further studied in future work.
- There are potential grey areas allowing developers running incentivized install campaigns to increase revenue that we did not explore in our study. The completion of some in-app tasks may require users to purchase gaming credits, virtual currencies, or premium subscriptions, which are not mentioned in the offer description. We would not be able to identify these tasks since our existing categorization of in-app tasks relies on reading offer descriptions. Understanding these behaviors would require more extensive and comprehensive use of advertised apps (e.g., automating interaction with advertised apps).

6 RELATED WORK

To the best of our knowledge, we are the first to study the characteristics of incentivized install services for mobile apps, and the possibilities they offer for manipulating in-app metrics such as user engagement or revenue. Prior work addressed similar, but distinct, areas of pay-per-install (PPI) services for distributing malicious or unwanted desktop apps, and reputation manipulation for online social networks and mobile app stores as discussed next.

App Install Campaigns on Desktop. In contrast to the mobile app install campaigns we observe in this paper, desktop app install campaigns appear to be predominantly used by attackers and malicious actors to disseminate malware and potentially unwanted software (PUP) through deceptive practices and compromised hosts [19, 36, 60]. Caballero *et al.* found that such campaigns installed malicious apps from 12 out of 20 prevalent malware families on compromised host machines [19]. Kotzias *et al.* studied the presence of potentially unwanted software distributed through app install campaigns and detected at least one sample of potentially unwanted software on 54% of 9 million hosts [36]. These services rely on deceptive practices such as drive-by downloads [41] or default opt-in to install potentially unwanted software.

The ecosystem of incentivized mobile app installs differs from desktop app installs in at least two ways. First, users participating in mobile app installs consciously install an app to receive a reward, whereas desktop app installs typically occur without the user’s awareness, such as through deceptive practices, or covertly on compromised hosts. Second, the apps distributed through the install campaigns in desktop settings tend to be PUP or malware. In contrast, the install campaigns in mobile advertise predominantly

benign apps, presumably due to the presence of gatekeepers such as Google’s Play Protect program to prevent the distribution of malware. Furthermore, incentivized install providers seem to base their business on artificially increasing the reputation of the apps with the goal of driving future business opportunities and revenues.

Reputation Manipulation. Prior work has studied reputation manipulation on online social networks to artificially inflate social metrics such as likes, followers, or retweets in Facebook, Twitter, and Instagram [25–27, 32, 57, 58, 61, 63]. Similar to install counts of mobile apps, manipulating these metrics can increase the perceived value of a social media profile. Incentivized mobile app installs differ from social network reputation manipulation in the ecosystem of key actors and the methods they employ, such as detailed task “offers” distributed to crowd workers through a network of affiliate apps instead of using automation. These tasks often go beyond the mere installation of apps, requesting for example that workers reach a certain level in a game, or make in-app purchases before they can get paid.

Manipulating App Store Charts. A different line of prior work focused on the detection of app ranking manipulation in marketplaces according to publicly available metrics such as reviews, ratings, and ranking patterns of apps in app stores’ top charts [23, 47, 49, 66]. For example, Zhu *et al.* [66] built a statistical behavioral model of historical ranking patterns and the rating score of apps to detect manipulation. In contrast, our work detects actual apps engaging in manipulation by monitoring “offer walls” rather than making statistical inferences. Our methodology reveals tasks aiming to manipulate app store metrics that are not publicly visible and not targeted at improving the app’s standing in the app store, such as inflating in-app revenue or user engagement. Prior work has also interviewed fraudsters selling fake installs, reviews, and ratings in Google Play Store to reveal their operational characteristics [31, 48]. While we do not directly interact with the people behind the incentivized install platforms, our work sheds light on the ecosystem at scale, and provides new insights into how these services operate, as well as how app developers may stand to benefit from using these services.

7 CONCLUSION

In this paper, we provided an in-depth analysis of the ecosystem of incentivized installs on Google Play Store by purchasing installs for our honey app and monitoring IIPs. We find that most of the users are potentially crowd workers installing apps advertised on IIPs to earn money. Our results also showed that IIPs can have negative effects on the app store and beyond (e.g., deceitfully raise funding and erode consumer trust) due to their ability to manipulate user engagement metrics. Finally, our results demonstrate a likely lack of enforcement from Google Play Store to detect the potential violations of their policies by incentivized installs.

Our methodology to investigate IIPs could be deployed to provide a semi-automated monitoring service of IIPs. This could be leveraged by the Google Play Store to act on potential violations of its terms and services. It could also be leveraged by apps conducting non-incentivized marketing campaigns to detect if a third-party

is deceiving them and delivering incentivized installs, and by investors during due diligence investigations. We believe this would help increase the trustworthiness of user engagement metrics in Google Play Store.

ACKNOWLEDGEMENT

We would like to thank our shepherd, Maria Papadopoulou, and the anonymous reviewers for their useful feedback on this paper. We would also like to thank Sojhal Ismail, Abubakar Aziz, and Muzammil Hussain for their help with data collection. This work is supported in part by the National Science Foundation (under grant numbers 1564329, 1715152, 1814816, and 1954224), the European Union’s Horizon 2020 Innovation Action program (grant agreement number 786741, SMOOTH Project), the Spanish Ministry of Science and Innovation (grant agreement number PID2019-111429RB-C22, ODIO project), and by unrestricted gifts from Facebook and Google.

REFERENCES

- [1] [n.d.]. 2019 Interim Report iGG. http://iis.quan.net.com/media/IRAnnouncement/799/EN_US/003588684-0.PDF.
- [2] [n.d.]. App Store Review Guidelines. <https://developer.apple.com/app-store/review/guidelines/>.
- [3] [n.d.]. Chi-Square Test for Independence. <https://stattrek.com/chi-square-test/independence.aspx>.
- [4] [n.d.]. Choosing a Business Model. <https://developer.apple.com/app-store/business-models/>.
- [5] [n.d.]. Do angel investors look for a certain number of active users when investing in an app startup that has recently launched, 5,000–10,000+? <https://www.startups.com/community/questions/1841/do-angel-investors-look-for-a-certain-number-of-active-users-when-investing-in>.
- [6] [n.d.]. Earn more revenue with the right monetization options. <https://developer.android.com/distribute/best-practices/earn/monetization-options>.
- [7] [n.d.]. Fyber Competitors, Revenues and Employees. <https://www.owler.com/company/fyber>.
- [8] [n.d.]. INCENTIVIZED INSTALL MISLABELING. <https://appsflyer.com/glossary/incentivized-install-mislabeling/>.
- [9] [n.d.]. Mobile App User Acquisition Cost (2019). <https://www.statista.com/statistics/185736/mobile-app-average-user-acquisition-cost/>.
- [10] [n.d.]. RootBeer. <https://github.com/scottyab/rootbeer>.
- [11] [n.d.]. Fyber User Ratings, Reviews, and Installs. <https://play.google.com/about/storelisting-promotional/ratings-reviews-installs/>.
- [12] [n.d.]. View app statistics. <https://support.google.com/googleplay/android-developer/answer/139628>.
- [13] 2018. A Healthy Tension Between Growth and Profit is a Common Phenomenon For Any Startup. <https://theaccountancycloud.com/resources/blogs/a-healthy-tension-between-growth-and-profit-is-a-common-phenomenon-for-any-startup>.
- [14] 2019. Fyber N.V. published results for Q1 2019. <https://investors.fyber.com/assets/media/fyber-nv-q1-2019-results-announcement-english.pdf>.
- [15] Appium. [n.d.]. Automation for apps. appium.io.
- [16] Artyom Dogtiev. 2020. Top App Install Ad Platforms (2019). businessofapps.com/guide/app-install-ad-networks/.
- [17] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The Menlo report. *IEEE Security & Privacy 10* (2) (2012), 71–75.
- [18] Alex Beutel, WanHong Xu, WenKatesan, Christopher, and Christos. 2013. Copy-Catch: Stopping Group Attacks by Spotting Lockstep Behavior in Social Networks. In *WWW*.
- [19] Juan Caballero, Chris Grier, Christian Kreibich, and Vern Paxson. 2011. Measuring pay-per-install: The commoditization of malware distribution. In *USENIX Security Symposium*.
- [20] Qiang Cao, Xiaowei Yang, Jieqi Yu, and Christopher Palow. 2014. Uncovering Large Groups of Active Malicious Accounts in Online Social Networks. In *ACM CCS*.
- [21] Caroline Banton. 2020. Public Company. <https://www.investopedia.com/terms/p/publiccompany.asp>.
- [22] Craig Chapple. 2019. Devs lose out on thousands of dollars as Apple cracks down on offerwall ads. <https://www.pocketgamer.biz/news/70610/devs-lose-thousands-of-dollars-as-apple-cracks-down-on-offerwall-ads/>.
- [23] Hao Chen, Daojing He, Sencun Zhu, and Jingshun Yang. 2017. Toward detecting collusive ranking manipulation attackers in mobile app markets. In *ACM*

[24] Crunchbase. [n.d.]. Crunchbase. www.crunchbase.com.

[25] Louis DeKoven, Trevor Pottinger, Stefan Savage, Geoffrey Voelker, and Nektarios Leontiadis. 2018. Following Their Footsteps: Characterizing Account Automation Abuse and Defenses. In *ACM Internet Measurement Conference (IMC)*.

[26] Shehroze Farooqi, Muhammad Ikram, Emiliano De Cristofaro, Arik Friedman, Guillaume Jourjon, Mohamed Kaafar, Zubair Shafiq, and Fareed Zaffar. 2017. Characterizing Key Stakeholders in an Online Black-Hat Marketplace. In *IEEE/APWG Symposium on Electronic Crime Research (eCrime)*.

[27] Shehroze Farooqi, Fareed Zaffar, Nektarios Leontiadis, and Zubair Shafiq. 2017. Measuring and Mitigating OAuth Access Token Abuse by Collusion Networks. In *ACM Internet Measurement Conference (IMC)*.

[28] Geri Terzo. 2020. The Impact of Earnings Announcements on Stock Prices. <https://finance.zacks.com/impact-earnings-announcements-stock-prices-4265.html>.

[29] Gurdeep Singh. 2018. 29 App Store Stats 2018-19 that Proves Apple's Uprising Growth Trend. <https://appinventiv.com/blog/apple-app-store-statistics/>.

[30] Hamed Haddadi. 2010. Fighting online click-fraud using bluff ads. *ACM SIGCOMM Computer Communication Review* 40, 2 (2010), 21–25.

[31] Nestor Hernandez, Mizanur Rahman, Ruben Recabarren, and Bogdan Carburan. 2018. Fraud de-anonymization for fun and profit. In *ACM CCS*.

[32] Muhammad Ikram, Lucky Onwuzurike, Shehroze Farooqi, Emiliano De Cristofaro, Arik Friedman, Guillaume Jourjon, Mohammed Ali Kaafar, and Zubair Shafiq. 2017. Measuring, Characterizing, and Detecting Facebook Like Farms. *ACM Transactions on Privacy and Security* 20, 4 (2017), 1–28.

[33] Meng Jian, Peng Cui, Alex Beutel, Christos Faloutsos, and Shiqiang Yang. 2014. CatchSync: catching synchronized behavior in large directed graphs. In *ACM KDD*.

[34] Joseph Cox. 2019. America's DIY Phone Farmers. https://www.vice.com/en_us/article/d3naek/how-to-make-a-phone-farm, Vice.

[35] Jay Kim. 2017. How To Catch The Attention Of One Of The Most Active Startup Investors In The World. <https://www.forbes.com/sites/kimjay/2017/11/02/how-to-catch-the-attention-of-one-of-the-most-active-startup-investors-in-the-world/>.

[36] Platon Kotzias, Leyla Bilge, and Juan Caballero. 2016. Measuring PUP Prevalence and PUP Distribution through Pay-Per-Install Services. In *USENIX Security Symposium*.

[37] Ziang Ma, Haoyu Wang, Yao Guo, and Xiangqun Chen. 2016. LibRadar: fast and accurate detection of third-party libraries in Android apps. In *International Conference on Software Engineering (ICSE) Companion*.

[38] Mashable. 2015. Google Play just hit a major milestone. <https://mashable.com/2015/09/29/google-play-1-billion-users/>.

[39] Mary McHugh. 2013. The chi-square test of independence. *Biochimia medica: Biochimia medica* (2013).

[40] mitmproxy. [n.d.]. Homepage. mitmproxy.org.

[41] Alexander Moshchuk, Tanya Bragin, Steven D Gribble, and Henry M Levy. 2006. A Crawler-based Study of Spyware in the Web. In *Network and Distributed Systems Security (NDSS) Symposium*.

[42] Kazushi Nagayama. 2016. Keeping the Play Store trusted: fighting fraud and spam installs. <https://android-developers.googleblog.com/2016/10/keeping-the-play-store-trusted-fighting-fraud-and-spam-installs.html>, Android Developers Blog.

[43] Kazushi Nagayama. 2017. Google Play's policy on incentivized ratings, reviews, and installs. <https://android-developers.googleblog.com/2017/06/google-plays-policy-on-incentivized.html>, Android Developers Blog.

[44] Paul Bankhead. 2018. Improving discovery of quality apps and games on the Play Store. <http://ndrdnwbs.blogspot.com/2018/06/improving-discovery-of-quality-apps-and.html>.

[45] Phil Pearson. 2015. Eyeballs vs Engagement: the Struggle of Digital Advertising. <https://medium.com/@phlgpearson/eyeballs-vs-engagement-the-struggle-of-digital-advertising-301ffd61c773>, Medium.

[46] Sarah Prez. 2019. Password manager Dashlane closes on \$30M, adds former Spotify CMO to board. <https://techcrunch.com/2019/04/12/password-manager-dashlane-closes-on-30m-adds-former-spotify-cmo-to-board/>.

[47] Mizanur Rahman, Nestor Hernandez, Bogdan Carburan, and Duen Horng Chau. 2018. Search Rank Fraud De-Anonymization in Online Systems. In *ACM Hypertext and Social Media (HT)*.

[48] Mizanur Rahman, Nestor Hernandez, Ruben Recabarren, Syed Ishtiaque Ahmed, and Bogdan Carburan. 2019. The Art and Craft of Fraudulent App Promotion in Google Play. In *ACM CCS*.

[49] Mahmudur Rahman, Mizanur Rahman, Bogdan Carburan, and Duen Horng Chau. 2016. Fairplay: Fraud and malware detection in google play. In *International Conference on Data Mining (ICDM)*.

[50] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, and Phillipa Gill. 2018. Apps, trackers, privacy, and regulators: A global study of the mobile tracking ecosystem. In *Network and Distributed Systems Security (NDSS) Symposium*.

[51] Abbas Razaghpanah, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Christian Kreibich, Phillipa Gill, Mark Allman, and Vern Paxson. 2015. Haystack: In situ mobile traffic analysis in user space. *arXiv preprint arXiv:1510.01419* (2015).

[52] Redfin. 2019. Redfin Second-Quarter 2019 Revenue up 39% Year-over-Year to \$197.8 Million. <https://www.prnewswire.com/news-releases/redfin-second-quarter-2019-revenue-up-39-year-over-year-to-197-8-million-300895226.html>.

[53] Scott Reyburn. 2019. The depreciation of incentivized download offers in Offer Wall Edge for iOS, explained. <https://blog.fyber.com/the-depreciation-of-incentivized-download-offers-in-offer-wall-edge-for-ios-explained/>.

[54] Saima Salim. 2020. Apple App Store and Google Play Users Spent over 83 Billion USD on Mobile Apps in the Last 12 Months, Globally. <https://www.digitalinformationworld.com/2020/01/global-consumers-spent-over-83-billion-on-mobile-apps-in-the-last-12-months.html>.

[55] Sarah Prez. 2012. ASO (App Store Optimization) Is The New SEO, And Here's A Tool To Do It. <https://techcrunch.com/2012/02/28/aso-app-store-optimization-is-the-new-seo-and-heres-a-tool-to-do-it/>.

[56] Shani Rosenfelder. 2020. Global App Install Ad Spend To Double By 2022 to Hit 118 Billion USD. <https://www.appsflyer.com/blog/app-install-ad-spend/>.

[57] Jonghyuk Song, Sangho Lee, and Jong Kim. 2015. CrowdTarget: Target-based Detection of Crowdurfing in Online Social Networks. In *ACM CCS*.

[58] Gianluca Stringhini, Gang Wang, Manuel Egele, Christopher Kruegel, Giovanni Vigna, Haitao Zheng, and Ben Y. Zhao. 2013. Follow the Green: Growth and Dynamics in Twitter Follower Markets. In *ACM Internet Measurement Conference (IMC)*.

[59] Sujata Sungwan. 2019. Online automobile marketplace Droom gets \$10M in Series F from its Singapore-based holding entity. <https://yourstory.com/2019/06/startup-funding-online-automobile-marketplace-droom>.

[60] Kurt Thomas, Juan A. Elices Crespo, Ryan Rasti, Jean-Michel Picod, Cait Phillips, Marc-Andre Decoste, Chris Sharp, Fabio Tirelo, Ali Tofigh, Marc-Antoine Courteau, et al. 2016. Investigating commercial pay-per-install and the distribution of unwanted software. In *USENIX Security Symposium*.

[61] Bimal Viswanath, M. Ahmad Bashir, Mark Crovella, Saikat Guha, Krishna P. Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2014. Towards Detecting Anomalous User Behavior in Online Social Networks. In *USENIX Security Symposium*.

[62] Kurt Wagner. 2017. What kind of apps catch the attention of Silicon Valley investors? <https://www.fox.com/2017/11/18/16673982/app-growth-investors-jeremy-liew-snapchat>.

[63] Janith Weerasinghe, Bailey Flanigan, Aviel Stein, Damon McCoy, and Rachel Greenstadt. 2020. The Pod People: Understanding Manipulation of Social Media Popularity via Reciprocity Abuse. In *The Web Conference*.

[64] Zack Whittaker. 2019. Password manager Dashlane raises \$110M in Series D, adds CMO. <https://techcrunch.com/2019/05/30/dashlane-series-d/>.

[65] Zhirong Liu Zhenhua Dong Jiebo Luo Yingtong Dou, Weijian Li and Philip S. Yu. 2019. Uncovering download fraud activities in mobile app markets. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*.

[66] Hengshu Zhu, Hui Xiong, Yong Ge, and Enhong Chen. 2013. Ranking fraud detection for mobile apps: A holistic view. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. ACM, 619–628.

[67] Tai Ziv. 2015. What's Better? Incentivized Or Non-Incentivized App-Install Campaigns. <https://techcrunch.com/2015/09/10/whats-better-incentivized-or-non-incentivized-app-install-campaigns/>.

[68] Ziv Bass Spektor. 2020. Attribution model explained. <https://support.appsflyer.com/hc/en-us/articles/207447053-Attribution-model-explained>.