

Spillover Effects and Freemium Strategy in the Mobile App Market

Yiting Deng

Anja Lambrecht

Yongdong Liu*

Working Paper

December 16, 2020

*The authors are listed alphabetically. Yiting Deng is an Assistant Professor at the UCL School of Management, University College London (email: yiting.deng@ucl.ac.uk). Anja Lambrecht is a Professor at London Business School (email: alambrecht@london.edu). Yongdong Liu is an Assistant Professor at the UCL School of Management, University College London (email: yongdong.liu@ucl.ac.uk). The authors would like to thank AppMonsta for providing data on App Store and Google Play, and thank Dave Verwer for providing data on now-defunct website AppReviewTimes.com. We would also like to thank Catherine Tucker, Prasad Vana, Rob Waiser, participants at the 2017 Marketing Science Conference, the 2017 Marketing Dynamics Conference, the 2018 Berlin IO Day, and the Workshop on Perception and Public Policy at Bar Ilan University, as well as seminar participants at Cornell, ESSEC Business School, HKUST, Johns Hopkins University, Peking University, SHUFE, Tsinghua University, Sun Yat-sen University, UIBE, University of California, Irvine, University of Southern California, University of Colorado Boulder, University College London, University of Cambridge, University of Mannheim, University of Toulouse, Western University, Wharton and Yale for useful comments. Elias Djurfeldt, Julian Hohlweg and Lucas Weidenholzer provided excellent research assistance. The usual disclaimer applies.

Abstract

“Freemium” whereby a basic service level is provided free of charge but consumers are charged for more advanced features has become a popular business model for firms selling digital goods. However, it is not clear whether the launch of a free version helps or hurts the demand of an existing paid version. The free version may allow consumers to sample the product before making a purchase decision and subsequently increase demand of the paid version but it may also cannibalize demand of the paid version. We use a comprehensive data set on game apps from Apple’s App Store that tracks the launch of both the paid and the free versions of individual apps on a daily level, to identify whether a freemium strategy stimulates or hurts demand of an existing paid version. We estimate the spillover effects between the free version and the paid version of the same app accounting for app-level product heterogeneity. We find that the launch of a free version increases demand of the paid version of the same app and present multiple robustness checks for these results. We then present evidence that the results are driven by consumers sampling the free version as well as enhanced app discovery.

Keywords: Freemium, mobile app, spillover effect, cannibalization, sampling

JEL Classification Codes: C21, D12, L15, M31

1 Introduction

For firms selling digital products, “freemium” – providing a basic service level for free but charging consumers for using their “premium” version – has in recent years become an increasingly popular business model. Well-known examples include Dropbox, Evernote, LinkedIn and Spotify. Freemium is a prevalent pricing model for mobile apps, a market that is predicted to total almost 935 billion U.S. dollars in worldwide revenues by 2023 (Statista, 2019). In 2014, about half of the apps downloaded from the Google Play store were freemium, and these apps garnered 98% of global Google Play revenue (AppAnnie, 2015). Similarly, 92% of revenue on Apple’s App Store (“App Store” hereafter) is generated by freemium apps (Kosner, 2015).¹ For example, the game Fruit Ninja is available for free with a separate paid version for a price of \$1.99.

Understanding whether and when offering a free version alongside a paid version is an effective strategy is important: the free version may allow consumers to sample the product before making a purchase decision and so ultimately increase demand of the paid version. But the free version may also cannibalize demand of the paid version if consumers who otherwise would have paid instead opt for the free version. As of yet, there is contradictory evidence on whether the launch of a free version helps or hurts the demand of the existing paid version. While some findings suggest that offering feature-limited free versions in the mobile app market negatively affects the paid app’s adoption speed (Arora et al., 2017), others suggest that a free version increases sales of the paid app (Liu et al., 2014).

The objective of the present paper is threefold. First, we aim to identify the effect of introducing a free version of an app on the demand of its existing paid counterpart by focusing on a short time window around the free version’s launch. Thus, we circumvent identification issues encountered by prior research that analyzed an extended time period where broader changes in the marketplace, in patterns of consumer demand or in the app itself may have affected the result. Second, we aim

¹During the period we study, apps with a freemium offering typically offered both a free and a paid app version of the same app. More recently, firms have started to adopt another variant of freemium: they offer free apps with in-app purchases. The data by AppAnnie (2015) and Kosner (2015) cover both types of freemium apps.

to provide evidence that the patterns we observe are indeed driven by the free version's launch and not by correlated marketing actions. Third, we explore the mechanisms behind the results in order to make recommendations to firms on when offering a free version can be an effective strategy.

We use daily data on individual game apps from App Store. We focus the analysis on apps that have both a paid version and a free counterpart, and on a tight four-week window around the launch of the free version. We find that the launch of the free version increases the demand of the paid version. The interpretation that this effect is causal holds under the assumptions that the firm cannot time precisely the date for the free version to be approved and launched by the platform and that, within this brief four-week time period around the launch, no other actions correlated with the launch affect demand. Both assumptions are supported by our data. We then conduct a battery of checks. We demonstrate the effect holds under alternative specifications and report a falsification test. Further, we provide evidence suggesting that the effect is not driven by other events correlated with the introduction of the free version, such as changes to the paid version's design or price or advertising and promotion.

We then explore two potential mechanisms. First, we explore the possibility that the free version allows consumers who are uncertain about the quality of the product to sample before purchasing the full version (Bawa and Shoemaker, 2004; Lee and Tan, 2013; Shoemaker and Shoaf, 1975; Wang and Zhang, 2009). We conceptualize the process of consumers sampling the free version and upgrading to the paid version. We then propose how observable characteristics of an app including its number of ratings, the average star rating and its size contribute to a consumer's likelihood to sample and upgrade. Perhaps most importantly, we suggest that sampling should be most effective for apps that have a moderate number of ratings and a moderate star rating. This is because consumers who expect a product might be of very low quality will find the expected benefits too low given the cost associated with sampling. At the same time, if the product's quality is guaranteed to be above a high threshold, consumers may prefer to purchase the paid version without sampling. The patterns in our data support this relationship. In addition, we suggest that sampling should be more effective for apps that are larger in size, because these apps generally

contain richer content and functionalities, and thus have higher potential of providing greater value to consumers. This relationship is also supported by the data. Last, we propose that the effect of sampling should be more pronounced when the paid version provides significantly greater utility for consumers to upgrade. Again, we find evidence consistent with these predictions, supporting further that sampling is one mechanism behind the findings.

Second, we examine whether the free version may have enhanced the paid version's visibility thus facilitating app discovery for consumers in a market where millions of options are available (Li et al., 2016).² In such a setting, having two versions of the same app can make the app more visible to consumers. We provide two pieces of evidence supporting that enhanced visibility can increase demand in the market for mobile apps. First, we document for apps that initially launched a free version and later added a paid version that the introduction of a new paid version increases the initial free version's demand. Second, we demonstrate that the effect of launching a second version decreases as the size of the category increases, suggesting that the effect on the marginal probability to discover an app becomes less pronounced as a category becomes more crowded. Both patterns are consistent with an app's additional version enhancing visibility of the original version but cannot directly be attributed to sampling.

Our work relates to three streams of academic research. First, it relates to a nascent stream of empirical research on freemium pricing when firms offer both a free and a paid version of a digital product. Existing research has demonstrated the effect of providing a free version in various contexts such as software (Chen et al., 2017; Lee and Tan, 2013; Runge et al., 2016), cloud storage (Lee et al., 2017), digital TV service (Foubert and Gijsbrechts, 2016) and on-demand dramas (Hoang and Kauffman, 2018). For apps on Google Play, Liu et al. (2014) find that paid apps with a free counterpart receive more sales than those that do not have a free counterpart. By contrast, in the same empirical context Arora et al. (2017) find that offering free versions reduces the paid app's adoption speed. However, this existing literature has focused on an overall effect of a free version, without being able to tie it directly to the actual timing of the introduction of the free

²See also <http://www.tbray.org/ongoing/When/201x/2010/07/18/How-to-run-an-App-store> and <https://www.cnet.com/news/the-mobile-app-discovery-problem/>.

version. Our work adds to this literature in three ways. First, in order to identify the effect of the free version's launch, we identify the precise launch date of the free version and focus our analysis on a tight time window around the launch. We demonstrate that launching a free version increases the demand for the paid version. Second, we discuss in detail that other marketing actions that could in theory be correlated with the launch are unlikely to cause this effect of interest. Third, to our knowledge we are the first to empirically identify mechanisms that lead to a spillover effect of a free version on an existing paid version. Specifically, we provide evidence for both sampling and app discovery as the mechanisms.

The second stream of research we contribute to is a broader literature on the pricing of digital goods. Pauwels and Weiss (2008) demonstrate for an online content provider for marketing professionals that moving from free to fee can be profitable, despite loss of advertising revenue. Chiou and Tucker (2013) find that paywalls decrease viewership of an online newspaper and, as a result, depress advertising revenues. Lambrecht and Misra (2016) show that an online content provider that also generates advertising revenues can benefit from offering more free content in periods of high demand than in periods of low demand. Kannan et al. (2009) explore how a firm that offers a print and a PDF version should set prices. We contribute to this literature by demonstrating that offering a free version of a product can generate additional demand for the paid version, potentially making this a profitable pricing strategy.

Third, our research relates to largely analytical work on sampling and versioning. Theory literature has long discussed the role of sampling of free products. The earlier literature focused on sampling of physical goods (Bawa and Shoemaker, 2004; Heiman et al., 2001; Jain et al., 1995). More recently, researchers have investigated free trials of digital goods that limit available features (Chellappa and Mehra, 2017; Dey and Lahiri, 2016; Halbheer et al., 2014; Lahiri and Dey, 2018; Niculescu and Wu, 2014; Wang and Zhang, 2009; Cheng and Liu, 2011) or limit the time of use (Dey et al., 2013; Cheng and Liu, 2011). For example, Halbheer et al. (2014) model an online content publisher's choice between offering paid content, sampling and free content. They find that it can be optimal for the publisher to generate advertising revenue by offering free samples

even when sampling reduces prior quality expectations and content demand. Appel et al. (2019) develop an analytical framework for app monetization, and demonstrate that even if advertisers do not pay for ads, an app developer can still profit from offering a free version of the app. Gu et al. (2018) and Li et al. (2019) conduct field experiments at a book publisher to investigate how to design free samples for books to increase revenue. We add to this literature by providing empirical evidence that consumers' sampling of digital products can indeed drive consumers' decision to purchase the paid version and so potentially increase firm revenue, and demonstrating conditions under which the effect is more prominent.

At the same time, our findings are relevant to managers. They demonstrate, at least in the context of mobile apps, that a freemium strategy can indeed increase demand for the paid version of a product. At the same time our results suggest that firms may not want to indiscriminately offer free versions but instead assess the likely benefits and cost given the specific product under consideration. First, firms should particularly consider investing into a free version when the paid product is of medium quality. If a product is of low quality, opportunity cost associated with the free version may prevent consumers from sampling while for high quality products, consumers may forgo sampling and instead directly purchase the paid version so in either case a free version may not be as effective in driving demand. Second, firms need to ensure that the paid version offers substantial additional benefits relative to the free version as only then will it be attractive for consumers to upgrade. Third, especially in crowded marketplaces such as the market for mobile apps – at the time of our data, App Store offered close to 200,000 game apps – a free version may be useful as such simply because it enhances product visibility.

2 Industry Background and Data

2.1 Industry Background

Since Apple and Google launched their respective mobile application stores in 2008, the number of apps has grown exponentially: App Store has now 3.42 million available apps as of July 2020 (Statista, 2020b) and Google Play offers about 3.04 million apps as of September 2020 (Statista,

2020a). Each platform offers a wide range of different apps developed by a large number of independently operating developers, the vast majority of which are small. Because the iOS and Android operating systems are incompatible, apps are not portable between the two platforms without a substantial change in the source code. In this paper, we focus on App Store, which, during the sample period (2011-2013), was the largest app platform.

During our sample period, a developer who adopted a freemium strategy typically offered both a free and a paid version of an app. The free version may have differed from the paid version along a number of dimensions but generally had a reduced set of capabilities. To access the full version, users were required to separately download the paid app. In recent years, a different variant of freemium pricing, in-app purchases, has become popular. While App Store started allowing for in-app purchase microtransactions in 2009, these were not yet common during the period of our data.

2.2 Data Collection

We use a comprehensive data set that was collected from App Store between September 2011 and October 2013. The data set is provided to us by AppMonsta, a company specialized in app data analysis and consulting. The company scraped the U.S. App Store on a daily basis and each day collected all data that were publicly available.

We focus on game apps as it is the most popular category in App Store and purchases tend to be for personal use, not business purposes. In addition, game apps are highly heterogeneous in terms of functionality and quality, thus, before using it, consumers are likely to have limited knowledge about how well an app fits their needs. In total, the data cover 197,884 mobile game apps, of which 95,990 (48.5%) are paid apps and the remaining are free apps. Among these paid apps, 55,320 were released during our data period. The data include the name of the developer, app characteristics such as price, version, size, genre, and release date, as well as user ratings, reviews and app rankings on various top charts.³ In the data, the paid version and the free version of the

³App Store provides three types of top charts: Top Free, Top Paid, and Top Grossing. The Top Free chart ranks free apps. The Top Paid chart ranks paid apps. The Top Grossing chart ranks all apps based on total revenue. For each genre (e.g., Action, Board, Trivia), there is a top free chart, a top paid chart, and a top grossing chart.

same app are marked as different products. Table 1 summarizes the relevant variables in the raw data, some of which such as app price, the number of ratings, the star rating (on a scale of 1 to 5) and ranking can change on a daily basis.

Table 1: Variables in Raw Data

Variable Category	Item	Description
App Identifier	App name	App name shown on mobile devices
	App ID	A string of numbers and/or characters that uniquely identifies an app
Developer Identifier	Developer name	Developer name shown on mobile devices
	Developer ID	A string of numbers and/or characters that uniquely identifies a developer
App Characteristics	App price	App price and currency unit
	App size	App size in bytes
	Release date	App launch date
	Genre	App category
	App description	One or several paragraphs outlining app functions
Ratings	Version	App version number
	Number of ratings	Both daily incremental and accumulative number of ratings received
Ranking	Star ratings	Daily and overall ratings
	Ranking on top charts	Rankings on each related genre's top charts

2.3 Freemium Apps

We define freemium apps as paid apps that have a free counterpart. We take three steps to identify among all available apps those apps that are freemium apps (see Online Appendix A for details). First, we identify apps launched by the same developer. Second, within these apps, we compare similarities in app names for all possible app pairs. The pairs that are associated with high similarity measures are identified as potential freemium pairs. Third, we manually check all pairs to make sure they are indeed different versions of the same app. The free version and the paid version can usually be distinguished by words such as “lite”, “free”, “pro”, and “premium”.

We identify 9,806 app pairs that adopted the freemium strategy. For 8,284 (84.5%) apps, the paid version was launched first. In only 1,522 (15.5%) cases, the free version was launched first. There are no instances when the two versions were launched at the same time. We focus on the

vast majority of cases where the paid version was launched first. Furthermore, for identification purposes, we drop from our sample 5,586 app pairs where the free version was launched before the start of our sample period. For 167 apps, data are missing for an extensive period of time after the release and are dropped from the sample. Thus, we retain a total of 2,531 app pairs which form the focus of our empirical analysis. For these apps, on average, the gap between the two versions' launch dates is 115 days (s.d. = 200 days). We will later use the 87,706 paid apps that did not have a free version during our sample period, as control apps.

For each app, App Store displays a description of the app on its download page. We use the descriptions of the two versions displayed on the day when the free version was launched to classify differences between the paid and the free versions. Research assistants manually coded the differences into six categories. Each app was coded by two coders independently. We manually checked inconsistent codings to make final decisions. Table 2 summarizes the six dimensions along which the two versions may differ. In 34.6% of cases, the paid version allows the user to progress to more game levels than the free version. For example, the description of the game Phase 10's free version indicates "Phase 10 Free lets you play 3 random phases. Update to the full version for \$.99 and play all 10 phases." Typically, this means a more complete, advanced or challenging game experience. 17.9% of paid versions offer more modes or themes than the free versions. For instance, Bean's Quest's paid version includes five themes (worlds) whereas the free version includes three themes. This means that while the ability to progress or the difficulty of the game does not change, the user experience can be customized. In 16.7% of cases, the paid version offers more functions or features (e.g., more powerful weapons, record keeping) than the free version. Further, in 11.7% of cases, the paid version allows for social interactions while the free version does not. For example, the paid version may be integrated with Game Center, Apple's social gaming network, which enables users to track their best scores on a leaderboard, compare their achievements, invite friends to play a game, and start a multiplayer game through auto-matching. The paid version may also allow the user to link to Facebook to post scores and share progress with friends. Lastly, in 11.0% of cases, the paid version may be ad-free while the free version is

ad-supported and in 1.5% of cases, the paid version provides better user support, for example an email contact to address user questions. These differences are not mutually exclusive, and there can be multiple differences between the paid version and the free version. For 43.3% of apps, there is no difference in the descriptions of the two versions.⁴

Table 2: Version Differences based on App Description

Difference	Percent
More levels	34.6%
More modes/themes	17.9%
More functions/features	16.7%
Social interactions	11.7%
Ad free	11.0%
Better support	1.5%
No difference	43.3%

2.4 Descriptive Statistics

For our analysis, we focus on App Store and retain two types of data: paid apps that also have a free version (2,531 apps), and paid apps that did not have a free version in our data period (87,706 apps). Table 3 provides descriptive statistics for these apps, measured on the last day of the sample period (October 31, 2013). Columns (1) and (2) respectively show the descriptive statistics for the paid version and the free version of the freemium apps, and Column (3) shows the descriptive statistics for apps that only had a paid version. The data suggest that the paid versions of freemium apps are broadly similar to apps that only exist as a paid version but tend to have a somewhat lower price, somewhat higher star rating, and appear somewhat more likely to be ranked highly. In addition, compared with its paid counterpart, the free version tends to receive somewhat fewer ratings and somewhat lower star rating.

⁴Note that for these apps, the two versions may still differ along one or more dimensions, but we were unable to identify the differences from our data.

Table 3: Descriptive Statistics of Free Apps

	(1) Paid version of freemium apps	(2) Free version of freemium apps	(3) Only paid
Price (\$)	1.2 (1.1)	0.0 (0.2)	1.8 (6.5)
If rated	39.6% (48.9%)	27.3% (44.6%)	37.0% (48.27%)
ln (No. ratings+1)	1.4 (2.0)	0.9 (1.6)	1.4 (2.2)
Average star rating	4.2 (0.7)	4.0 (0.7)	3.6 (0.9)
Ranked on top 10	0.04% (1.99%)	0.00% (0.00%)	0.02% (1.43%)
Ranked on top 11-20	0.04% (1.99%)	0.04% (1.99%)	0.01% (1.22%)
Ranked on top 21-50	0.00% (0.00%)	0.16% (3.97%)	0.04% (2.00%)
Ranked on top 51-100	0.24% (4.86%)	0.28% (5.25%)	0.09% (2.94%)
Ranked on top 101-	1.82% (13.36%)	5.53% (22.86%)	2.87% (16.70%)
Age (Days)	515 (291)	400 (220)	691 (459)

Note: Each entry is the average value of the corresponding variable for the group, with standard deviations in parentheses

2.5 Google Play as an Alternative Data Source

Our main analysis focuses on data from App Store. We collect similar data for the same data period (September 2011 to October 2013) from Google Play. In total, the data cover 150,822 game apps of which 124,004 are free. The sample of freemium apps in Google Play is significantly smaller, with a total of 405 cases where a free version of a pre-existing paid version was launched during our data period. To allow for a meaningful analysis, we need to observe the paid version's number of ratings for two weeks before the launch of the free version as well as for two weeks after the launch of the free version, reducing the data to 217 pairs of freemium apps.⁵ The significantly smaller sample size means that we focus the majority of our analysis on App Store but demonstrate the robustness of results using the data from Google Play.

2.6 Proxying App Demand

Ideally, we would want to explore the effect of launching a free version on the number of downloads of the existing paid version. As download data are not publicly available,⁶ we use the number

⁵This means we exclude apps where the gap between launches of the free and paid versions was less than two weeks as well as those where the free version was launched close to the start or the end of the data period. More generally, two reasons for the smaller number of freemium apps are that at the time of our data, Google Play was still smaller than App Store and had a greater focus on free apps.

⁶Download data are confidential and only held by the platforms (App Store or Google Play) and individual developers, typically small firms. Third-party firms offer so-called "download data" for sale but these firms do not access

of ratings an app has accumulated as a proxy for demand. Using ratings as a proxy for demand is possible because the free and paid versions are treated as separate products in App Store meaning that each has a separate measure for the number and valence of ratings. In this section, we first explain why the number of ratings of an app is a good proxy for demand. We then discuss alternatives proxies for demand.

Number of ratings: Our approach to proxy demand by an app's number of ratings builds on previous literature. Kummer and Schulte (2019) use the number of ratings as the preferred demand variable for Google Play apps, and demonstrate that the number of incremental ratings correlates with download ranges.⁷ Boudreau (2018) reports that for apps sold in App Store, sales are highly correlated with the number of user ratings. He then approximates the rank order of app demand by the rank order of the numbers of user ratings.

We test whether the number of ratings and the number of downloads are indeed highly correlated as we assume. Such a test is difficult to implement using data from App Store as no download information is provided. However, Google Play reports for each app and day a range for the cumulative number of downloads an app has reached as of yet.⁸ These download ranges allow us to broadly test the relationship between the number of ratings and downloads. To do so, we use all available apps (i.e., not limited to freemium apps) in Google Play data and test the relationship between the number of ratings and the upper, respectively lower, bound of the download range. We include in the regression app fixed effects to control for the variation in the propensity across apps to obtain their first rating. Table 4 displays the estimation results. We use as dependent

true download data at scale. Instead, they typically build models to estimate downloads based on apps' rankings, sometimes supplemented with information on release date and ratings and calibrate the models based on downloads of a selection of apps. They then estimate downloads only for days on which an app was ranked. Thus, these firms only offer a proxy for downloads, the data are skewed to more successful apps on days when they performed particularly well and the scope of such data would be too narrow as very few apps are ranked consistently across a 28- or 14-day time period, which we will require for identification. As a result, such data are not useful for our purpose.

⁷They also use change in the number of installations as an alternative measure but point out that this variable is less preferred than changes in the number of ratings, because the number of installations is observed only in discrete step-size form which leads to low intertemporal variation.

⁸Google Play reports data for 18 download ranges: [1, 5], [5, 10], [10, 50], [50, 100], [100, 500], [500, 1,000], [1,000, 5,000], [5,000, 10,000], [10,000, 50,000], [50,000, 100,000], [100,000, 500,000], [500,000, 1,000,000], [1,000,000, 5,000,000], [5,000,000, 10,000,000], [10,000,000, 50,000,000], [50,000,000, 100,000,000], [100,000,000, 500,000,000], and [500,000,000, 1,000,000,000].

variables the logarithm of the upper bound in Column (1), respectively lower bound in Column (2), of the download range, and as independent variable the logarithm of the number of ratings.⁹ In both cases, $R^2 = 0.97$, indicating a strong relationship between the logarithm of the bounds of download ranges and the logarithm of the number of ratings. To rule out that the results are driven by highly popular apps, Columns (3) and (4) replicate Columns (1) and (2), excluding app-day observations that have more than 492 ratings (95th percentile). The results are robust and together give us confidence that the number of ratings is indeed a good proxy for demand.

Table 4: Relationship between the Number of Ratings and Download Range Bounds on Google Play

Variables	All observations		Exclude observations with ratings over 95th percentile	
	(1) ln(Upper bound+1)	(2) ln(Lower bound+1)	(3) ln(Upper bound+1)	(4) ln(Lower bound+1)
ln (No. Ratings+1)	1.245*** (0.0001)	1.240*** (0.0001)	1.298*** (0.0001)	1.294*** (0.0001)
Constant	4.583*** (0.0002)	3.282*** (0.0002)	4.507*** (0.0002)	3.204** (0.0002)
App FE	Yes	Yes	Yes	Yes
Number of Observations	279,547,860	279,547,860	265,555,690	265,555,690
Number of Apps	1,169,351	1,169,351	1,158,507	1,158,507
Adj. R-squared	0.97	0.97	0.96	0.96

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Despite the strong relationship between download range bounds and number of ratings, relying on any proxy for sales could introduce bias. In our context, we expect that if such a bias existed, it would lead us to underestimate, not overestimate, the effect of the launch of the free version on the sales of the paid version. Prior research suggests that the number of new consumer reviews decreases in the total number of reviews available (Wu and Huberman, 2008; Godes and Silva, 2012). By implication, the fraction of paid app downloads that lead to app ratings should have a downward trend. Thus, if the free version's launch does not affect the paid version's downloads, we would expect the paid version to accumulate fewer daily ratings after the free version has been

⁹We use the logarithm of these variables as both the download range bounds and the number of ratings are skewed.

launched. If instead we observe an increase in the paid version's daily incremental number of ratings after the free version's launch, this would strongly suggest an increase in the paid version's downloads. We further acknowledge the possibility that consumers might delay rating an app after downloading, which would again lead to an underestimation of the effect immediately following adoption.

Download ranges: As an alternative to the number of ratings as a proxy for demand, previous research has used download ranges reported by Google Play to proxy for demand (Liu et al., 2014; Arora et al., 2017). As these ranges vary in size and can be very broad (the lowest download range is 0 to 5, the highest 500,000,000 to 1,000,000,000; see footnote 8), this approach requires apps to have a significant variation in downloads, for example because apps quickly become very popular or data cover a long time period.

Proxying demand using download ranges is difficult for two reasons. First, App Store does not report download ranges. Second, there are significantly fewer freemium apps on Google Play, and in the Google Play data, apps rarely change download ranges, likely owing to the large width of most ranges. For an individual app, we can proxy the number of downloads by the upper bound of the previous interval or the lower bound of the new interval at the time of a download range change. Observing two changes would thus allow us to infer the speed of download accumulation. However, our identification approach requires us to observe such changes within a short time window around the launch of a free version. Out of the 217 freemium apps, 54 apps have at least one change in download ranges around the time of the launch of the free version, 16 apps have two or more download range changes and no app has two download range changes both before and after the launch of the free version. Thus, such data would not provide sufficient variation to identify the effect of launching a free version.¹⁰

¹⁰Liu et al. (2014) and Arora et al. (2017) circumvent these constraints because of three characteristics of their empirical approach. First, their approaches do not require them to observe the free version's launch. Second, they extend their observation period until the end of their data period, even if this lies many weeks or months after the launch of the free version. As a result, they observe apps over a longer time period and thus have more variation in download ranges. As we discuss in more detail in §3, in our case, we are concerned that such an approach would make identification of a causal effect difficult. Third, Liu et al. (2014) focus only on apps that were ranked at least once. These are obviously popular apps, suggesting that if they are followed over a long time period, one would observe

Ranking: Since Chevalier and Goolsbee (2003) demonstrated for books a relationship between actual sales and sales rank at Amazon and Barnes & Noble, researchers have tried using ranking to proxy for sales or downloads.¹¹ In our context, however, there are multiple challenges with using ranking as a proxy for sales. First, ranks are only available for a small fraction (15%) of freemium apps during the period of two weeks before and two weeks after the launch of the free version.¹² Second, relying on ranks would require excluding unranked apps and so bias the sample towards more successful apps. Third, ranking is a relative measure as a change in rank of one app could be caused by a change in performance of another app and not the focal app. Fourth, app rankings not only reflect the number of downloads, but also account for other factors such as the average rating, freshness (e.g., recency, updates) of the app or app retention (Liu et al., 2014).

Notwithstanding, we will report a robustness check in §3.3 where we use ranking in App Store as an alternative outcome variable.

3 Empirical Strategy and Estimation

3.1 Model-Free Evidence

To identify the effect of launching a free version of a mobile game app on the demand for the existing paid version, proxied for by its number of ratings, we focus the analysis on a four-week window around the free version's launch.¹³ The assumption is that within this brief time period, changes to the incremental number of ratings can be attributed to the launch of the new version only. Figure 1 shows the average of the logarithms of the number of ratings accumulated by the paid version in the 2-week windows before and after the free version's launch. Clearly, the paid version received more incremental ratings after the introduction of the free version, suggesting

a significantly greater variation in downloads than for the majority of apps that have never been ranked. Instead, we want to ensure that our results apply to a broad range of freemium apps.

¹¹Garg and Telang (2013) build on this approach and calibrate the relationship between app ranks and sales. Ghose and Han (2014) calibrate the relationship between app ranks and sales using additional data that contain rank and download information from an app store.

¹²Kummer and Schulte (2019) demonstrate a close relationship between the number of ratings and demand measures based on app rankings, and suggest that using app rankings as the demand measure limits the scope of the analysis because rankings are available only for the most successful apps.

¹³In the remainder of the paper we simplify the language for ease of exposition. When we refer to the demand for the paid version, we refer to demand as proxied for by the number of ratings, even if we do not explicitly spell this out.

a relationship between the free version's launch and the paid version's demand. The figure also indicates that there is significant variation in the number of ratings for apps.



Figure 1: $\ln(\text{Number of Incremental Ratings} + 1)$ in a 2-Week Window

This analysis does, of course, not control for other potentially confounding factors. It is possible that the relationship in Figure 1 simply reflects a time trend. To address these concerns, we control for app fixed effects as well as a large number of covariates. We then conduct a falsification test that demonstrates that the results are not merely a reflection of a time trend. Further, we present a battery of robustness checks that show that the results are robust to alternative definitions of the time window and other marketing actions related to product, price, and promotion.

One might also be concerned that a developer would introduce a free version strategically at the time when they expect a rise in demand. There are two reasons why we believe this selection effect is limited. First, many developers in the data are small developers, who are likely to have little knowledge on whether and when it would be effective to introduce a free version for the paid app. Second, even if the developers had enough knowledge to strategically decide whether to introduce a free version, this would only influence the selection of apps but not the timing of launch, because Apple first reviews and approves each individual app before making it available.

In the data period, there was considerable uncertainty for developers over an app’s review time and in turn the precise launch date. While app review times are not usually publicly available, a now-defunct website <https://appreviewtimes.com/> tracked App Store review times crowdsourced from developers from 2011 to 2018. We obtained from the owner all the data he collected when the website was active. 9,617 review times were reported in our data period, ranging between 0 and 49 days with a mean of 7.05 days and a standard deviation of 4.05 days (see Figure 2 for boxplots of the review times by month). This high variation makes it unlikely that a developer could determine with precision the day or week of a launch to coincide with a time when demand is expected to be higher.




Figure 2: Distribution of App Review Times by Month. Each box is drawn from the 25th percentile to the 75th percentile and the horizontal line in the middle denotes the median.

3.2 Empirical Strategy and Estimation

Ideally, to estimate the effect of the launch of a free version on the existing paid version's number of ratings, we would compare the paid version's number of ratings in the presence of a free version to its number of ratings in the absence of a free version, holding everything else constant. However, we cannot simultaneously observe for the same paid version of an app the number of ratings in the presence and in the absence of a free version. As a result, we focus our analysis on a tight time window around the launch and compare the paid version's incremental number of ratings before the launch of the free version to its incremental number of ratings after the launch. We include app fixed effects to control for time-invariant app features that can potentially influence rating accumulation.

We specify the econometric model as follows:

$$\ln(N_{it} + 1) = \alpha + \beta T_{it} + X_{i,t-1}\gamma + \delta_i + \varepsilon_{it}, \quad (1)$$

where i denotes a paid app, t denotes day. N_{it} denotes the number of ratings accumulated by app i on day t which we use to proxy for demand. We specify our dependent variable in log-form as the data are skewed. T_{it} denotes treatment: $T_{it} = 0$ if t is in the window before the launch of the free version and $T_{it} = 1$ if t is in the window after the launch of the free version. $X_{i,t-1}$ denotes the vector of control variables measured on day $t - 1$, including ranking on top charts, price, the total number of ratings app i had accumulated, and the average star rating of the app. δ_i is an app fixed effect and captures unobservable app characteristics. ε_{it} is an idiosyncratic error term.

Table 5 summarizes the estimation results when focusing on a time window of two weeks before the launch of the free version and two weeks after the launch. We include in this analysis 1,710 apps as we need to ensure the paid version's number of ratings is observed in the entire time window before and after the free version's launch. This means that the gap between the release dates of the two versions must be at least two weeks and the free app must be launched at least two weeks before the end of our data period. We cluster standard errors on app level.

Table 5: Window Analysis

Variables	ln (Number of Ratings of Paid Version+1)			
	OLS			Type I Tobit
	(1)	(2)	(3)	(4)
<i>T</i> (After new version launch)	0.0184*** (0.0042)	0.0179*** (0.0036)	0.0290*** (0.0036)	0.2109*** (0.0520)
Lag (Ranked on top 10)		1.4530*** (0.3002)	1.4786*** (0.2824)	2.8645*** (0.3754)
Lag (Ranked on top 11-20)		1.5997*** (0.2834)	1.5914*** (0.2732)	2.6918*** (0.5283)
Lag (Ranked on top 21-50)		1.3159*** (0.3292)	1.3365*** (0.3369)	3.0978*** (0.2990)
Lag (Ranked on top 51-100)		0.7745*** (0.1666)	0.7700*** (0.1679)	2.4258*** (0.2957)
Lag (Ranked on top 101-)		0.1142*** (0.0266)	0.1102*** (0.0271)	1.2627*** (0.0971)
Lag (Price (\$))			-0.0027 (0.0025)	-0.0201 (0.0329)
Lag (ln (Total # ratings+1))			0.0139 (0.0499)	0.7354*** (0.0403)
Lag (If rated)			-0.6287*** (0.2245)	-2.1235*** (0.4539)
Lag (Average star rating)			0.0939** (0.0387)	0.2355** (0.0938)
App FE	Yes	Yes	Yes	
App RE				Yes
Number of Observations	49,205	49,205	49,186	49,205
Number of Apps	1,710	1,710	1,710	1,710
Adj. R-squared	0.5002	0.5265	0.5315	-

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Columns (1)-(3) report OLS estimates. Column (1) controls only for app fixed effects. It indicates that the paid app receives more ratings in the window after the launch of the free version. Column (2) additionally controls for rankings. The result holds and it appears that the number of ratings increases for higher ranked apps. Column (3) shows that the result is robust to including the full set of controls. Not surprisingly, the number of ratings increases with the average star rating. To deal with the concern that many apps receive no rating on a given day, we replicate Column (3) with a Type I Tobit model. The results reported in Column (4) are qualitatively similar to those in Column (3). As discussed, the interpretation that this effect is causal holds under the assumptions

that the firm cannot time precisely the date for the free version to be approved and launched by App Store and that, within this brief time period, no other factors correlated with the launch affect demand. The first assumption is supported by Figure 2. We discuss the validity of the second assumption in §3.3.

Table 5, Column (3) indicates that the introduction of a free version increases the paid version's log number of number of ratings by 0.029. The mean value of the dependent variable in the data is 0.068, so that on average, the dependent variable increased by 42.6% following the free version's introduction, which is not insubstantial.

3.3 Robustness Checks

In this subsection, we present a battery of robustness checks. First, we demonstrate that the results are robust to alternative specifications, and are not driven by a general trend (§3.3.1). Then, we explore whether the effect could arise from developers systematically coordinating other marketing activities (e.g., changes to existing product's design or price, advertising and App Store promotion) with the launch of the free version (§3.3.2).

3.3.1 Existence of Effect

Window length: Table 5 focuses on a 2-week window before and after the introduction of the free version. We check that our results are robust to the specification of the observation window around the time when the free version was launched. Table 6 replicates the results presented in Column (3) of Table 5 for alternative windows before/after the launch.¹⁴ We find the effect holds. The effect sizes under 1-week (Column 1) and 3-week (Column 2) windows are not statistically different from the effect size under a 2-week window, and the effect size under a 4-week window (Column 3) is smaller than that under a 2-week window, suggesting the effect is stronger in the shorter term.

¹⁴Note that to conserve space, in Table 6 and subsequent tables, we do not report the estimates associated with other control variables, which are qualitatively similar to those in Column (3) of Table 5.

Table 6: Robustness Checks on Window

Variables	ln (Number of Ratings of Paid Version+1)		
	(1) 1-week window before/after	(2) 3-week window before/after	(3) 4-week window before/after
<i>T</i> (After new version launch)	0.0313*** (0.0037)	0.0245*** (0.0037)	0.0174*** (0.0038)
Other control variables (ranking, price, the total number of ratings, the average star rating)	Yes	Yes	Yes
App FE	Yes	Yes	Yes
Number of Observations	31,878	60,057	68,050
Number of Apps	2,159	1,404	1,199
Adj. R-squared	0.5287	0.5458	0.5811

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Alternative dependent variables: While our main analysis uses the number of ratings as a measure of demand, we replicate the result with an app’s ranking, instead of the number of ratings, as an alternative outcome variable. We focus on the 28 apps for which within a window of one week before and one week after the launch we observe a ranking on each of the fifteen days.¹⁵ Column (1) of Table 7 illustrates that ranking improves after the launch of the free version, consistent with our prior finding.

Next, we replicate the analysis with data from Google Play, estimating Equation (1) with the 217 freemium apps in Google Play. Column (2) of Table 7 demonstrates that the results are consistent, and the effect size is not statistically different from the effect size for App Store.

¹⁵As described in §2.6, there is a great variation in the ranking of apps. This means that many apps are ranked among the top 1000 apps on one day and drop out of the list on the subsequent day. But for the purpose of the estimation, we need a complete ranking history within the window. If we were to use the window of two weeks prior to and after the launch, the data set would reduce to only 18 apps.

Table 7: Robustness Checks on Alternative Dependent Variables

Variables	(1) ln (ranking)	(2) ln(Number of Ratings of Paid Version+1) in Google Play
<i>T</i> (After new version launch)	-0.0961* (0.0512)	0.0242** (0.0096)
Other control variables (ranking, price, the total number of ratings, the average star rating)	Yes (except ranking)	Yes
App FE	Yes	Yes
Number of Observations	420	6,233
Number of Apps	28	217
Adj. R-squared	0.8373	0.6217

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Falsification test We wish to rule out that we measure a general trend in the number of ratings an app accumulates over time rather than the effect of the launch of the free version. For this purpose, we run a falsification test using as control group the 87,706 paid apps that did not have a free counterpart during the observation period. For each control app, we generate a random date as the “placebo free version launch date.” This date is evenly drawn from the time period starting from 14 days after the app’s launch date or the first day of our data period, whichever is later, to 14 days before the last day of our data period. We then estimate Equation (1) using the control group data for a window of two weeks before and two weeks after the placebo launch date.¹⁶

Column (1) in Table 8 shows a negative coefficient, suggesting that the number of daily ratings an app accumulated after the placebo launch date is smaller than those accumulated prior to that date. Results are robust if we remove apps that have no rating in either window, or when removing apps that appear in the top charts in either window. This result is consistent with prior research that has pointed to a decreasing trend in rating accumulation over time (e.g., Wu and Huberman, 2008; Godes and Silva, 2012). When adding the full set of controls in Column (2), the main coefficient

¹⁶Alternatively, one could use a matching procedure to match the freemium apps with other paid apps, narrowing down the group of control apps further to more closely reflect the observable characteristics of freemium apps. We follow this approach in the Online Appendix B and find similar results.

of interest becomes positive though its magnitude is much smaller than that of the corresponding coefficient for freemium apps in Table 5. We then pool the control group data with data on the treatment group (freemium apps) to obtain a difference-in-difference estimator and present the results in Column (3). The binary variable “treatment group” takes the value of 1 if the app has a free counterpart. The interaction term between “after ‘new version launch’” and “treatment group” therefore measures the treatment effect above and beyond the general trend. The results indicate that indeed an app is rated more frequently following the introduction of a free version and the effect is not driven by a time trend, providing further evidence that the addition of a free version increases demand of its paid counterpart.

Table 8: Robustness Checks with Control Apps

Variables	ln (Number of Ratings of Paid Version+1)		
	(1) Paid control apps	(2) Paid control apps	(3) Paid control apps + Freemium apps
T (After “new version launch”)	-0.0005* (0.0003)	0.0028*** (0.0003)	0.0027*** (0.0003)
$T \times$ (Treatment group)			0.0325*** (0.0044)
Lag (Ranked on top 10)		1.1992*** (0.1263)	1.2254*** (0.1195)
Lag (Ranked on top 11-20)		0.9408*** (0.0989)	0.9782*** (0.0967)
Lag (Ranked on top 21-50)		0.6348*** (0.0571)	0.6766*** (0.0588)
Lag (Ranked on top 51-100)		0.3326*** (0.0361)	0.3544*** (0.0356)
Lag (Ranked on top 101-)		0.0604*** (0.0041)	0.0629*** (0.0041)
Other control variables (price, the total number of ratings, the average star rating)	No	Yes	Yes
App FE	Yes	Yes	Yes
Number of Observations	2,479,511	2,479,511	2,528,697
Number of Apps	85,684	85,684	87,394
R-squared	0.5678	0.5809	0.5786

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

3.3.2 Effect is Not Driven by Correlated Marketing Actions

As discussed in §3.1, while the decision to introduce a free version may be endogenous, the precise timing of the launch of the app cannot be determined directly by the developer. Therefore, it would be difficult for the developer to coordinate other actions within the relatively short time period that we consider to precisely fit the short window following the launch. Nevertheless, we provide additional evidence supporting that this effect is due to the launch of the free version and unlike an outcome of other marketing actions (e.g., changes to existing product's design or price, advertising and App Store promotion) correlated with the launch.

Change of existing version's design: One alternative explanation for the positive spillover effect is that app developers simultaneously improve the existing version when launching the new version, thus enhancing the existing version's demand. If this were the case, our estimates might reflect an improvement in the paid version rather than an effect of the free version. To alleviate this concern, we replicate the analysis in Column (3) of Table 5 with a subset of apps that updated neither in the two-week window before nor in the two-week window after the free version's launch. Results are presented in Column (1) of Table 9 and the effect holds. Note that the effect size is lower than in Column (3) of Table 5, possibly because apps that update more frequently also differ from other apps along other characteristics (such as star rating, number of ratings, and size of the app) that can potentially affect the size of the effect. We discuss how these characteristics can matter in §4.1.

Change of existing version's price: We also check that the results hold for apps where there is no price change for the paid app. If the developer lowered the price of the paid version at the same time that they launched a free version, we might observe increased sales of the paid version coinciding with the launch of the free version simply due to the price change of the paid version. During the period of two weeks before and two weeks after the launch of the free version, 27.6% of apps had at least one price change. To ensure that price changes did not influence our results, we replicate the analysis in Column (3) of Table 5 excluding any apps that had a price

change within the four-week window around the launch. Table 9, Column (2) shows that the results hold and the effect size is not significantly different from that in Column (3) of Table 5 ($p=0.47$). Table 9, Column (3) shows the results hold when focusing on apps that had neither updates nor price changes within the four-week window around the launch, and the effect size is lower than in Column (3) of Table 5 due to the same reason as in Column (1) of Table 9.

Table 9: Robustness Checks on Changes on Existing Version

Variables	In (Number of Ratings of Paid Version+1)		
	(1) Exclude apps with updates on existing version	(2) Exclude apps with price changes	(3) Exclude apps with updates or price changes
<i>T</i> (After new version launch)	0.0125*** (0.0041)	0.0254*** (0.0034)	0.0135*** (0.0036)
Other control variables (ranking, price, the total number of ratings, the average star rating)	Yes	Yes	Yes
App FE	Yes	Yes	Yes
Number of Observations	21,637	35,620	14,737
Number of Apps	751	1,238	509
Adj. R-squared	0.6757	0.5151	0.7161

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Advertising: If developers typically advertise the free version when it was launched, this could likewise lead to awareness of the paid version. We provide five pieces of evidence from the industry and our data that suggest our results are not an outcome of a change in advertising levels.

First, during the period of our data, advertising was not yet a popular means for app developers to promote apps. In October 2012, Facebook was the first firm to enter the app install ad market.¹⁷ Google and Twitter both entered this market only in 2014 and Pinterest in March 2017.¹⁸ eMarketer did not start monitoring app install ad spending until 2013¹⁹ while Sensor Tower, a leading

¹⁷An app install ad is a mobile ad that offers a direct click to an app's page in the app store. When people click on the ad on a mobile device, they will be directed to the app store to install the app.

¹⁸<https://techcrunch.com/2014/11/30/like-advertising-a-needle-in-a-haystack/>; <http://www.adweek.com/digital/pinterest-is-now-offering-app-install-ads-to-all-marketers-will-it-see-sales-growth-like-facebook-and-google/>

¹⁹<http://www.emarketer.com/Chart/US-Mobile-App-Install-Ad-Spending-2013-2015/167985>

company in app intelligence, did not start monitoring app advertising until March 2015.²⁰ Together these facts suggest that the app-install ad market was insignificant in the period we study and it appears extremely unlikely that it played a significant role in the effect we observe.

Second, we analyze the period before Facebook, the first company to enter the app-install ad market, entered this market in October 2012 separately from the subsequent period. If advertising contributed significantly to the effect, we should observe a significantly higher effect size between October 2012 and October 2013 than from September 2011 to September 2012. In Table 10, Column (1) we interact our coefficient of interest with an indicator for whether the observation falls into the first or second year of the data and find that there is no significant difference. Similarly, if we interact the coefficient of interest with a linear weekly time trend, we find this to be insignificant ($p=0.14$). The fact that there was little change in the effect size over time further supports that advertising is unlikely to have played a significant role.

Table 10: Robustness Checks on Advertising

Variables	ln (Number of Ratings of Paid Version+1)	
	(1) Add time indicator	(2) Smaller developers
T (After new version launch)	0.0307*** (0.0049)	0.0258*** (0.0037)
$T \times$ First year of data	-0.0041 (0.0075)	
Other control variables (ranking, price, the total number of ratings, the average star rating)	Yes	Yes
App FE	Yes	Yes
Number of Observations	49,186	30,071
Number of Apps	1,710	1,045
R-squared	0.5315	0.2523

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Third, we used the website <https://www.spyfu.com/> to check whether the top 50 developers in our data (measured by their number of apps) that had launched a free version of at least one paid app did search advertising during our data period. We found that none of these developers did

²⁰Based on email correspondence with Sensor Tower's Head of Mobile Insights from June 13 to June 21, 2017.

search advertising during that time period.

Fourth, to rule out the possibility that the increase in the number of ratings of the paid version comes from advertising of the free version that subsequently affected the sales of the paid version directly, we replicate the analysis in Column (2) of Table 10 with a subset of developers that we expect to have particularly limited resources for advertising. Specifically, we focus on apps by developers who at the time of the launch of the free version either had only one app, or had other apps but none of these apps had any ratings. These app developers are presumably small in size, had limited budget, and thus had little capacity for cross promotion. Column (2) of Table 10 illustrates a similar positive spillover effect to that of Column (3) in Table 5. We additionally limit the analysis to apps from developers that have a single app only and again find very similar results ($est=0.0293$, $se=0.0047$). The effect sizes are not significantly different from that in Column (3) of Table 5.

Fifth, as a final piece of evidence, we contacted 14 developers of apps in our sample. They all indicated they did little or no advertising during the period of our data.

Together, these results lead us to conclude that it is highly unlikely that developers systematically coordinated advertising strategies to coincide with the launch of the free version, leading to an effect of the free version's launch on the paid version's demand.

Featuring in App Store: If the launch of the free version would lead the app to be featured in App Store, then it would enhance the app's visibility and in turn increase demand (Liang et al., 2019). In our data period, App Store had a section entitled "New and Noteworthy". This section is distinct from rankings and Apple chooses the apps to be featured there. As a result of being featured, apps typically receive a lot of attention and a significant increase in demand. If the launch of a free version would motivate Apple to feature the app, then we might observe an increase in the number of ratings that is correlated with the launch of the free version but not a result of a spillover effect from users of the free version. We do not observe whether apps were featured. However, to be selected by Apple for this section, an app needs to be original and outstanding and receive media

attention.²¹ Because the free version is inferior to the existing paid version, it is extremely unlikely that Apple would ever feature the free version and not the paid version or that Apple would feature the paid version simply because of the release of the free version. It is therefore not plausible that the effect we observe would be a result of the free version being featured in App Store. Note also that anecdotal evidence suggests that being featured as “New and Noteworthy” typically leads to immediate big jumps in demand larger than those we observe following the launch of a free version.²²

Solicited Ratings: Another possible driver to our result would be that concurrently with the launch of the free version, developers systematically solicit consumers to rate apps. Such a strategy could lead to an increase in the number of ratings, holding constant the number of downloads. We check this possibility. First, if developers were more likely to solicit ratings shortly after the launch of the free version, then we might expect different types of consumers to rate the app than previously. If users with a different usage experience of the app evaluate it, the valence of ratings should change. We estimate Equation (1) with the star rating as the dependent variable, and the estimated coefficient of free version launch is -0.0294 (s.e.=0.0511), indicating there is little difference in the valence of ratings before and after the launch of the free version. In addition, we tested whether the distribution of high (4 or 5)/low (1 to 3) ratings is the same before and after the launch of the free version. The average share of high ratings is 81.2% before the launch and 82.4% after the launch. A t-test indicates that the null hypothesis that the two shares are equal cannot be rejected ($p=0.31$). Overall we find that there is little difference in the valence of ratings before and after the launch of the free version.

Further, we were unable to find any direct evidence suggesting that a developer would ben-

²¹See for example: <http://www.appfreak.net/how-to-get-your-app-into-the-new-and-noteworthy/>, <https://growthtower.com/7-ways-to-get-featured-on-the-app-store/> or <https://www.quora.com/How-do-you-get-onto-new-noteworthy-on-the-App-Store>

²²One developer stated that “My app, called Citrus, was featured once on the German iOS AppStore... Our App Store page got around 50,000 views and 4,000 downloads for the week that it was up. Before that we were getting about 2–4 downloads a day.” (see <https://www.quora.com/What-is-it-like-to-have-your-app-featured-on-the-App-Store>). For similar evidence, see <http://blog.anylist.com/2012/08/app-store/>. We acknowledge that while such anecdotal evidence supports the notion that our results are unlikely driven by app being featured in App Store shortly after the free version’s introduction, it does not completely rule out this possibility.

efit from requesting consumers to rate an app specifically around the point in time when a new version is launched. Few consumers are likely to follow the number of reviews posted over time. Indeed, recommendations in online forums that suggest strategies when firms should request reviews largely relate to usage patterns of a specific user and not to the broader product strategy of the developer. None of the recommendations relates to the launch of a free version.²³

In sum, it seems unlikely that increased promotion or advertising at the time of the launch of the free version can explain the effect we observe across a large number of developers and apps.

4 Exploring the Mechanism

The result that the introduction of a free version increases the demand of the paid version is consistent with prior literature on sampling: the free version allows consumers who are uncertain about the quality of the product to sample before purchasing the full version. At the same time, the availability of a free version may enhance visibility of the paid version in App Store, and thus facilitate product discovery. We explore sampling and enhanced product discovery by consumers as potential mechanisms that can explain why the launch of the free version increases demand for the paid app.

4.1 Sampling

Before purchasing a product, consumers are often uncertain about product attributes and the product's quality. Such uncertainty is a major impediment to sales in online markets (Appel et al., 2019; Hong and Pavlou, 2014). One way to resolve this fit uncertainty is for consumers to try the product and learn about the product quality prior to purchasing, thus reducing the uncertainty around the product value.

Several patterns in our data are consistent with sampling as a mechanism behind the spillover

²³For example, it is suggested to consider how long and how extensively a user has been using an app, the moment when a user has finished a positive interaction, or when updating the app but not immediately when an app is opened. See, for example <https://www.quora.com/What-analytics-are-iOS-app-developers-using-to-identify-the-best-time-to-prompt-for-a-user-rating-in-order-to-generate-positive-results>; <https://www.quora.com/Whats-the-best-way-to-ask-an-app-user-to-rate-this-app-without-berating-them>.

effect on the paid version. First, if consumers upgrade to the paid version of an app after sampling and enjoying the free version, then we should find a positive correlation between the free version's incremental star ratings and the paid version's incremental number of ratings as the latter proxies for demand. We indeed find a correlation of 0.17 ($p < 0.0001$). In addition, we also find a correlation of 0.24 ($p < 0.0001$) between the paid version's daily incremental number of ratings and the free version's daily incremental number of ratings, a pattern which is consistent with users purchasing the full version after sampling the free version. Note that if the free version would solely aid product discovery – that is help consumers find the paid version in the store but without consumers sampling the free version – then we would not expect these relationships.

Next, we replicate the analysis in Column (3) of Table 5 with a subset of apps that have identical descriptions for the paid version and the free version. The point estimate is 0.0261 with a standard error of 0.0054, which is very close to that in Column (3) of Table 5. For these apps, consumers cannot observe the difference between the two versions before downloading, and we would expect them to first download the free version. While playing the game, they would be prompted to download the paid version for additional features. The fact that for this sample we still find a similar effect further supports that sampling is a mechanism at work.

We next conceptualize when a consumer samples and the effect of sampling on purchase of the paid version. We focus here on a representative consumer and abstract away from potential heterogeneity across consumers. We then take this framework to our data.

4.1.1 Conceptual Framework

While sampling may help consumers learn the value of a product, sampling is not without cost. Consumers may incur cost of time in acquiring the free sample (e.g. downloading the app) and in getting familiar with its functionalities (Cheng and Liu, 2011). Further, downloading even a free app uses data allowance and storing it uses space on a consumer's device, introducing further opportunity cost for a consumer. As a result, a consumer needs to weigh the opportunity cost associated with acquiring, setting up and learning about the product with the expected utility from

sampling (Lee and Tan, 2013). This expected utility of sampling the free version accounts for utility that may arise from upgrading to the paid version.

As sampling is not costless, consumers' decision on whether or not to sample depends on the expected benefits. Assume the cost for the consumer to sample is c which captures any opportunity cost, such as cost related to time or storage requirements. Also, assume the cost of purchasing the paid version to be p , which includes the price of the product as well as any opportunity cost, such as that associated with time or storage. Assume that a consumer knows the value of the free version, v_f , but prior to sampling has uncertainty about the paid version's value, v_p .²⁴ Instead, the consumer only knows that v_p is uniformly distributed on $[\underline{v}_p, \bar{v}_p]$ where $p > \underline{v}_p > v_f$, and $\bar{v}_p > v_f + p$. After the consumer has used the free version, the value of v_p is revealed. We assume utility is a linear function of the valuation, and for simplicity assume the consumer obtains utility $U_f = v_f - c$ from the free version and $U_p = v_p - p$ from the paid version.

The overall effect of sampling is the difference between the paid version's demand when the firm offers a free version and the paid version's demand when the firm does not offer a free version. We now analyze how the values of \underline{v}_p and \bar{v}_p affect a consumer's decision to sample.

In the absence of a free version, a consumer will purchase the paid version if the expected utility of the paid version is larger than zero, that is, $E(U_p) = E(v_p) - p > 0$. $E(U_p)$ is also the expected utility of purchasing the paid version directly in the presence of a free version.

When the firm offers a free version, a consumer can sample the free version to learn about the value of v_p . After sampling, the consumer upgrades to the paid version if the difference between the value of the paid version and the value of the free version exceeds the price of the paid version, $v_p - v_f > p$. Thus, the probability of upgrading is $Prob(v_p - v_f > p) = \frac{\bar{v}_p - (p + v_f)}{\bar{v}_p - \underline{v}_p}$. Upon upgrading, the consumer obtains utility $v_p - p$. If the difference between the value of the paid version and the free version is smaller than the price of the paid version, $v_p - v_f < p$, then the consumer does not upgrade and only receives utility v_f from the free version.

²⁴The consumer knowing the value of the free version, v_f , but not the value of the paid version, v_p , is a simplification for ease of analysis. The underlying assumption is that the consumer has a more accurate belief of v_f than of v_p , that is a tighter distribution around v_f than around v_p . By definition, the free version is a simplified version of the paid version and would have a lower maximum value it could possibly provide, thus leading to a tighter distribution.

Thus, the expected utility from sampling is:

$$E(U_s) = v_f \text{Prob}(v_p - v_f \leq p) + E(v_p - p) \text{Prob}(v_p - v_f > p) - c \quad (2)$$

For consumers to sample the free version, two conditions need to be satisfied. First, the expected utility of sampling needs to be greater than zero ($E(U_s) > 0$). Second, the expected utility of sampling should exceed the expected utility from purchasing the paid version directly ($E(U_s) > E(U_p)$). For consumers to upgrade to the paid version after sampling, v_p needs to satisfy $v_p - v_f > p$. The condition $v_p - v_f > p$ will be met with a positive probability because we assume $\bar{v}_p > v_f + p$. Thus, the introduction of a free version can increase the paid version's demand through sampling if $E(U_s) > 0$ and $E(U_s) > E(U_p)$.²⁵

Result: It can be shown that $E(U_s)$ increases with \underline{v}_p and \bar{v}_p , and $E(U_s) - E(U_p)$ decreases with \underline{v}_p and increases with \bar{v}_p (for details, see Online Appendix C). Therefore, the conditions for sampling, $E(U_s) > 0$ and $E(U_s) > E(U_p)$, are likely to be met simultaneously when two conditions hold: (1) \underline{v}_p is neither too large nor too small; and (2) \bar{v}_p is large.

If $E(U_s) > 0$ but $E(U_s) < E(U_p)$, then consumers would directly purchase the paid version regardless of whether there is a free version. The introduction of a free version would not affect the demand for the paid version.

4.1.2 App Quality

One factor that plausibly affects whether a consumer expects sampling to be beneficial is the quality information a consumer receives about a product. Such quality information about a product may shift the lower bound (\underline{v}_p) or upper bound (\bar{v}_p) of the consumer's belief about the distribution of v_p . We discuss in turn how publicly available information on app quality in the marketplace shifting either bound may affect a consumer's decision to sample.

²⁵Note that it is also possible that a consumer downloads the free version without the intention to upgrade and then is enticed to upgrade. This would happen if $v_f - c > 0$. Because $E(U_s) > v_f - c$, when $v_f - c > 0$, the condition $E(U_s) > 0$ is also met. Therefore, this can be considered as a special case which is in line with our discussion.

Quality information shifting lower bound: Our result in §4.1.1 indicates that sampling is more effective when \underline{v}_p is neither too large nor too small. The intuition is as follows: When a consumer receives information on app quality indicating a very low value of \underline{v}_p , they expect to receive a low utility, $E(U_p)$, from the product and the marginal benefit of sampling is low as the cost of sampling may exceed the marginal benefit of sampling. It may then not be optimal for the consumer to sample. Put simply, the consumer has such low expectations of product quality that it is not worthwhile to incur the cost of sampling, and as a consequence, offering a free sample will have little effect on the consumer's purchase of the paid version.

A similar pattern occurs when the consumer receives information that indicates a very high value of the lower bound \underline{v}_p . If \underline{v}_p is very high, then the consumer's expected utility, $E(U_p)$, is very high and the marginal benefit of sampling may be below the marginal cost of sampling. Put differently, if the consumer expects a product to be of high quality, they may prefer to purchase the paid version directly and avoid the opportunity cost associated with sampling. Again, there will be little effect of offering a free sample on the consumer's purchase of the paid version.

Two pieces of public information on app quality that are readily available to consumers in the mobile app market are an app's number of reviews (Schoenmueller et al., 2020) and its star rating. If an app received very few ratings or is poorly rated, consumers may conclude that its quality is too low to be a good fit and that sampling is not worthwhile. On the other hand, a large number of past ratings or a very high star rating indicate that many users found the app useful, making it more efficient for consumers to directly purchase the paid version and avoiding the cost of sampling. In consequence, sampling is more beneficial for apps that have a moderate number of ratings and/or a moderate star rating.

We argue that the number of ratings, as well as the average star rating of an app, shift the lower bound of the value distribution, \underline{v}_p , but not its upper bound, \bar{v}_p . This is because a large number of ratings would inform a consumer that the app was sufficiently attractive for many other consumers. However, this information would not allow the consumer to infer the upper bound of the value distribution, \bar{v}_p , as the number of ratings provides little information about the likely

maximum value a consumer could possibly derive from the product. A similar logic applies to the average star rating. Note, however, that while we argue that both the number of ratings and the average star rating shift the consumer's lower bound of the value distribution, \underline{v}_p , our results are robust if they would likewise shift the upper bound, \bar{v}_p (see Online Appendix C).

We thus expect an inverted-U shaped relationship between the effect of sampling and the number of ratings – and as an alternative measure the average star rating – of an app. We examine empirically the moderating effects of the paid version's number of ratings, or the average star rating, prior to its free counterpart's introduction on the treatment effect. That is, we estimate Equation (1) adding interaction terms between the free version's introduction and the quality measure and its squared term. The estimation results in Column (1) Table 11 demonstrate indeed an inverted U-shaped relationship between the number of ratings and the effect size. Column (2) demonstrates a similar effect for the star rating. This result has a somewhat lower significance level, possibly because there is less variation in the average star rating across apps as star ratings are bounded between 1 and 5.

Table 11: Moderating Effect of Number of Ratings and Star Rating

Variables	ln (Number of Ratings of Paid Version+1)	
	(1)	(2)
T (After new version launch) \times ln(Number of ratings+1)	0.0750** (0.0335)	
$T \times$ ln(Number of ratings+1) ²	-0.0077** (0.0038)	
$T \times$ Star rating		0.1490* (0.0770)
$T \times$ Star rating ²		-0.0198* (0.0106)
Other control variables (ranking, price, the total number of ratings, the average star rating)	Yes	Yes
App FE	Yes	Yes
Observations	49186	49186
Number of Apps	1710	1710
Adj. R-Squared	0.5319	0.5315

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Quality information shifting upper bound: Our result in §4.1.1 indicates that the value of sampling to the consumer increases in \bar{v}_p , that is as the upper bound of the value distribution increases, the consumer becomes more likely to sample. However, unlike for \underline{v}_p , the value of sampling does not decrease for very large values of \bar{v}_p : as the consumer still faces the risk of a low valuation, \underline{v}_p , it continues to be optimal to sample, even for large values of \bar{v}_p .

We argue that the size of an app provides information to the consumer about the upper bound of the value distribution, \bar{v}_p . A larger app, that is an app that has a greater volume of code and data, may have richer content, more functionalities, more features or better graphics and thus may have the potential of providing greater value to the consumer than a small app which is likely more limited. Thus the size of the app is correlated with \bar{v}_p . At the same time, for larger apps, there is more room to differentiate the free version from the paid version, simply because the code may include a larger number of additional characteristics (e.g. levels of a game). Note though that the size of the app sheds little light on the lower bound of the value distribution, \underline{v}_p . This is because any app may be implemented in ways that are difficult to use and may have features that are not valuable to the consumer. It may contain code that makes the app large but provides little value.²⁶ Thus, the size of an app is unlikely to be correlated with the lower bound of the value distribution, \underline{v}_p .

We estimate the moderating effect of app size on the effect of the free version's introduction. Column (1) of Table 12 indicates that indeed, the effect of the free version's introduction increases with the size of the paid app. To be consistent with Table 11, Columns (2) and (3) further allow the effect to vary respectively with the number of ratings and the star rating as well as their squared terms. The results continue to hold.²⁷

Together, the results in Tables 11 and 12 suggest that consumers respond to public information on app quality which in turn lead to behavior that is consistent with sampling.

²⁶See, for example, https://developer.apple.com/documentation/xcode/reducing_your_app_s_size.

²⁷We also estimate a regression adding a quadratic term of app size. There is no evidence supporting a non-linear relationship between app size and effect of launching a free version.

Table 12: Effect by App Size

Variables	ln (Number of Ratings of Paid Version+1)		
	(1)	(2)	(3)
T (After new version launch)	-0.0980* (0.0533)	-0.0729 (0.0517)	-0.0904* (0.0521)
$T \times \ln(\text{App Size})$	0.0077** (0.0033)	0.0061* (0.0032)	0.0072** (0.0032)
$T \times \text{Rated}$		-0.1370** (0.0594)	-0.2350* (0.1320)
$T \times \ln(\text{number of rating}+1)$		0.0708** (0.0338)	
$T \times \ln(\text{number of rating}+1)^2$		-0.0073* (0.0038)	
$T \times \text{Star Rating}$			0.1320* (0.0768)
$T \times \text{Star Rating}^2$			-0.0176* (0.0105)
Other control variables (ranking, price)	Yes	Yes	Yes
App FE	Yes	Yes	Yes
Observations	49186	49186	49186
Number of Apps	1710	1710	1710
Adj. R-Squared	0.5316	0.5319	0.5316

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

4.1.3 Difference between Free and Paid Versions

Prior research has demonstrated that the effect of sampling depends on the quality difference between the free and the paid products. For instance, Li et al. (2019) find that when a publisher provides free digital samples to print copies of books, it is beneficial to increase the functionality difference between the high-quality digital sample and the print format. Cheng and Liu (2011) find that offering a feature limited free sample is more suitable for software products that exhibit strong network effect.

In our conceptual framework, the consumer will upgrade to the paid version after sampling if the difference between the paid version and the free version ($v_p - v_f$) is larger than the price of the paid version. Thus, the paid version may provide sufficiently greater utility for consumers to upgrade once consumers are satiated with the free version (Appel et al., 2019). To see if this

holds empirically, we examine the moderating effects of version design on the treatment effect. Specifically, we estimate Equation (1) adding interaction terms between treatment and the app description-based design differences described in §2.3: more levels, more modes/themes, more functions/features, social interactions, ad free, and better support. We group apps with the same combination of design differences, and compute robust standard errors clustered at the group level. Table 13 shows the results. Consistent with the sampling explanation, compared with apps with identical descriptions for the paid version and the free version, this effect is more pronounced for apps where the paid version appears to offer greater additional benefits. It is greatest when the paid version offers “more levels” or “social interactions”. It seems plausible that the free version provides consumers an opportunity to try the game on initial levels while additional levels provide sufficient utility for consumers to upgrade. Similarly, the ability to compete, compare scores and communicate with others may give consumers who found they like the initial version significant extra utility, thus making a purchase of the full version attractive.

Interestingly, we find a negative moderating effect for “more modes/themes”, indicating that consumers value having more modes/themes (horizontal differentiation) relatively less than additional levels in the game (vertical differentiation). This appears plausible as additional modes/themes (e.g., a different background picture) do relatively little to change the gaming experience, so that the introduction of a free version with reduced modes/themes cannibalizes the demand on the paid version to some degree. The moderating effect of “ad free” is also negative. It suggests that for the average consumer, the negative utility associated with advertising is small, so the introduction of the free version with ads cannibalizes the demand on the paid version to some degree. The result seems plausible in view of the broader evidence that many consumers prefer ads rather than paying for content (Papies et al., 2011; Chiou and Tucker, 2013).²⁸ Note that despite the negative effects associated with “more modes/themes” or “ad free”, the relatively larger size of the main effect continues to imply that on average, the launch of the free version still positively affects

²⁸A study by the Digital Advertising Alliance (DAA) finds more than 85% of respondents prefer an ad-supported Internet instead of having to pay for the content they now get for nothing (see <https://www.consumeraffairs.com/news/survey-finds-85-of-consumers-prefer-an-ad-supported-internet-051216>).

demand of the paid version. Finally, the moderating effect of “support” is negative and its effect size exceeds that of the main effect. The estimate suggests that the average consumer does not place much value on support offered by the developer, and the introduction of a free version that provides less support may cannibalize demand of the paid version.

Table 13: Moderating Effects of Design Differences based on App Description

Variables	ln (Number of Ratings of Paid Version+1)
<i>T</i> (After new version launch)	0.0237*** (0.0031)
<i>T</i> × More levels	0.0138*** (0.0035)
<i>T</i> × More modes/themes	-0.0117* (0.0066)
<i>T</i> × More functions/features	0.0100 (0.0059)
<i>T</i> × Social	0.0275*** (0.0082)
<i>T</i> × Ad free	-0.0122* (0.0064)
<i>T</i> × Support	-0.0478*** (0.0076)
Other control variables (ranking, price, the total number of ratings, the average star rating)	Yes
App FE	Yes
Number of Observations	49,186
Number of Apps	1,710
Adj. R-squared	0.5317

Robust standard errors clustered at app group level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

In sum, this set of results supports the explanation that consumers’ sampling of the free version leads to the purchase of the paid variant. The strategy of releasing a free version to let consumers sample is effective when apps have a moderate number of ratings, a moderate level of the average star rating, as well as for apps that are larger in size. In addition, consumers are more likely to upgrade from the free to the paid version when the two versions differ substantially in the value they provide to consumers.

4.2 App Discovery

With the rapidly increasing number of apps on app stores, “the gap between search and find is getting bigger and bigger”,²⁹ making app discovery a problem as it is difficult for any individual app to be discovered by a consumer (Li et al., 2016). At the time of our data, App Store offered close to 200,000 game apps, most of which were by small independent developers. In such a setting, having two versions instead of one increases visibility of the app and makes it more likely to be discovered by a consumer. Thus, the availability of a free version may increase the likelihood of discovery of the paid version. This can impact demand for the paid version in two ways: First, the enhanced visibility may lead consumers to discover the free version, sample the free version and later purchase the paid version; that is discovery may lead to sampling. Second, the availability of the free version may increase a consumer’s probability of discovering the paid version and then purchasing the paid version directly, without first sampling the free version.³⁰

Ideally, one would explore app discovery based on consumer browsing activity in App Store. However, such data are not publicly accessible. Instead, we use two features of our data to shed light on whether app discovery is a likely mechanism. First, we turn to a setting where we would expect app discovery to matter but not sampling. We focus on apps that initially offered a free version and later added a paid version. If app discovery matters, then the introduction of a paid version should increase the probability for the app to be discovered, and thus demand for the original free version. However, we would not expect sampling to matter in this context. For this analysis, we focus on the 288 apps where the free version was launched first and the paid version was launched later during our sample period. We then estimate the effect of introducing the paid version on the free version’s number of ratings. Table 14, Column (1) indeed indicates a positive

²⁹<https://www.forbes.com/sites/allbusiness/2013/08/15/app-discovery-why-cant-anyone-figure-this-out-yet/>

³⁰Note that on a high level, the question of discovery is related to research which documented that music streaming services serve as an alternative channel for music discovery, which increases downloads at alternative consumption channels (Aguiar, 2017) or recorded sales (Kretschmer and Peukert, 2020), or increases volume and diversity of music consumption (Datta et al., 2017). Both their research and ours argue that a wider availability may facilitate discovery and thus increase sales. However, in our context we argue that a greater number of similar products promotes discovery while prior research focused on the number of channels that made a single product available for purchase or consumption.

effect, suggesting that the availability of an additional version may enhance discovery.

Second, the app discovery mechanism suggests that when the number of apps in a category increases, providing an additional version of an app has increasingly less effect on the probability for an app to be discovered, simply because the “haystack” gets larger. Thus, the marginal benefit of providing an additional version of an app is smaller. In our data, the number of apps per genre varies significantly between 1,405 (Dice) and 40,704 (Puzzle). We estimate the moderating effect of the number of apps in the same genre. Column 2 of Table 14 shows the results for apps where the paid version was launched first and Column 3 shows the corresponding results for instances where the free version was launched first. In both instances, the estimates indicate that adding an additional version of an app has a lower effect on the demand for the original version when a genre comprises more apps. This is in line with the mechanism related to app discovery that we outline.

In sum, our results support that offering a free version in addition to a paid version can enhance visibility of the app and thus increase demand for the paid version.

Table 14: App Discovery

Variables	(1) ln (Number of Ratings of Free Version+1) for free-first apps	(2) ln (Number of Ratings of Paid Version+1) for paid-first apps	(3) ln (Number of Ratings of Free Version+1) for free-first apps
<i>T</i> (After new version launch)	0.0216** (0.0101)	0.0430*** (0.0083)	0.0462*** (0.0150)
<i>T</i> (After new version launch)× Number of Apps in the Genre (100,000)		-0.0591** (0.0285)	-0.1080** (0.0520)
Other control variables (ranking, price, the total number of ratings, the average star rating)	Yes	Yes	Yes
App Fixed Effects	Yes	Yes	Yes
Observations	8,282	49,186	8282
Number of Apps	288	1,710	288
Adj. R-Squared	0.6702	0.5347	0.7107

5 Conclusion

Freemium pricing has become an increasingly popular business model in the digital economy. While the free version may tempt consumers to purchase the paid version, it may likewise can-

nibalize the sales of the paid version. In this paper, we aim to measure the overall impact of a freemium strategy using granular data on game apps available in App Store when a developer offers both a free and a paid version of the same game. As data on downloads of apps are not publicly available, we use the number of ratings available for an app to proxy app demand. We focus on apps that introduced a free version after an initial paid version and estimate how this free version's introduction changed the number of ratings of the paid version focusing on a tight time window around the launch of the free version. We find that, on average, demand for the paid version increased following the launch of the free version. This result holds in a battery of robustness checks and is unlikely to be an outcome of other marketing actions correlated with the launch.

We then explore the mechanism at work. First, we develop a conceptual framework on sampling of the free version and the effect of sampling on purchase of the paid version, and put forward predictions on when sampling is more effective. We then show evidence from the data that is in line with these predictions, supporting that consumers use free versions to sample and later upgrade to the paid versions. Second, we demonstrate that a new version of an app may increase the initial version's demand by enhancing the visibility (and thus chance to be discovered) of the app.

Our results are relevant for app developers and, more broadly, digital firms who offer freemium pricing. First, they confirm that a freemium strategy can indeed increase demand for the paid version of a product. At least on average, this spillover effect outweighs any possible cannibalization in the market we study. Second, our results indicate that a freemium strategy would be more effective for products that prior users evaluated as moderately good. Third, our results suggest that to truly benefit from a freemium strategy, firms need to ensure that the difference in utility for the consumer between the free product and the paid product is sufficiently high to induce upgrades. At least in the market for game apps, this appears to be the case when the user can have a substantially enhanced usage experience by playing more levels or through social interaction. Fourth, we document that simply the fact that an additional version is available can increase demand of the paid version by making it easier for consumers to discover the app.

While the focus of our empirical analysis is on a setting where the firm offers two separate

versions of an app, we expect that many of our results translate to other forms of freemium pricing such as when in-app purchases can be made within a free app. While the technology to implement freemium pricing differs across domains and even within domains changes over time, we believe that the underlying economics and implications for consumers are similar. In all instances, consumers can access a certain set of features or levels for free, or use the product for free for a certain amount of time. Full access to all features or further usage requires payment. We therefore believe that our results regarding consumers decision to upgrade are indicative of consumer response to freemium pricing more broadly.

There are several limitations to the present study that represent opportunities for future research. First, we focus on a specific empirical application – the market for game apps during a time period when developers marketing freemium apps offered both a free and a paid version, and in-app purchases within the same app were less common. Our estimation results most directly apply to this specific empirical setting and data period. Exploring in-app purchases may widen our understanding of how consumers respond to freemium offerings and of how firms should best design freemium products. Second, as we lack data on app downloads, we use an app’s number of ratings as a proxy for its demand. Given the downward trend in rating accumulation, this approach appears conservative as it likely underestimates incremental downloads as the number of cumulative downloads increases. Nonetheless, using a more precise measure of demand could yield more precise results. Third, we focus on the short-term effect of the free version introduction, as identification of a long-term effect is more difficult and would be significantly less precise due to changes in the app itself, in the market for apps and in patterns of consumer demand over time. Properly measuring the long-term effect of a freemium strategy while accounting for such potential shifts over time may be interesting to explore in future research. Finally, while we provide ample evidence that correlated marketing activities are unlikely the driver of the spillover effect as a rule, we cannot fully rule out that in some instances firms may still have tried to support the launch of a free version with marketing activities, which might have had some effect on our estimates.

References

- Aguiar, Luis. 2017. Let the music play? free streaming and its effects on digital music consumption. *Information Economics and Policy* **41** 1–14.
- AppAnnie. 2015. The state of play: A look at the growth of Google Play. Available at <https://s3.amazonaws.com/files.appannie.com/reports/The-State-of-Play-Google-I0-2015-EN.pdf>.
- Appel, Gil, Barak Libai, Eitan Muller, Ron Shachar. 2019. On the monetization of mobile apps. *International Journal of Research in Marketing* .
- Arora, Sandeep, Frenkel ter Hofstede, Vijay Mahajan. 2017. To free, or not to free: the implications of offering free versions for the performance of paid mobile apps. *Journal of Marketing* **81**(6) 62–78.
- Bawa, Kapil, Robert Shoemaker. 2004. The effects of free sample promotions on incremental brand sales. *Marketing Science* **23**(3) 345–363.
- Boudreau, Kevin J. 2018. Amateurs crowds & professional entrepreneurs as platform complementors. Working Paper 24512, National Bureau of Economic Research.
- Chellappa, Ramnath K., Amit Mehra. 2017. Cost Drivers of Versioning: Pricing and Product Line Strategies for Information Goods. *Management Science* .
- Chen, Hong, Wenjing Duan, Wenqi Zhou. 2017. The interplay between free sampling and word of mouth in the online software market. *Decision Support Systems* **95** 82–90.
- Cheng, Hsing Kenneth, Yipeng Liu. 2011. Optimal Software Free Trial Strategy: The Impact of Network Externalities and Consumer Uncertainty. *Information Systems Research* **23**(2) 488–504.

- Chevalier, Judith, Austan Goolsbee. 2003. Measuring prices and price competition online: Amazon.com and BarnesandNoble.com. *Quantitative Marketing & Economics* **1**(2) 203–222.
- Chiou, Lesley, Catherine Tucker. 2013. Paywalls and the demand for news. *Information Economics and Policy* **25**(2) 61–69.
- Datta, Hannes, George Knox, Bart J. Bronnenberg. 2017. Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery. *Marketing Science* **37**(1) 5–21.
- Dey, Debabrata, Atanu Lahiri. 2016. Versioning: Go vertical in a horizontal market? *Journal of Management Information Systems* **33**(2) 546–572.
- Dey, Debabrata, Atanu Lahiri, Dengpan Liu. 2013. Consumer learning and time-locked trials of software products. *Journal of Management Information Systems* **30**(2) 239–268.
- Foubert, Bram, Els Gijbrecchts. 2016. Try it, you'll like it - or will you? the perils of early free-trial promotions for high-tech service adoption. *Marketing Science* **35**(5) 810–826.
- Garg, Rajiv, Rahul Telang. 2013. Inferring app demand from publicly available data. *MIS Quarterly* **37**(4) 1253–1264.
- Ghose, Anindya, Sang Pil Han. 2014. Estimating demand for mobile applications in the new economy. *Management Science* **60**(6) 1470–1488.
- Godes, David, Jose C. Silva. 2012. Sequential and temporal dynamics of online opinion. *Marketing Science* **31**(3) 448–473.
- Gu, Xian, PK Kannan, Liye Ma. 2018. Selling the premium in freemium. *Journal of Marketing* **82**(6) 10–27.
- Halbheer, Daniel, Florian Stahl, Oded Koenigsberg, Donald R. Lehmann. 2014. Choosing a digital content strategy: How much should be free? *International Journal of Research in Marketing* **31**(2) 192–206.

- Heiman, Amir, Bruce McWilliams, Zhihua Shen, David Zilberman. 2001. Learning and forgetting: Modeling optimal product sampling over time. *Management Science* **47**(4) 532–546.
- Hoang, Ai-Phuong, Robert J. Kauffman. 2018. Content sampling, household informedness, and the consumption of digital information goods. *Journal of Management Information Systems* **35**(2) 575–609.
- Hong, Yili (Kevin), Paul A. Pavlou. 2014. Product fit uncertainty in online markets: Nature, effects, and antecedents. *Information Systems Research* **25**(2) 328–344.
- Jain, Dipak, Vijay Mahajan, Eitan Muller. 1995. An approach for determining optimal product sampling for the diffusion of a new product. *Journal of Product Innovation Management* **12**(2) 124–135.
- Kannan, P. K., Barbara Kline Pope, Sanjay Jain. 2009. Practice prize winner - pricing digital content product lines: A model and application for the National Academies Press. *Marketing Science* **28**(4) 620–636.
- Kosner, Anthony Wing. 2015. Apple App Store revenue surge and the rise of freemium app pricing. Available at <https://www.forbes.com/sites/anthonykosner/2015/01/11/apple-app-store-revenue-surge-and-the-rise-of-the-freemium>.
- Kretschmer, Tobias, Christian Peukert. 2020. Video killed the radio star? online music videos and recorded music sales. *Information Systems Research* .
- Kummer, Michael, Patrick Schulte. 2019. When private information settles the bill: Money and privacy in Google’s market for smartphone applications. *Management Science* **65**(8) 3470–3494.
- Lahiri, Atanu, Debabrata Dey. 2018. Versioning and information dissemination: A new perspective. *Information Systems Research* **29**(4) 965–983.
- Lambrecht, Anja, Kanishka Misra. 2016. Fee or free: When should firms charge for online content? *Management Science* **63**(4) 1150–1165.

- Lee, Clarence, Vineet Kumar, Sunil Gupta. 2017. Designing freemium: Strategic balancing growth and monetization. SSRN Scholarly Paper ID 2767135, Social Science Research Network, Rochester, NY.
- Lee, Young-Jin, Yong Tan. 2013. Effects of different types of free trials and ratings in sampling of consumer software: An empirical study. *Journal of Management Information Systems* **30**(3) 213–246.
- Li, Hongshuang (Alice), Sanjay Jain, P. K. Kannan. 2019. Optimal design of free samples for digital products and services. *Journal of Marketing Research* **56**(3) 419–438.
- Li, Xing, Timothy Bresnahan, Pai-Ling Yin. 2016. Paying incumbents and customers to enter an industry: Buying downloads. Available at SSRN 2834564 .
- Liang, Chen, Zhan (Michael) Shi, T. S. Raghu. 2019. The spillover of spotlight: Platform recommendation in the mobile app market. *Information Systems Research* **30**(4) 1296–1318.
- Liu, Charles Zhechao, Yoris A. Au, Hoon Seok Choi. 2014. Effects of freemium strategy in the mobile app market: An empirical study of Google Play. *Journal of Management Information Systems* **31**(3) 326–354.
- Niculescu, Marius F., D. J. Wu. 2014. Economics of Free Under Perpetual Licensing: Implications for the Software Industry. *Information Systems Research* **25**(1) 173–199.
- Papies, Dominik, Felix Eggers, Nils Wlömert. 2011. Music for free? How free ad-funded downloads affect consumer choice. *Journal of the Academy of Marketing Science* **39**(5) 777–794.
- Pauwels, Koen, Allen Weiss. 2008. Moving from free to fee: How online firms market to change their business model successfully. *Journal of Marketing* **72**(3) 14–31.
- Runge, Julian, Stefan Wagner, Jorg Claussen, Daniel Klapper. 2016. Freemium pricing: Evidence from a large-scale field experiment. SSRN Scholarly Paper ID 2888471, Social Science Research Network, Rochester, NY.

- Schoenmueller, Verena, Oded Netzer, Florian Stahl. 2020. The polarity of online reviews: Prevalence, drivers and implications. *Journal of Marketing Research* 0022243720941832.
- Shoemaker, Robert W., F. Robert Shoaf. 1975. Behavioral changes in the trial of new products. *Journal of Consumer Research* 2(2) 104–109.
- Statista. 2019. Worldwide mobile app revenues in 2014 to 2023. Available at <https://www.statista.com/statistics/269025/worldwide-mobile-app-revenue-forecast/>.
- Statista. 2020a. Number of available applications in the Google Play store from December 2009 to September 2020. Available at <https://www.statista.com/statistics/266210/number-of-available-applications-in-the-google-play-store/>.
- Statista. 2020b. Number of available apps in the Apple App Store from 2008 to 2020. Available at <https://www.statista.com/statistics/268251/number-of-apps-in-the-itunes-app-store-since-2008/>.
- Wang, Chong (Alex), Xiaoquan (Michael) Zhang. 2009. Sampling of information goods. *Decision Support Systems* 48(1) 14–22.
- Wu, Fang, Bernardo A. Huberman. 2008. How public opinion forms. Christos Papadimitriou, Shuzhong Zhang, eds., *Internet and Network Economics*. No. 5385 in Lecture Notes in Computer Science, Springer Berlin Heidelberg, 334–341.

Online Appendix

A Identify Freemium Pairs

We use Python to identify freemium pairs in three general steps described below.

Step 1: Match

For each publisher, we generate a group comprising of all apps developed by this publisher.

For each app, we compute its similarity to all other apps in the same groups, using app name as the basis of comparison. For each app pair, we first compute a Jaro-Winkler distance between the full strings of the names. After this, we tokenize app names, leading to two lists of tokens (words) for each app pair. For instance, “Angry Birds” and “Angry Birds Lite” are tokenized into [“angry”, “birds”] and [“angry”, “birds”, “lite”]. We then remove paid/free keywords (e.g., “lite”, “pro”, “free”, “premium”). In the the previous example, the lists become [“angry”, “birds”] and [“angry”, “birds”]. For each pair of words in order (meaning “angry” will only compare to “angry”), we compute a new Jaro-Winkler distance which is added to the previous measure of the full string distance. The aggregated Jaro-Winkler distance is then normalized to a scale of 0-1.

Using this established Jaro-Winkler distance, each app pair gets sent into a series of heuristics using cut-off values to determine the quality of match. Very confident matches (≥ 0.9) are added as a “NoFlag” match indicating that they are very close matches. Matches that score a distance between 0.9-0.8 are added but with an added “Flag” variable indicating that they are closely related but may contain some false positives. After these two cut-offs, two more samples are added at 0.8-0.75 and 0.75-0.8 respectively. These last two cut-offs add an “OtherSample” variable to the match.

Step 2: Remove False Matches

We remove the following types of false matches:

- We remove matches between an iPad app and an iPhone app using a set of iPad specific keywords (e.g. “Angry Birds” and “Angry Birds iPad” or “Angry Birds HD”).

- Sequels such as “angry birds” and “angry birds 2” would have a very high similarity score but are not a true freemium pair. We remove sequels by comparing suffix numbers of app names. If app1 has a “2” at the end whilst app2 has nothing, app1 is considered a sequel to app2 or vice versa. Sequels up to the number 4 are considered.
- If a free (paid) app matched with another free (paid) app respectively (i.e. paid-paid or free-free matches only), this pair is marked for removal (e.g. “angry birds lite” and “angry birds free”).

Step 3: Manual Check

We order all the identified freemium pairs so that app1 is the paid version and app2 is the free version. We then manually check to ensure that they are indeed a freemium pair.

B Matching Results

As an alternative to Table 8 where we use all paid apps as controls, we narrow down the group of control apps further to more closely reflect the observable characteristics of freemium apps using matching. Specifically, for each paid version of the freemium app (treated app), we look for control paid apps that meet the following criteria:

1. Control app is in the same genre as the treated app
2. Difference between control app’s launch date and treated app’s launch date is smaller than 30 days
3. On the launch date of the free version, if the treated app is not rated, then the control app is not rated either; if the treated app is rated, then the control app is also rated

For each candidate app, we compute the Mahalanobis distance between this app and the treated app on the free version’s launch date, based on number of ratings, average star rating, and price. We then choose the app with the smallest distance as control app for the corresponding treated app. If multiple candidate apps have the same distance to the treated app, then we randomly pick one

as the control app. Once a control app is selected for a treated app, then we drop this control app from the candidate pool. Hence, each control app can be matched to at most one treated app.

We replicate Table 8 with the new control group. Table A.1 displays the estimation results. We find that the results are very similar to those obtained when using all apps as a control group.

Table A.1: Robustness Checks with Control Apps from Matching

Variables	ln (Number of Ratings of Paid Version+1)		
	(1) Paid control apps	(2) Paid control apps	(3) Paid control apps + Freemium
T (After new version launch)	0.0024 (0.0029)	0.0065** (0.0028)	0.0065** (0.0028)
$T \times$ (Treatment group)			0.0220*** (0.0048)
Lag (Ranked on top 10)		0.5170 (0.4629)	0.9790*** (0.3088)
Lag (Ranked on top 11-20)		1.2358** (0.5051)	1.3883*** (0.2942)
Lag (Ranked on top 21-50)		0.5953*** (0.2269)	1.0811*** (0.2505)
Lag (Ranked on top 51-100)		0.6928*** (0.1797)	0.7347*** (0.1251)
Lag (Ranked on top 101-)		0.0689*** (0.0228)	0.0941*** (0.0181)
App FE	Yes	Yes	Yes
Other control variables (price, the total number of ratings, the average star rating)	No	Yes	Yes
Number of Observations	49,437	49,437	98,835
Number of Apps	1,708	1,708	3,416
Adj. R-squared	0.4028	0.4273	0.4918

Robust standard errors clustered at app level in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

C Analyses on the Effect of Sampling

Effect of \bar{v}_p and \underline{v}_p on $E(U_s)$:

$$\begin{aligned}
 E(U_s) &= v_f \text{Prob}(v_p - v_f \leq p) + E(v_p - p) \text{Prob}(v_p - v_f > p) - c \\
 &= v_f \int_{\underline{v}_p}^{p+v_f} f(v_p) d(v_p) + \int_{p+v_f}^{\bar{v}_p} f(v_p) (v_p - p) d(v_p) - c
 \end{aligned}$$

$$\begin{aligned}
&= v_f \frac{p + v_f - \underline{v}_p}{\bar{v}_p - \underline{v}_p} - p \frac{\bar{v}_p - (p + v_f)}{\bar{v}_p - \underline{v}_p} + \frac{1}{\bar{v}_p - \underline{v}_p} \int_{p+v_f}^{\bar{v}_p} v_p d(v_p) - c \\
&= \frac{(p + v_f)^2 - 2v_f \underline{v}_p - 2p\bar{v}_p + (\bar{v}_p)^2}{2(\bar{v}_p - \underline{v}_p)} - c
\end{aligned}$$

Thus, we have:

$$\frac{dE(U_s)}{d\underline{v}_p} = \frac{(p + v_f)^2 + \bar{v}_p(\bar{v}_p - v_f) + \bar{v}_p(2p - v_f)}{2(\bar{v}_p - \underline{v}_p)^2} > 0$$

$$\frac{dE(U_s)}{d\bar{v}_p} = \frac{v_f^2 + \bar{v}_p^2 - p^2 + 2p(\underline{v}_p - v_f)}{2(\bar{v}_p - \underline{v}_p)^2} > 0$$

Therefore, $E(U_p)$ increases with \underline{v}_p and \bar{v}_p .

Effect of \bar{v}_p and \underline{v}_p on $E(U_s) - E(U_p)$:

$$\begin{aligned}
E(U_s) - E(U_p) &= v_f \int_{\underline{v}_p}^{p+v_f} f(v_p) d(v_p) + \int_{p+v_f}^{\bar{v}_p} f(v_p)(v_p - p) d(v_p) - c - \left[\int_{\underline{v}_p}^{\bar{v}_p} v_p f(v_p) d(v_p) - p \right] \\
&= \frac{(p + v_f)^2 - 2v_f \underline{v}_p - 2p\bar{v}_p + (\bar{v}_p)^2}{2(\bar{v}_p - \underline{v}_p)} - \frac{\underline{v}_p + \bar{v}_p}{2} + p - c \\
&= \frac{(p + v_f)^2 - 2v_f \underline{v}_p - 2p\bar{v}_p - \underline{v}_p^2}{2(\bar{v}_p - \underline{v}_p)} + p - c
\end{aligned}$$

Thus, we have:

$$\begin{aligned}
\frac{dE(U_s - U_p)}{d\underline{v}_p} &= \frac{-2v_f \bar{v}_p - 2\bar{v}_p \underline{v}_p + \underline{v}_p^2 + (p + v_f)^2 - 2p\bar{v}_p}{2(\bar{v}_p - \underline{v}_p)^2} \\
&< \frac{-2v_f^2 - 2\bar{v}_p \underline{v}_p + \underline{v}_p^2 + (p + v_f)^2 - 2p^2}{2(\bar{v}_p - \underline{v}_p)^2}
\end{aligned}$$

$$\begin{aligned}
&= \frac{v_p (v_p - 2\bar{v}_p) - (p - v_f)^2}{2(\bar{v}_p - v_p)^2} \\
&< 0
\end{aligned}$$

$$\begin{aligned}
\frac{dE(U_s - U_p)}{d\bar{v}_p} &= \frac{2p(v_p - v_f) + (v_p^2 - v_f^2) - p^2 + 2v_f v_p}{2(\bar{v}_p - v_p)^2} \\
&> \frac{2p(v_p - v_f) + (v_p^2 - v_f^2) + v_f v_p}{2(\bar{v}_p - v_p)^2} \\
&> 0
\end{aligned}$$

Therefore, $E(U_s) - E(U_p)$ decreases with v_p and increases with \bar{v}_p .

Effect of \bar{v}_p and v_p on the sampling effect: Sampling is effective when $E(U_s) > 0$ and $E(U_s) > E(U_p)$. Assume $E(U_s) = 0$ when $v_p = v_p^*$, $\bar{v}_p = \bar{v}_p^*$; and $E(U_s) - E(U_p) = 0$ when $v_p = v_p^{**}$, $\bar{v}_p = \bar{v}_p^{**}$.

Because both $E(U_s)$ and $E(U_s) - E(U_p)$ increase with \bar{v}_p , and $E(U_s)$ increases with v_p but $E(U_s) - E(U_p)$ decrease with v_p , the consumer both samples and also upgrades to the paid version when $\bar{v}_p \geq \max\{\bar{v}_p^*, \bar{v}_p^{**}\}$ and $v_p^* \leq v_p \leq v_p^{**}$.

For a piece of information on app quality that only increases \bar{v}_p but does not shift v_p , condition $\bar{v}_p \geq \max\{\bar{v}_p^*, \bar{v}_p^{**}\}$ is more likely to meet as the information indicates higher app quality and thus the likelihood that a consumer both samples and upgrades to the paid version increases with this quality signal. For a piece of information on app quality that shifts v_p only, condition $v_p^* \leq v_p \leq v_p^{**}$ is more likely to meet when the signal indicates that app quality is neither too low nor too high, thus the consumer is most likely to sample and to upgrade when this quality signal is moderate. If the information on app quality shifts both v_p and \bar{v}_p , the consumer would sample and upgrade as long as both $\bar{v}_p \geq \max\{\bar{v}_p^*, \bar{v}_p^{**}\}$ and $v_p^* \leq v_p \leq v_p^{**}$ are satisfied.