International Journal of Research in Marketing 36 (2019) 185–199

Contents lists available at ScienceDirect

IJRM
International Journal of Research in Marketing

journal homepage: www.elsevier.com/locate/ijresmar

Full Length Article

Measuring and forecasting mobile game app engagement

Oliver Rutz [a,*], Ashwin Aravindakshan [b], Olivier Rubel [b]

[a] University of Washington, Seattle (WA), USA
[b] University of California, Davis (CA), United States

ARTICLE INFO

Article history:
First received on June 1, 2017 and was under review for 8 months
Available online 29 January 2019

Area Editor: Andrew Stephen

Keywords:
Mobile
Games
Usage behavior
Apps

ABSTRACT

Monetization strategies in the large and strongly growing mobile app space must go beyond traditional purchase revenue as most mobile apps are now free to download. One key marketing innovation allows mobile app publishers to monetize ongoing engagement — in-app advertising. We study ongoing user engagement with a mobile app after the initial download decision using the $40 billion mobile gaming industry as an example. Our study investigates and forecasts user engagement after the initial download aiming to help publishers to monetize their engagement via in-app advertising. We leverage a novel dataset containing user-level engagement for 193 mobile games and propose a hierarchical Poisson model on a mobile-game level. We find significant usage heterogeneity across the mobile games studied and generate forecasts publishers can use when trying to monetize engagement via pre-sold contracts.

1. Introduction

In 2009, Apple was awarded a trademark for "There's an App for That." Today, the Apple App store alone has over 2.2 million apps available and passed the 180 billion download mark in June 2017 (Statista, 2017a). An analysis of the download data reveals that there is an ever increasing number of apps downloaded every year across all mobile operating systems — 197 billion in 2017 alone, an increase of roughly 33% over 2016 (Statista, 2017b). The number of mobile apps available keeps growing and current estimates put the number of smartphone users over 2.3 billion in 2017 (Statista, 2018). An even more significant development is that daily smartphone usage surpassed TV viewing in 2014 with users, on average, spending over 300 min on their smartphones (Andrews, Ellis, Shaw, & Piwek, 2015). While the average consumer uses about 30 mobile apps per month, he/she habitually accesses only about ten apps per day (AppAnnie, 2017). Given the large number of mobile apps downloaded, it is not surprising that, on average, ongoing engagement, i.e., repeat usage, is low for the large majority of mobile apps. While 77% of users give a mobile app a second try (Localytics, 2017), after 30 days, only about 3% of the downloaded apps are still active (Apptentive, 2017).

According to Comscore, in 2017, mobile game apps constitute the fourth largest segment of the mobile app market in terms of time spent with a share of about 10%, after Social Media (20%), Music (18%) and Multimedia (10%). The mobile gaming industry is expected to generate $52 billion in revenues in 2018 with a share of 45% of the global games market (Newzoo, 2017). This expected growth is driven not only by platforms, e.g., cellphones and tablets, but, more importantly, by marketing innovations specific to the mobile industry. First, in-app purchasing allows users to improve the gameplay by, for example, buying items to be used in the game or saving time needed for certain achievements. Second, in-app advertising allows monetizing users'

* Corresponding author.
E-mail address: orutz@uw.edu. (O. Rutz).

https://doi.org/10.1016/j.ijresmar.2019.01.002
0167-8116/

engagement with the game by showing ads, for example, before a user can continue using the game. Thus, unlike traditional games (i.e., non-mobile), a mobile game app can generate revenues above and beyond the initial purchase by monetizing the ongoing engagement with users. These new mobile game revenue models (e.g., freemium[1]) contrast with traditional revenue models of non-mobile games played on consoles where revenues come first and foremost from sales.

The marketing literature on product usage provides only partial insights with regard to mobile game usage behaviors and monetization strategies. Existing related studies focus on how purchase behaviors are based on anticipation of product usage (e.g., Lambrecht, Seim, & Skiera, 2007) as well as how pricing models affect product usage in contractual settings (e.g., Gopalakrishnan, Iyengar, & Meyer, 2015; Iyengar, Jedidi, Esseghaier, & Danaher, 2011). Despite the fact that engagement generated by mobile game is a critical component of the monetization strategy, the extant marketing literature on mobile apps focuses mostly on the purchase or download decision rather than on engagement (usage) behavior *after* the app was initially downloaded (e.g., Ghose & Han, 2014; Lambrecht et al., 2014). A notable exception is the work of Dew and Ansari (2018) focusing on in-app purchasing behavior after the initial download. However, their study does not investigate ongoing engagement and does not leverage usage data. Next to in-app purchasing, the second key dimension of a successful app monetization strategy is to show advertising to users as a “price” to keep using the free app. Critical to that strategy is the ability to understand and forecast engagement (usage) on an app-level to assess the opportunities for selling this engagement to interested advertisers.

As fewer and fewer apps charge for the initial download, monetizing ongoing engagement is becoming more critical for any app to succeed. Often, in-app purchasing is not a strategy that generates much additional revenue. In the mobile game space, it is well known that a very small percentage of users, 3.5% according to one source (AppsFlyer, 2016), spend any money in a game in terms of in-app purchasing. Thus, leveraging in-app advertising seems a more promising strategy and is an always feasible alternative to monetize ongoing engagement. Our paper aims to study what drives such ongoing engagement (usage) with mobile game apps after the initial download that could be monetized via in-app advertising. We also show how engagement (usage) can be forecasted on the app-level to allow assessment of the engagement monetization potential through in-app advertising. To do so, we leverage a novel dataset consisting of mobile-game level usage information in the form of (repeat) usage counts over a period of eleven months.

From our perspective, the primary reason for a lack of studies on mobile app engagement in terms of ongoing usage is the absence of easily available metrics that relate to engagement with a mobile game as well as the novelty of the engagement-based business models now prevalent in the app space. Potential proxies available via app stores (e.g., Google's Play Store) are, for example, the number of downloads or the rank of the app. These metrics do not necessarily predict engagement beyond the initial download. For example, 49% of consumers install at least one app per month while 29% install between 2 and 8 apps and 5% install more than eight apps per month (Statista, 2017a), but about 75% of their app usage time is restricted to about four apps (Comscore, 2017). Thus, freely available metrics seem inadequate to determine the level of engagement an app can create. Our data are coming from a large panel of consumers that have installed a tracking app on their mobile devices allowing us access to their usage behaviors after they have downloaded apps. We use these data to create app-level usage data to investigate and forecast mobile game app usage.

The remainder of the paper evolves as follows. We first discuss the mobile app space before we present our novel mobile game usage data. We then detail our modeling approach and discuss our empirical findings and forecast. We finish with the conclusion and questions for future research.

2. Mobile app monetization strategies

In a traditional marketing setting, value for the firm is created when a customer purchases the good or service the firm is selling. Thus, most of the value in a customer–firm relationship is captured at the point of the transaction between the two players — which often occurs before the customer realizes how much they actually use the product. However, in the case of apps, and mobile games in particular, from the customer's perspective, the majority of the value is not realized at the point of download, either paid for or free, but through (ongoing) engagement with the app (i.e., continuing usage). While this is true for many traditional products, e.g., books or music, firms usually cannot extract more value from the ongoing use of these products. In the case of apps, developers have the unique opportunity to continually capture value while the customer uses the app. The two key methods to capture this value are in-app purchasing and in-app advertising. Using an in-app purchasing strategy allows users to upgrade to the paid version, unlock paid features, and buy special items for sale. In-app advertising follows a more traditional advertising model, e.g., TV advertising, where users need to watch advertising before they can continue to use the app for free (Fig. 1).

There is ample evidence that in traditional customer acquisition settings, monetary incentives, e.g., “free toaster”, can lead to adverse selection in terms of attracting deal-prone consumers who are likely leave for the next deal offered. For example, Datta, Foubert, and Van Herde (2015) find that for digital TV subscriptions, customers who joined using a free trial have a 59% lower customer lifetime value. It is not clear whether the freemium strategy can be seen as a mobile version of a free trial, especially given that apps that are positioned as freemium are not targeted as freemium to some audiences and are costly to others. Investigating one key monetization dimension of freemium apps, in-app purchasing, Dew and Ansari (2018) propose a novel approach on how firms can leverage in-app purchasing data to evaluate customer lifetime value and generate a dashboard for managers to understand the lifecycle of an app. However, to the best of our knowledge, no work currently investigates what drives ongoing engagement, which is necessary to successfully leverage in-app adverting as a monetization strategy.

[1] Freemium is a pricing strategy by which a product or service (typically a digital offering or application such as software, media, games or web services) is provided free of charge, but money (premium) is charged for proprietary features, functionality, or virtual goods.

Fig. 1. In-game advertising — a sponsored level in "Angry Birds".

In 2017, global mobile app revenue is still driven by paying for the app (39%) closely followed by in-app purchases (38%) and in-app advertising (23%, AppsFlyer, 2016). Mobile app advertising is either pre-sold by a so-called "eyeball" contracts or sold using auction algorithms, so called real-time bidding or RTB. The "eyeball" contracts mirror how traditionally banner ads have been sold while the auctions are similar to methods used to sell search engine advertising. While RTB is increasing, the majority of mobile ads are still sold via "eyeball" contracts (Fisher & Liu, 2016). To successfully monetize a mobile game's available in-game inventory, i.e., the to be generated engagement by the mobile game that allows to show ads, a game publisher benefits from understanding engagement as well as forecasting such engagement. As a large part of the in-game inventory is sold by contracts, it is beneficial to the publisher to be able to forecast the engagement to be sold.

We propose an approach to model and forecast the "eyeball" supply of games (or the in-game inventory) that can be monetized via in-app advertising leveraging novel data on app usage available. Our work complements that of Dew and Ansari (2018) focusing on in-app purchasing as an app monetization strategy. Our data are obtained from a market research firm in India. They contain roughly one thousand unique consumers per month for a span of 11 months and detail each consumer's complete smartphone usage in the month this consumer has been in the panel. The data are collected via a 'metering app' installed on the panel members' smartphones and captures the usage behavior in a non-invasive manner. While a user-level model might be desirable, in our application as in many others, the data are sparse on a user-game level. We are aiming to understand engagement and ultimately in-game inventory for in-app advertising on a game-level. We aggregate user engagement on the game-level and create a dataset of ongoing engagement that we measure as repeat usage counts for a large number of popular games. We employ a Poisson model to understand how usage differs across these mobile games based on their characteristics, enabling us to better understand the engagement a game generates as well as forecast usage (i.e., the "eyeball" supply available to publishers). Our empirical findings reveal that there is significant heterogeneity across mobile games with respect to the level of engagement they generate. We quantify how observable characteristics (e.g., from the app store) inform estimates of both the mean and the variance of mobile game usage.

As discussed above, one of the available monetization strategies is in-app advertising. Typically, in-app advertising is sold in one of two ways. One way is more aligned with traditional advertising where an advertiser purchases advertising exposure in a specific app or content domain via an exposure contract. The second way in-app advertising is sold is via real-time bidding (RTB). Currently, roughly 30% of mobile ads are sold via programmatic advertising of which RTB is one type. RTB, introduced in 2009, has quickly become one of the main methods of buying and selling mobile ads. It currently accounts for over a third of the mobile advertising market in the US, and is estimated to exceed $14 billion in 2017 (Fisher & Liu, 2016). It is interesting to consider the future of "eyeball" contracts (or reservation contracts) vs. RTB. Recent work shows that despite RTB's higher efficiency in ad allocation and advertisers' higher willingness to pay for RTB impressions, publishers should not phase out reservation contracts as these allow publishers to set higher reserve prices in RTB and achieve higher total revenue (Sayedi, 2018). Consistent with this finding and despite its rapid growth in the past several years, industry reports show that the growth of RTB is slowing down. The spending on RTB, as a fraction of total mobile advertising in the US, is estimated to be 31%, 34%, 36% and 36% in 2015, 2016, 2017 and 2018, respectively (Fisher & Liu, 2016; Liu, 2016).

To be able to sell ads via the non-RTB route, a publisher needs to be able to forecast its "eyeball supply" to enable writing contracts that are based on exposure and a cost per exposure. In practice, publishers do not know how many users (or usage) their apps will have in the future. As such, publishers are uncertain about the size of their supply when selling impressions in reservation contracts. Before RTB, this uncertainty, along with the fact that reservation contracts include penalties for under-delivery, made the decision of how many impressions to sell in reservation contracts a major challenge for many publishers, and publishers were often left with unsold impressions. Our model addresses these issues in understanding engagement (or usage) of mobile apps and providing a method to forecast future "eyeball supply".

In sum, in-app purchases are currently the dominant monetization strategy. However, about 70% of in-app purchases are driven by 10% of the gaming population. The vast majority, 90% of users, only generates about 30% of the in-app revenue (AdWeek, 2016). To better monetize these remaining 90%, mobile game publishers have begun to rely on in-app advertising.

From a manager's perspective, it is therefore critical to understand and forecast how much *engagement* each game will create to inform in-game inventory supply. Apart from the marketing literature discussed above, several studies in computer science also model engagement with apps in various scenarios, ranging from quantifying the diversity in smartphone usage (Falaki et al., 2010), using app-usage behavior to predict the next app that will be used (Liao, Pan, Peng, & Lei, 2013), predicting how certain types of app-usage leads to brief bursts of interaction with multiple apps (Ferreira, Goncalves, Kostakos, Barkhuus, & Dey, 2014), the influence of contextual information like location or time of day on app usage (Böhmer, Hecht, Schöning, Krüger, & Bauer, 2011; Shin, Hong, & Dey, 2012), the relationship between trust and app-usage (Kim, Shin, & Lee, 2009), to understanding how consumers transition between apps based on usage (Gouin-Vallerand & Mezghani, 2014). However, to the best of our knowledge, no prior study in the marketing or computer science literature focuses on modeling and forecasting ongoing engagement itself as a supply for in-app advertising.

While downloads are necessary for usage, they are not sufficient to understand engagement and, in turn, the value of the engagement when it comes to monetizing it through in-app advertising. We posit that understanding customer usage can allow determining how valuable an app might become as an advertising vehicle. This value, from a publisher's perspective, should inform the supply this app can provide and potentially inform on the price to be charged. For example, does an app generate mostly new users (very limited repeat usage but large number of new downloads) that will be exposed to an ad only one or very few times? Or does an app have a stable user base that will allow exposure to the same ad multiple times potentially leading to ad wear-out or ad fatigue? Clearly, depending on the scenario, the price to be charged for the advertising exposure should be different.

A caveat for such a usage-based approach is that usage information is hard to collect. However, the app is a true online app, i.e., cannot be played offline and requires a connection to the publisher's server, such information would be available. But typically data across apps are not available if they are not by the same publisher. We circumvent these data availability issues by collaborating with a market research firm that employs a large panel of mobile users allowing the collection of all mobile usage behavior as it related to apps installed on the panel members' mobile devices. The data collected are in count form and represent how many users in a given period of time have used a large set of apps at least once or a repeat number of times. For example, we observe that a user has used the mobile game app “Angry Birds” 15 times over a time period of 4 weeks. Additionally, we collected publicly available data from the app store on app characteristics. For example, the rank of “Angry Birds” in the app store is 12. We aim to show how publicly available app store information can be leveraged to predict the customer engagement for an app and generate usage forecasts. In the next section, we describe the data collection process and the data used in our study.

3. Data description

The data are obtained from market research firm situated in India.[2] The firm collects data on mobile usage from a panel of consumers. They provided us with data on mobile usage for approximately one thousand unique consumers per month over the period of 11 months from July 2011 to May 2012. We note that the firm only tracks a consumer for a period of 30 consecutive days in a given year. The panelists are not aware of the exact month they will be tracked. The firm follows this practice in order to preclude any change in behavior due to the panelist knowing that their mobile behavior is being tracked. The identity of the firm and exact location of the data are confidential. The firm collects data from its panel members via a ‘metering app’ installed on the panel members' smartphones and the metering application captures the usage behavior of panel members in a non-invasive manner. This allows the firm to not only accurately monitor usage, but also precludes the necessity of primary data collection via surveys as usage behavior is directly observed and recorded. The firm aims to randomly choose consumers to invite to join the panel. Prior to joining the panel, the firm records the panelist's demographic, cell phone and mobile plan information. The firm updates this information on an annual basis to record any changes that might occur. Upon inclusion in the panel, consumers install the metering app on their smartphone allowing the firm to monitor smartphone usage. Behaviors monitored include the usage of apps. More precisely, the metering app records the name of the app the consumer uses, its category (e.g., game), hour and date the app was opened, and length of use. For the purposes of this analysis, we restrict our attention to the mobile game apps category. In our panel, we observe 1014 unique mobile game apps played by the panelists.

While we have data on mobile gameplay at the user-level (i.e., detailed usage information on which games each panelist engages with), we do not have enough information per mobile game on a user-level to investigate engagement within any given mobile game leveraging a user-level modeling approach. As discussed in the Introduction, most users do not stick with any given game (or app) for long and we find evidence of this behavior also in our data. User-level mobile game usage data are sparse for any given game that is not a highly successful game such as Angry Birds. As such, we could either study only every games on a user-level or aggregate across users and study engagement on a mobile game-level. As we aim to understand and forecast mobile game engagement for in-app advertising, we aggregate across users and create mobile game-level usage data as described below.

In order to study mobile game engagement, we aggregate mobile usage of 1014 mobile games over a period of eleven months across consumers. In total, we have data on 49,292 usage occurrences. While this seems to indicate roughly 50 usage occurrences per mobile game, the usage data are highly skewed and the median number of mobile game usage occurrences is only 5. Due to the large number of low use mobile games (again, the majority of mobile games is used very few times if at all), we only include mobile games that account for about 90% of the observed usage in our data, which results in a total of 193 mobile games being included.[3] The average usage for the 193 mobile games is 168.1 with a standard deviation of 484.6. The distribution of the number

[2] The name of the company and further identifying details of the data are confidential.
[3] Similar to studies of brand choice only focusing on the large market share brands we can also not assess engagement with niche mobile games.

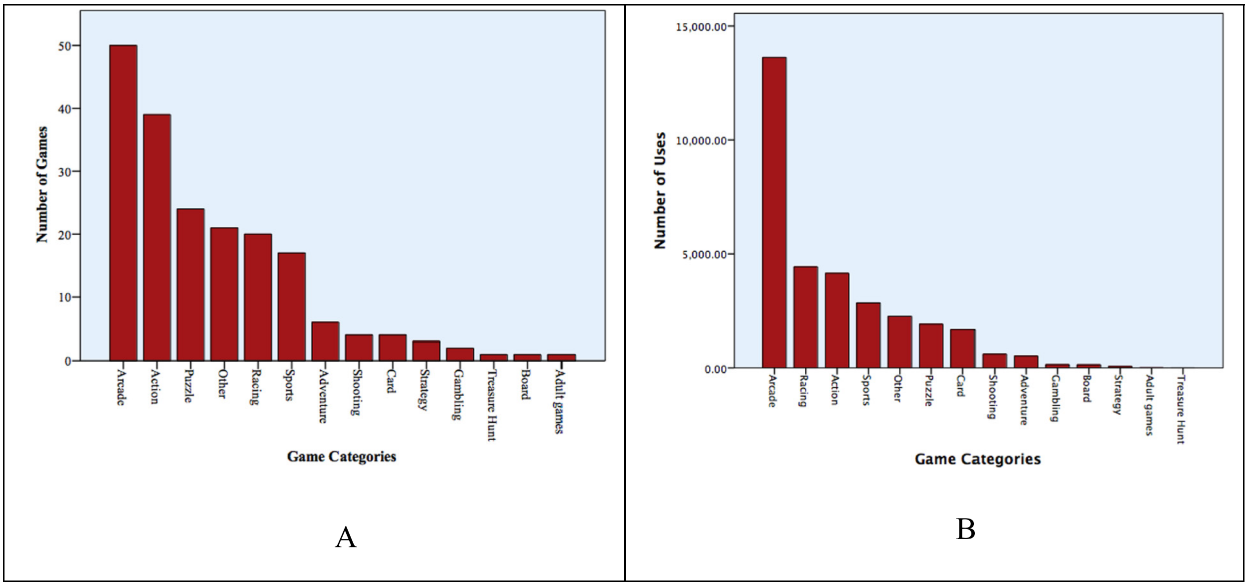

Fig. 2. Number of games and number of uses by game category.

of mobile games by category (as defined by the firm) is shown in Panel A of Fig. 2. Note that the top 5 mobile game categories account for almost 80% of the mobile games observed in the data and the average number of mobile games in a category is roughly 14. Panel B of Fig. 2 displays the distribution of usage by the mobile game category.

We augment the data with mobile game-specific information from the mobile game's page from the Play Store or iTunes Store. Specifically, we collect the average rating, the number of reviews, the number of screen shots displayed and the text describing the mobile game. Fig. 3 provides an example of a mobile game's page from the Play Store. We provide a summary of the data collected in Table 1.

We find that the mobile games in our sample rate highly with an average rating of 4.3 (out of 5). We note that this is an artifact of our sampling approach. Recall we are focusing on the top 193 mobile games that generate about 90% of the usage out of 1014 mobile games that were used at least once by a panel member. Indeed these are also the mobile games that are rated highest by other users assuming our panel is representative. Similarly, the number of reviews for the focal 193 mobile games is high with a mean of 42K and a median of 8.1K. The difference between the mean and median is mostly driven by a few top selling and very popular mobile games such as Angry Birds, which alone has 986K reviews. We also find that publishers vary extensively in the amount of text they use to describe their games, with a minimum of 19 words for the shortest description and a maximum of 585 words for the longest description. We also note that the number of screenshots that allow a potential customer to "see" the game varies greatly from 2 to 8 with a mean of 5. Additionally, Panels A, B, C and D in Fig. 4 graphically display the distributions of the average ratings, number of reviews, word counts and the number of screenshots across all 193 mobile games.

Next, we introduce our model and detail the estimation procedure.

4. A model of mobile game engagement

Our goal is to determine what drives engagement, measured by repeat usage, of mobile games. We propose a mobile game-centered approach and model usage across mobile games conditional on purchase (i.e., download). Our approach is in line with existing research, which so far has investigated why certain apps are purchased and while others are not (e.g., Ghose & Han, 2014). We add to this literature by extending the view beyond the initial purchase occasion to engagement we measure by repeat usage.

As described above, for each app (in our case, mobile game) we have user-level mobile game usage data. While an individual-level model is desirable, the data become sparse for the large majority of mobile games in our dataset when trying to construct user-level and game-level panel data. We are interested in understanding engagement with a mobile game to value the engagement that mobile games can create. As such, we model engagement on a mobile game-level and aggregate over individuals. As discussed above, we focus on the top 193 mobile games out of 1014 total mobile games played by our panelists due to data sparseness for the majority of the mobile games. While it is an interesting question to study engagement of a long-tail (low usage) mobile game, our data are not rich enough to allow us to do so, which is not untypical for marketing research in general.

Given our mobile game-level approach, we describe in detail next how we construct the dataset to implement our proposed approach. We observe, for each mobile game, engagement in terms of the unique total usage counts[4] as well as the number of users with these usage counts. For example, say we would observe the following usage count data for an exemplary mobile game over our observation period: (2 2 2 7 7 15). Usage count refers to the observed number of repeat usage observations per

[4] For the ease of exposition, we will use usage count(s) instead of total usage count(s) for the remainder of the paper.

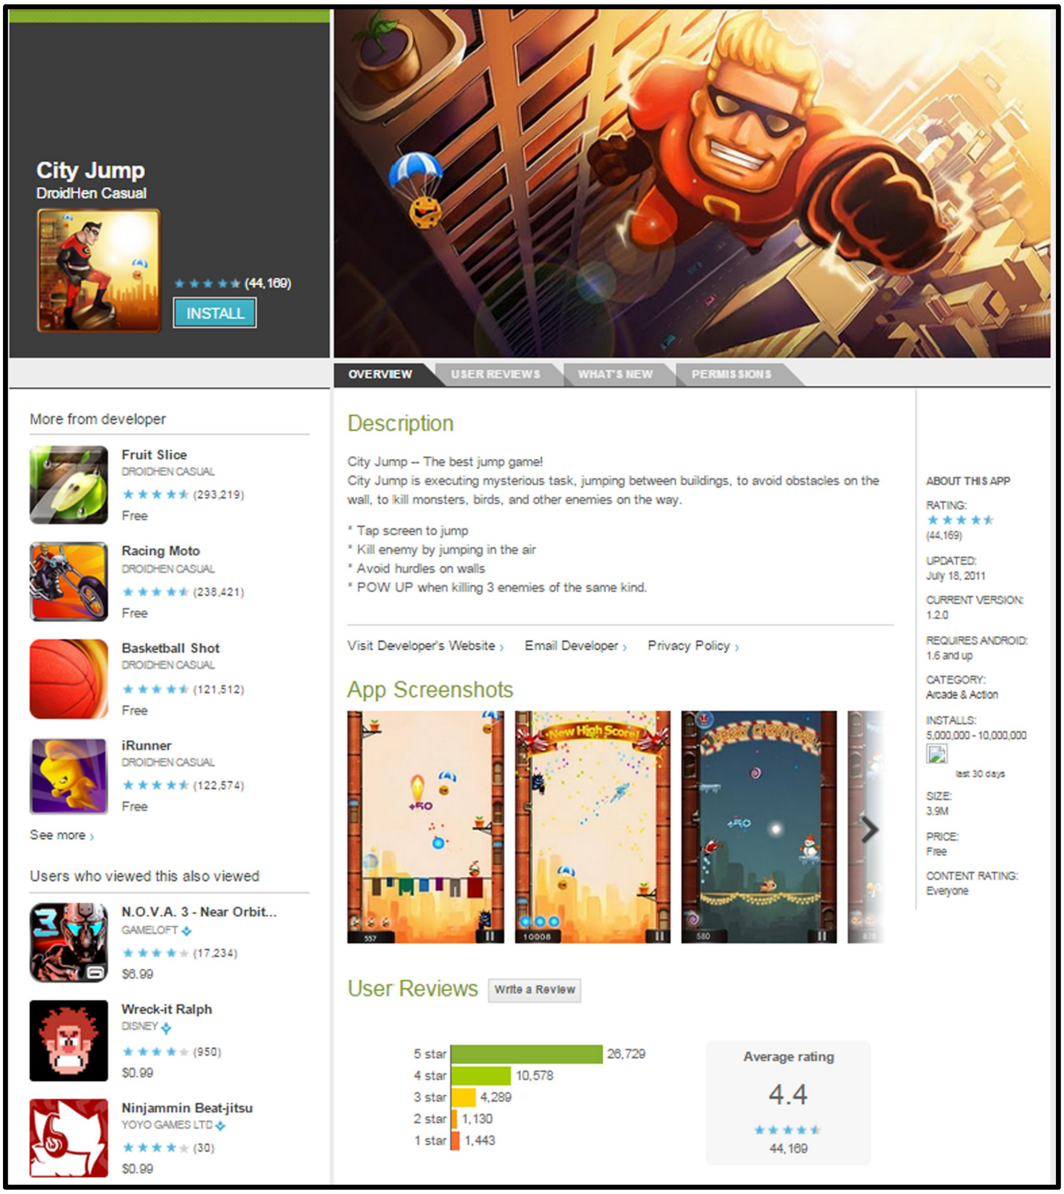

Fig. 3. Example of a Google Play Store page — City Jump.

user for a given mobile game for one month (recall each user is in the panel for 1 month). We call the unique usage counts y_i for mobile game i. For example, "2" is a unique usage count for the exemplary mobile game. Continuing with our example, we observe that 3 users have used the exemplary mobile game "2" times. Note that we calculate total usage for each user only one time per mobile game (i.e., no double counting) over the user's observation period and use it as their unique usage count, e.g., "2" in our running example.

Table 1
Overview of game characteristics metrics.

	Mean	Std. dev.	Median	Minimum	Maximum
Number of uses	168.0	484.6	32.0	1.0	5760.0
Average rating	4.3	0.4	4.4	2.5	4.9
Number of reviews	42,237.0	102,036.0	8140.0	8.0	986,000.0
Number of words	187.4	113.1	173.0	19.0	585.0
Number of screenshots	5.0	2.0	5.0	2.0	8.0

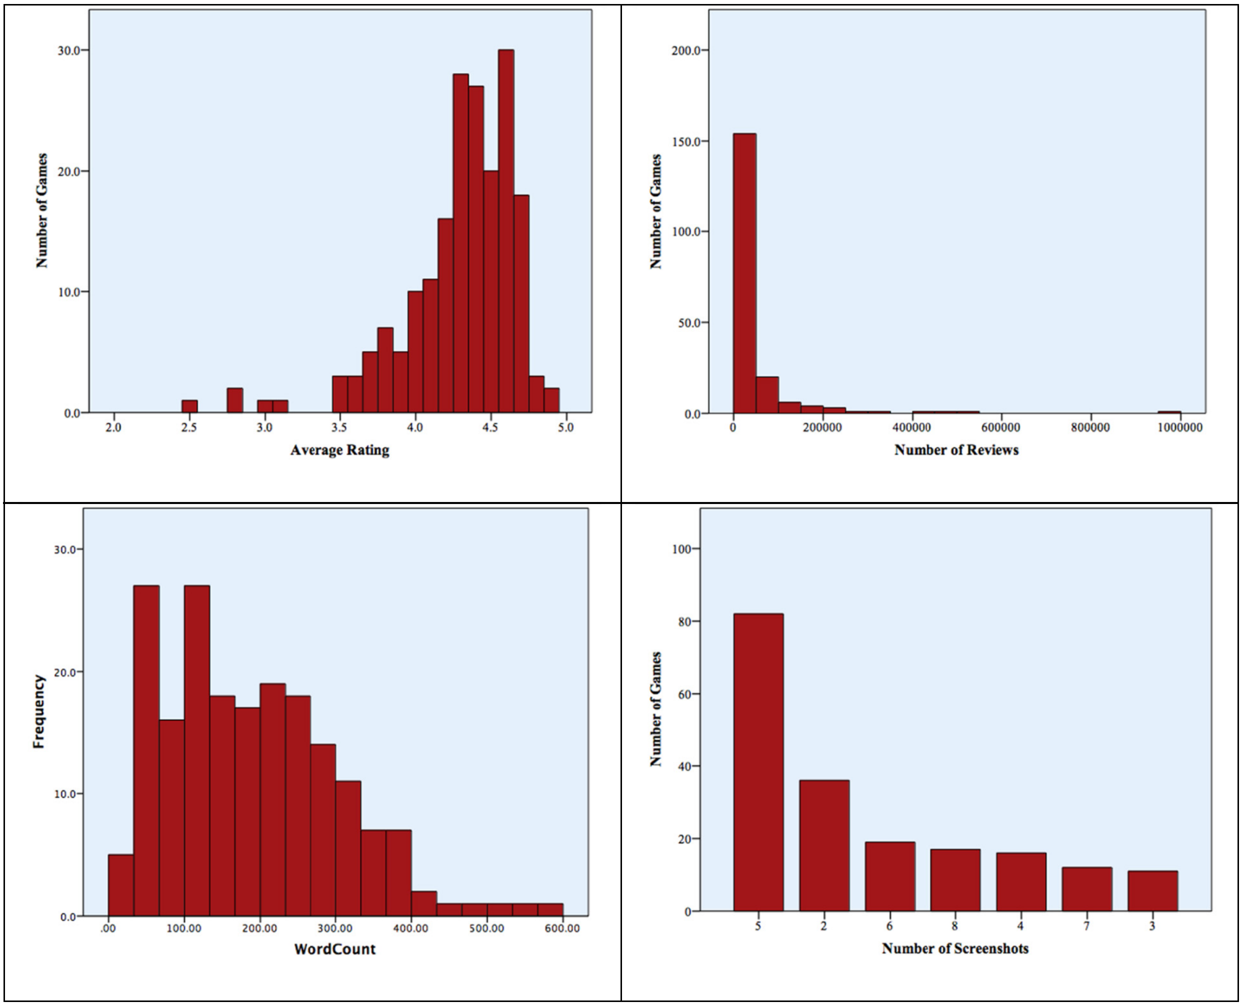

Fig. 4. Distributions of game characteristics metrics.

Given that we have data across mobile games, we simplify the data set-up by modifying the likelihood function given that we have a homogenous model on the user-level and address heterogeneity on the mobile game-level. Instead of having the usage count for "2" enter 3 times in the likelihood, we create an additional variable called d (in our example above $d = 3$) allowing to account for the number of users who have used the mobile game the same amount of time in the likelihood.[5] Summing up, for each mobile game i, we have usage count data for usage occasion j, called y_{ij}, as well as d_{ij} to account for the number of times usage count y_{ij} has been observed in our data.

Next to the mobile game usage counts, we collect, for each mobile game, additional characteristics, e.g., category and textual description by the company, as well as consumer feedback in the form of mobile game ratings.

We model the data using a zero-truncated Poisson model. The usage data are modeled as

$$y_{ij} \sim TruncPoisson\left(\lambda_{ij}\right), \tag{1}$$

where i denotes the mobile game and j denotes the j-th usage count observation, e.g., "2" from our example,and λ_{ij} is to be estimated.

Our goal is to model ongoing mobile game engagement, i.e., repeat usage. As such, we model the mobile game- and usage-count specific Poisson rate λ_{ij} using a hierarchical Poisson regression (Christiansen & Morris, 1997). In order to deal with the constraint $\lambda_{ij} > 0$ we define λ_{ij} as follows:

$$\lambda_{ij} = \exp(x_i\beta), \tag{2a}$$

where x_i are mobile game-specific covariates,and β is a parameter vector to be estimated.

[5] For example, instead of having $p(y_1 = 2) * p(y_2 = 2) * p(y_3 = 2)$ enter the likelihood we have $p(y = 2)^d$ enter the likelihood where $d = 3$.

It is well known, and our data are no exception, that the Poisson distribution suffers from the problem of over-dispersion. For models with few parameters like the Poisson model, theoretical predictions may not match empirical observations for higher moments. Over-dispersion occurs when the observed variance is higher than the variance of a theoretical model (in the Poisson model mean and variance are the same by definition). We follow Chib and Winkelmann (2001) and address this issue adding a random effect that increases the models flexibility. As they note "… the model allows for over-dispersion, a variance in excess of the expectation…" (Chib & Winkelmann, 2001). We modify the basic hierarchical Poisson regression model given by Eqs. (1) and (2a) as follows:

$$\lambda_{ij} = \exp\left(x_i\beta + \varepsilon_{ij}\right), \tag{2b}$$

where x_i are app-specific covariates,

β is a parameters to be estimated,

and ε_{ij}~$N(0,\sigma^2)$.

This is the Poisson LogNormal model. Typically, an uninformative prior is chosen for the variance of ε_{ij}, e.g.,

$$\sigma \sim gamma(c, c), \tag{3a}$$

with c chosen to make the prior uninformative.

In typical applications of this model, the data come from one underlying process with one assumed underlying variance to be estimated. Our data do not come from one underlying process as we have different mobile games (193) for which we observe usage counts. As such, the assumption of one shared variance over the 193 mobile games is potentially untenable. Similar to models on consumer choice, we also do not wish to run 193 separate models (one for each mobile game) but to pool across the 193 games to share information efficiently. In our case, a number of mobile games still have relatively low usage and a model based on such limited data might be questionable. We leverage the fact that we have 193 different mobile games by introducing a hierarchy on the variance. In essence, our model assumes that the mobile game-level variances arise from a common distribution of variances similar to individual-level effects capturing consumer heterogeneity in choice models. Our data has a panel-like structure in that we observe multiple usage counts for each mobile game in the dataset. This allows us to add additional flexibility into our model and we allow the over-dispersion to vary across mobile games by adding a level of hierarchy on the over-dispersion parameter, i.e., σ. We introduce an app-specific variance, σ_i, arising from a gamma distribution as follows:

$$\sigma_i \sim gamma(c_i, c_i), \tag{3b}$$

where c_i are prior parameters.

We select an informative prior to leverage observed mobile-game heterogeneity in over-dispersion. In our case, the parameter c_i (again, $c_i > 1$) of the gamma distribution depends on mobile game-specific covariates as follows:

$$\log(c_i) = z_i\gamma, \tag{4}$$

where z_i are mobile game-specific covariates,and γ is a parameter to be estimated.

To summarize, the model hierarchy is as follows

$$\begin{aligned}
&\beta|y_{ij}, x_i, d_{ij}, \mu_\beta, \sigma_\beta, \varepsilon_{ij} \\
&\varepsilon_{ij}\Big|y_{ij}, x_i, d_{ij}, \mu_\beta, \sigma_\beta, \sigma_i^2 \\
&\sigma_i^2\Big|\varepsilon_{ij,} z_i, \lambda \\
&\lambda|\sigma_i^2, z_i, \mu_\lambda, \sigma_\lambda
\end{aligned} \tag{5}$$

We estimate our model in a Bayesian framework using MCMC sampling (please, see Appendix I for details).

Note that in a standard Poisson regression model ε_i is an observation-specific shock arising from a distribution characterized by the variance (Chib & Winkelmann, 2001). Our data allows us to relax this assumption and introduce a hierarchical Poisson model for panel-style data. We have 193 mobile games and leverage these mobile game panel-style data to pool across the mobile games. We introduce an observation- and mobile game-specific error shock ε_{ij} from a mobile game-specific error distribution with variance σ_i. Instead of estimating 193 separate models, each with its own variance, we use a Bayesian approach to pool over our 193 mobile games and leverage information across mobile games to better assess mobile game-level variances. In the presence of a large dataset and similar amount of data per mobile game, the pooling approach would essentially result in similar estimates to estimating 193 separate models. However, because some mobile games have many more observations than other mobile games, pooling allows us to more precisely estimate variances of the mobile games with lower number of observations. The model also shares information across mobile games and it is dependent on observable mobile game characteristics, in contrast to the variances simply arising from a common distribution alone. We have roughly 4500 data points over which we pool the 193 mobile games and estimate 25 beta parameters and 25 gamma parameters based on these data. We also include a simulation study to detail recovery (see Appendix II). Via a simulation study, we also investigate the issue of aggregation bias.[6] Potentially, one could calculate the average usage per mobile game and investigate differences in average usage per mobile game based on these 193 data points. Typically, such an approach is untenable in the face of richer data and referred to as aggregation bias. In a simulation study, we find that our proposed model correctly recovers the true parameters while the aggregate model is not able to do so (see the Appendix II). Next, we discuss the empirical results of the proposed model and its performance compared to other approaches.

5. Empirical results

We estimate our proposed model and other, competitive, models on data of 193 mobile games. For these 193 mobile games, we have 1359 unique usage counts and 4500 total consumer usage observations. Note that a specific usage count for a certain mobile game is often observed to occur for more than one consumer, please see the Data description section for details. The minimum observed usage count is 1, i.e., the mobile game gets used only one time and the maximum observed usage count is 264. Recall that a usage count would be, for example, that a certain mobile game has been used 5 times by at least one consumer.[7] For each mobile game, we also observe information provided by the mobile game publisher in terms of the text as well as the screen shots used to describe the mobile game in the app store. We have used text mining to generate predictors that capture information in the text as well as content from the actual website in the app store (see Fig. 7). We capture the content of the web presentation of the mobile game by the count of screen shots in the description, the number of words in the description, as well as the counts of the following words: fun, enjoy, play, paced, "?", "!", "*", best, addictive, realistic, Free, Million, Award(s), Winner. Next, we have information on the category as provided by the market research firm that collected the data. We use their secondary category indicator that allows segmenting the game category into finer sub-categories. Our sample includes 17 secondary category indicators. Based on the variation in the data, we estimate 10 secondary category indicators.[8] We have indicators for the following secondary categories: Action, Adventure, Arcade, Card, Other Games, Puzzle, Racing, Shooting, Sports, and Strategy. Lastly, we have user-based covariates for the mobile games. We have collected information on users' response to the mobile game. For each mobile game, we have the average rating (on a 1 to 5 star scale where 5 stars is the best and 1 star the worst rating) as well as the number of ratings. For our 193 mobile games, the mean rating is 4.3 with a minimum rating of 2.5 and a maximum rating of 4.9. In terms of number of ratings, we find that, on average, the mobile games have about 42K ratings. While some mobile games have few ratings (the minimum is 8), some mobile games have a very large number of ratings (the maximum is 986K).[9] All together, we have 25 variables describing every mobile game. We use these mobile-game specific data in our mean, i.e., x_i, as well as our variance, i.e., z_i, model. Note that while the marketing research firm has collected additional information on the user-level, e.g., demographics or details on the smartphone, we cannot use these data in our proposed approach as we need to aggregate over users to circumvent the sparseness in the mobile game usage data. As more data becomes available, a fascinating topic for future research would be investigating mobile game usage on a user-level.

Next to estimating our proposed model, we have estimated a base model assuming no over-dispersion, the model given by Eqs. (1) and (2a), as well as a model that allows for over-dispersion but no heterogeneity in the over-dispersion parameter, the model given by Eqs. (1), (2b) and (3a). We find that our proposed model, i.e., given by Eqs. (1), (2b) and (3b), fits the data best with a marginal log-density of −6462.3 compared to a model not accounting for over dispersion (marginal log-density −30,584.2) and a model with a homogenous over-dispersion parameter (marginal log-density −6584.5).[10] We also compute mean average percentage error (MAPE) and find that our proposed model has a MAPE of 0.29 while the model with homogenous over-dispersion has a MAPE of 0.39, meaning our proposed model improves MAPE by roughly 25%. The model not accounting for over-dispersion fares poorly with a MAPE of 1.61. We also detail the recovery of the usage data for one randomly chosen mobile game in Fig. 5. As we can see, a model not accounting for over-dispersion in the data fares poorly in fitting the data (as indicated by the very poor performance with regards to the log-marginal density and MAPE). The proposed model with heterogeneous over-dispersion fits the data the best while the model with homogenous over-dispersion seems to typically overestimate the usage counts, at least for the exemplary mobile game we use to detail model performance.

First, we discuss the results of the mean component in the model, β. Recall that we model $\lambda_{ij} = \exp(x_i\beta + \varepsilon_{ij})$, where β describes how mobile game covariates, x_i, e.g., text or user ratings, affect the mean usage. In terms of the mean component, we find, on average, a baseline usage of 1.22 times[11] based on the intercept alone (please see Table 2 for parameter estimates). The secondary category as defined by the market research firm who provided the data enables differentiating across mobile game in terms of different categories. Compared to the baseline, Adventure and Other Games do not have a significant different usage. Only Strategy usage is higher than the baseline usage by 0.35. Action, Arcade, Card, Puzzle, Racing, Shooting, and Sports usage is

[6] We thank the review team for these suggestions.

[7] Please note that we are not using one consumer multiple times within a mobile game app. For example, a consumer has used a mobile game a total of 5 times. Naturally, the consumer has also used the mobile game 4 times. But we do not count this consumer as a 4-time user, we only use the total usage of the consumer per mobile game.

[8] For seven secondary category indicators, one or less mobile games are observed. Thus, we collapse these seven secondary categories into the intercept. The seven secondary categories not estimated are: Adult, Ball, Board, Gambling, Quiz, Treasure Hunt, and Word or Number games.

[9] The Angry Bird app has the highest number of ratings in our data followed by the Paradise Island (515,836 ratings) and the Tap Fish app (471,736 ratings).

[10] Note that the difference in marginal log-density between our proposed model and the model with a homogenous over-dispersion parameter is roughly 22. A log-marginal density difference is significant above the 6–8 points in difference typically referred to in Bayesian model selection, also called Bayes factor (Newton and Raftery, 1994).

[11] Recall that mean usage is λ_i which is modeled as $\log(\lambda_i) = x_i\beta$.

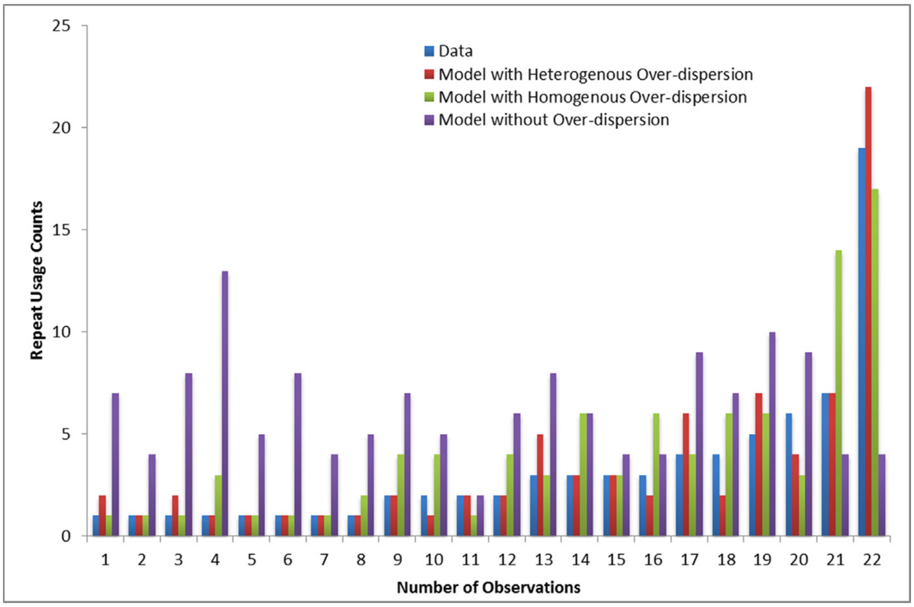

Fig. 5. Exemplary model fit.

lower than the baseline usage. Respectively, the usage is lower by 0.52 (Action), 0.42 (Arcade), 0.17 (Card), 0.51 (Puzzle), 0.43 (Racing), 0.33 (Shooting) and 0.14 (Sports). Next, we find that average rating as well as number of ratings is, as expected, positively related to usage. An increase in star rating of 1 star (say from 3.5 stars to 4.5 stars) increases usage by about 0.38 compared to the baseline. We also find that the number of ratings is positively related to usage. In order to account for decreasing marginal effects of the number of ratings (some apps have very large ratings), we use the log of the number of ratings in our analysis. We find that an increase in the number of ratings by 1000 increases usage by about 0.22. In terms of the text components (or the marketing actions by the firm), we find that the number of screen shots used as well as the text components "paced", "addictive", "realistic", and "million" coincide with the higher than average usage, while the text components "fun", "enjoy", "award(s)" and the log number of words coincide with lower than average usage.

Table 2
Results — mean component (β).

	Mean	Coverage interval	
Intercept	**0.20**	**0.07**	**0.32**
Action (category indicator)	**−0.56**	**−0.62**	**−0.49**
Adventure (category indicator)	0.05	−0.03	0.19
Arcade (category indicator)	**−0.43**	**−0.56**	**−0.31**
Card (category indicator)	**−0.15**	**−0.28**	**−0.07**
Other games (category indicator)	−0.06	−0.17	0.04
Puzzle (category indicator)	**−0.55**	**−0.60**	**−0.49**
Racing (category indicator)	**−0.43**	**−0.52**	**−0.32**
Shooting (category indicator)	**−0.32**	**−0.40**	**−0.25**
Sports (category indicator)	**−0.12**	**−0.20**	**−0.06**
Strategy (category indicator)	**0.25**	**0.18**	**0.30**
Average rating	**0.27**	**0.23**	**0.32**
Number of ratings	**0.16**	**0.13**	**0.20**
Number of screen shots	**0.03**	**0.01**	**0.04**
Number of words	**−0.21**	**−0.27**	**−0.16**
Fun (word indicator)	**−0.10**	**−0.16**	**−0.07**
Enjoy (word indicator)	**−0.14**	**−0.21**	**−0.09**
Play (word indicator)	0.01	−0.01	0.03
Paced (word indicator)	**0.34**	**0.28**	**0.43**
Best (word indicator)	0.02	−0.02	0.05
Addictive (word indicator)	**0.10**	**0.03**	**0.18**
Realistic (word indicator)	**0.17**	**0.08**	**0.25**
Free (word indicator)	−0.03	−0.08	0.02
Million (word indicator)	**0.09**	**0.01**	**0.19**
Award(s) (word indicator)	**−0.17**	**−0.24**	**−0.10**

Figures in bold are significant on a 95% level.

Table 3
Results — variance component (γ).

	Mean	Coverage interval	
Intercept	**0.96**	**0.61**	**1.30**
Action (category indicator)	−0.01	−0.18	0.11
Adventure (category indicator)	−0.12	−0.31	0.12
Arcade (category indicator)	**−0.14**	**−0.34**	**−0.01**
Card (category indicator)	0.15	−0.06	0.36
Other games (category indicator)	0.15	−0.06	0.29
Puzzle (category indicator)	−0.09	−0.21	0.05
Racing (category indicator)	−0.11	−0.29	0.06
Shooting (category indicator)	**0.27**	**0.04**	**0.43**
Sports (category indicator)	0.06	−0.10	0.30
Strategy (category indicator)	**0.34**	**0.07**	**0.54**
Average rating	**0.18**	**0.05**	**0.31**
Number of ratings	**0.05**	**0.02**	**0.10**
Number of screen shots	−0.02	−0.05	0.02
Number of words	0.03	−0.03	0.10
Fun (word indicator)	−0.02	−0.10	0.06
Enjoy (word indicator)	0.01	−0.20	0.17
Play (word indicator)	0.01	−0.05	0.05
Paced (word indicator)	−0.01	−0.26	0.23
Best (word indicator)	−0.02	−0.13	0.07
Addictive (word indicator)	−0.04	−0.11	0.05
Realistic (word indicator)	0.09	−0.14	0.31
Free (word indicator)	**−0.15**	**−0.25**	**−0.02**
Million (word indicator)	−0.17	−0.35	0.04
Award(s) (word indicator)	**0.08**	**0.02**	**0.15**

Figures in bold are significant on a 95% level.

Second, we discuss the results of the variance component in the model (please see Table 3 for parameter estimates). Recall that we model the prior of the variance, $\sigma_i \sim gamma(c_i, c_i)$, that allows for over-dispersion as $\log(c_i) = z_i\gamma$. We use the same covariates, z_i, for the variance component that we have used for the mean component. We find that two secondary categories exhibit a significantly higher over-dispersion compared to the baseline (as captured by the intercept). These categories are Shooting (0.27) and Strategy (0.34). Only for the secondary category, Arcade, we find a negative effect of 0.14. We find that the average rating as well as the log number of ratings increases the over-dispersion parameter. In other words, mobile games with a higher rating and/or a higher number of ratings have a higher variation in repeat usage.

This indicates, potentially, that a higher number of ratings, even very good ones, might not be predictive of improved fit of the mobile game to an individual. For example, very popular mobile games do not need to match everybody's tastes. Our findings are in line with previous literature on the effects of mean ratings as well as the variance of ratings. Sun (2012) finds that a high average rating indicates a high quality product, whereas a high variance of ratings is associated with a niche product, one that some consumers love and others hate. The author finds that given its informational role, a higher variance would correspond to a higher subsequent demand if and only if the average rating is low. We note that Sun's (2012) findings relate to the initial purchase decision where uncertainty exists in terms of the fit of the product with the consumer's needs. In our scenario, these findings could be used to inform on the user's initial download decision but not to the repeat usage decision that we study. In the case of repeat usage of games, it is hard to know whether a game with a higher variance and everything else equal might be better or worse in creating engagement/value for the mobile game publisher.

Summing up, we find evidence for our proposed model in terms of in-sample fit. Our proposed model that allows for mobile game-level over-dispersion is preferable over a basic model and, more importantly, a model that restricts over-dispersion to be the same for all mobile games — repeat usage (or engagement) is a mobile-game level phenomenon and as such should be modeled on the mobile game-level. Our proposed model introduces an approach that allows doing so even for mobile games with relatively low numbers of observations. Substantively, we find that publicly available mobile game-level characteristics and categories are useful in understanding usage and engagement not only when it comes to understanding mean usage (often the focus of many marketing models) but also the variance of usage. In terms of mean usage, existing mobile game categories allow to correctly differentiate across games. Crowdsourced information like user ratings and number of reviews is valuable in understanding usage and engagement. While we have not leveraged the content of the reviews, it would be interesting to use the mined review content to see whether it would allow an even better fitting model. Similarly, some of the same information is useful when it comes to understanding the variance of usage. Compared to their effect on the mean of usage, mobile game categories seem less relevant when it comes to the variance of usage. Interestingly, for Strategy mobile games, we find a higher mean and a higher variance of usage. So while a Strategy mobile game seems to draw users in more and generate more engagement this engagement is more extremely distributed than for other mobile game categories, i.e., more low engagement and high engagement users compared to the other categories. As we hypothesized above, this could all be about fit. Typically, Strategy mobile games are more involved and require more time to "learn" them. This could mean that users, after download, have more extreme usage patterns as some users very quickly drop out but some users get drawn in and engage at a much higher level compared to

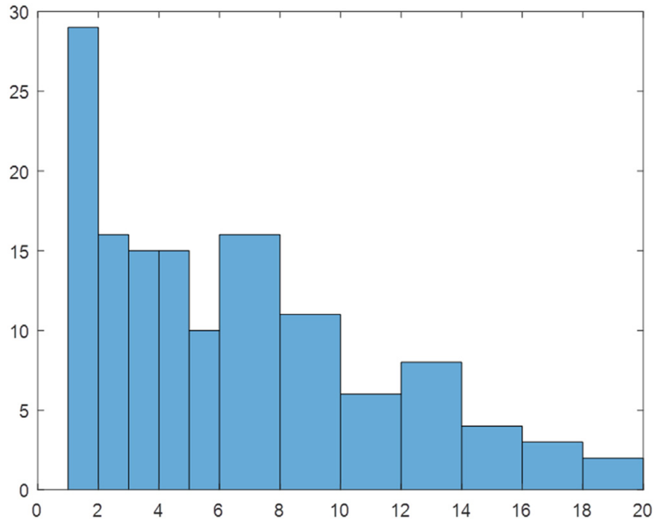

Fig. 6. Histogram of variance/mean for 193 game forecast.

other mobile game categories. This is an interesting topic for future research if more user-level data become available. Lastly, we find that while ratings drive mean usage ratings also drive the variance of usage. This poses an interesting managerial challenge in the mobile game monetization question. Typically, if higher ratings lead to more purchases (downloads), managers might not be as concerned with the downstream consequences of users not liking the product, i.e., the increase in the variance of engagement, if the product is monetized during purchase (or download). As we discuss above, this is not true anymore for most mobile games. So higher ratings might lead to not only more downloads and higher mean usage, but also to higher variance in usage making forecasting of usage to sell as in-game advertising inventory harder. It would be interesting to investigate this trade-off on a user-level to understand the potential mismatch between a user and a game based on the user's reading of the ratings and the download as well as engagement decisions.

6. Forecasting usage

Next, we aim to forecast usage to enable publishers to sell ad impressions through reservation contracts as “eyeball supply”. As discussed above, most games require a connection to a server today. This allows tracking of user behavior and the opportunity to create data similar to the data we use. While it is possible for games with large replay counts to forecast each user's behavior, in reality, this would be a very complex undertaking. More essential, it is not as important as to forecast who will see an ad but more important to forecast how many ads can be shown.

Our model is built at the mobile game-level but leveraging user data. As such, our model can forecast how many repeat usage occasions an additional user can generate. The publisher can observe how many new users download the mobile game, and, more importantly, use it at least one time.[12] We first provide a forecast for 100 additional users across all 193 games. We note that our model forecasts a much larger variance compared to the mean, correctly accounting for over-dispersion in the Poisson approach (see Fig. 6). We find that for only 28 games (roughly 14%), the mean and the variance are close, i.e., the standard Poisson approach would have been appropriate. For the majority of the games, the variance of the forecasted repeat usage counts is significantly larger than the mean. We detail some more properties of the forecast in Fig. 7. We can see that for most games, the minimum forecast is 1, which is in line with the fact that even in our sample of 193 popular mobile games, average gameplay is still low. However, a few highly played mobile games have minimum forecasts significantly above 1. Again, our model is flexible and allows accommodating small as well as large mobile games and generates forecasts that are in line with the mobile game's observed behavior.

We next show a more realistic scenario for a randomly chosen mobile game. We assume that the mobile game generates a random number of new sign-ups every day for 30 days. We generate random sign-ups with a mean of 50. For this mobile game, on average, a user generates 12 repeat usages observations. This means that after 30 days of sign-up, the manager would expect 1500 new users with an expected usage count of roughly 18,000 or, in other words, at least 18,000 occasion to show ads. We replicated our simulation 1000 times and find that while, on average, indeed roughly 1500 new users did sign up (1530) we found some significant noise as the forecast moves further out (see Fig. 8). On average, the manager can expect 18,832 usage occasions. On the lower end, the 2.5% confidence band, only 14,550 usage occasions realize. On the higher end, the 97.5% confidence band, 22,915 usage occasions realize.

12 There is a difference between downloading a mobile game and opening it for the first time. Our approach requires usage and not only the download.

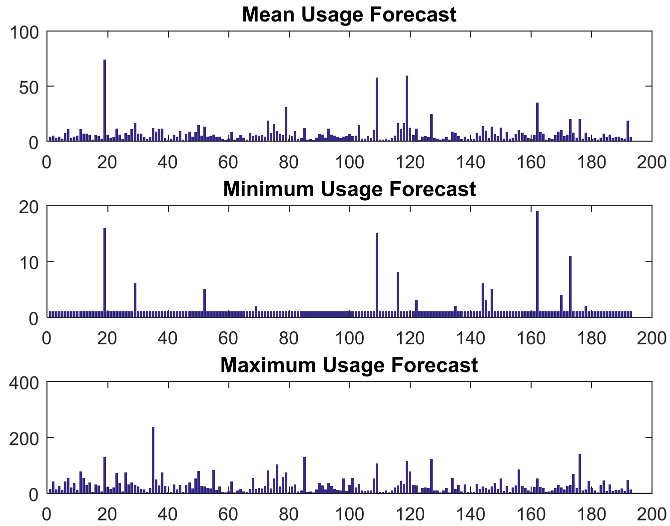

Fig. 7. Usage forecast details.

From a manager's perspective, the ability to build forecast scenarios and, in expectation, determine the overall number of usage occasions can be a useful resource. This helps determine the number of usage occasions, i.e., "eyeball supply", for which an ad can be displayed and hence, the overall value that can be derived from engagements with the mobile game via ads shown during gameplay.

7. Conclusion

Mobile apps are large and strongly growing space. With this increase in importance and strong growth predications, managers grapple with monetizing their mobile content and offerings. Marketing academia, so far, focused on the more traditional revenue models of the mobile app market and investigated issues related to the initial app purchase or in-app purchases. Next to these two monetization strategies, app publishers are increasingly leveraging in-app advertising as the Fremium model has effectively curtailed the ability to charge for the app purchase (download). As such, app publishers are trying to monetize the ongoing engagement a mobile app can generate by selling it to interested advertisers. In this paper, we focus on ongoing engagement with a mobile app

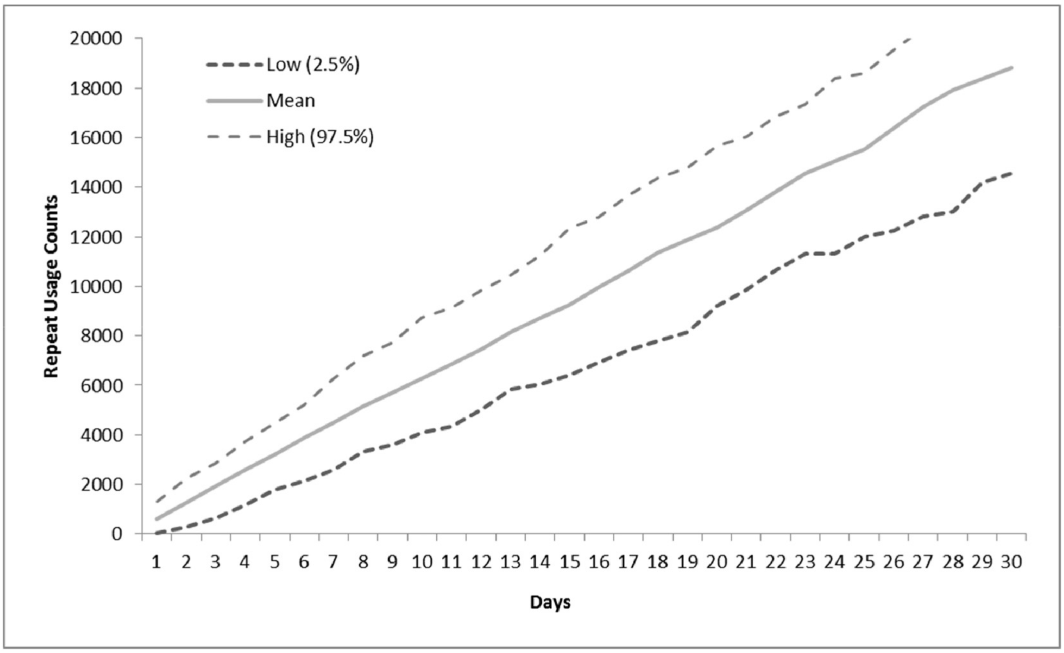

Fig. 8. Usage forecast interval.

beyond the initial purchase. To the best of our knowledge, no research as of yet has investigated the engagement generated by mobile apps and how to forecast this engagement. Our goal is to study engagement and detail what could be drivers of differences in engagement across apps. We aim to also forecast user engagement to allow publishers to better forecast their "eyeball" supply to enable selling advertising via "eyeball" contracts next to RTB. We use a novel data set that contains repeat mobile app usage counts across a large number of mobile apps over an eleven-month period. The data was collected by a collaborating market research firm using a large panel of consumers.

In our study, we focus on engagement in the mobile game space, a growing multi-billion dollar market. As most mobile games are now played online, publishers should have access to very similar usage data and be able to implement a model like ours to measure and forecast the engagement that their mobile games are generating even when the number of observations for individual games is small. We model mobile game-level engagement, measured as repeat usage occasions, in a framework that allows us to investigate the mean as well as the variance of usage on a mobile game-level. We find that the mean and the variance of usage are significantly different over 193 mobile games that we study. Using publicly available data from the app stores, we find that, as expected, average rating and number of ratings increases mean usage of a mobile game. Interestingly, we also find that average rating and number of ratings increase the variance of usage. Our approach accounts for mobile game characteristics as well as marketing tools, e.g., the number of screenshots used by the app developer on the mobile game's app store page. We also forecast mobile game usage based on new sign-ups that a publisher can observe.

At this stage of the industry development, very few games generate value at the point of sale (i.e., download). Per the website ThinkGaming, in January 2018 there was not a single paid game in the 50 top grossing games in the Apple App store. A major source of revenue are in-app purchases, e.g., *Candy Crush Saga*, generates $1,329,309 on average per day in revenue from in-app purchases.[13] The second major source is in-app advertising. Our paper focuses on monetizing user engagement via in-app advertising. We note that other studies have focused on purchase (download and in-app purchase). While in-app purchases are a topic of key interest, it is hard to imagine a dataset across apps that would allow an investigation into the effectiveness of in-app purchasing strategies.

In sum, we have investigated the engagement of users with mobile games leveraging novel user-level usage information. Our study focuses on the aggregate levels of usage and engagement as a first step to understand the monetization opportunities of apps that are significantly different from products traditionally studied in marketing that are either sold or have usage service contracts. Going forward, future research could investigate how consumers switch between apps types (i.e., game vs. social media) and allocate their time across apps or how certain types of usage occur at certain times of the day. Another interesting topic could be "wear-in" and "wear-out" of apps to detail how consumers use apps over the lifetime of the app. Last but not the least, studies into pricing models such as freemium are needed to understand when app developers should opt for traditional upfront models and when they should leverage the power of mobile devices to use a freemium or a tiered pricing strategy.

Acknowledgements

Marketing Science Institute (MSI) Grant # 41821, "Valuing Mobile App Engagement: The Case of Mobile Games" in 2015.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.ijresmar.2019.01.002.

References

AdWeek (2016). Infographic: 'whales' account for 70% of in-app purchase revenue. http://www.adweek.com/digital/infographic-whales-account-for-70-of-in-app-purchase-revenue/.
Andrews, M., Luo, X., Zhang, D., & Ghose, A. (2015). Mobile ad effectiveness: Hyper-contextual targeting with crowdedness. *Marketing Science*, *35*(2), 218–233.
AppAnnie (2017). Spotlight on consumer app usage. http://files.appannie.com.s3.amazonaws.com/reports/1705_Report_Consumer_App_Usage_EN.pdf.
AppsFlyer (2016). The state of in-app spending. http://hub.appsflyer.com/hubfs/IAP_Guide/The_State_of_In-app_Spending_AppsFlyer.pdf.
Apptentive (2017). How many mobile apps are actually used? https://www.apptentive.com/blog/2017/06/22/how-many-mobile-apps-are-actually-used.
Böhmer, M., Hecht, B., Schöning, J., Krüger, A., & Bauer, G. (2011). Falling asleep with Angry Birds, Facebook and Kindle: A large scale study on mobile application usage. *Proceedings of the 13th international conference on human computer interaction with mobile devices and services* (pp. 47–56). ACM.
Chib, S., & Winkelmann, R. (2001). Markov chain Monte Carlo analysis of correlated count data. *Journal of Business & Economic Statistics*, *19*(4), 428–435.
Christiansen, C., & Morris, C. (1997). Hierarchical Poisson regression modeling. *Journal of the American Statistical Association*, *92*(438), 618–632.
Comscore (2017). The 2017 US mobile app report. https://www.comscore.com/Insights/Presentations-and-Whitepapers/2017/The-2017-US-Mobile-App-Report.
Datta, H., Foubert, B., & Van Herde, H. J. (2015). The challenge of retaining customers acquired with free trails. *Journal of Marketing Research*, *52*(2), 217–234.
Dew, R., & Ansari, A. (2018). Bayesian nonparametric customer base analysis with model-based visualizations. *Working Paper*. Columbia University.
Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R., & Estrin, D. (2010). Diversity in smartphone usage. *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services* (pp. 179–194). MobiSys.
Ferreira, D., Goncalves, J., Kostakos, V., Barkhuus, L., & Dey, A. K. (2014). Contextual experience sampling of mobile application micro-usage. *MobileHCI* (pp. 91–100).
Fisher, L., & Liu, C. (2016). US programmatic ad spending forecast. *eMarketer Report, September*.
Ghose, A., & Han, S. (2014). Estimating demand for mobile apps in the new economy. *Management Science*, *60*(6), 1470–1488.
Gopalakrishnan, A., Iyengar, R., & Meyer, R. (2015). Consumer dynamic usage allocation and learning under multipart tariffs. *Marketing Science*, *34*(1), 116–133.
Gouin-Vallerand, C., & Mezghani, N. (2014). An analysis of the transitions between mobile application usages based on Markov chains. *UbiComp Adjunct* (pp. 373–378).
Iyengar, R., Jedidi, K., Esseghaier, S., & Danaher, P. (2011). The impact of tariff structure on customer retention, usage and profitability of access services. *Marketing Science*, *30*(5), 820–836.

[13] https://thinkgaming.com/app-sales-data/2/candy-crush-saga/.

Kim, G., Shin, B., & Lee, H. G. (2009). Understanding dynamics between initial trust and usage intentions of mobile banking. *Information Systems Journal*, *19*(3), 283–311.
Lambrecht, A., Goldfarb, A., Bonatti, A., Ghose, A., Goldstein, D., Lewis, R., ... Yao, S. (2014). How do firms make money online. *Marketing Letters*, *25*(3), 331–341.
Lambrecht, A., Seim, K., & Skiera, B. (2007). Does uncertainty matter? Consumer behavior under three-part tariffs. *Marketing Science*, *26*(5), 698–710.
Liao, Z. X., Pan, Y. C., Peng, W. C., & Lei, P. R. (2013). On mining mobile apps usage behavior for predicting apps usage in smartphones. *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management* (pp. 609–618).
Liu, C. (2016, November). US ad spending: eMarketer's updated estimates and forecast for 2015–2020. *Industry report*.
Localytics (2017). App user retention improves in the U.S., but declines internationally. http://info.localytics.com/blog/app-user-retention-improves-in-the-us.
Newton, M. A., & Raftery, A. E. (1994). Approximate bayesian inference with the weighted likelihood bootstrap. *J. R. Stat. Soc. Ser. B Methodol.*, *1*(1), 3–48.
Newzoo (2017). The global games market will reach 4108.9B in 2017 with mobile taking 42%. https://newzoo.com/insights/articles/the-global-games-market-will-reach-108-9-billion-in-2017-with-mobile-taking-42/.
Sayedi, A. (2018). Real-time bidding in online display advertising. *Marketing Science*, *37*(4), 507–684.
Shin, C., Hong, J. H., & Dey, A. K. (2012). Understanding and prediction of mobile application usage for smart phones. *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (pp. 173–182). *ACM*.
Statista (2017a). Number of app downloads per month of smartphone users in the United States as of June 2017. https://www.statista.com/statistics/325926/monthly-app-downloads-of-us-smartphone-users/.
Statista (2017b). Number of mobile app downloads worldwide in 2016, 2017 and 2021 (in billions). https://www.statista.com/statistics/271644/worldwide-free-and-paid-mobile-app-store-downloads/.
Statista (2018). Number of smartphone users worldwide from 2014 to 2020 (in billions). https://www.statista.com/statistics/330695/number-of-smartphone-users-worldwide/.
Sun, M. (2012). How does the variance of product ratings matter? *Management Science*, *58*(4), 696–707.