

iPhone's Digital Marketplace: Characterizing the Big Spenders

Farshad Kooti
Facebook
Menlo Park, USA
farshadkt@fb.com

Mihajlo Grbovic
Airbnb
San Francisco, USA
tua95067@temple.edu

Luca Maria Aiello
Bell Labs
Cambridge, UK
luca.aiello@nokia-bell-labs.com

Eric Bax
Yahoo Research
Playa Vista, USA
ebax@yahoo-inc.com

Kristina Lerman
USC Information Sciences
Institute
Marina del Rey, USA
lerman@isi.edu

ABSTRACT

Mobile devices have become ubiquitous, and people are spending increasing amounts of money to purchase digital goods on their phones. Despite the surging popularity, this phenomenon has not been studied at the large-scale. In this paper we analyze a large data set consisting of more than 776M digital purchases made on Apple mobile devices that include songs, apps, and in-app purchases. We found that 61% of all the spending is on in-app purchases and that the top 1% users are responsible for 59% of all the spending. We characterize these big spenders, finding that they are more likely to be male, older, and less likely to be from the US. We study how they adopt and abandon a specific app, and find that, after an initial phase of increasing daily spending, users gradually lose interest: the delay between their purchases increases and the spending decreases with a sharp drop towards the end. Finally, we model the in-app purchasing behavior in multiple steps: 1) we model the time between purchases; 2) we train a classifier to predict whether the user will make a purchase from a new app or continue purchasing from the existing app; and 3) based on the outcome of the previous step, we attempt to predict the exact app, new or existing, from which the next purchase will come. The results bring important insights and conclusions about spending habits in the mobile digital marketplace.

Categories and Subject Descriptors

H.4.3 [Information Systems]: Information systems applications

Keywords

Phone purchases, demographics, prediction

1. INTRODUCTION

Consumer spending is an essential component of the economy, accounting for 71% of the total US gross domestic product (GDP) in 2013 [5]. In spite of representing just over 10% of the entire consumer spending [3], online shopping is rapidly growing as people are becoming more comfortable with the online payment systems, security and delivery of the purchased goods. Over the last three years, online sales grew over 45% [11] and are showing signs of exponential growth.

One of the largest and fastest-growing online markets is Apple's iOS market, where people can make digital purchases from different categories. Apple's revenue from digital purchases surpassed \$20 billion in 2015 [4] and more than doubled in just three years [1, 2].

Mobile digital markets offer a wide variety of digital goods, including apps, songs, movies and digital books. They also include items that users can buy within an app, called *in-app purchases*, such as virtual currencies, bonuses, extra game lives and levels, etc. Despite the popularity of the iOS market, there has not been a large-scale study characterizing user's spending on different types of content and apps. For example, it is not known how much is spent on in-app purchases compared to songs or how disproportionate the spending is. Learning how people spend their money in this context has direct practical implications on the business of several stakeholders, including app developers and managers of online app stores, but it also has important theoretical implications for the understanding of the consumer behavior in an emerging market whose dynamics are still poorly known.

We study a longitudinal dataset extracted from hundreds of millions of email receipts for digital purchases on iPhones, iPads, and other iOS devices (which we refer to as "iPhone purchases" for brevity). Besides its large scale, our data has two unprecedented advantages over other data collected in the past. First, besides recording the purchase history, it also includes *demographic information* such as people's age, gender, and country of residence. This enables us to not only characterize the consumer population, but also study how spending differs by income and location. For example, we find that the average spending is not correlated with income but strongly depends on the country of residence. And second, in contrast to the application-centric view of previous works, our data is user-centric and allows for the observation of user spending behavior across *multiple apps*. This allows us to study how users abandon

an app and start using a new one, and to compare the behavior of users who make purchases from a single app with users who make purchases from multiple apps simultaneously.

We analyze how users spend money across different categories and show that a small fraction of users are responsible for the majority of spending. We call these users *big spenders*. Moreover, our analysis informs predictive models of digital spending behavior. We first show that the time between purchases is best described by a Pareto distribution. We then build a supervised classifier that accurately predicts if a user will make a purchase from a new app, or an existing app that he/she previously consumed. Finally, based on the outcome of the previous step we predict the exact app from which the user will purchase.

In summary, our main findings and contributions are as follows:

- In-app purchases account for 61% of all money spent on digital purchases on iPhones, followed by songs (23%) and app purchases (7%).
- The spending is highly heterogeneous: the top 1% of spenders account for 59% of all money spent on in-app purchases.
- Big spenders tend to be 3-8 years older, 23% more likely to be a man, and 31% less likely to be from the US, compared to the typical spender. Interestingly, income is comparable between the big spenders and the rest of the users.
- Big spenders become slower to repurchase from an app as time passes, but their rate of spending within an app initially increases, then decreases.
- From the perspective of app developers and ad networks big spenders are the most valuable targeting segment. Even if they abandon an app they are frequently buying from, they are 4.5x more likely to be a big spender in a new app compared to a random user.
- We model the entire purchasing process in several steps: modeling time between purchases, predicting purchase from a new app or an app from which the user has already purchased, and predicting the exact app of an in-app purchase.

Both consumers and producers in the mobile app market might benefit from this study. Our results can inform the deployment of better app recommendation systems that can lead people to download apps they are more likely to enjoy, thus creating higher revenue opportunities for app developers.

2. DATASET AND MARKETPLACE

Shortly after each digital purchase on an iPhone (or any other iOS device), the user receives a confirmation email with details of the purchase. This email contains information about the purchase, including the amount of money spent and the type of purchase. The email has a specific format, making it easy to parse automatically. We obtained information about digital purchases of Yahoo Mail users, using an automated pipeline that hashes the names and IDs of users to preserve anonymity. All the analyses have been performed in aggregate and on non-personally unidentifiable data. We gathered data covering 15 months, from March 2014 to June 2015. Our dataset includes 26M users who together made more than 776M purchases totaling \$4.6B.

There are six main categories of iPhone purchases: applications (apps), songs, movies, TV shows, books, and in-app purchases (purchases within an app, e.g., bonuses or coins in games). These

categories differ vastly in the numbers of purchasers: 16M people purchased at least one song, but only 671K people purchased a TV show. The number of purchases in each category varies greatly as well: there are 430M song and 255M in-app purchases, while movies, books, and TV shows have fewer than 40M purchases all together. The total money spent in each category varies even more: in-app purchases account for \$2.8B, or 61%, of all money spent on digital purchases; 23% of the money is spent on songs, 7% on app purchases (purchasing apps themselves, not the purchases within apps), 6% on movies, 2% on books, and only 0.7% on TV shows. Even though there are considerably fewer total in-app purchases compared to songs (60% fewer), the money spent on in-app purchases is 2.7 times more than on songs, showing that on average in-app purchase is much more expensive than a song purchase. Figure 1 shows the number of users, purchases, and the money spent on each of these six categories.

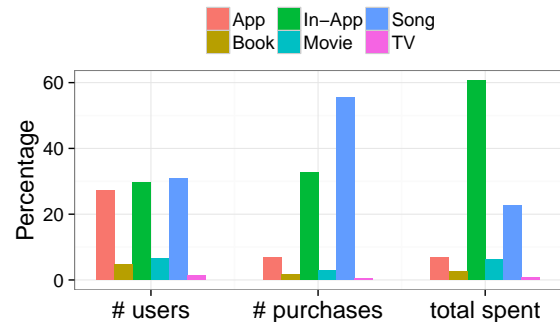



Figure 1: Percentage of users, purchases, and money spent on each category.

Our dataset also includes user age, gender, and zip code as provided by the users at the time of sign up. Spending varies significantly based on demographics. Figure 2(a) shows the cumulative distribution function (CDF) of the spending for men and women. Men spend more money on purchases than women; the median spending for women is \$31.1, and for men it is \$36.2, which is 17% higher. Age also affects spending: Figure 2(b) shows that the peak age for iPhone spender is the mid-30s, and after that the spending level decrease quickly. For US residents, we use the median income of the zip code they declared as their residence as an estimate of their income. This estimate has shown meaningful trends in earlier work [18]. Surprisingly, income only affects the spending of people with less than \$40K annual income (Figure 3). This is in contrast with online shopping, where users with higher income tend to spend more money shopping online [18].



(a) Spending and gender (b) spending and age
Figure 2: Effect of gender and age on iPhone purchases.

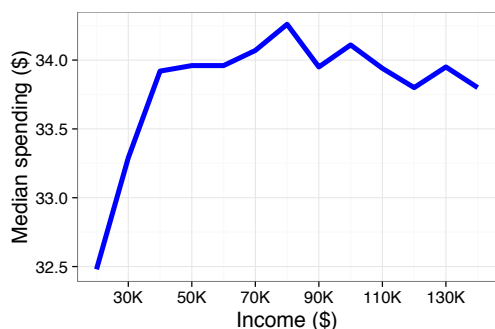


Figure 3: Effect of income on spending. There are more than 10k users for each income category.

Table 1: Top 10 apps by in-app earnings, with demographics.

App Name	Earnings	# of Purchases	Avg. Age	% Women
Clash of Clans	\$356.5M	13.6M	33.0	29%
Candy Crush Saga	\$168.0M	45.3M	40.3	70%
Game of War	\$159.8M	1.9M	35.0	25%
Boom Beach	\$60.9M	1.7M	34.4	18%
Hay Day	\$52.0M	3.2M	36.8	67%
Farm Heroes Saga	\$40.4M	6.8M	42.3	81%
Candy Crush Soda	\$37.5M	8.8M	41.0	77%
Big Fish Casino	\$32.9M	1.1M	44.5	57%
DoubleDown Casino	\$27.2M	1.1M	49.4	65%
Pandora Radio	\$24.3M	5.7M	37.8	54%

Spending on iPhones varies considerably by geography. Figure 4(a) shows that European countries, especially Scandinavian countries, have the highest spending per Yahoo mail user. This higher spending could be explained by higher income and fewer people making app purchases. Canada, Mexico, and Australia have higher spending per person than the US, while most African and Asian countries have lower levels of spending.

There are also more than 154K applications in our dataset with at least one user purchase or in-app purchase. As shown above, the earnings from in-app purchases are considerably higher than from app purchases themselves (almost 9 times higher). Table 1 shows the top 10 apps by in-app earnings, along with their numbers of purchases, average purchaser age, and percentage of female purchasers. We make some observations about these data. First, there are considerable differences among the earnings of top apps. Second, the average price of in-app purchases varies widely across different apps: while there are more than 3.3 times as many purchases in *Candy Crush* compared to *Clash of Clans*, the earnings for *Clash of Clans* is 2.1 times higher than for *Candy Crush*. Third, there is a significant difference in the demographics of purchasers of different apps. The average age of buyers is 33 years for *Clash of Clans* and 49.4 years old for *DoubleDown Casino* and 81% of the users making purchases from *Farm Heroes Saga* are women, compared to only 18% for *Boom Beach*. Other apps, such as *Pandora Radio* and *Netflix* (not shown in the table), have a balanced audience. Knowing the audience of an application could be useful for both the advertisers and game designers. Advertisers can target the particular population and the game designers can make changes to their apps to make them more appealing to the audience that is not currently engaged.

We also collected the category information for each application, e.g., puzzle game or travel, from Apple’s iTunes. Then, for each

Table 2: Top 5 gender-biased categories.

Top Categories for Men	% Men	Top Categories for Women	% Women
Sport magazines	84.6%	Board games	70.5%
Sports	74.9%	Word games	64.9%
Racing games	69.9%	Puzzle games	63.6%
Navigation	68.1%	Family games	62.1%
Sports games	67.4%	Educational games	61.5%

Table 3: Top 5 age-biased categories.

Categories for Youth	Avg. Age	Categories for Older Users	Avg. Age
Photo & Video	32.2	Food & Drink	47.7
Strategy games	33.2	Weather	45.4
Racing games	33.6	Travel magazines	45.2
Trivia games	33.6	Board games	44.7
Social networking	34.1	Business	42.9

category, we calculated the percentage of people who purchased an app or made an in-app purchase from an app belonging to that category. The top five categories by gender are shown in Table 2. Men are more likely to make a purchase from apps relating to sports, and women prefer games, especially brain games. Similarly, we found the top 5 categories with the youngest and oldest average age. Younger buyers are interested in photo and video apps, racing games, and social networking applications, while older buyers are interested in more general applications for food, weather, business, and travel (Table 3).

Limitations. Our dataset only includes purchases from people who are Yahoo Mail users, there may be a selection bias in the subset of users being studied. While this might occur to some extent, but given the popularity of Yahoo Mail (over 300M users¹), we believe our data includes a somewhat representative sample and the findings can be generalized to other users. Moreover, the email receipts are sent within a day after the purchase, so our dataset does not include the exact time of the purchase and we have to conduct our analysis in granularity of a day.

3. BIG SPENDERS

In this section we focus on the category of in-app purchases since it is the largest spending category in the digital marketplace. We first show that a small number of users are responsible for the majority of spending. Then, we characterize these users demographically, by age, gender, country of residence, and income. Finally, we study how these buyers discover a specific app, start spending money within it, and how they stop making purchases within it as their interest in that app diminishes.

In-app spending patterns vary significantly across different users. In Figure 5 we show the PDF and CDF of spending on in-app purchases. It demonstrates that the spending has a heavy-tailed distribution, with most users spending \$20 or less, and a small minority (2.4% of users) spending more than \$1000 over the studied 15 months.

To better demonstrate the disparities in spending, we plot the Lorenz curve, which shows the percentage of the total spending by different percentiles of the population when ordered by spending (Figure 6). The diagonal line represents a perfect equality of spending (i.e., when each person spends the same amount). The

¹www.comscore.com

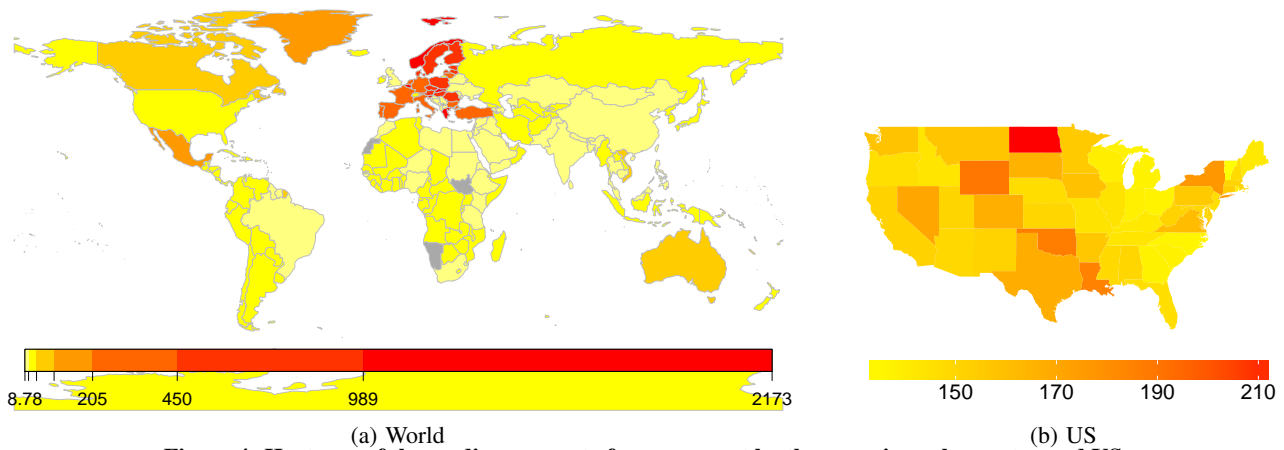


Figure 4: Heatmap of the median amount of money spent by the users in each country and US.




Figure 5: PDF and CDF of user's spending on in-app purchases.

larger the distance from the diagonal, the larger the inequality in spending. The figure shows very high inequality: the bottom half of buyers spend less than 2% of the total amount of money, while the top 10% are responsible for 84% of all spending. In fact, *just the top 1% is responsible for 59% of all the money spent on in-app purchases*. This inequality can be captured by the Gini coefficient that summarizes the distance from equality in a single number, which turns out to be 0.884, representing an extremely high inequality. Interestingly, if we consider the earnings of the apps, the inequality is even higher, with Gini coefficient 0.989, and 0.1% of the apps earn 71% of all the in-app purchase income (Figure 7). As a comparison, the Gini coefficient for the income of the US population is 0.469, which is the highest among Western industrialized nations (according to census data).

As mentioned above, the 1% of buyers, that represent 154K users, are responsible for the majority of in-app purchases. In the rest of this section, we focus on this set of users, and since most of the apps with high earnings were games, we call this set of users *big spenders*. We also calculated the top 1% of spenders in each month separately. Among the big spenders, who are in the top 1% of spenders over the entire 15 month period, 68.4% are in the top 1% of spenders for half of these months or fewer. This shows that there are bursts and pauses in individual spending levels for each big spender.

3.1 Characteristics of big spenders

We start by comparing the demographics of big spenders with the rest of the users in our dataset. Understanding the differences in demographics could be useful for advertisers and app stores to




Figure 6: Lorenz curve of the spending of the users on in-app purchases, showing high disparity among users.




Figure 7: Lorenz curve of the earning of the apps, showing extremely high inequality in the earning of the apps.

better target the population that is more likely to be a big spender.

Big spenders are 23% relatively more likely to be men (59% men vs. 48% women). Regardless of gender, big spenders tend to be older. Men who are big spenders have a median age of 37 years, while the median age in our dataset is 34 years. The difference

is even larger among women: 43 vs. 35 years. Moreover, there are considerable differences in the country of residence statistics for big spenders compared to the typical user. For some countries, like the US, a random user is less likely than average to be a big spender, but for the other countries users are much more likely to be big spenders. For example, Greek, Turkish, and Romanian users are respectively 50, 33, and 29 times more likely to be big spenders than users from our study population at large.

We also consider the role of income for the people from the US, by calculating the fraction of people with a given income who are big spender. Figure 8 shows that income has a very small effect on users being big spenders, except for the users with less than \$20K or more than \$140K annual income. Note that the percentage is almost always smaller than the expected 1%, which is how we define big spenders, because this analysis is conducted only on users from the US, and big spenders are less prevalent in the US.




Figure 8: Fraction of big spenders, given the income of the users.

3.2 App adoption and abandonment

To better understand the behavior of big spenders, we focus on how they start making in-app purchases within apps, and how they abandon them. In order to analyze the behavior of big spenders over a long period of time within apps they use frequently, we only considered pairs of big spenders and apps that had more than 50 in-app purchases. Furthermore, we were interested in the entire time span of the user’s app usage, i.e. from the first time they make a purchase to the last time they make a purchase, so we filtered out the cases in which the usage started before or ended after the period of data collection. This was done by considering only the (user, app) pairs for which the first purchase happened after the first month in our dataset, and the last purchase was before the last month.

We start our analysis by looking at the time delay between consecutive purchases. Because our data has one-day granularity, we count multiple purchases by a user in one day as a single, more expensive, purchase. To account for the large heterogeneity in time delays between purchases by different users, we normalize the values for each user individually. Figure 9(a) shows the 2nd to 9th delays, normalized by the first delay. On average, both the time between purchases and the spending per purchase increase (Figure 9(b)). This means that even though people make fewer in-app purchases, they spend more money after couple of transactions. Since a considerable fraction of the spendings are for virtual game coins and bonuses, this suggests that users start by buying small packages of coins and bonuses and move on to the larger ones as they progress in the game. Also, as they buy more and more bonuses and/or coins per purchase, it takes longer for them to re-




Figure 9: Normalized change in time between purchases (i.e. delay) and spending in the first 10 purchases from an app, showing that the delay between consecutive purchases increases, while more money is spent on each purchase.

plenish their supplies by making a follow-up purchase.

Similarly, when focusing on the user’s last 10 purchases within an app, we find that users’ delays still get longer, but now at a much higher rate. The very last delay is six times longer than the first delay, on average (Figure 10(a)). This long delay is a strong indicator of app abandonment. Finally, in Figure 10(b) we show that as users get closer to their last purchase, they start spending less and less money on their daily purchases.

Switching to other apps. Next, we investigate what big spenders do after they abandon an app. More precisely, we wanted to find out what fraction of those users switch to another app in which they again become a frequent buyer. We conduct this analysis on the same data as above, i.e., user-app pairs that have more than 50 purchases. We find that 8.6% of big spenders who stop making purchases from an app will start making purchases from another app and will become a big spender in the new app (i.e., making more than 50 purchases). This number may seem small, but consider that the big spenders who abandon an app are 2.1x more likely to be a big spender in another app when compared to a random user from our entire dataset (because only 4.1% of users make 50 or more purchases from at least one app). Consequently, from a marketing perspective, it makes much more sense to advertise the new apps to the big spenders of existing apps. Furthermore, if we consider a more restrictive definition for big spender and only examine user-app pairs with is more than 100 purchases, the difference becomes even larger, and the big spenders who have abandoned an app are 4.5x more likely to become a big spender in another app.




Figure 10: Normalized change in delay and spending in the last 10 days with a purchases from an app, showing that the delays become significantly longer and less money is spent on each purchase.

4. PURCHASE MODEL

Table 4: AIC for different distributions. Lower AIC scores are preferred.

Distribution	AIC
Pareto	59.55M
Log Normal	60.86M
Weibull	61.87M
Gamma	62.22M

In this section, we model the sequence of purchases people make in order to understand the purchasing behavior better. The insights into the future purchasing behavior could be used by both the gaming companies and the app store to increase user engagement and provide better app recommendation.

Following the prior work on user consumption sequences [10, 8] we model the in-app purchases in 3 main steps: 1) modeling time between purchases, 2) predicting whether the next in-app purchase will come from an app that the user already purchased in the past or a completely new app, and finally 3) predicting the exact app that the user will purchase from, given the output of the previous step. The output of each step is used in the next step; the estimated time interval is one of the main indicators for predicting if the user will purchase from a new app, and we can predict the next app to be consumed much more accurately if we know whether the app is a new app for that user or not.

4.1 Temporal model

First, we investigate a set of parametrized distributions to see which one best describes the distribution of inter-purchase times. We considered Weibull, Gamma, Log normal, and Pareto and find that Pareto best fits the data. We used the Akaike Information Criterion (AIC) [7] and the P-P and Q-Q plots, such as the ones shown in Figure 11 for Pareto, to compare different distributions. The AIC values were fairly close for all distributions as shown in Table 4, but the plots showed that the Pareto distribution with $shape = 3.21$ and $scale = 20.17$ matches the data better than other distributions. Figures 11(a), 11(b) show that the modeled distribution fits the probability density function and cumulative density function very well. There are some deviances in the Q-Q plot (Fig 11(c)), where the empirical and theoretical quantiles are matched, that shows distribution failing to capture very large values in the data. However, the Pareto distribution still is the best fit for our data distribution, considering both the AIC and the density distributions.

4.2 Novelty prediction

Next, we predict whether the user will purchase from a new app or from an app he or she has purchased in the past. We approach the problem as supervised learning at the time that the user will make the purchase and use the following features: age, gender, time since previous purchase, average time between purchases, average time between re-purchases, total number of purchases, day of the current purchase, percentage purchases from new vs. past apps, whether last three purchases are from a new app, and the number of apps from which the user has purchased in the past. We use the first year of our dataset for training and the last three months for testing, such that we do not use any future information in our predictions.

We tested a collection of different classification algorithms, including several types of decision tree algorithms and SVMs. Finally, the C5.0 algorithm in R achieves the best result [20]. Our classifier achieves a high accuracy, predicting the right class in 84.5% of cases with precision of 0.862, recall of 0.965, and F-score of 0.964. This accuracy is slightly higher than the result reported in [10] for a similar problem on music and video re-consumption.

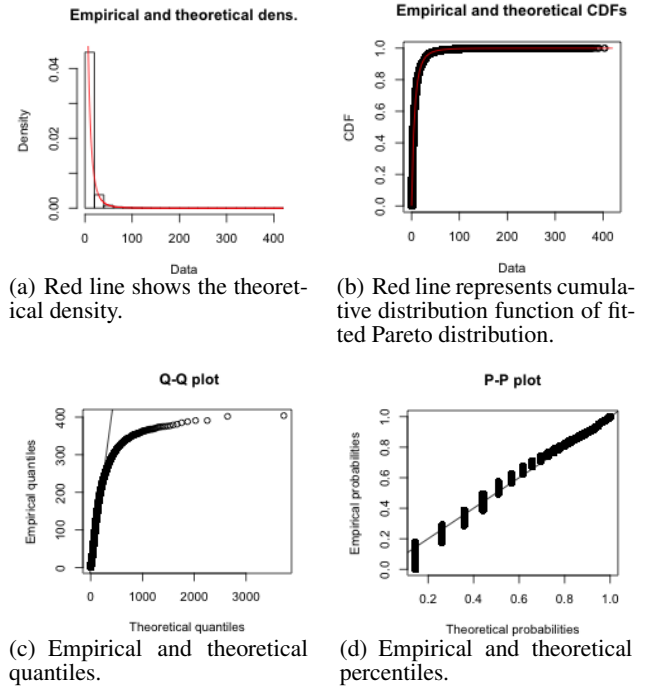


Figure 11: Results of fitting the time between purchases to a Pareto distribution.

Table 5: Results of logistic regression on the independent variables for abandonment prediction. * p -value < 0.001.**

Variable	Coeff.
% of re-purchase	6.236e+00***
Previous class (re-purchase)	2.878e-01***
2nd to the last class (re-purchase)	1.624e-01***
3rd to the last class (re-purchase)	7.878e-02***
Gender (m)	7.375e-02***
Mean inter-purchase time	4.764e-02***
Time since last purchase	-2.782e-02***
Total number of re-purchases	2.232e-02***
Day of the purchase	1.236e-03***
Age	1.069e-03***

To better understand the importance of each feature, we also fit a Logistic Regression model to our dataset after removing the correlated features. Figure 12 shows the pairwise correlation coefficient between the features. We removed one of the features from pairs with correlation coefficient higher than 0.7. Table 5 shows the result of the Logistic Regression: the percentage of re-purchases that the user has made is the most important feature; it captures the tendency of the user to re-purchase from an app. The three next most important features capture user’s recent history of re-purchases and purchases from new apps. These are followed by gender, with a positive correlation, showing men are more likely to have a repurchase. This is in tune with our earlier findings that men are more likely to be big spenders, and big spenders make many purchases from the same app.

4.3 App prediction

In two previous steps, we modeled the time between purchases and whether the user’s next purchase will come from a new app, with no previous purchases by that user, or from an existing app, with past purchases by that user. If the outcome of the model in-




Figure 12: Pairwise correlation coefficient among the features for predicting the purchase from new apps. Crossed cells do not have a statistically significant coefficient.

indicates that the purchase will come from a new app, then our task is to predict the most likely new app, given the previous apps that user made purchases from. If, on the other hand, the outcome of the classifier indicates that the purchase is from an existing app, then we use the sequence of all previous purchases to select the most likely existing app.

4.3.1 New app prediction

Given the apps that user purchased from in the past our goal is to predict the most likely new app the user will purchase from next. Similar problems have been studied extensively in the area of the recommendation systems [6, 13].

Motivated by recent success of embedding models in a number of natural language processing tasks [19], we propose to use a language model to learn app vectors in a low-dimensional space, trained from sequences of user in-app purchases such that apps that appeared in similar context reside nearby in the embedding space. Following the embedding step, we propose to use a k -nearest neighbor approach in the learned vector space to predict the most likely new app given the existing apps consumed by the user.

More formally, let us assume we are given a set of apps $A = \{a_j | j = 1 \dots M\}$, each identified by a unique identifier a_j . In addition, in-app purchase times for N users over a time period T from our dataset are also known. For the n^{th} user we collect data in a form $d_n = \{(a_j, t_i), i = 1, \dots, K_i, t_1 < t_2 < \dots < t_{K_i}\}$, where d_n denotes the user’s in-app purchase sequence, K_i is the total number of in-app purchases user made, and t_i is time of i^{th} in-app purchase from app a_j .

Given a set \mathcal{D} of N user in-app purchase sequences, where sequence $d_n \in \mathcal{D}$ is defined as an in-app purchase from K apps, the objective is to maximize log-likelihood of the training data \mathcal{D} ,

$$\mathcal{L} = \frac{1}{N} \sum_{d \in \mathcal{D}} \left(\sum_{a_j \in d} \sum_{-b \leq i \leq b, i \neq 0} \log \mathbb{P}(a_{j+i} | a_j) \right), \quad (1)$$

where b is the context widths for in-app purchase sequences and probability $\mathbb{P}(a_{j+i} | a_j)$ of observing a neighboring in-app purchase given the current in-app purchase is defined using a softmax function [19] expressed using app vectors.

Once we learn a vector representation for each app, we can lever-

Table 6: Top 5 closest apps by cosine similarity for 3 apps.

Kim Kardashian West Official	
Top 5 closest apps	Cosine similarity
Khloe Kardashian Official	0.907
Kourtney Kardashian Official	0.866
Kylie Jenner Official	0.863
Kendall Jenner Official	0.805
kimoji	0.733
Homework	
Top 5 closest apps	Cosine similarity
Smart Studies	0.502
iStudy Pro	0.500
Barrons Hot Words	0.491
Physics 101	0.480
PSAT Preliminary SAT Test Prep	0.479
Checkbook Pro	
Top 5 closest apps	Cosine similarity
Accounts 2 Checkbook	0.678
Checkbook Spending	0.657
Checkbook HD Personal Finance	0.641
My Check Register	0.617
My Checkbook	0.608

age vector cosine similarities to calculate similarities between apps. In Table 6 we show 5-nearest neighbor apps for several randomly picked apps along with the cosine similarities between the corresponding app vectors. As demonstrated in the table, the proposed approach can accurately capture similarities between apps.

We used the first year from our dataset to train the app embeddings and leveraged them to predict the new app user will purchase from in the remaining three months. The prediction was done using cosine similarity between the apps that user has already purchased from and all the remaining apps. All predictions were made per user. Specifically, if the user made purchases from k apps in the first year, we used these k apps to predict the new $k/4$ apps user will likely purchase from in the next three months ($k/4$ was chosen because the test period is one fourth of the training period). The $k/4$ apps were predicted by considering the most similar app to each of the k previous apps and picking one fourth of them randomly.

We predict the exact new app that the user is going to make a purchase from in 4.7% of purchases. This seems very low, but considering that there are more than 216K apps the user can choose from, in the context of the recommendation systems the approach is working considerably well. To further quantify the accuracy of the proposed embedding approach, we compared our recommendation strategy to several baseline models: 1) Non-negative matrix factorization (NMF) approach trained using the matrix of in-app purchases formed from \mathcal{D} ; 2) LDA [12] applied on the app description text; 3) ranking the apps by popularity and always predicting top $k/4$ apps users still have not made a purchase from.

The results, presented in Table 7, show that the app embedding approach outperforms the considered baselines. Poor performance of LDA agrees with previous research [16] that found that this method perform poorly when trained on short text documents. Top apps also achieves a low accuracy, due to users experimenting with many less popular apps. Better performance of app embeddings over the NFM approach can be explained by the fact that NFM model loses the notion of time and sequence order once it transforms the dataset \mathcal{D} into the matrix. We also experimented with other ways of selecting the most similar apps, such as considering larger and smaller number of candidates, and in all cases our

Table 7: New app prediction accuracy.

Method	Accuracy
App embeddings	4.7%
NMF	4.1%
Top apps	2.2%
LDA	1.7%

method outperformed all three other baselines with very similar margins to the reported one.

4.3.2 Existing app in-app re-purchase

In the case the model from Section 4.2 predicts that a re-purchase is most likely to happen, we use the frequency and recency of user’s previous app consumption to predict the app from which the re-purchase will occur. This may appear to be an easy prediction problem, as one might think that users almost always purchase from the last app they purchased from, or the app from which they made the majority of purchases. However, in case of in-app re-purchases, only 46.5% of the time the re-purchase comes from the latest app they purchased from, and only 45.3% of them come from the app from which the user made most of the purchases. This justifies the need for a more involved re-purchase model.

We follow a similar approach to the one from [10, 8] and use both recency and popularity of previous apps to predict from which app the user’s next in-app purchase will come from. We use a weight function and a time function that maps the frequency of the usage and time since previous usage to the learnt values. This repeat consumption model could be used for the i th consumption:

$$\mathbb{P}(x_i = e) = \frac{\sum_{j < i} I(x_j = e) s(x_j) T(t_i - t_j)}{\sum_{j < i} s(x_j) T(t_i - t_j)} \quad (2)$$

In this equation, function s represents the frequency of the purchase from the app, and function T represents the time between the purchases. These functions are optimized jointly by calculating the negative log-likelihood over the equation. The negative log-likelihood is not convex in s and T , but is convex in each function when the other one is fixed. Thus, we use a standard gradient descent to maximize the likelihood with respect to s and T , separately. After learning the weight functions, we are able to predict the correct app from which the user is going to make a purchase with 54.8% accuracy, which is considerably higher than the baselines mentioned above, i.e. 46.5% and 45.3% accuracy by always predicting the latest or the most consumed app, respectively.

5. RELATED WORK

Online shopping is becoming more popular as people learn to trust online payment systems, which was not the case in the past [11]. Multiple studies, aimed at profiling online shoppers, found that online shoppers tend to be younger, wealthier, and more educated compared to the average Internet user [28, 26, 27, 15]. A more recent work showed that while women are more likely to be online shoppers, men spend more money per purchase and make more purchases overall [18]. In our work, we focus on a particular subset of online purchases, iPhone digital purchases. There are considerable differences in characteristics of iPhone purchases and purchases of physical goods. One of the main differences is that people are much more likely to purchase the same item multiple times.

Similar to online shopping, spending on mobile digital goods is increasing, and people have spent more than \$20 billion dollars in

the Apple App Store in 2015 [4], which is four times more per user than in Android App Stores². This might be due to the different demographics of iPhone users. Given this high level of spending, understanding the market would help us to more effectively target apps towards users who are likely to become regular users and frequent spenders. Despite the popularity of the iPhone digital market, there has not been any large-scale study of how people are spending money on this platform. In this work, we show that most of the money is spent on in-app purchases, and we present a demographic and prediction analysis of spending.

Usage and purchases from apps have been the subject of a few studies. Sifa et al. studied the purchase decisions in free-to-play mobile games [25]. They built a classifier that predicts whether a user is going to make any purchase in the future and also build a regression model to estimate the amount of money that will be spent by each user. The models are moderately accurate. Schoger studied the monetization of popular apps in the global market, identifying growing markets and that in-app purchases are increasingly accounting for larger fraction of total purchases [24]. Our study, unlike those studies, includes the full history of iPhone purchases by the users and considers that many users make purchases from multiple apps. Moreover, the large scale of the dataset allows us have enough big spenders to analyze their behavior accurately.

We also study changes in user purchases over time, how users become frequent buyers in a particular app, and how their purchases evolve over time. The abandonment of a service is called *consumer attrition* or *churn*. The importance of consumer attrition analysis is driven by the fact that retaining an existing consumer is much less expensive than acquiring a new consumer [21]. Thus, prediction of consumer churn is of great interest for companies, and has been studied extensively. For example., Ritcher et al. exploit the information from users’ social networks to predict consumer churn in mobile networks [22]. Braun and Schweidel focus on the causes of churn rather than when churn will occur [14]. They find that a considerable fraction of churn in the service they studied happens due to reasons outside of the companies’ control, e.g., the consumer moving to another state. In context of mobile games, Runge et al. study user churn for two mobile games and predict it using various machine learning algorithms [23]. They also implement an A/B test and offer players bonuses before the predicted churn. They find that the bonuses do not result in longer usage or spending by the users. Kloumann et al. study the usage of apps by users who use a Facebook login and model the lifetime of apps using the popularity and sociality of apps, showing that both of these affect the lifetime of the app [17]. Baeza-Yates et al. addressed the problem of predicting the next app the user is going to open through a supervised approach [9]. In our work, we model the whole sequence of purchases that users make, including adoption, churn, and prediction of the next app.

Our work is the first work that studies the details of all iPhone purchases made by a large number of users. This allows us to better understand the interplay between usage of multiple apps that are competing for the same users, their attention and purchasing power.

6. CONCLUSION

Mobile devices have grown wildly in popularity and people are spending more money purchasing digital products on their devices. To better understand this digital marketplace, we studied a large dataset of more than 776M purchases made on iPhones, includ-

²<http://fortune.com/2014/06/27/apples-users-spend-4x-as-much-as-googles/>

ing songs, apps, and in-app purchases. We find that, surprisingly, 61% of all the money spent is on in-app purchases, and a small group of users are responsible for most of this spending: the top 1% of users are responsible for 59% of all spending on in-app purchases. We characterize these users, showing that they are more likely to be men, older, and less likely to be from the US. Then, we focus on how these big spenders start and stop making purchases from apps, finding that as users gradually lose interest, the delay between purchases increases. The amount of money spent per day on purchases initially increases, then decreases, with a sharp drop before abandonment. Nevertheless, from the perspective of app developers these big spenders are a valuable user segment as they are 4.5x more likely to be a big spender in a new app than a random app user. In the last part of our study we model the purchasing behavior of users by breaking it down in three different steps. First, we model the time between purchases by testing a variety of different distributions, and we find the Pareto distribution fits the data most accurately. Second, we take a supervised learning approach to predict whether a user is going to make purchase from a new app. Finally, if the purchase is from a new app, we use a novel approach to predict the new app based on the previous in-app purchases. If the purchase is from an app that user purchased from in the past, we combine the earlier frequency of the purchases and the time between the purchases to predict from which app the re-purchase will come from. The findings and the models proposed in our study can be leveraged by app developers, app stores and ad networks to better target the apps to the corresponding users.

7. REFERENCES

- [1] Annual apple app store revenue in 2013 and 2015 (in billion u.s. dollars). <http://www.statista.com/statistics/296226/annual-apple-app-store-revenue/>. Accessed: 2016-07-11.
- [2] Apple's cut of 2015 app store revenue tops \$6b. <http://www.computerworld.com/article/3019716/apple-ios/apples-cut-of-2015-app-store-revenue-tops-6b.html>. Accessed: 2016-07-11.
- [3] Online shopping is still a pathetically tiny fraction of all shopping. <http://goo.gl/ecTKxT>. Accessed: 2016-07-11.
- [4] People spent an insane amount of money on apps this year. <http://time.com/4169153/apple-app-store-stats-2015>. Accessed: 2016-07-11.
- [5] Personal consumption expenditures. <https://research.stlouisfed.org/fred2/series/PCE/>. Accessed: 2016-07-11.
- [6] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering*, 17(6):734–749, 2005.
- [7] H. Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected Papers of Hirotugu Akaike*, pages 199–213. Springer, 1998.
- [8] A. Anderson, R. Kumar, A. Tomkins, and S. Vassilvitskii. The dynamics of repeat consumption. In *Proceedings of the 23rd international conference on World wide web*, pages 419–430. ACM, 2014.
- [9] R. Baeza-Yates, D. Jiang, F. Silvestri, and B. Harrison. Predicting the next app that you are going to use. In *WSDM '15: Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pages 285–294, New York, NY, USA, 2015. ACM.
- [10] A. R. Benson, R. Kumar, and A. Tomkins. Modeling user consumption sequences. In *Proceedings of the 25th International Conference on World Wide Web*, pages 519–529. International World Wide Web Conferences Steering Committee, 2016.
- [11] A. Bhatnagar, S. Misra, and H. R. Rao. On risk, convenience, and internet shopping behavior. *Commun. ACM*, 43(11):98–105, Nov. 2000.
- [12] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. In *The Journal of machine Learning research*, 2003.
- [13] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutiérrez. Recommender systems survey. *Knowledge-Based Systems*, 46:109–132, 2013.
- [14] M. Braun and D. A. Schweidel. Modeling customer lifetimes with multiple causes of churn. *Marketing Science*, 30(5):881–902, 2011.
- [15] S. Farag, T. Schwanen, M. Dijst, and J. Faber. Shopping online and/or in-store? a structural equation model of the relationships between e-shopping and in-store shopping. *Transportation Research Part A: Policy and Practice*, 41(2):125–141, 2007.
- [16] L. Hong and B. D. Davison. Empirical study of topic modeling in twitter. In *In Proceedings of the First Workshop on Social Media Analytics*, 2010.
- [17] I. Kloumann, L. Adamic, J. Kleinberg, and S. Wu. The lifecycles of apps in a social ecosystem. In *Proceedings of the 24th International Conference on World Wide Web*, pages 581–591. ACM, 2015.
- [18] F. Kooti, K. Lerman, L. M. Aiello, M. Grbovic, N. Djuric, and V. Radosavljevic. Portrait of an online shopper: Understanding and predicting consumer behavior. In *9th ACM International Conference on Web Search and Data Mining (WSDM'16)*, 2015.
- [19] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119, 2013.
- [20] R. Quinlan. Data mining tools see5 and c5. 0. 2004.
- [21] F. Rechinheld and W. Sasser. Zero defections: Quality comes to service. *Harvard Business Review*, 68(5):105–111, 1990.
- [22] Y. Richter, E. Yom-Tov, and N. Slonim. Predicting customer churn in mobile networks through analysis of social groups. In *SDM*, volume 2010, pages 732–741. SIAM, 2010.
- [23] J. Runge, P. Gao, F. Garcin, and B. Faltings. Churn prediction for high-value players in casual social games. In *2014 IEEE Conference on Computational Intelligence and Games*, pages 1–8. IEEE, 2014.
- [24] C. Schoger. How the most successful apps monetize globally. Retrieved January, 15:2015, 2014.
- [25] R. Sifa, F. Hadiji, J. Runge, A. Drachen, K. Kersting, and C. Bauckhage. Predicting purchase decisions in mobile free-to-play games. In *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*, 2015.
- [26] W. R. Swinyard and S. M. Smith. Why people (don't) shop online: A lifestyle study of the internet consumer. *Psychology & Marketing*, 20(7):567, 2003.
- [27] W. R. Swinyard and S. M. Smith. Activities, interests, and opinions of online shoppers and non-shoppers. *International Business and Economics Research Journal*, 3(4):37–48,

2011.

[28] M. Zaman and M. Y. W. Meng. Internet shopping adoption:

A comparative study on city and regional consumers. In *ANZMAC 2002*, pages 2421–2428. Deakin University, 2002.