

Matching Mobile Applications for Cross Promotion^{*†}

Gene Moo Lee
Department of Computer Science
University of Texas at Austin
gene@cs.utexas.edu

Joowon Lee
McCombs School of Business
University of Texas at Austin
jwlee1015@gmail.com

Andrew B. Whinston
McCombs School of Business
University of Texas at Austin
abw@uts.cc.utexas.edu

October 13, 2014

*We thank Kangwoo Lee, Seungbeom Kim, Seung-guk Yang, and Young-gil Kim from IGAWorks for providing data. We also thank John W. Hatfield from University of Texas at Austin and JaeJoon Han from Inha University for the suggestions on modeling, and Shu He from University of Texas at Austin and Liangfei Qiu from University of Florida for the discussions on empirical analysis. We are responsible for all the problems in the paper.

†Patent pending

Abstract

Mobile ecosystem is one of the most successful markets in recent years with billions of mobile device users and millions of mobile applications (apps) in multiple app platforms. As the market grows, the challenge is how app developers advertise their apps to the right users and how customers search the right apps that fit their needs. Cross promotion, advertising a mobile app in another established app, is introduced as a new promotion framework. The performance of this emerging ad framework has not been well studied in the literature. Using data from 1,011 cross promotions that ran between September 2013 and May 2014 in Korean app markets involving one million consumers and 325 mobile apps, we evaluate the effectiveness of cross promotion in comparison with existing ad channels such as mobile display ads. While cross promotion, on average, is still suboptimal as compared with display ads, we find evidence that a careful matching of mobile apps can significantly improve the effectiveness of cross promotions. We model the ad placement in cross promotions as a matching problem and identify significant factors that contribute to better app matching. Results show that app similarity, measured by topic models from apps' text descriptions, is a significant factor that increases the user engagement in cross promotions. With the observations, we propose a matching mechanism to generate app matches with stability and improved effectiveness.

Keywords: mobile applications, cross promotion, matching, topic modeling, market design, mobile targeting

1 Introduction

Mobile ecosystem is one of the most successful markets in recent years (Petsas et al. 2013; Breshnahan and Greenstein 2014; Yin, Davis, and Muzyrya 2014). Millions of mobile applications (apps) are developed in multiple mobile app markets such as Apple’s App Store, Google’s Play Store, and Microsoft’s Windows Phone Store. Billions of people are adopting smartphones and tablets as their main Internet devices, so the demand for mobile apps keeps increasing. This successful two-sided market is opening up a post-PC era in the computing industry.

Product diversity is one of the key success factors in the mobile app market. In addition to the first-party apps developed by the platform builders, open application programming interface (API) allows third-party developers to bring innovative products to the market. Of note is that a significant number of third-party apps are developed by relatively small-sized startups with the support of various platforms. New mobile apps can reach the global market through well-established distribution channels, and new app services can support large user demands with cloud services without large investments on infrastructure. As a result, we are experiencing a huge growth in mobile app markets.

Our expectation of this market is that the mobile app popularity follows a long-tail distribution (Anderson 2008): many apps with small user bases contribute to a significant portion of the total market share. However, recent studies have found evidence that mobile app markets are actually experiencing a “winner-take-all” phenomenon (Petsas et al. 2013; Zhong and Michahelles 2013). A recent TechCrunch report indicated that 54% of total app store revenue goes to only 2% of the developers and that almost half of the developers earn less than \$500 a month¹. This is a sharp contrast to other online markets such as video streaming (Anderson 2008), auctions (Hu and Bolivar 2008), retail (Linden, Smith, and York 2003), and even music stores. Actually, many independent app developers have already switched to more stable positions in established firms². Norumra recently reported that even the Chinese mobile game market shows signs of slowdown because no killer apps emerge in the market³. We argue that this phenomenon can compromise the vitality of the mobile app markets.

It is believed that this market inefficiency is due to the fact that app advertising (ad) heavily relies on app marketplaces’ in-house ranking systems, which provide lists of popular and growing apps in different ranking criteria. Hence the developers’ primary goal is to somehow get into the rankings, rather than to produce high-quality software. Without an efficient app search mechanism, customers are mainly exposed to the top ranked apps, which cover only a small fraction of the whole market. This trend calls for better marketing strategies to promote mobile apps to potential active customers and to enable users to search the right apps that fit their needs.

Cross promotion has recently emerged as a way to recommend new apps to the users who are already using related established apps. For example, game app developers can promote their new products to the active users playing other games of a similar genre. For new app developers, this is an effective ad channel to reach potential customers. For the established app publishers, cross promotion provides a way to monetize their visibility. Potentially established apps may even improve their reputations by providing good app recommendations to their customers. Cross promotions incentivize users to install and use new apps by providing credits (e.g., free game

¹<http://techcrunch.com/2014/07/21/the-majority-of-todays-app-businesses-are-not-sustainable/>

²<http://apple.slashdot.org/story/14/07/30/1838203/is-the-app-store-broken>

³<http://blogs.barrons.com/asiastocks/2014/09/08/nomura-tencent-qihoo-may-see-pressure-on-mobile-gaming/>

items) in the apps they use.

Figure 1 shows a screenshot of a cross promotion event from IGAWorks, a Korean mobile ad company. In this promotion, an app introduces a list of other apps along with the rewards to give if users participate the event by installing or using the apps. There are many active cross promotion networks including AppFlood⁴, Chartboost⁵, Tapjoy⁶, and LeadBolt⁷. In a broader sense, Facebook and Twitter also provides cross promotions by providing their real-estates in news feeds to the app publishers. Despite the pervasiveness of cross promotion, this new ad framework has not been studied in the literature.

This paper sheds light on the cross promotion platform in mobile app markets. The contribution of the paper is sixfold.

First, we empirically evaluate the ad effectiveness of cross promotion using data with 1,011 cross promotions conducted from September 2013 to May 2014 in Korean app markets, involving with one million consumers and 325 mobile apps. We compare this emerging ad framework with other user acquisition channels such as organic growth and mobile display ads. While data shows that cross promotion is still suboptimal in terms of the acquired users' engagement, we also find evidence that careful ad placements can significantly improve the ad effectiveness of cross promotions. Based on the observations of successful campaigns, we hypothesize that the effectiveness of a cross promotion depends on pairwise app similarity as well as individual apps' characteristics.

Mobile targeting is the one of the most important agenda items in both academia and industry. There is a growing literature on various user targeting strategies (Goldfarb and Tucker 2011; Luo et al. 2013; Ghose, Goldfarb, and Han 2013; Baker, Fang, and Luo 2014; Bart, Stephen, and Sarvary 2014). The industry is also actively experimenting with different approaches to place the ads to the right customers at the right time and location. Facebook is trying to leverage their strong social graph in mobile app ads market⁸. Google recently announced a new technology to track mobile app usages along with mobile web behaviors for better ad targeting⁹. Existing approaches target users according to locations, times, and social relationships. Our approach is to target potential active app users by selecting the right apps where cross promotions are conducted. In doing so, we leverage topic model based app similarity between apps hosting the promotions and those to be promoted.

The second contribution of the paper is to model ad placement in cross promotion as a matching problem. Given the apps to promote and those where ads can be placed, the cross promotion platform should arrange the most effective matchings between apps to meet the requirements of the stakeholders. Matching markets have been well studied in the economics literature with many applications such as marriage and dating (Gale and Shapley 1962; Hitsch, Hortaçsu, and Ariely 2010), labor market (Roth 1984; Roth 1991; Roth and Peranson 1999), and school admission (Abdulkadiroglu and Sönmez 2003; Ergin and Sönmez 2006). To the best of our knowledge, our work is the first to frame a matching problem in mobile app markets.

Third, we propose a novel app similarity measure constructed with apps' text descriptions. Specifically, we apply latent Dirichlet allocation (LDA) topic modeling (Blei, Ng, and Jordan

⁴<http://appflood.com/>

⁵<https://www.chartboost.com/en/platform#cross-promotion>

⁶<http://home.tapjoy.com/>

⁷<http://www.leadbolt.com/developer-tools/>

⁸<https://developers.facebook.com/docs/ads-for-apps>

⁹<http://adage.com/article/digital/google-tie-mobile-web-app-trackers-ad-targeting/294502/>

2003; Blei 2012) on the app description texts. The resulting topic model gives the trending topics in the current app market and also transforms individual apps into topic vectors. Then the app similarity is calculated by the cosine similarity between topic vectors.

Next, we empirically estimate our model to identify the variables that improve the ad effectiveness in cross promotion. Specifically, we are interested in similarity between source apps (where the ads are placed) and target apps (which are the ones to promote in the campaign). We find evidence that the proposed app similarity has significantly positive effects to improve the ad effectiveness. In other words, a cross promotion is likely to be successful if source and target apps are closely related. This can be a basis for a recommender system for app markets.

Based on the empirical results, we design a matching mechanism for cross promotions. Using the learned model, a linear programming (LP) based algorithm is used to provide stable matchings. Our counterfactual analysis shows that the matching obtained from the LP can improve the ad effectiveness by 260%.

Lastly, this work can serve as an example of “Big Data” approach to bring machine learning techniques and economic theory into the marketing literature. Many ad frameworks can be modeled as matching problems as done in the present paper. Also, an unprecedented large amount of unstructured text information about products can be analyzed with machine learning algorithms, as shown in this work.

The remainder of the paper is organized as follows. In Section 2, we describe the data on mobile apps and promotions, then compare the ad effectiveness of different ad channels. In Section 3 we model ad placements in cross promotion as a matching problem, and overview the independent variables in the model with the introduction on the novel app similarity measure in Section 4. Empirical results are given in Section 5. Based on the observations, a stable matching algorithm is designed in Section 6. Section 7 concludes the paper with future directions.

2 Data

We first describe data on mobile app markets, then compare the effectiveness of three ad channels – organic growth, mobile display ads, and cross promotions – in terms of user engagements.

2.1 Data Description

We use data from IGAWorks, a leading mobile advertising company in Korea. The product line includes a mobile app analytics tool called Adbrix and a mobile app monetization platform supporting various promotions such as mobile display ads and cross promotions. It has the largest mobile ad network in Korea, including hundreds of mobile apps and 2.4 million users. The data was shared by the company using a secure channel. All personally identifiable information (PII) is anonymized to preserve user privacy.

The data consists of three parts: app meta data, usage data, and funnel data. The meta data includes descriptive information about 383,896 mobile apps in three major app markets in Korea: Apple’s App Store, Google’s Play Store, and SK Telecom’s T-store. Play store and T-store provide Android apps, whereas the App Store serves iOS apps. Each app record contains the app name, text description, screenshots, developer, registration time, last update time, price, number of ratings, average rate, and file size. Note that this information is publicly available in the app markets.

Usage data includes detailed information about user engagements. This user level data includes daily app session times (i.e., how long a customer uses an app), daily connection counts (i.e., how many times a customer executes an app), and daily buy activities (i.e., how many times a customer makes in-app purchases). Usage data is available for 501 apps that adopted the Adbrix analytic tool and a total of 1.1 million users' activity data is captured over a six-month period in our data. Note that buy activity is available only for apps with in-app purchase options.

Lastly, funnel data provides information on promotions that IGAWorks has executed with its clients (app developers). The promotions were conducted from September 2013 to May 2014, involving 310,183 user participations and 325 mobile apps. Ad types include cross promotions and mobile display ads. The data keeps track of user acquisition channels for each app. In other words, we observe how and when a given user installed the promoting app, which is the basis to evaluate the effectiveness of promotions.

2.2 Effectiveness of Ad Channels

We measure the effectiveness of a given ad campaign by combining funnel and usage data. We divide user groups according to the acquisition channels: organic growth, mobile display ads, and cross promotions. A user is organic with respect to a mobile app if the app installation is not associated to any ad campaigns. Users are associated to mobile display ads if they installed the app by clicking the banner ads placed in mobile websites or mobile apps (Bart, Stephen, and Sarvary 2014). Lastly, a user is in cross promotion group if he or she installed the app through a reward-based cross promotion conducted in another app. Note that reward is the differentiator of cross promotion as compared with mobile display ads placed in other mobile apps.

Ad effectiveness can be measured with various user engagement metrics such as session times, connection counts, or buy activities. In our study, we focus on session times and connection counts because buy activities are only available in mobile apps with in-app purchase options. We say an ad channel is effective if the users acquired through the channel show active engagements (e.g., longer session time). We argue that the number of app downloads is not a good metric of ad effectiveness because the users acquired from promotions may not end up being active users.

Figure 2 shows the average user session times in three user acquisition channels: organic (red line), mobile display ads (green line), and cross promotions (yellow line). The X-axis lists mobile apps sorted by the average session times of total users (blue line) and the Y-axis shows the average session times in each channel. We observe that organic users are the most active group. This finding is intuitive because an app installation without any external inputs indicates the user's strong motivation to use the app. User groups from display ads show 50% less engagement than organic user groups. Lastly, we clearly observe that users acquired by cross promotions are the least active group. Since the app installation in cross promotion is incentivized by the rewards, users may install the promoting apps but do not use them afterwards. This is an issue for both the promotion platform and participating apps because the promotion yields a low return on investment.

Next, we conduct an in-depth analysis within cross promotions. For a given app to advertise (we call it *target* app), there are multiple apps where the ads can be placed (we call them *source* apps). For a given target app, we divide its users according to the specific acquisition subchannel (e.g., the source app). Then for each source-target pair, we calculate the average user engagement levels, then identify 1% and 10% best pairs for each target app. Figure 3 shows two additional lines for the top 1% (purple dotted line) and 10% (sky-blue dotted line) pairs. We find that top

1% matches are 690% more effective than the average ones and that the top 10% are 130% more effective than the average. Results also show that the top 1% matches outperform the display ads in almost half of the target apps (48%), and they even outperform organic acquisitions in 22% of the samples. Based on these observations, we argue that the app matches in cross promotion should be optimized so that the ads are targeted to the right source apps which users are likely to be active in the target apps.

Given the large impact of source-target matching on the ad effectiveness, the question is what makes a good match. We compare the list of good matches with that of bad ones to find that a pair of apps with similar genres and topics makes a good match. For example, a new poker game is actively used by the users acquired from other similar gambling games. On the other hand, bad matches involve two unrelated apps such as a celebrity photos app and a utility app. Based on these observations, we hypothesize that app similarity positively contributes to the ad effectiveness of cross promotions. In the next section, we build a model of ad effectiveness in cross promotions. Then we operationalize the app similarity measure in Section 4.2.

3 Modeling Cross Promotion

A cross promotion involves with three groups of entities: source app, target app, and the promotion platform. App publishers who want to promote their (target) apps make contracts with the platform to launch a campaign with the specific number of app installations to acquire. Then the cross promotion platform places the ads in the (source) apps that agreed to conduct cross promotions. Note that source apps are mostly popular ones that already have large user bases, whereas targets are usually new apps with limited awareness in the market. Thus we assume no overlaps in source and target apps.

Source apps are paid by the targets according to the number of target app installations they achieved and the promotion platform gets a cut on each installation. Essentially, this is a cost-per-action (CPA) pricing model. A campaign is finalized when the number of app installations reaches the goal. One thing to note is that the utility of source apps and the platform is based on app download counts, where the objective of target apps is to acquire *active* users. This misalignment of these two objectives may explain the suboptimal ad effectiveness of the current cross promotion data shown in Section 2.2. In order for the promotion market to sustain, the objectives of sources, targets, and the platform should be harmonized.

Another economic insight about cross promotion is that the platform acts as an intermediary match maker to match source and target apps. Thus cross promotion framework creates a two-sided matching market rather than a commodity market. In a commodity market, it is assumed that sellers (source apps in our case) and buyers (target apps in our case) have perfect information about each other and that sellers and buyers can switch their roles in different situations. Also, prices and transactions can be determined without any intermediary. However, the cross promotion market has information asymmetry issues: Source apps have superior information about the customers than do target apps and they may only want to reveal private information to the matched counterparts. Also, the platform has extensive knowledge about the whole market. Thus the existence of the promotion platform as a match maker is essential.

Matching markets have a strong theoretical foundation established in the economics literature (Gale and Shapley 1962; Roth 1984; Roth 1991; Roth and Peranson 1999; Abdulkadiroglu and

Sönmez 2003; Ergin and Sönmez 2006; Hitsch, Hortacısu, and Ariely 2010; Hatfield et al. 2013). The theory has been applied to many empirical studies involving with marriage (Gale and Shapley 1962), online dating (Hitsch, Hortacısu, and Ariely 2010), labor market (Roth 1984; Roth 1991; Roth and Peranson 1999), and school admission (Abdulkadiroglu and Sönmez 2003; Ergin and Sönmez 2006).

We frame the ad placement in cross promotion as a matching problem. Let S be the set of source apps where ads can be placed and T be the set of target apps to be advertised. Then let $G = \langle V, E \rangle$ be the bipartite graph where $V = S \cup T$ and $S \cap T = \emptyset$. For a given target app $t \in T$, the platform should select a source app $s \in S$, creating an edge $(s, t) \in E$. Note that an edge is not created within the same subset (S or T) under our assumption.

The effectiveness of an app match $u(s, t)$ is measured by the user engagement levels in target t . Our hypothesis is that the effectiveness depends on the individual characteristics of s and t and the pairwise similarity between s and t . Thus the effectiveness of an app match is given by a linear functional form:

$$u(s, t) = \alpha_0 + \alpha_1 X_s + \alpha_2 X_t + \alpha_3 P_{s,t} + \epsilon_{s,t} \quad (1)$$

where X_s and X_t represent individual characteristic vectors of apps s and t (e.g., popularity, quality, age). $\epsilon_{s,t}$ is the individual heterogeneity of a match s and t , and is independent across all pairs (s, t) . Then $P_{s,t}$ is the symmetric app similarity between apps s and t ($P_{s,t} = P_{t,s}$) and parameter α_3 measures the tendency that users engage in similar apps. In our context, the similarity measure is operationalized by apps' text descriptions. Details on the independent variables are described in Section 4.

4 App Characteristics and Similarity

In this section, we describe mobile apps' individual characteristics considered in the model, then propose a novel pairwise app similarity measure by applying a machine learning technique to apps' text descriptions.

4.1 Individual App Characteristics

Recent empirical studies on app markets have shown that various app characteristics (e.g., popularity, quality, age, complexity) affect the user preference (Ghose and Han 2014; Lee and Raghu 2014; Yin, Davis, and Muzyrya. 2014). To capture app popularity in our model, we use number of ratings (Num_Rates) reported in app markets. It is worth noting that the number of app downloads is not publicly available in most markets (Ghose and Han 2014). Thus we use rate count as a proxy for app popularity. Then we use the average rate (between 1 and 5) to capture the latent app quality observed by the existing app users (Avg_Rate). We also consider two age-related variables: number of days since the initial app registration (Days_Regist) and number of days since the last update (Days_Update). One may argue that old apps are likely to lose attention as people search for new things (Feinberg, Kahn, and McAlister 1992; Xu et al. 2011). On the other hand, we may expect that apps that have survived a long time have some compelling features that keep consistent user engagements. Recent update time reveals the developer's engagement level in the product: If an app does not have update for a long period, it may indicate that developers lost

interest in adding new features. The last individual app characteristic is the file size in megabytes (File_Size). Large file size may indicate that the developer made significant efforts and that the app has complex functionalities.

4.2 Topic Models and App Similarity

Besides individual app characteristics, we argue that app similarity can positively affect the ad effectiveness in the model. Studies show that people usually stick to a certain taste when they select products in online shopping (Linden, Smith, and York 2003), music streaming (Hariri, Mobasher, and Burke 2012) and mobile app usage (Natarajan, Shin, and Dhillon 2013). Essentially, customers’ tendencies to choose similar products is the basis for online recommender systems. One may argue that app genre can be used to measure app similarity. However, this method can only provide binary relationships between apps, which is not sufficient for our purpose to measure the degree of closeness.

App similarity is operationalized by processing apps’ text descriptions. Developers provide detailed app descriptions in the app market so that potential users can understand the features provided by the apps. A pair of apps with similar descriptions is supposed to share common features such as game genres, usage scenarios, and so on. The issue is how we process unstructured text descriptions in a principled way to quantify the pairwise closeness.

Our approach is to use latent Dirichlet allocation (LDA) topic modeling on the app description corpus (Blei, Ng, and Jordan 2003; Blei 2012). LDA is a natural language processing technique that allows a set of documents to be explained by hidden “topics,” which are sets of related keywords. LDA has been successfully used to analyze documents in various domains such as scientific articles (Griffiths and Steyvers 2004; Wang and Blei 2011; Blei 2012), music (Hariri, Mobasher, and Burke 2012), social media (Ramage, Dumais, and Liebling 2010; Weng et al. 2010; Lee, Qiu, and Whinston 2014), and firms (Shi, Lee, and Whinston 2014). In our context, each app description is a mixture of a small number of app features and each word in the description is a realization of the app features. For details on LDA see Blei (2012).

We run LDA on the text descriptions of 195,956 mobile apps in Korean market. We vary the number of topics to find that 100-topic model gives the best result. Table 1 shows a partial list of 100-topic model¹⁰. The keywords in each topic are translated into English for readability. We believe that the topics give a reasonable overview of the app market. Topics in the Korean app market include music (topics 0, 27), social networks (topics 1, 14, 25, 41, 89), kids (topics 6, 34), religion (topic 11), games (topics 16, 27), sports (topic 76), online dating (topic 96), foreign language education (topics 19, 33, 81, 93), e-commerce (topics 18, 29), and utilities (topics 10, 13, 48, 49, 97).

Once the topic model is built, an app i ’s description can be represented by a topic vector $V_i = \langle V_{i,1}, V_{i,2}, \dots, V_{i,K} \rangle$, where K is the number of topics, $V_{i,k}$ is the weight on the k -th topic, and the sum of weights is 1 ($\sum_{k=1}^K V_{i,k} = 1$). Given a pair of source s and target t and their topic vectors V_s and V_t , we define the app similarity $P(s,t)$ (Topic_Similarity) to be the cosine similarity of the two topic vectors as follows:

¹⁰For full list of topics and keywords, see <http://diamond.mcombs.utexas.edu/app.topic.keywords.txt>

$$P(s,t) = \frac{V_s \cdot V_t}{\|V_s\| \|V_t\|} = \frac{\sum_{k=1}^K V_{s,k} V_{t,k}}{\sqrt{\sum_{k=1}^K (V_{s,k})^2} \sqrt{\sum_{k=1}^K (V_{t,k})^2}} \quad (2)$$

where the resulting values range from 0 to 1. For the extreme cases, $P(s,t) = 0$ if two apps do not share any common topics and $P(s,t) = 1$ if two apps have identical topics. Similar approaches are used to measure user similarity in social networks (Lee, Qiu, and Whinston 2014) and firms’ business proximity in high tech industry (Shi, Lee, and Whinston 2014).

5 Empirical Analysis

In this section, we present the estimation results on the ad effectiveness of cross promotions. We collect the list of target apps that have conducted cross promotion campaigns along with the list of corresponding source apps where the ads were placed. The cross promotion data includes 1,011 app matches and 310,183 user participations. An app match in a promotion is said to be effective if the promotion acquires active users with longer session times and higher connection counts.

Table 2 shows the estimation results on user session times and Table 3 gives those on user connection counts. For a robustness check, we estimate four different models by including and excluding various app characteristics. Characteristics can be divided into two groups: customer-given and developer-given. Customer-given variables include number of ratings (for popularity) and average rates (for quality), and developer-given ones are registration time (for age), update time (for responsiveness), and file size (for complexity).

We find strong evidence that the effect of app topic similarity, `Topic_Similarity`, on ad effectiveness is significantly positive. The results are consistent with all models in both dependent variables. This result validates our hypothesis that people tend to like target apps that are highly similar to sources. It means that the user preference on app adoption is to some extent predictable based on the current apps they are using. This result can be a basis for a recommender system to introduce new apps to users according to the topic similarity.

Empirical results also show that various individual app characteristics have significant impacts on app engagement. First, the effects of average ratings of both source (`Avg_Rate_Source`) and target (`Avg_Rate_Target`) apps are significantly positive. This finding indicates that apps with better quality are more attractive to the customers, which follows intuition. An interpretation on the source app quality effect can be that promotions from high quality apps are perceived to be more reliable to the customers, which leads to high user engagements. A similar phenomenon can be found in job markets: applicants recommended by well established people are more likely to be accepted by the recruiters.

We do not observe consistent effects of app popularity on the ad effectiveness (`Num_Rates_Source` and `Num_Rates_Target`). Target apps are usually new in the market, so the rate counts may not matter. However, it is interesting that even the source app’s popularity does not have consistent effects. This may indicate that ads should be placed with the “right” apps, not the “popular” ones.

Next we consider developer-given variables. The target app’s age (`Days_Regist_Target`) has a significantly positive impact on user engagement. An interpretation can be that apps that have survived in the market for a long time have intrinsic values in them. The number of days since last update (`Days_Update_Target`) has a significantly negative impact on engagement. In

other words, target apps with infrequent updates are less likely to keep the customer’s attention. This may suggest that app developers should actively respond to their customers’ feedback and add new features to their products. Results show that source apps’ age-related variables do not have consistent effects. Lastly, the file size of target apps (`File_Size_Target`) has a significantly positive effect in all the models, indicating that well-made apps are more likely to increase user engagements.

6 Matching Mechanism Design

We design a matching mechanism for cross promotions, followed by the model introduced in Section 3. Given the set of target apps that want to be advertised and the set of source apps who can provide real-estate for cross promotions, the platform should decide an assignment to meet the requirements from sources and targets. We leverage the model on ad effectiveness to calculate the expected utility of each app pair. There are three main issues to consider in designing the matching mechanism: utility transferability, information structure, and monogamy.

We first discuss the utility of matchings. In the literature on marriage matching market (Gale and Shapley 1962), the utility of each side is separated as compensating transfers are not allowed. However, in the cross promotion market, utility can be transferred from targets to sources according to the performance of the promotions. This is similar to the model from Shapley and Shubik (1972). A target app’s gained utility of a match can be interpreted as the engagement levels of the users achieved by the matched cross promotions. The utility of a source app is the reward it gets when one of its users installed the target. Based on the empirical results in Section 5, we define the utility of a potential app match to be the ad effectiveness given by Equation 1.

The next design issue is about the information structure. We assume that perfect and costless information about potential matches is available to all participants. In other words, each target (source) app is aware of the potential utility achievable from all possible source (target) apps. This is a reasonable assumption because all the variables (text descriptions, ratings, ages, etc.) needed to estimate the ad effectiveness are public information available in the app markets.

Lastly, we assume monogamous matching in cross promotions: one target (source) can be assigned to at most one source (target). In most cases, the platform should perform one-to-one matchings. However, some promotions involve multiple target apps where a popular source app hosts multiple cross promotions simultaneously. This scenario can be modeled as many-to-one matchings as in job markets, where multiple employees can work for a single company (Kelso and Crawford 1982).

In summary, the app matching problem can be considered a frictionless one-to-one matching with transferable utilities.

Now we formally design the matching mechanism. Let S be the set of source apps where ads can be placed and let T be the set of target apps to advertise. Then let $u_{s,t}$ be the utility of a match between source s and target t . Note that the utility is transferred so the gained utility value is given by a pair of apps. Then let $u_{0,t}$ be the utility that target t receives if no ads are placed in any source app. We assume that apps get zero utility if they are not matched with any other apps ($u_{0,t} = 0$ and $u_{s,0} = 0$). We define the match assignment indicator, $m_{s,t}$, such that $m_{s,t} = 1$ if and only if source s is advertising target t and $m_{s,t} = 0$ otherwise. Then, following Gale (1960) and Shapley and Shubik (1972), a *stable* assignment can be obtained by solving an integer linear programming

(LP) problem as below:

$$\max_{m_{s,t}} \sum_{s \in S} \sum_{t \in T} m_{s,t} u_{s,t} \quad (3)$$

subject to

$$\sum_{t \in T} m_{s,t} \leq 1, s = 1, 2, \dots, S \quad (4)$$

$$\sum_{s \in S} m_{s,t} \leq 1, t = 1, 2, \dots, T \quad (5)$$

The solution of this LP can serve as a recommended matching for cross promotions. Note the inequality in the constraints (4) and (5): As the number of sources and that of targets can be different, some apps may not be matched for cross promotions.

There are a few remarks about the problem. The first issue is about stability of the matching. An assignment is said to be *stable* if there is no app that would rather not be matched and if there are no two apps that would prefer to form a new matching for cross promotion. From Shapley and Shubik, it is shown that the assignment obtained by solving the LP is stable. In other terms, this app match assignment has the *core* property from cooperative game theoretic perspective (Chapter 9 in Myerson 1991; Sorenson, Tschirhart, and Whinston 1978)¹¹. The core is the set of assignments that cannot be improved by the deviation from any subset of players. In other words, there are no source or target app developers who can achieve better utility by deviating from the assignment proposed by the platform. This property secures the authority of the platform.

One can actually assume that the assignment indicator, $m_{s,t}$, can be real numbers, instead of integers. Intuitively, $m_{s,t}$ can be interpreted as the probability of source s being matched to target t . However, it is shown that the constraint matrix of the LP assignment problem is totally unimodular, thus all extreme points are integers (Nemhauser and Wolsey 1988). In other words, the solution of the LP always gives the results with all $m_{s,t}$ being zero or one.

The next remark is that the assignment problem is defined as a standard LP, where we want to find a vector that maximizes the objective function (3) with the constraints (4) and (5). Therefore, we can use a standard tool of LP: duality theory, which says that every maximization problem, called primal, can be converted into a dual minimization problem. Aggregate utility maximization that decides the assignments is a dual cost minimization problem that determines the set of possible divisions of the gained utility. Specifically, we define a dual variable x_s for each constraint (4) and a dual variable y_t for each constraint (5). Then the dual program is given as follows:

$$\min_{x_t, y_s} \left(\sum_{s \in S} x_s + \sum_{t \in T} y_t \right) \quad (6)$$

subject to

$$x_s + y_t \geq u_{s,t}, s \in S, t \in T \quad (7)$$

$$x_s \geq 0, y_t \geq 0 \quad (8)$$

¹¹In cooperative game theory, a subset of players form a *coalition* and the payoff of each player is decided by the coalition. Mobile apps form coalitions in the cross promotions. Side payments are also possible within the matched app developers, which means that the utility is transferable. These properties are different from the non-cooperative games where it is assumed that the players in the game cannot directly communicate each other and do not share the utility.

The optimal values of x_s and y_t can be interpreted as the *prices* of the constraint in the original maximization problem (the primal). Then $x_s + y_t = u_{s,t}$ if the match is formed, and $x_s + y_t \geq u_{s,t}$ otherwise. This dual LP can serve as a mechanism to recommend the prices of app matches according to their competitive advantage. In other words, x_s can be the price to pay the source app in order to conduct a cross promotion and y_t can be the price for the target. Note that payments from targets to sources are conditional on the number of downloads achieved, which is different from the fixed price case in Kelso and Crawford (1982).

With the proposed LP based matching mechanism, we conduct a counterfactual analysis to produce optimal matching. From the empirical analysis from Section 5, we learn the parameters for Equation 1 in Section 3. We use this model to calculate the predicted utility values for all possible matches ($u_{s,t}$). Using the GNU Linear Programming Kit (GLPK), we run the primal LP to find the optimal assignment ($m_{s,t}$). It turns out the assignment obtained from the LP gives much higher predicted utility value than the current matching in the promotion data: The existing matching in the data gives an average utility of 0.189 for each app pair. As a comparison, the average utility of all possible app pairs is 0.204, which shows the suboptimality of the current matches. Furthermore, the matching obtained by the LP achieves an average predicted utility value of 0.679, which is a 260% improvement from the baseline. This counterfactual analysis shows that the proposed matching algorithm can achieve both stability and improved effectiveness. One may argue about the accuracy of the predicted utility values. Thus we plan to conduct a randomized field experiment to compare the performance of different matchings.

7 Conclusion and Future Directions

In this paper, we study cross promotion in the mobile app market. As compared with other user acquisition channels such as organic growth and mobile display ads, cross promotion shows sub-optimal ad effectiveness in terms of user engagement. However, it has also shown that carefully matches source and target apps can significantly improve the ad effectiveness. We built a model to identify significant factors that contribute to better app matching. Empirical results show that app similarity, measured by app descriptions' topic model, has a significantly positive effect to improve ad effectiveness. Lastly, we proposed a matching mechanism for cross promotions to achieve stable app matching with improved ad effectiveness.

From the modeling perspective, we assume a frictionless one-to-one matching in cross promotion markets. We plan to extend our studies by relaxing some assumptions. For the information structure, some variables related to matching effectiveness can be privately shared. Also, source apps can host multiple targets simultaneously, thus we may extend the model to the many-to-one matching market. Eventually, we may consider many-to-many matching markets as one target app can perform promotions on multiple source apps and a single source app may advertise multiple targets.

Mobile app market is highly dynamic: new apps enter the market, existing ones disappear or update themselves with new features, and app demands change rapidly. Thus our matching model can be extended to capture the dynamics of the market (Anderson and Smith 2010; Akbarpour, Li, and Gharan 2014).

References

- [1] Abdulkadiroglu A. and Sönmez T. (2003). “School choice: A mechanism design approach.” *American Economic Review* 93(3): pp. 729-747.
- [2] Akbarpour M., Li S., Gharan S. O. (2014). “Dynamic Matching Market Design.” *Proceedings of the ACM Conference on Economics and Computation*: p. 355.
- [3] Anderson A. and Smith L. (2010). “Dynamic Matching and Evolving Reputations.” *The Review of Economic Studies* 77(1): pp. 3-29.
- [4] Anderson C. (2008). “The Long Tail: Why the Future of Business is Selling Less of More.” Hyperion.
- [5] Baker B. J., Fang Z., and Luo X. (2014). “Hour-by-Hour Sales Impact of Mobile Advertising.” Working paper, Temple University.
- [6] Bart Y., Stephen A. T., and Sarvary M. (2014). “Which Products Are Best Suited to Mobile Advertising? A Field Study of Mobile Display Advertising Effects on Consumer Attitudes and Intentions.” *Journal of Marketing Research* (51): pp. 270-285.
- [7] Blei D. B. (2012). “Introduction to Probabilistic Topic Models.” *Communications of the ACM*, 55(4): pp. 77-84.
- [8] Blei D. B., Ng A. Y., and Jordan M. I. (2003). “Latent Dirichlet Allocation.” *Journal of Machine Learning Research*, 3: pp. 993-1022.
- [9] Breshahan T. and Greenstein S. (2014). “Economics of the Internet and Mobile Computing: The Next Platform Rivalry.” *American Economics Review: Papers & Proceedings*, 104(5): pp. 475-480.
- [10] Ergin H. and Sönmez T. (2006). “Games of School Choice under the Boston Mechanism.” *Journal of Public Economics*, 90: pp. 215–237.
- [11] Feinberg F. M., Kahn E. B., and McAlister L. (1992). “Market Share Response When Consumers Seek Variety.” *Journal of Marketing Research*, 29(2): pp. 227-237.
- [12] Gale D. (1960). “The Theory of Linear Economic Models.” New York: McGraw-Hill.
- [13] Gale D. and Shapley L. S. (1962). “College Admission and the Stability of Marriage.” *American Mathematical Monthly*, 69(3): pp. 9-15.
- [14] Ghose B. Goldfarb A., and Han S. P. (2013). “How is the mobile Internet different? Search costs and local activities.” *Information Systems Research*, 24(3): pp. 613-631.
- [15] Ghose A. and Han S. P. (2014). “Estimating Demand for Mobile Applications in the New Economy.” *Management Science*, Forthcoming.
- [16] Goldfarb A. and Tucker C. (2011). “Online display advertising: Targeting and obtrusiveness.” *Marketing Science*, 30(3): pp. 389-404.

- [17] Griffiths T. L. and Steyvers M. (2004). "Finding Scientific Topics." *Proceedings of the National Academy of Sciences*, 101: pp. 5228-5235.
- [18] Hariri N., Mobasher B., and Burke R. (2012). "Context-Aware Music Recommendation Based on Latent Topic Sequential Patterns." *Proceedings of the ACM Conference on Recommender Systems*: pp. 131-138.
- [19] Hatfield J. W., Kominers S. D., Nichifor A., Ostrovsky M., and Westkamp A. (2013). "Stability and Competitive Equilibrium in Trading Networks." *Journal of Political Economy*, 12(5): pp. 966-1005.
- [20] Hitsch G. J., Hortaçsu A., and Ariely D. (2010). "Matching and Sorting in Online Dating." *American Economic Review*, 100(1): pp. 130-163.
- [21] Hu W. and Bolivar A. (2008). "Online auctions efficiency: a survey of eBay auctions." *Proceedings of the International Conference on World Wide Web*: pp. 925-934.
- [22] Kelso A. S. and Crawford V. P. (1982). "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50(6): pp. 1483-1504.
- [23] Lee G. M., Qiu L., and Whinston A. B. (2014). "Strategic Network Formation in a Location-based Social Network: A Topic Modeling Approach." Working paper, University of Texas at Austin.
- [24] Lee G. and Raghu T. S. (2014). "Determinants of Mobile Apps Success: Evidence from App Store Market." *Journal of Management Information System*, Forthcoming.
- [25] Linden G., Smith B., and York J. (2003). "Amazon.com Recommendations: Item-to-Item Collaborative Filtering." *IEEE Internet Computing*, 7(1): pp. 76-80.
- [26] Luo X., Andrews M., Fang X., Phang C. W. (2013). "Mobile Targeting." *Management Science*: pp. 1-20.
- [27] Myerson R. B. (1991). "Game Theory." Harvard University Press.
- [28] Natarajan N., Shin D., and Dhillon I. S. (2013). "Which app will you use next? Collaborative filtering with interactional context." *Proceedings of the ACM Conference on Recommender Systems*: pp. 201-208.
- [29] Nemhauser G. L. and Wolsey L. A. (1988). "Integer and Combinatorial Optimization." Vol. 18, Wiley, New York.
- [30] Pan W., Aharony N., and Pentland A. (2011). "Composite Social Network for Predicting Mobile Apps Installation." *Proceedings of Association for the Advancement of Artificial Intelligence*.
- [31] Petsas T., Papadogiannakis A., Polychronakis M., Markatos E. P., and Karagiannis T. (2013). "Rise of the Planet of the Apps: A Systematic Study of the Mobile App Ecosystem." *Proceedings of the Internet Measurement Conference*, pp. 277-290.

- [32] Ramage D., Dumais S., and Liebling D. (2010). "Characterizing Microblogs with Topic Models." *Proceedings of International AAAI Conference on Weblogs and Social Media*: pp. 130-137.
- [33] Roth A. E. (1984). "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy*, 92(6): pp. 991-1016.
- [34] Roth A. E. (1991). "A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom." *American Economic Review*, 81(3): pp. 415-440.
- [35] Roth A. E. and Peranson E. (1999). "The Redesign of the Matching Market for American Physicians." *American Economic Review*, 89(4): pp. 748-780.
- [36] Shapley L. and Shubik M. (1972). "The Assignment Game 1: The Core." *International Journal of Game Theory*, 1: pp. 111-130.
- [37] Shi Z., Lee G. M., and Whinston A. B. (2014). "Towards a Better Business Proximity: Topic Modeling for Analyzing M&As." *Proceedings of the ACM Conference on Economics and Computation*: p. 565.
- [38] Sorenson J., Tschirhart J., and Whinston A. B. (1978). "A Theory of Pricing under Decreasing Costs." *American Economic Review*, 68(4): pp. 614-624.
- [39] Wang C. and Blei D. B. (2011). "Collaborative topic modeling for recommending scientific articles." *Proceedings of ACM Conference on Knowledge Discovery and Data Mining*: pp. 448-456.
- [40] Weng J., Lim E. P., Jiang J., and He Q. (2010). "TwitterRank: Finding topic-sensitive influential twitterers." *Proceedings of the ACM International Conference on Web Search and Data Mining*, pp. 261-270.
- [41] Xu Q., Erman J., Gerber A., Mao Z. M., Pang J., and Venkataraman S. (2011). "Identifying diverse usage behaviors of smartphone apps." *Proceedings of the ACM Internet Measurement Conference*: pp. 329-344.
- [42] Yin P.-L., Davis J. P., and Muzyrya Y. (2014). "Entrepreneurial Innovation: Killer Apps in the iPhone Ecosystem." *American Economic Review: Papers & Proceedings*, 104(5): pp. 255-259.
- [43] Zhong N. and Michahelles F. (2013). "Google Play Is Not A Long Tail Market: An Empirical Analysis of App Adoption on the Google Play App Market." *Proceedings of the Annual ACM Symposium on Applied Computing*: pp. 499-504.

Figures and Tables

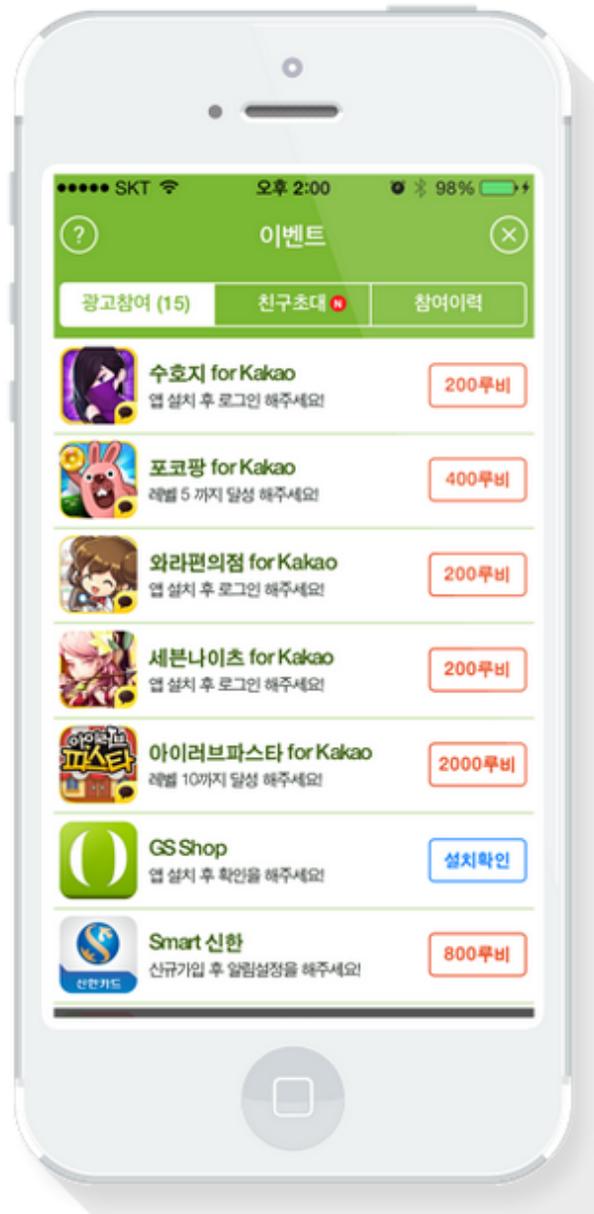


Figure 1: Cross promotion screenshot

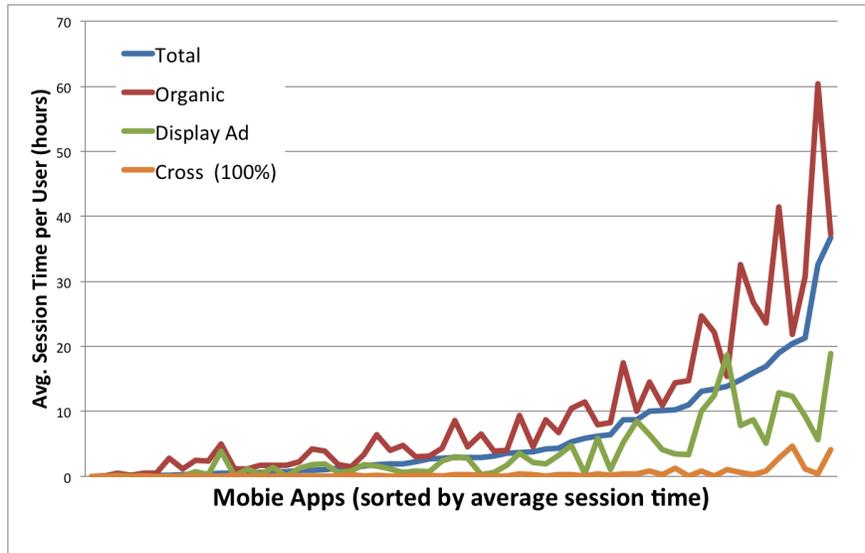


Figure 2: Comparison on average session time by user acquisition channels.

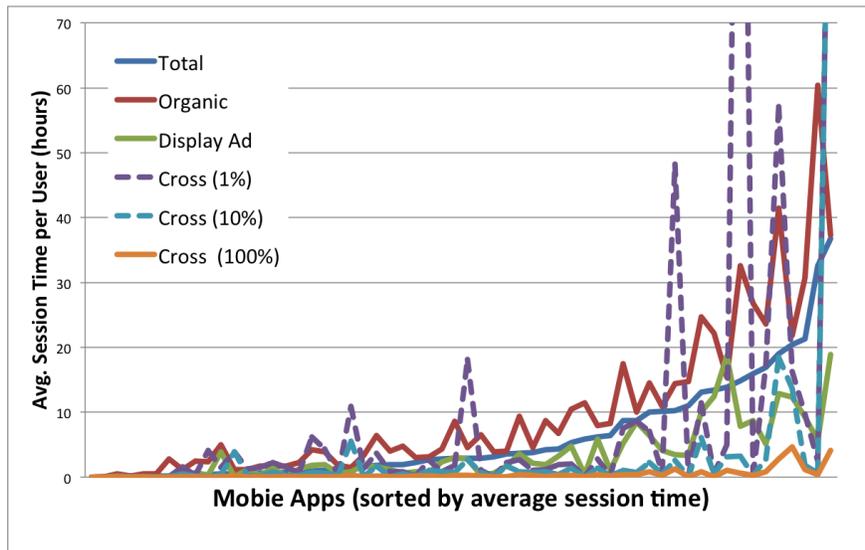


Figure 3: Comparison on average session time app usage by acquisition channels with two additional lines on cross promotion.

Topic ID	Keywords
0	piano, sskin, classic, flipfont, sound, symphony, Beethoven, Mozart
1	Naver, Kakaotalk, subway, radio, radion, radic, developer, DMB
3	color ring, background, service provider, copyright, go launcher
6	kids, Cocomong, animation, hearts, master, fun, Cocomong2
8	icon, Hello Kitty, atom, screen, game, cute, diary
10	LTE, contract, content, SK Telecom, SKT, promotion, free call
11	hymn, copyright, bible, the Lord's prayer, the Apostle's Creed, Ten testaments
13	series, galaxy, final, system, fantasy, wifi, network, player
14	friends, facebook, play, graphics, developers, upgrade, twitter
16	car, racing, simulation, parking, bicycle, place, driver, graphic, simulator
18	point, gift card, reference, cookie run, content, kakaotalk, convenience store
19	Chinese, maker, content, foreign language, kids, Korean, HSK, mp3
25	camera, image, frame, emoticon, sticker, gallery, twitter, facebook, email, friends
27	music island, epilus, mr karaoke, karaoke, hellip, pop, musicsum, romance, sound
28	lotto, tethering, seller, lottery, lottery number, round
29	social commerce, shopping mall, gifts, brand, style, emart store, category
33	English listening, smart teps, ted, smart, player, vocabulary, movie
34	Pororo, friends, animation, sing, kids, adventure, content
36	what's the number, poweramp, go locker, bull, phone number, dotemu, voice phishing
41	naver, dodol launcer, dodol home, blog, icon, dodol, installation
42	kakao talk, alert, kakao story, passrod, theme, developer, copyright
45	recruiting, job korea, resume, check card, part-time job, saramin, job posting
48	calendar, anniversary, diary, point, day, time management
49	subway, bus stop, guide, public transporation, offline, etips, GPS, restaurant
51	Korean language, Korea, travel, tourism, smart wallet, travel information
53	fortune telling, 2014, love, money, content, new year, health, star sign
56	drama, vod, content, rate, youtube, high resolution, story, animation, streaming
67	NFC, touch, USIM, smart, sd card, app, record
68	diet, calory, receipe, stetching, fitness, trainer, graph, weight
76	sports, baseball, NBA, wordcup, score, KBO, Spain, France, Brazil, Italy
80	book 21, story, series, shw, homepage, email, twitter
81	title, YBM, CNN, TOEIC, YFS, word, Japanese, network, Korean-English
85	mp3, battery, 50 songs, series, recorder, ebooks
89	naver, blog, post, mail, diary, NHN, content, navercc
93	Korean, Spanish, Chinese, French, German, Japanesse, Italian, Russian
96	blind date, date, ideal, profile, social dating, random chatting, single, people
97	wall paper, 7days, subway, love, image, background image
99	voca, megabox, vocabulary bible, traffic information, text, highway

Table 1: Subset of Topic Modeling Results (100 topics). Korean keywords translated into English.

User session time of target apps (minutes)				
	(1)	(2)	(3)	(4)
Topic_Similarity (0~1)	25.4915*** ($<2e-16$)	5.801e+01*** ($<2e-16$)	54.846372*** ($<2e-16$)	6.116e+01*** ($<2e-16$)
Num_Rates_Source		1.538e-02*** (0.000128)		2.778e-03 (0.7313)
Num_Rates_Target		-1.302e-03 (0.268803)		-2.218e-03* (0.0625)
Avg_Rate_Source (1~5)		1.689e+01*** ($<2e-16$)		2.280e+01*** ($<2e-16$)
Avg_Rate_Target (1~5)		1.162e+01*** (4.44e-05)		1.434e+01*** (7.34e-07)
Days_Regist_Source			-0.087131*** ($< 2e-16$)	2.156e-02 (0.1919)
Days_Regist_Target			0.073222*** ($< 2e-16$)	6.567e-02*** ($< 2e-16$)
Days_Update_Source			0.074570*** (0.00919)	-3.822e-02 (0.2231)
Days_Update_Target			-0.230001*** (4.14e-13)	-2.405e-01*** (4.29e-14)
File_Size_Source			-0.108862 (0.12014)	-5.627e-01*** (1.12e-09)
File_Size_Target			0.253022*** ($< 2e-16$)	2.338e-01*** ($< 2e-16$)
Intercept	15.1479*** ($<2e-16$)	-1.117e+02*** ($< 2e-16$)	8.535493** (0.28598)	-1.588e+02*** ($< 2e-16$)
Observations	310,183	310,183	310,183	310,183

Table 2: Multivariate linear regression results on user session time.

a

^aNote: This table shows the estimation result on ad effectiveness in an app match. Results show that the effect of app similarity is significantly positive. * indicates statistical significance at the 10% level, ** at the 5% percent level, and *** at the 1% level.

	User connection count of target apps			
	(1)	(2)	(3)	(4)
Topic_Similarity (0~1)	5.1517*** ($<2e-16$)	9.255e+00*** ($<2e-16$)	8.018898*** ($<2e-16$)	8.939e+00*** ($<2e-16$)
Num_Rates_Source		-1.393e-03* (0.0525)		2.271e-03 (0.116667)
Num_Rates_Target		-4.128e-04** (0.0500)		-3.627e-04* (0.088721)
Avg_Rate_Source (1~5)		3.134e+00*** ($<2e-16$)		3.650e+00*** ($<2e-16$)
Avg_Rate_Target (1~5)		3.999e+00*** (3.94e-15)		4.259e+00*** ($<2e-16$)
Days_Regist_Source			-0.008145*** (5.65e-07)	1.146e-02*** (0.000107)
Days_Regist_Target			0.006517*** (2.38e-09)	6.057e-03*** (7.60e-08)
Days_Update_Source			0.029420*** (9.27e-09)	9.692e-03* (0.084234)
Days_Update_Target			-0.053952*** ($<2e-16$)	-5.477e-02*** ($<2e-16$)
File_Size_Source			0.091901*** (2.26e-13)	3.489e-02** (0.034796)
File_Size_Target			0.022574*** (3.20e-06)	2.057e-02*** (2.41e-05)
Intercept	4.0839*** ($<2e-16$)	-2.586e+01*** ($<2e-16$)	2.164251*** (0.00144)	-3.417e+01*** ($<2e-16$)
Observations	310,183	310,183	310,183	310,183

Table 3: Multivariate linear regression results on user connection count.

a

^aNote: This table shows the estimation result on ad effectiveness in an app match. Results show that the effect of app similarity is significantly positive. * indicates statistical significance at the 10% level, ** at the 5% percent level, and *** at the 1% level.