Marketing Science Institute Working Paper Series 2016
Report No. 16-112

# An Integrated Analysis of Mobile Application Usage and In-App Advertising Response

Liye Ma and Baohong Sun

**Report Summary**

Mobile applications account for 86% of consumers' time spent on mobile devices, and mobile advertising, projected to exceed $100 billion in revenue in 2016, is increasingly shifting to in-app advertising. To successfully engage with consumers, marketers need to understand more about their mobile application usage and in-app advertising response across time and locations.

In this study, Liye Ma and Baohong Sun develop an integrated model of consumers' mobile app usage and advertising response. Their study is the first to link consumers' usage of mobile applications with response to in-app advertisements, and the first to evaluate the effects of consumers' underlying involvement in different activities and the application context on advertising response in a real world setting.

Their dataset, obtained from a large mobile advertising platform company, includes panel data on both application usage and advertising responses across four categories of mobile applications: information, entertainment, utility, and social. Their dataset includes 424 mobile applications and the impressions and clicks of 14 advertising campaigns.

Overall, their analysis shows that consumers' usage of different types of mobile applications and their propensity to click in-app advertisements vary significantly over time. Involvement in information activities is more pronounced in the morning, while involvement in utility and entertainment activities peak later in the day, and involvement in entertainment activities persists into evening. Involvement in social activities varies less dramatically over time, with a double-peak in the morning and early afternoon, remaining stable into early evening.

Further, earlier usage of entertainment, information, and utility applications leads to reduced use later on, while earlier use of social applications leads to higher use later.

On advertising response, repeated delivery of a product trial advertisement reduces the probability to click, while repeated delivery of a promotion advertisement increases the click probability.

Their analysis also shows that consumers' ad clicks are related to not only the specific application in which the ad is displayed, but also to their usage of other applications at around the same time. Higher involvement in entertainment activities significantly reduces a consumer's interest in advertisements. Higher involvement in utility and information activities also reduces click propensity, although to lesser extents. Higher involvement in social activities, in contrast, increases a consumer's interest in advertisements.

At the same time, estimates of the contextual effect show that information applications provide the most favorable context for clicking advertisements, while social applications are least conducive to clicks. This contrast between the contextual effect and the effect of underlying involvement underscores the importance of analyzing application usage and in-app advertising response in an integrated framework.

Through simulation, the authors show that targeted ad delivery strategies derived from their model yield significantly higher click-through rates than the benchmark strategy, and the advantage is greater for lower ad impression quotas which require more precise targeting. The targeting strategies deliver more ad impressions in times that are more conducive to clicks, thus better aligning the two than does the benchmark strategy. Ad impressions delivered based on inference of individual consumer level involvements are also shown to be more effective than those based on population level estimates or only time effects, suggesting there is much potential to individual consumer based targeting.

*Liye Ma is Assistant Professor of Marketing, Robert H. Smith School of Business, University of Maryland. Baohong Sun is Dean's Distinguished Chair Professor of Marketing, Cheung Kong Graduate School of Business.*

## Introduction

Consumer activities are rapidly shifting to mobile devices. With more than two billion smartphone users worldwide, mobile advertising spending is projected to exceed $100 billion in 2016 (eMarketer 2015). Firms are quickly ramping up their capabilities to engage consumers on mobile platforms, on advertising, sales, service, and many other functions. Recently, consumer activities on mobile devices have been shifting from mobile web to *mobile applications*, with more than 100 billion mobile application downloads in 2014 and growing (Forbes 2014a).[1] In 2013, mobile applications account for 80% of the time consumers spent on mobile devices, and the proportion grew to 86% in 2014 (TechCrunch 2014). Accompanying this trend, advertising expenditure on mobile devices is also shifting to mobile apps from mobile web. *Mobile in-app advertising* spending is projected to almost triple mobile web advertising spending in 2016, accounting for more than half of overall mobile application revenues (VentureBeat 2015).[2] This shift seems well justified, as the click-through rate of mobile in-app advertisements is shown to be significantly higher than that of mobile web advertisements (Forbes 2014b).

Understanding consumers' mobile application usage and their in-app advertising response thus becomes imperative for successfully engaging with consumers on this new platform. However, while a rapidly growing literature has begun to answer many important questions about mobile consumers (Shankar et al. 2010, Ghose and Han 2011, Luo et al. 2013, Bart et al. 2014, Andrews et al. 2015, etc), little is yet known about mobile application usage and in-app advertising response. Adding to the challenge is that unlike traditional media channels such as TV or desktop computer, mobile device has the distinguishing feature of ubiquity in time and location. With a smartphone in pocket, and countless mobile applications providing all kinds of functionalities a few taps away, a consumer can use mobile applications virtually anytime, anywhere, and for any purpose, on a 24/7 basis. A consumer may set up schedules using a calendar application during morning rush, check stock prices using a financial application at lunch, and play casual games in the evening. In addition, the ease with which consumers can switch between multiple applications makes multi-tasking commonplace on mobile devices.

[1] Mobile applications are software programs that run on mobile devices that provide certain functionalities, such as map, casual game, news, ebook, etc. A large variety of mobile applications can be readily installed from online app stores.
[2] A mobile in-app advertisement is an advertisement displayed when a consumer is using a mobile application, usually at the top or bottom region of the mobile phone screen.

Juggling between several applications, a consumer may check for news update periodically, engage in a conversation with a friend back and forth, and read an ebook, all within the block of time. In other words, consumers are likely involved in various activities on mobile phones simultaneously, with such involvements evolving dynamically over time and situation. Understanding consumers' usage of different applications over time is thus not a trivial task. However, while preliminary analyses in the industry have confirmed the complexity of mobile application usage (Salesforce 2014, Rosenstein 2015), little in-depth research exists to date on this behavior. The ubiquitous nature of application usage poses further challenge to advertisers. While the large amount of time consumers spend on mobile applications in many settings affords advertisers ample opportunities to access consumers, consumers may not be equally interested in advertisements in all these occasions. A consumer who is playing a game on mobile phone at night, for example, may be more interested in an advertisement than a consumer who is balancing her account book while rushing to work. Literature has shown that advertising response is a complex phenomenon, depending both on consumers' underlying involvements and on the context in which the advertisement is shown (Krugman 1965, Zaichkowsky 1985, Park and Young 1986, Pavelchak et al. 1988, Kamins et al. 1991, Howard and Barry 1994, Sharma 2000, Lord et al. 2001, etc). Recent field experiment shows that consumers in a crowded environment are more likely to respond to SMS ads, which confirms the importance of understanding mobile consumers' context (Andrews et al. 2015). Little systematic knowledge exists, though, on how consumers' involvement in different activities affects their response to mobile in-app advertisements in the real world. Many questions thus remain open: 1. How does consumers' involvement in different activities drive their use of mobile applications, and how does the involvement evolve over time? 2. How does consumers' involvement in various types of activities affect their interest in mobile advertisements? 3. How does the context of mobile applications affect consumer responses to mobile advertisements? 4. How do consumers respond to repeated deliveries of different advertisements? 5. How should firms perform targeted advertising delivery to optimize the result? All these questions are crucial for firms to successfully conduct mobile marketing operations, and are the focus of this study.

In this study, we develop an integrated model of consumer's usage of mobile applications and response to mobile in-app advertisements. Drawing from the activity consumption literature, consumers' usage of mobile applications is driven by their underlying involvements in different types of activity: *entertainment*, *information*, *utility*, and *social*. To account for the ubiquitous and multitasking nature of mobile application usage, the model explicitly incorporates consumer's underlying involvement in multiple types of activity simultaneously, and accounts for the dynamic evolution of these involvement levels over the time of day. Meanwhile, drawing from the advertising response literature, our model also accounts for the effect of these involvement levels on consumer's decisions to click mobile in-app advertisements. The involvement levels thus connect application usage and advertising response in a unified framework. Furthermore, the model also incorporates the contextual effect of mobile applications on advertisements, and the sequential effect through repeated exposure to ad impressions. Potential endogeneity concerns are addressed in the model through explicitly modeling targeted delivery of mobile advertisements. The key feature that enables model identification, especially the distinction between the effect of involvements and the contextual effect, is the multitasking nature of mobile devices, as consumers routinely use different mobile applications concurrently. The integration of application usage and advertising response through the underlying involvements is especially important, since it both provides deeper insights on consumers' response to mobile advertisements and enables the creation of effective ad targeting strategies based on mobile application usage.

We estimate our model using a unique dataset obtained from a large mobile advertising platform company. Coming from the platform instead of an individual application provider or advertiser, the dataset has 24/7 coverage of both consumers' usage of a large set of mobile applications, and the impressions and clicks of advertisements delivered in those applications. This comprehensive coverage makes the dataset well suited for our study. Rich variations over time, applications, advertisements, and consumers, are present in the data. Specifically, the data shows that consumers' usage of different types of mobile applications and their propensity to click in-app advertisements vary significantly over time. Furthermore, model free data patterns show that the clicks of in-app advertisements are related not only to the specific applications in which such ads are displayed, but also to the other applications the consumers are using at the same time. Our model setup is thus not just grounded in existing literature, but empirically justified by the data patterns as well.

Model estimates reveal a strong temporal pattern in consumers' involvements in different types of activities throughout a day. While involvement in Information activities is more pronounced in the morning, involvement in Utility and Entertainment activities peak later in the day, and the latter persists into evening. Involvement in Social activities, in contrast, varies less dramatically over time, with a double-peak in the morning and early afternoon, remaining stable into early evening. Meanwhile, we find that the usage of Entertainment, Information and Utility applications are inter-temporal substitutes, i.e. earlier use of the applications leads to reduced use later on, while the usage of Social application is inter-temporal complement, where earlier use leads to higher use later. Model estimates also show that the involvement levels in Entertainment, Utility, and Information activities are highly persistent. On advertising response, we find that repeated delivery of a product trial advertisement reduces the probability to click, while that of a promotion advertisement increases the click probability. More importantly, we find that higher involvement level in Entertainment activities significantly reduces a consumer’s interest in advertisements. Higher involvement in Utility and Information activities also reduce click propensity, although to lesser extents. Higher involvement in Social activities, in contrast, increases a consumer’s interest in advertisements. Different from the effect of involvement, estimates of the contextual effect show that Information applications provide the most favorable context for clicking advertisements, while Social applications are least conducive to clicks. The contrast between the contextual effect and the effect of underlying involvement is particularly notable, and underscores the importance of analyzing application usage and in-app advertising response in an integrated framework. In addition to adding to the knowledge of this growing phenomenon, these findings are also crucial for managers to devise ad targeting strategies, and our simulations show that targeted advertisement delivery strategies based on our model generate significantly higher click-through rates than do benchmark strategies.

We contribute to the literature in the following ways. First, we are the first to provide an integrated framework for jointly modeling consumers’ usage of mobile applications and response to mobile in-app advertisements. We show that consumers’ underlying involvement is a key driver of both decisions, and it is important to analyze these two actions together rather than in isolation. Second, on mobile advertising response, we are the first to distinguish the effect of consumers’ underlying involvement in different activities from the contextual effect of the mobile applications in a real world setting. Third, we are among the first to investigate the dynamic evolution of consumer’s involvement in different activities over time of day. This provides insights into how consumers shift their focus over time to fulfill different needs, and advances our understanding of consumers' complex application usage behavior on the mobile platform. Finally, we show the managerial implications of the knowledge on mobile application usage and in-app advertising response, by demonstrating their potential for improving ad targeting effectiveness. Marketing research on consumer's use of mobile devices is still at an early stage. All these advance our understanding of this important phenomenon.

## Literature Review

Our study falls into the small but rapidly growing literature on mobile marketing. Banerjee and Dholakia (2008) study the effect of location based advertising. Shankar et al. (2010) provide an early summary of and discuss opportunities for mobile marketing research in the retailing environment. Ghose and Han (2011) investigate content generation and consumption on mobile Internet. Ghose et al. (2013) compare the internet browsing activities on mobile web with those on personal computers. Luo et al. (2013) investigate the effect of temporal and geographical mobile targeting through a field experiment. Bart et al. (2014) investigate what type of products benefit most from mobile display advertising. Andrews et al. (2015) analyze how the crowdedness of the environment affects consumer's response to mobile advertisements. These studies have expanded our understanding of mobile marketing. However, to date no studies have investigated the effect of consumers' underlying involvements in different activities on their responses to mobile in-app ads, nor the contextual effect of the applications in which such ads are shown. Furthermore, existing studies also have not studied consumers' usage of mobile applications and the change of the usage over time. Our study bridges this gap, which is particularly important given the increasing market share of both mobile applications and mobile in-app advertisements. From another perspective, the distinctive feature of mobile is time and location ubiquity. While existing literature provides valuable insights on the location dimension (Luo et al. 2013, Andrews et al. 2015), relatively less is known about the time dimension, a gap this study fills.

Our modeling of consumer's usage of mobile applications and response to mobile in-app advertisements draws from the rich literature of advertising response, and the literature on activity consumption and time use. Consumers spend a great deal of time on a variety of activities throughout the day, and the use of time has been studied in different fields (Jacoby et al. 1976). The allocation of time on different activities has been modeled as driven by different underlying needs (Bhat 2005, Kamakura 2009, Luo et al. 2013, Lin et al. 2013). Furthermore, studies also show that consumers' states of mind change throughout the day, with important implications on the activities they engage in and decisions they make (Yoon et al. 2007, Danziger et al. 2011). A large variety of mobile applications exist to satisfy consumers' different needs, and consumers are involved in using these applications throughout a day. Drawing from the activity consumption and time use literature, we model the usage of different categories of mobile applications as driven by consumers' involvements in different type of activities, and model the dynamic evolution of such involvements.

The effect of advertising, a key topic of marketing research, has been studied extensively from different perspectives. Our study is closely related to individual consumer's response to advertisements and to Internet and mobile advertising. Numerous studies have shown that consumer's underlying state of mind, such as level of involvement, mood, emotion, etc. has significant implications on the response to commercials (Krugman 1965, Clancy and Kweskin 1971, Krugman 1983, Zaichkowsky 1985, Park and Mittal 1985, Park and Young 1986, Goldberg and Gorn 1987, Kamins et al. 1991). Meanwhile, the effect of contextual factors, particularly the program-commercial congruity for TV commercials, has also been investigated, with certain studies explaining the contextual effects by linking the program and commercial through the former's effect on the underlying emotions, which in turn affects the latter (Pavelchak et al. 1988, Howard and Barry 1994, Sharma 2000, Lord et al. 2001, Cho 2003). Although these studies show that both the underlying involvement and the context affect consumer's advertising response, they differ in their findings on the effect of such factors. These studies also do not apply directly to the real world setting of mobile marketing. With the growth of Internet advertising, studies have also analyzed the effect of such advertisements from perspectives such as the context congruity (Moore et al, 2005), repeated exposures (Manchanda et al. 2006), and multiple creatives (Braun and Moe 2013). Recent studies have also started to investigate the effect of mobile advertising, focusing on how the ad effectiveness depends on product types and location (Luo et al. 2013, Bart et al. 2014, Andrews et al. 2015). Drawing from this literature, in our model consumers' responses to mobile in-app ads are driven by both their underlying involvements in different types of activities and the context of the mobile applications, while we also account for the effect of repeat deliveries.

## Data

### Data overview

The data used in this study is obtained from a major advertising platform company in an Asian country. The platform company runs an advertising engine that delivers in-app advertisements to applications on mobile phones. The industry setup and the technical aspects of in-app advertising delivery are discussed in Technical Appendix 1. The dataset is a rich panel dataset which contains the application usage and advertising response information of 3,988 randomly selected mobile phone users over a 7-day period, from June 25$^{th}$ 2012 to July 1$^{st}$ 2012. Usage information of 424 mobile applications and the impressions and clicks of 14 advertising campaigns are included. Both types of information have precise timestamps. Specifically, the data contains individual records of ad impressions and responses. Each record contains the following information: the time of the ad impression, the unique identifier of the mobile device for the ad impression, the mobile application the user is using (in which the ad impression is displayed), whether an ad impression is successfully delivered, the identifier of the advertisement that is displayed, and whether the user clicked on the advertisement. Coming from the platform company, the dataset covers multiple advertising campaigns and a large set of mobile applications, on a 24/7 basis. In addition to the mobile in-app ad information, the data contains detailed application usage information throughout the day, which is the key to understanding the evolution of consumers' simultaneous involvements in various activities. Uncovering these underlying involvement levels, and analyzing how they affect consumers' application usage and advertising response, are the focus of our study. The dataset does have a weakness, in that it does not contain subsequent website visitation and conversion information after ad clicks. Consequently, in our evaluation of advertising response we focus on the click-through behavior, instead of the eventual purchase decisions.

### Descriptive statistics

The advertising platform company classified the mobile applications in the dataset into eight different categories: Game, IT/Digital, Entertainment, News, Finance, Sports, Social, and Fashion. Following the typology of Gupta (2013), we further consolidate these into four categories: *Entertainment*, *Utility*, *Information*, and *Social*.[3] Application usage is measured in the number of 20-second segments during which the application is open.

Table 1 reports the per-user mobile applications usage information by category. The table shows that Entertainment is the category with highest average amount of usage, followed by Utility and Information, while the category Social has the lowest amount of usage.[4] For all categories, the usage across users is positively skewed, with both the standard deviation and the maximum much larger than the mean. Overall, this shows a group of active mobile users with considerable diversity in their usage of mobile applications both across users and across different categories.

Application usage varies significantly over time. Figures 1 reveals a clear time-of-day pattern of application usage: the usage in early morning is fairly low; it starts to ramp up quickly after 7AM; it reaches a active level after 9AM, and remains roughly at that level into evening and night; after midnight the activity level drops quickly. Around noon time there is also a spike in the usage, which recedes after noon. This pattern corresponds well with people's normal daily routines. While this temporal pattern is generally shared across categories, each category also has its own characteristics. The usage of Utility and Information applications, for example, peaks earlier in the day than Entertainment and Social applications; the usage of Information applications drops steadily in the afternoon into early evening; finally, there is an uptick in the usage of Entertainment and Information applications in late night, while the usage of the Utility and Social applications remain stable. Taken together, the application usage data shows rich heterogeneity in application usage across users and time, both in general and at individual category level. The data gives us the opportunity to understand how consumers' underlying involvement in different activities evolves throughout a day.

On the advertising side, the impression and click information of 14 advertising campaigns, coming from a diverse set of industries, are reported in Table 2. The advertising

[3] Gupta (2013) classified mobile apps into five categories: games and entertainment, social networks (such as Facebook), utilities (such as maps, clocks, and calendars), discovery (such as Yelp and TripAdvisor), and brand. Our dataset does not contain brand apps, and our Information category corresponds to the discovery category in Gupta (2013). The mapping from original categories to the new ones is as follows: the original Game, Entertainment, Sports, and Fashion categories are mapped to the new Entertainment category; the original IT/Digital and Finance categories are mapped to the new Utility category; the original News category is mapped to (renamed as) the new Information category; the Social category remains the same.

[4] We note that the dataset was collected in 2012, and recent data may show higher social activity level on mobile applications.

company further classified the advertising campaigns into three types: Promotion, Product Trial, and Product Launch, a classification which we follow in this study. Ad campaign 5 has the highest number of impressions (336,744), and ad campaign 13 has the least impressions (5,454). These ad campaigns also have different click-through rates, ranging from 0.34% (ad campaign 9) to 2.34% (ad campaign 2). This large variation across ad campaigns on impressions and clicks provides rich information for investigating consumers' advertising response behavior.

On mobile devices, consumers are routinely shown the same advertisements repeatedly. This is also true in our dataset. The data demonstrates a sequential pattern of click-throughs: When a user sees an advertisement for the first time, she has relatively high probability of clicking on it (above 3%). As the same advertisement is repeatedly delivered to the user, however, the click probability decreases gradually.

**Model-free patterns**

We now discuss two notable data patterns that both inform and motivate our model and analysis. The first pattern shows that among significant variations of application usage and ad clicks over time, there is strong evidence that the current ad delivery practice is suboptimal. Figure 2 plots the statistics of ad impression requests, delivered ad impressions, and clicks, for each hour of day. The *requests* line shows the total number of times the mobile phones sought delivery of ad impressions. This represents the maximum numbers of ad impressions that could be displayed, as determined by the amount of time users spend on using mobile applications. The *impressions* line shows the total number of actual ad impressions served, while the *delivery failures* line shows the number of times an ad impression was not displayed, even though the users were using the applications so ads can be served. Finally, the *CTR* line shows the average click-through-rate for each hour. As the figure shows, although almost a million impressions were delivered, the advertising engine was actually not running at full capacity, especially later in a day. In the hour immediately after midnight, most ad impression requests were successfully filled (in about 89% of the time ad impressions were displayed, while 11% of ad impression requests were unfilled). During morning and noon times, 58% to 75% of ad impression requests were filled, also fairly high although lower than the early morning hour. Starting from early afternoon, though, the percentage of unfilled ad impression requests increased. And in the evening, around 90% of the ad impression requests were left unfilled. In other words, in the evening hours, in around 90% of the time when users were using mobile applications, the display ad areas were left empty instead of showing ad impressions. According to the advertising platform company, this was caused by two factors. First, the company did not get enough ad campaign purchases to fill all the advertising slots at the mobile applications. Second, product firms often ask the advertising company to display a target number of impressions per day, and ask the advertising company to display the impressions as soon as possible. To illustrate, if a client purchased 10,000 impressions every day for its ad, and all impressions are displayed by 2pm, then subsequent ad impression requests from mobile applications will be left unfilled. These two factors combined lead to large vacancies in the evening.

The click-through rates (CTRs) also vary greatly over time, as the figure shows. CTRs are highest around noon. The number of ad impressions displayed is also fairly high for this time. This suggests that the advertising company captures the most effective time slots rather well. However, closer examination suggests the current practice is still suboptimal. For example, CTR is consistently lower in early morning than in the evening, yet a much larger portion of ad impression requests in the former time slots are filled than the latter. The hour after midnight has 89% of ad requests filled, yet these ad impressions have a CTR of only 0.55%. Compared with that, the CTRs during evening hours are much higher, but only about 10% of ad impression requests in those hours were filled. The last hour of day has a click-through rate of 1.49%, approaching the midday levels, yet less than 6% of ad impression requests in that hour were filled. In other words, over 90% of opportunities to display ad impressions in the evening were missed, when they could have led to more clicks, while many more ad impressions were served in other times slots with lower click-throughs. Taken together, this temporal pattern of ad request, delivery, and response confirms the rich dynamics in consumers' advertising response throughout the day, and shows that the current ad impression delivery practice leaves a lot of room for improvement. In-depth analysis of response to mobile in-app advertisements thus is both academically interesting and practically important.

While the first data pattern focuses on the time dimension, the second data pattern focuses on the multitasking nature of mobile usage. This data pattern shows that consumers' ad clicks are related to not only the specific application in which the ad is displayed, but also their usage of other applications around the same time. Table 3 reports how the CTRs are related to the applications the consumers are using at the time. The first column shows the average CTRs of advertisements displayed in each category of applications. The average CTR of advertisements shown in Utility applications is 1.14%, while the CTR is higher for Information and Social applications, at 1.81% and 1.96%, respectively. This shows that the application context has significant implications on consumers' click-through behaviors. Furthermore, the next four columns report the average CTRs, also for the advertisements displayed in each category of applications, but conditional on the consumer having above-average use of another category of applications in the hour the advertisement is displayed.[5] These statistics reveal additional complexities in consumers' click-through behavior. For example, if a consumer has been using Utility applications heavily (above-average) in a specific hour, then the CTR of advertisements is much lower, even for those advertisements *displayed in other categories of applications* in that hour. Similarly, the CTR of advertisements is also lower when the consumer has been using Entertainment applications heavily. In contrast, the CTR is much higher if the consumer has been using Social applications heavily, regardless of the category of the specific application in which the advertisement is displayed. These statistics show that both the consumers' underlying levels of involvement in different types of activities, and the context of the application in which an advertisement is delivered, are key drivers of advertising response, and these are two separate factors with distinct impacts.

These preliminary data analyses reveal rich variations of consumers' usage of mobile applications and response to mobile in-app advertisements, across category and time. Meanwhile, it points to the complex and close connections between these two actions. For in-depth analysis, we next set up the formal model and discuss the empirical results. The whole dataset contains 3,988 consumers. Since the focus of the study is to analyze consumers' mobile application usage across different categories and advertising response under those varied circumstances, we used the subset of consumers who used applications in at least two categories. This subset contains 571 consumers. There are seven days in the dataset. We use the first six days for calibration and the last day as hold-out sample.

[5] That is, the consumer's usage of the applications of the specific category at that hour is higher than the average hourly usage rate of that category of applications.

## Model

### Conceptual framework

Our model consists of two major components, as depicted in the conceptual framework in Figure 3. The first component characterizes the usage of mobile applications, while the second models the decisions to click the in-app advertisements displayed when users use mobile applications. Our modeling of mobile application usage is grounded in the activity consumption and time use literature, while that of clicking of in-app advertisements is grounded in the advertising response literature. At the core of the model, connecting both components, is consumer's *underlying involvement* in different types of activities. Following the activity consumption literature, we model consumers' *usage of mobile applications* as driven by these underlying involvement levels. Meanwhile, the advertising response literature shows that consumer's involvement is a key driver of response to advertisements. Data patterns discussed in the previous section also show that consumer's click propensity is related to both the specific application in which an ad is displayed, and other applications the consumer is using at the same time. This points to the necessity of accounting for the effect of involvement levels on advertising response. Thus grounded in the literature and motivated by the data pattern, in our model consumers' *response to in-app advertisements* is also driven by these underlying involvement levels. The underlying involvement effectively unifies both the application usage and the ad response activities.

Consumers use mobile applications throughout different times of day, at different locations, and for different purposes. Involvements in different activities would change depending on time and circumstances. For example, a consumer might be more involved in mobile applications in general when she is not at work. As another example, a consumer might be more involved in information gathering activities earlier in the day, while more involved in entertainment later in the evening. The ubiquity and the multitasking nature of mobile application usage make it necessary to model consumer's involvements in different types of activities simultaneously, and to account for their change over time. Considering this, our model incorporates the dynamic evolution of involvement levels in a flexible manner, accounting for both time specific effects and persistence of involvement levels for individual consumers.

The usage of mobile applications may also depend on history. For example, if a consumer uses an application to set up schedules in the morning, then she might not use the application in the afternoon, as the task is already complete. Accordingly, we model consumers' application usage as also depending on their previous usage. Meanwhile, our model also accounts for the contextual effect of the applications in which advertisements are displayed, as the advertising response literature shows that such context also has important effects. Finally, consumers' response to advertisements is likely dependent on the history of ad exposures, and there may also be direct time effects on ad response. When analyzing the involvement and contextual effects, it is important to control for these additional factors. Both are also incorporated in our model.

Formally, there are $I$ consumers, or mobile phone users, each indexed by $i$, $i = 1, \dots, I$. Consumers use mobile applications and click in-app advertisements in a period of $D$ days, indexed by $d, d = 1, \dots, D$. We partition each day into $T$ equal-length time intervals, and denote the time of day using $t, t = 1, \dots, T$. For our study, each time period is one hour, and there are 24 such time periods in a day, i.e. $T = 24$.[6]

**Application usage**

In any time period, a consumer may use one or more categories of mobile applications, driven by her underlying involvement levels in the corresponding types of activities. Following the typology of Gupta (2013), we account for consumer's involvements in four types of activities: Entertainment, Utility, Information, and Social, which drive the usage of mobile applications of the corresponding categories. Thus there are $K = 4$ categories of applications, each indexed by $k$.[7] Since our focus is on consumers' involvements in different types of activities, we focus only on the applications' categories, but not on the identities of individual applications. The multitasking nature of mobile application usage dictates us to account for the involvements in different types of activities simultaneously, and for the usage of multiple applications at the same time. Accordingly, we denote consumer $i$'s amount of application usage at time $t$ of day $d$ as:

$$\vec{x}_{idt} = (x_{idt,0}, x_{idt,1}, \dots, x_{idt,K}) \tag{1}$$

[6] We also examined 15-minute time periods, which shows similar data patterns.
[7] We note that the model structure is general enough to account for different classifications of activities.

In the equation, $x_{idt,k}$ is usage amount of category $k$ applications, in the number of time units that are distinguishable from the data (e.g. a second or a minute). The first element, $x_{idt,0}$, denotes the amount of usage of outside activities, i.e. when the consumer is not using mobile applications. The total usage amount of all applications, combined with the outside activity, is a constant that equals the number of time units a time period has, $\sum_{k=0}^{K} x_{idt,k} = C$, where $C$ is the total number of time units. In our data, the use of an application is recorded every 20 seconds, so each time unit represents 20 seconds. In other words, $x_{idt,k}$ represents the number of 20-second units the consumer used applications of category $k$ in the time period. The total number of time units in each time period is $C = 180$.

Application usage is thus represented using multi-dimensional count data, which we naturally model as following a multinomial distribution:

$$\vec{x}_{idt} \sim Multinomial(C; \vec{p}_{idt}) \tag{2}$$

In the equation, $\vec{p}_{idt} = (p_{idt,0}, p_{idt,1}, \dots, p_{idt,K})$ is the probability vector of each time unit being spent on mobile applications of different categories. We model the probabilities using the standard multinomial logit form:

$$\begin{cases} p_{idt,0} = \dfrac{1}{\sum_{k'=0}^{K} \exp\left(\bar{U}_{idt,k'}\right)} \\ p_{idt,k} = \dfrac{\exp\left(\bar{U}_{idt,k}\right)}{\sum_{k'=0}^{K} \exp\left(\bar{U}_{idt,k'}\right)}, k = 1, \dots, K \end{cases} \tag{3}$$

In equation (3), $\bar{U}_{idt,k}$ is the mean latent utility of using a category $k$ application ($\bar{U}_{idt,0} = 0$ for identification), which we now detail. Application usage is driven by and reflects a consumer's underlying levels of involvement in different activities. Throughout the course of a day, these involvement levels would change depending on time and circumstances. Accordingly, we model the mean latent utility of consumer $i$ using application of category $k$ at time $t$ of day $d$ as:

$$\bar{U}_{idt,k} = \beta_{idtk0} + \beta_{k1} \ln\left(H_{idt,k} + 1\right) \tag{4}$$

In the equation, $\beta_{idtk0}$ represents consumer $i$'s underlying involvement in category $k$ activities at time $t$ of day $d$. A higher $\beta_{idtk0}$ indicates that the consumer is more involved in activities of type $k$, and derives higher utility from using the applications of this category.

Uncovering this underlying involvement level is a key focus of our study. The modeling of the dynamic evolution of these underlying involvement levels, which allows for salient time-of-day patterns and is also flexible enough to admit inter-temporal dependence, is discussed in detail in the next subsection.

The variable $H_{idt,k} = \sum_{\tau=1}^{t-1} x_{id\tau,k}$ is the cumulative amount of usage of category $k$ applications since the beginning of the day, and $\beta_{k1}$ is the corresponding coefficient.[8] This term captures potential inter-temporal substitution or complementarity. A positive $\beta_{k1}$ means that the more a consumer uses category $k$ application earlier in a day, the more she will use it later, i.e., usage of this category of applications is inter-temporally complementary. In contrast, a negative $\beta_{k1}$ means there is an inter-temporal substitution for the category, as higher earlier usage leads to lower later usage. The coefficients are specific to each category, as different application categories may not have the same degree of complementarity or substitution.

**Evolution of involvement levels**

Throughout the course of a day, consumers' involvements in different activities evolve simultaneously. Flexible modeling of this dynamic evolution is necessary to provide insight into the underlying driver of consumers' mobile application usage. The evolution of consumers' involvement levels depends on both history and time. For example, a consumer who is highly involved in Entertainment activities at a time, e.g. playing a casual game, is also likely involved in the activity in the next time period. Different activities may also be more salient at different times (Lin et al. 2013). For example, the involvement in Utility based activities may be higher during work time, while that for Entertainment may be higher in the evening. Accordingly, we model the dynamics of underlying involvement in different activities as follows:

$$\beta_{idtk0} = \beta_{ik} + \delta_{tk} + \tilde{\beta}_{idtk0} \tag{5}$$

$$\tilde{\beta}_{id,t,k0} = \phi_{ik}\tilde{\beta}_{id,t-1,k0} + \varepsilon_{idtk} \tag{6}$$

In Equation (5), $\beta_{ik}$ is the baseline involvement level for consumer $i$ of category $k$, which reflects the general inclination of the consumer to this type of activity. This baseline involvement

[8] We note that our model accounts for the dynamics mostly on an intra-day basis, implying that there is a "reset" at daily level. This helps highlight salient time-of-day patterns, as consumers are expected to use mobile phones to handle many daily routines. Extending the model to account for cross-day dependence will be straightforward, although a dataset covering a longer period of time will be needed for estimation.

level may differ across consumers. To account for the factor that different times of day may be more conducive to different types of activities, we include a time-specific fixed effect term, $\delta_{tk}$, in Equation (5). These time-specific parameters capture any salient time-of-day patterns. A higher value of $\delta_{tk}$ at noon time, for example, would indicate that consumers in general are more highly involved in category $k$ activities at that time.

The term $\tilde{\beta}_{idtk0}$ in Equation (5) is an individual and time-specific term that captures the fluctuation of the involvement level over time at the individual consumer level. The parameter is serially correlated, as shown in Equation (6), to admit persistence over time ($\phi_{ik} \in (-1,1)$). Higher $\phi_{ik}$ indicates higher degree of persistence of the involvement level, which may also be interpreted as this type of activities being addictive. A negative $\phi_{ik}$, on the other hand, would indicate a substitution effect across adjacent time periods. A value of 0 would indicate that the involvement levels across time periods are not related. Finally, $\varepsilon_{idtk} \sim N(0, \sigma_\varepsilon^2)$ is an i.i.d random term that creates the fluctuation of $\tilde{\beta}_{idtk0}$ over time.

**Advertising response**

When a consumer is using a mobile application, an in-app advertisement can be displayed every 20 seconds. An in-app advertisement brings to consumer's attention a certain product or service offering. Should the consumer be interested, she may click the advertisement, which will take her to the corresponding website where more information is shown. There are $J$ advertisements contained in our dataset, each belonging to one of three types: Promotion, Product Trial, and Product Launch. An advertisement is indexed by $j, j = 1, \ldots, J$. An advertisement can be delivered repeatedly to the same consumer over time.

The dependent variable of our advertising response model is the click decision for each impression.[9] We posit that the click behavior reflects consumer's interest in the content of the advertisement. We use a binary logit model, where the click decision is driven by perceived utility of clicking the ad:

$$Prob(Click_{ijn}) = \frac{1}{1 + \exp(\bar{U}_{ijn}^A)} \tag{7}$$

[9] Research shows that in addition to click-through, which is a crucial first step toward purchase, the impression of an ad itself may also have effects in the absence of clicks. Since our dataset does not contain outcome variables such as purchase, effects other than click-through cannot be evaluated, and are left for future research.

In the equation, $Click_{ijn}$ represents the event of consumer $i$ clicking the advertisement $j$ at the $n$-th time of its delivery to the consumer. The term $\bar{U}_{ijn}^{A}$ is the deterministic part of the latent utility that represents the attractiveness of the advertisement to the consumer in that occasion. This utility contains four components: the baseline value of the advertisement to the consumer, the effect of the underlying involvement in different activities on the consumer's interest in advertisement at the time, the sequential effects, and the contextual effect of the mobile application in which the advertisement is displayed. Specifically, the latent utility of consumer $i$ clicking advertisement $j$ the $n$-th time the advertisement is displayed (when the time of the display is $dt_{ijn}$), is:

$$U_{ijn}^{A} = \bar{U}_{ijn}^{A} + \epsilon_{ijn} = \theta_i q_j + \sum_{k=0}^{K} \theta_k^{\beta} \beta_{idt_{ijn}k0} + \sum_{a \in \{S,D\}} \theta_{aT(j)}^{S} A_{ijna} + \sum_{k=0}^{K} \theta_{kT(j)}^{C} C_{ijn}^{k} + \theta_{t(ijn)}^{T} + \epsilon_{ijn} \tag{8}$$

In the equation, $q_j$ is an ad-specific intercept term, which captures the intrinsic value or quality of advertisement $j$.[10] The coefficient $\theta_i$ is the consumer $i$'s intrinsic interest in advertisement. The first term $\theta_i q_j$ is thus the baseline utility of consumer $i$ clicking on ad $j$ irrespective of other factors.

More importantly, in the second term, each $\theta_k^{\beta} \beta_{idt_{ijn}k0}$ captures the effect of consumer's underlying involvement in category $k$ activities on her advertising response. A positive $\theta_k^{\beta}$ means that the more a consumer is involved in activities of category $k$, the more interested she is in advertisements. In contrast, a negative $\theta_k^{\beta}$ indicates that when a consumer is more involved in category $k$ activities, she is less interested in ads. Literature has shown that consumer's involvement significantly affects their response to commercials, and different types of involvement have been analyzed. However, not much is known about how involvements in different activities in the real world setting affect advertising response. It is thus necessary to account for the effect of consumer's underlying involvement levels in the model, and the empirical findings will provide insight in this important connection.

This third term captures the sequential effect of advertising response which is shown to have implications on internet display advertising (Manchanda et al. 2006, Braun and Moe 2013).

[10] Since we do not have detailed information about the advertisement itself, we use this fixed-effect term to capture the intrinsic quality of the advertisement.

In the term, $A_{ijnS} = \ln(N_{ijn}^{j} + 1)$ is the number of times the consumer has been shown the advertisement $j$ for the day, at the $n$-th time the advertisement $j$ is displayed, log-transformed.[11] Similarly, $A_{ijnD} = \ln(N_{ijn}^{-j} + 1)$ is the log-transformed number of times the consumer has been shown other advertisements on the day. $\theta_{ST(j)}^{S}$ and $\theta_{DT(j)}^{S}$ are the corresponding coefficients, where $T(j)$ is the type of advertisement $j$. A positive $\theta_{ST(j)}^{S}$ means there is a positive stocking effect of seeing the same advertisement repeatedly, while a negative value indicates a wear-out effect. A positive $\theta_{DT(j)}^{S}$means seeing other ads has a refreshing or restoration effect on the focal advertisement, while a negative value means seeing other ads distracts the consumer from the focal ad. The coefficients are type-specific, as different types of advertisements may have different sequential properties. For example, a consumer may decide whether to click a casual advertisement at the first look, but may be moved by a more serious advertisement only after several repetitions.

In the fourth term, $\theta_{kT(j)}^{C} C_{ijn}^{k}$ captures the contextual effect of category $k$ application in which the advertisement is displayed. Specifically, $C_{ijn}^{k}$ a dummy variable that indicates whether the application in which the advertisement is displayed is of category $k$. A positive $\theta_{kT(j)}^{C}$ means when a consumer is using the application of category $k$, she will be more interested in an advertisement of type $T(j)$, and will be more likely to click it. An important contrast needs to be made between this contextual effect term and the second term on the effect of involvements. Although both are classified according to the application categories, the second term captures the effect of the consumer's underlying involvement in every type of activities, while this fourth term captures only the contextual effect of this specific instance of ad impression. Both the effect of involvement and the contextual effect are recognized in the literature as affecting advertising response, although they are separate effects. For example, consumers may be more interested in advertisements when their state of mind is socially oriented, even if they do not like to click an ad when they are actively messaging a friend. The two effects can be separately identified in our model because of the multitasking nature of mobile device usage – consumers constantly juggle between various mobiles applications – which allow us to recover the underlying involvements in different types of activities simultaneously. As discussed in the data section, whether

[11] $N_{ijn}^{j}$ can be different than $n - 1$, as it counts only the number of times the ad has been shown on the same day.

consumers click on an ad not only depends on the category of the application in which the ad is delivered, but also is related to the other applications they are using at the same time. Understanding how the underlying involvement and the application context drive advertising response is thus both made possible and necessary.

The fifth term $\theta^T_{t(ijn)}$ controls for direct time fixed-effect of advertising response. Similar to mobile application usage, where different time of day may be conducive to different types of activities, consumers may also respond to mobile in-app ads differently at different time of day. To analyze the involvement and contextual effects, it is important to control for such direct time effect. Finally, the error term $\epsilon_{ijn}$ is assumed to follow an extreme value distribution, which leads to the binary logit probability as shown in equation (7)

**Control for ad targeting**

Marketing response modeling should account for the possibility that the marketing mix variables are not independent from response parameters (Manchanda et al. 2004). In our context of mobile in-app ads, it is possible that advertisers target their ad delivery based on certain knowledge of consumer responses. To the categories of applications or the time periods which are more likely to generate clicks, an advertiser may deliver more ad impressions. To address this endogneity concern, our model explicitly accounts for the targeted delivery of ads. Detailed discussion of the modeling of this component and the corresponding estimation result is provided in Technical Appendix 2.

**Heterogeneity, identification, and estimation**

Consumers will likely differ in their tendency to use different types of mobile applications, and in their interests in in-app advertisements. We account for unobserved heterogeneity in the standard hierarchical Bayesian fashion, by treating every individual-consumer specific coefficient (coefficients which have subscript $i$) as a random draw from the corresponding population level distribution. The model is estimated using MCMC, where the likelihood for application usage is according to equation (2), the likelihood for advertising response is according to equation (7), and the likelihood for ad delivery is according to equation (TA.2-1) in Technical Appendix 2. Identification of the model parameters primarily rests on the temporal and cross-sectional variations of the application usage and advertising response data, and certain parameters need to be normalized. The detailed discussion of identification and normalization of individual parameters is provided in Technical Appendix 3.

## Results

### Model comparison

We first compare our model with a few alternative model configurations. A key aspect of our model is incorporating the effect of consumers' involvements on advertising response. Considering this, the first benchmark model is otherwise identical to the proposed model, except that it excludes the effect of the involvement levels on ad response. In other words, in the ad response utility function, all $\theta_k^\beta$s are constrained to be zero. The second benchmark model excludes both the effect of underlying involvement and the contextual effect (i.e. all $\theta_k^\beta$s and $\theta_{kT(j)}^C$s are constrained to be zero). Finally, in the third benchmark model, the dynamic evolution of involvement levels is also excluded from the application usage equation (i.e. all $\tilde{\beta}_{idtk0}$ are constrained to be zero).

We compare the models based on both in-sample log marginal density (LMD) and the out-of-sample log-likelihood (LL), which are reported in Table 4. Comparing the proposed model with benchmark model 1 which excludes the effect of involvements, the proposed model has both higher in-sample LMD and higher out-of-sample LL. This confirms the importance of accounting for the effect of consumers' underlying involvements in various activities on their ad responses. Benchmark model 1 also slightly outperforms benchmark model 2 on in-sample LMD and out-of-sample LL, suggesting that accounting for mobile applications' contextual effect on ad response also improves model fit. Finally, benchmark model 2 significantly outperforms benchmark model 3 on both measures. This shows that it is crucial to incorporate the dynamic evolutions of consumers' underlying involvements. The dynamics of consumer's involvement in different activities, the effect of such involvements on responses to advertisements, and the contextual effect of mobile applications, are all key model components. The model comparison result confirms the necessity to incorporate all of them. The subsequent discussion of the results is based on the estimates of the proposed model, which account for all these effects.

## Application usage parameter estimates

Table 5 reports the parameter estimates for mobile application usage. The first region reports the population level mean estimates of consumers' baseline involvement levels in different activities ($\beta_{ik}$ in equation 5). Consumers on average are more involved in Entertainment activities on mobile phones (posterior mean estimate of $\bar{\beta}_k$ is -3.985 for Entertainment, highest among the four categories). This is followed by Utility activities, which is in turn slightly higher than and Information activities, although the difference is not statistically significant (posterior means are -7.946 and -8.193).[12] Social activities have the lowest baseline involvement level (posterior mean is -11.769).[13]

The second region of the table reports the effect of application usage history ($\beta_{k1}$ in equation 4). The coefficients for Entertainment, Utility and Information categories are negative and statistically significant. This suggests that these activities are inter-temporal substitutes. This inter-temporal substitution is especially pronounced for the Utility and Information categories (posterior means are -0.167 and -0.123, respectively). This is consistent with the nature of such activities: Consumers likely use Utility applications to perform certain tasks, such as balancing an account book or booking appointments on calendar. If a task is completed earlier, later usage may no longer be needed. Similarly, if a consumer acquires information, e.g. checks the news, earlier in the day, the need for information would be lower later. In contrast, the coefficient for the Social category is positive and statistically significant, suggesting that higher usage of such applications early on also leads to higher usage later. That usages of social applications are inter-temporal complements can be attributed to a potential stimulating effect of social interactions. For example, if a consumer starts an interaction with friends on a certain topic, the interaction may continue over time, with back and forth communications, leading to more subsequent usage of the application.

The last region of Table 5 reports the estimates of the persistence of involvement levels ($\phi_{ik}$ in equation 6). The coefficients for Entertainment, Utility, and Information categories are all positive and statistically significant (posterior means are 0.865, 0.884, and 0.768, respectively),

[12] Throughout the discussion, we consider a parameter estimate statistically significant if the 95% credible interval does not include zero, and consider the difference between two parameter estimates to be statistically significant if their 95% credible intervals do not overlap.

[13] As noted earlier, the dataset was collected in 2012, and recent data may show higher social activity level on mobile applications.

showing that all three activities are highly persistent. Comparatively, involvement in social activities is also positively serial correlated, but the extent of persistence is lower (posterior mean is 0.298). This again can be attributed to the nature of such activities: Entertainment activities, e.g. playing a game on mobile phone, can be quite addictive, so usage in one time period can easily stretch to the next. Similarly, Utility and Information activities may be task focused, so consumers may keep using the application until the task is complete. In contrast, social activities may be subject to external dependencies (e.g. the response from a friend), and may be more scattered through time.

Furthermore, Figure 4 plots the time fixed-effect of the involvement levels ($\delta_{tk}$ in equation 5). Consumers' involvements in all types of activities change significantly during the course of a day. The involvement levels across categories share certain common patterns: They are low in the early morning hours; they then increase rapidly during morning hours and peak around noon; after that the involvement levels decline but remain stable into the evening. Meanwhile, there are notable differences across the categories. The involvement in Information activities is more salient in the morning, and it peaks around 10AM. Comparatively, involvements in both Utility and Entertainment activities ramp up somewhat later, and both peak in early afternoon. After reaching the peak level, the involvement in Utility activities decline rapidly, while that in Entertainment activities decline at a slower pace and stabilizes. In contrast to all three categories, involvement in Social activities does not change as significantly. It has double peaks, one around 9AM and the other in early afternoon. Even after the second peak, the involvement level does not decline as much as the other categories, until late in the evening. Taken together, these variations in involvement levels paint a picture of consumers changing activity focus throughout a day. In the morning, they are more involved in acquiring information. The focus then shifts to utility-based activities and activities of entertainment nature, with the former wearing off rapidly but the latter persisting into evening. Spreading throughout a day is the involvement in Social activities, which starts and reaches a peak early in the day, followed by another peak in early afternoon, and remains relatively stable after that. Recovering these involvement levels provides crucial insight into the underlying drivers of consumer activities using mobile phones. Below, we show that the involvement levels also have significant implications for advertising response.

### Advertising response parameter estimates

Table 6 reports the parameter estimates for consumers' response to mobile in-app advertisements. The first region shows the sequential effect coefficients ($\theta^{S}_{aT(j)}$ in equation 8). For promotion advertisements, the coefficient for the same ad is positive and statistically significant (posterior mean is 0.260). This suggests that as a consumer sees a promotion advertisement repeatedly, the likelihood of clicking the ad increases. In contrast, for product trial advertisements, the same coefficient is negative and statistically significant (posterior mean is -0.176), suggesting that seeing a product trial advertisement repeatedly reduces click probability. The coefficient for product trial ads is consistent with the wear-out effect that has been documented in the literature, while that for promotion ads points to the opposite direction and indicates an ad stock effect. One way to understand this contrast is to consider a promotion ad as *purchase* based, which appeals to consumer's serious deliberation, and a product trial ad as *information* based that may be taken more casually by consumers. Consumers may slowly make up their mind about a promotional offering, so repeated deliveries gradually lead to conversion. Whereas for a more casual product trial ad, they may make up their mind early, such that repeated deliveries do not help. Meanwhile, the coefficient for different advertisements is negative and statistically significant for promotion ads (posterior mean is -0.104). This is also consistent with the ad stock effect discussed above: While repeated delivery of a purchase-based ad builds up the stock, seeing other ads in between may distract a consumer and reduce the stock effect. Both coefficients for product launch advertisements are close to zero and are not statistically significant, potentially due to the relatively small number of impressions of such ads in the dataset.

More important is the effect of a consumer's underlying involvements in the different types of activities on her in-app ad response ($\theta^{\beta}_{k}$ in equation 8), reported in the second region of Table 6. The coefficient is negative and statistically significant for Entertainment, Utility, and Information involvement levels (posterior means are -0.174, -0.107, -0.076, respectively). This suggests that the more a consumer is involved in these types of activities, the less interested in advertisement she is. The effect is most pronounced for Entertainment, while it is smaller for Utility and Information involvement levels. Meanwhile, the coefficient is positive and statistically significant for the Social category, suggesting that the more involved a consumer is in social activities, the more interested in advertisement she is. Theoretical literature on advertising response has established involvement as an important construct, and different types of involvement, e.g. cognitive and affective, have been analyzed (Park and Mittal 1985). However, little is known about how consumers' involvements in various activities affect their ad response in a real world setting. The findings here provide empirical insights in the mobile marketing context. These findings also have managerial implications – they suggest that for better effectiveness, firms should deliver an ad impression when the consumer is not heavily involved in Entertainment, Utility, and Information activities, especially for the former two types. High involvement in Social activities, in contrast, is more conducive to clicks.

The third region of Table 6 reports the contextual effect of the mobile application in which an in-app advertisement is displayed ($\theta_{kT(j)}^{C}$ in equation 8). The effect for Entertainment applications is normalized to 0, and the coefficients for the other categories represent the difference from Entertainment applications. For product trial advertisements, the coefficient is positive and statistically significant for both Utility and Information applications, with the latter larger than the former (posterior means are 0.281 and 0.622, respectively). The coefficient is negative and statistically significant for Social applications for the same type of advertisements (posterior mean is -2.108). This suggests that for product trial advertisements, Information applications provide the most favorable context, whereas Social applications present the least favorable context. The coefficients for the other two types of advertisements do not differ significantly from zero. As discussed earlier, promotion advertisements may appeal to consumers' more serious considerations. With more attention paid to the ad, a consumer may be less affected by the context, in contrast to a more casual product trial ad where context is a more salient factor. For product launch advertisements, this lack of contextual effect may be due to the relatively small number of impressions in the dataset.

An interesting contrast can be made between the effect of underlying involvement levels and the contextual effect of the mobile applications. The coefficient estimates show that the more involved a consumer is in Social activities, the more interested she is in advertisements. However, the context of a Social application itself is not favorable for ad clicks. Consumer's engagement in Social activities thus presents a mixed picture to advertisers. On one hand, the heightened underlying involvement suggests that the consumer is interested in ads, potentially because the involvement in social activities puts the consumer in an interactive mode, making her open minded to ads. On the other hand, if the consumer is using the Social application itself, the context is not favorable to clicks, possibly because the consumer does not want to be interrupted while actively communicating with friends. Instead, the advertiser will be better off delivering the ad to another type of application, if the consumer is known to be highly involved in social activities at the time. In contrast, involvement in Information activities slightly negatively affects consumers' interests in advertisements, yet an Information application provides the most favorable context among the four types of applications. Entertainment activities are the least conducive to ads, as involvement in Entertainment activities has the most negative effect on click propensity, and the context of an Entertainment application is also less favorable than a Utility or Information application. Distinguishing the effect of the underlying involvement levels from the contextual effects of application categories is a key aspect of this study, and doing so is made possible by the multi-tasking nature of mobile application usage.

Furthermore, Figure 5 plots the time fixed-effect of advertising response ($\theta_t^T$ in equation 8). Consumers' interest in mobile in-app ads also has a meaningful time pattern. The interest level is low throughout early hours of day, and increases rapidly from late morning to noon. The interest level remains stable after that, with additional increases in the final few hours.

**Simulations of ad delivery optimization**

Advertising firms seek to generate more consumer clicks with fewer number of ad impressions – the more clicks, the more leads generated from advertising; the fewer impressions, the lower the cost. Operationally, an ad impression can be displayed whenever a consumer uses an application. However, it may not be optimal to deliver the ad impression at every opportunity, as the consumer may not be in a state of mind to respond positively. Displaying an ad at the right time and circumstance is the key to effectiveness.

By connecting consumers' usage of mobile applications with their responses to mobile in-app advertisements, the model developed in this study not only advances our understanding of consumers' behavior on mobile phones, but also sets the foundation for effective ad targeting. For example, knowing that higher involvement in Entertainment activities reduces a consumer's interest in ads, and seeing a consumer is using Entertainment applications heavily at the time, an advertiser would decide that delivering the ad in this occasion is unlikely to be effective. As another example, knowing that higher involvement in Social activities increases a consumer's interest in ad and that an application of the Utility category provides a good context, an advertiser would infer that the ad should be delivered to a Utility application when the consumer is heavily involved in Social activities. More generally, using the model of our study, an advertiser can predict the probability of a consumer clicking an ad at each specific instance, and can deliver the ad only when the click probability is high, in order to improve targeting effectiveness.

To illustrate this point, we conduct simulation analysis of targeted ad delivery strategies based on our model. To recap, several model components provide opportunities for targeting. To begin, the time fixed-effects of ad click decisions already enable time-based targeting that is easy to implement. Taking this to the next step, two key model components are the effects of consumers' underlying involvement and the contextual effect of mobile applications. Accounting for the salient time patterns of involvement levels as well as the application categories, a firm can further improve targeting effectiveness by delivering ads when consumers' involvement levels are appropriate, and to the favorable context. Finally, consumers' application usage history and the involvement persistence enable the inference of involvement at individual consumer level, which can further enhance the accuracy of targeting. Accordingly, we simulate three targeting strategies. In the first, which we call the *time-only* strategy, at each time a consumer uses an application, the probability of the consumer clicking an ad is calculated based only on the time fixed-effect estimates. An ad impression is then delivered if the predicted click probability exceeds a certain threshold (which yields a predetermined target number of ad impressions). This represents the simple time-based targeting mechanism. The second is the *population-level targeting* strategy. In this strategy, the click probability is calculated using the population level parameter estimates, and an ad impression is then delivered if the predicted click probability exceeds a certain threshold. This strategy accounts for the contextual effect of mobile applications and the effect of consumers' underlying involvement, although the underlying involvement is inferred using the population level time patterns only, and no individual consumer based targeting is used. The third strategy, called the *individual-level targeting* strategy, extends from the second strategy by inferring the involvement levels for individual consumers. The posterior of a consumer's involvement levels is estimated form the application usage up to the point of the ad impression delivery. We compare all strategies with an *even-distribution* benchmark strategy, where a quota of the same number of ad impressions is given to each hour of day.[14]

One thousand consumers are randomly drawn, and their application usages for a day are simulated, both based on the model parameter estimates. Several scenarios with different ad campaign sizes, i.e. ad impression quotas, are simulated. Each scenario corresponds to a pre-determined number of ad impressions to be displayed to the consumers, ranging from 5,000 to 100,000 per advertisement. The result is reported in Table 7. For all the scenarios, all three targeted delivery strategies achieve significantly higher click-through rates than the benchmark strategy. Targeting based on time effects alone already generates 50%-100% increase in CTR from the benchmark strategy, and targeting based on population-level or individual-level estimates improves CTRs even further. The individual-level targeting strategy leads to almost three-fold increase in CTRs for 5,000 impressions. Meanwhile, the improvements of all the targeting strategies are higher when the number of ad impressions to be displayed is lower. As the number of ad impressions to be displayed decreases, all the targeting strategies become more selective in delivering ad impressions. Consequently, the CTRs become higher. However, since the benchmark strategy simply distributes ad impressions evenly across hours but does not perform targeting, the CTRs do not change meaningfully as the number of ad impressions changes. Comparing the three targeting strategies themselves, the individual-level targeting strategy achieves better CTR than does the population-level targeting strategy, which in turn performs better than the time-only targeting strategy. This shows that all the major model components – the time specific effects, the effect of involvements and the contextual effects, and the evolution of consumers' underlying involvement levels – can help significantly improve ad targeting performance. Furthermore, the difference among the three targeted delivery strategies is also higher for smaller ad impression quotas, especially between the population-level and individual-level targeting strategies. This shows that the more selective the ad impression delivery is, the more important it is to incorporate all the factors, especially the individual consumer level information. When the advertiser can deliver only 5,000 impressions, population-level targeting yields a CTR 50% better than time-based targeting, while individual-level targeting improves the CTR from population-level targeting by a further 30%.

[14] Advertisers typically deliver ad impressions using either such an even distribution strategy or a "greedy" strategy to deliver impressions whenever possible. The latter leads to more ad impressions being delivered earlier with no obvious benefits to click-through rates. Therefore, we use the former as the benchmark.

To take a closer examination of the targeting strategies in terms of their allocation of ad impressions over time, Figure 6 plots the proportions of ad impressions delivered in each hour of day (i.e. the number of impressions of delivered in that hour divided by the total number of impressions) according to the individual-level targeting strategy, together with the click-through rate achieved by the strategy in each hour. For comparison, the figure also plots the actual proportions of ad impressions delivered in each hour based on the actual dataset. The figure shows that both the actual proportions and those generated by the targeting strategy are generally consistent with the click-through rates: the time periods around mid-day have high click-through rates and account for high proportions of ad impressions, while the early hours of day with lower click-through rates also have fewer ad impressions. Detailed comparison, however, shows that the individual-level targeting strategy is clearly better. The proportions based on the actual data peak in mid-morning, while those based on the targeting strategy peak around noon, more closely matching the hours of higher click-through rates. More importantly, for afternoon and evening hours, the proportions based on the actual data are much lower than those based on the targeting strategy. The click-through rates of those hours are quite high, however, and are noticeably higher than morning hours. The targeting strategy improves the overall click-through rates partly by allocating more ad impressions to these hours, making the amount of ad impressions better aligned with the click-through rates. Also worth noting is that the click-through rates achieved by the targeting strategy in the afternoon and evening hours are as high as those in mid-day hours. This is because the individual-level targeting strategy infers consumers' involvement levels based on the application usage history, and can make better inference in later hours when a longer history is available.

In summary, the simulation confirms the managerial importance of the knowledge gained from our model. Note that by using a threshold approach, the strategies simulated here simply function as proofs of concept. Using the model estimates, more sophisticated decision support system can be crafted to optimize different performance criteria under different constraints. If the targeting strategy incorporates dynamic programming, by predicting the possibility of delivering ad impressions in future and anticipating the effect of current delivery on future clicks, click-through rates can potentially be improved even further. Developing such optimized strategies is itself an important research question, which we leave for future study.

## Discussion and Conclusion

The rapid migration of consumer activities to mobile phones brings about many new phenomena that are not yet well understood. Chief among them are consumers' usage of mobile applications and response to mobile in-app advertisements. With consumers spending more than 80% of their mobile phone time on mobile applications, and with billions of advertising dollars poured into this area, it is imperative for managers to understand the application usage and advertising response behaviors, and to optimize targeted delivery of in-app advertisements to the right consumers at the right circumstances. However, extant literature offers only limited insight in this new arena. The prevalence of 24/7 ubiquity and the multi-tasking nature of mobile phone usage present additional challenges to practitioners.

Drawing from the extant research on activity consumption and on advertising response, we develop an integrated model for application usage and in-app advertising response. The model sheds light on how consumers' usage of mobile applications are driven by their underlying involvements in different activities that evolve over time, and on how the involvements in different activities affect their propensity to click mobile in-app advertisements. Contextual effects and the effect of repeated deliveries are also accounted for in the model. Empirical estimates of the model using a unique dataset, which contains comprehensive information on consumers' usage of mobile applications and on the impressions and clicks of mobile ads delivered in those applications, show rich and intriguing findings. They show a salient temporal pattern of consumers' involvement in different types of activities, where involvement in Information activities peaks earlier in the day, while those in Utility and Entertainment activities peak later. They show that involvements in Entertainment, Utility, and Information activities are highly persistent. Equally importantly, the analysis shows that consumers' involvements in different activities have significant implications on their response to in-app advertisements. Higher involvement in Entertainment activities strongly reduces consumers' propensity to click ads. Higher involvements in Utility or Information activities also reduce the click propensity, although to a lesser extent. In contrast, higher involvement in Social activities increases a consumer's interest in ads and the likelihood of clicking them. The effect of these involvement levels is further contrasted to the contextual effects of mobile applications of different categories, which show that the context of a Social application is actually least favorable for clicks, while that of an Information application is the most favorable.

These findings not only advance our understandings of consumer behaviors on mobile devices, but also offer practical guidance to managers seeking to improve the effectiveness of mobile in-app ads. Through simulation, we show that targeted ad delivery strategies derived from our model yield significantly higher click-through rates than the benchmark strategy, and the advantage is greater for lower ad impression quotas which require more precise targeting. The targeting strategies deliver more ad impressions in times which are more conducive to clicks, thus better aligning the two than does the benchmark strategy. Ad impressions delivered based on inference of individual consumer level involvements is also shown to be more effective than those based on population level estimates or only time effects, suggesting there is much potential to individual consumer based targeting.

Our study contributes to the literature by being the first to jointly model consumers' usage of mobile applications and response to mobile in-app advertisements. The analysis confirms that these two key activities are closely connected. It is the first to show how consumer's underlying involvement in various activities affect their advertising response in a real world setting, and to distinguish the effect of the underlying involvement from that of the application context. It is also among the first to investigate the temporal patterns of mobile applications usage throughout the times of day, which assists managers in better gauging consumers' interests on mobile devices. The managerial importance of such knowledge is demonstrated from simulations, which show that it can help significantly improve the effectiveness of targeted ad delivery. All these contribute to the nascent yet rapidly growing literature on mobile marketing.

Several limitations of the study call for future research. First, although the dataset is rich with detailed, precisely time-stamped application usage and ad impressions and clicks information, it covers only a short period of seven days. A dataset that covers a longer period of time will enable the analysis of potential change in consumer behavior over time. For example, as consumers become more familiar with different types of mobile applications, the nature and extent of their response to advertisements may change. Second, the dataset does not contain consumer purchase information, so our analysis of ad response is restricted to the first step, i.e. clicking of the ads. While this is an important first step, not all clicks are the same, and it will be interesting to see clicks generated in what circumstances will lead to higher subsequent purchase. Meanwhile, literature has shown that even in the absence of clicks, ad impressions can still change consumers' brand perception and influence their purchase decisions. This also calls for richer datasets which enable the analysis of the ad effect on purchase in this context. Third, while the dataset contains around the clock information which reflects the ubiquity of mobile usage, it does not contain detailed time-stamped location information. Time-stamped location and other information about the consumers' activities will enable more in-depth analysis of the nature and extent of consumers' underlying involvement. With the rapid growth of mobile wearable devices, such data may soon become available for research. Finally, although the simulation shows that targeted ad delivery based on the model estimates can significantly improve click-through rates, the threshold-based targeting strategy is still rather primitive, and functions more as a proof of concept. Crafting optimal ad targeting strategies to maximize different performance measurement criteria under different constraints, which fully take advantage of the ubiquitous nature of the mobile channel, will likely call for dynamic programming under incomplete information. This is an important research question in its own right, and is an exciting topic for future study which can provide direct guidance to managers.

## Technical Appendix 1: Mobile In-App Advertising Overview

The data used in this study is obtained from a major mobile advertising platform company in a large Asian country. In this section, we briefly discuss the industry structure, which is illustrated in the top half of Figure TA.1-1. The mobile advertising company works as a platform, or two-sided market, in the mobile advertising ecosystem. On one side, the company contracts with mobile application developers. The application developers create and operate mobile applications, which are software programs that run on mobile devices that provide certain functionalities. The mobile applications span across many categories. Some are games that consumers play on mobile phones for leisure; some are social networking applications to connect with friends; some are tools for managing personal finances, etc. A large number of mobile applications are available for download from online app stores, and they account for more than 80% of the time consumers spend on mobile devices (TechCrunch 2014). In many mobile applications, a small area on the mobile phone display, usually at the top or bottom of the screen immediately above or below the content of the application, can be used to display advertisements while consumers are using the application. These display areas constitute the mobile display advertising "inventory". The advertising platform company contracts with application developers to fill these inventories with advertisements, and pays the application developers based on the amount of advertisements that are delivered.

On the other side of the market, the advertising platform company contracts with firms that seek to run advertising campaigns about their product or service offerings. These firms hire the platform company to conduct the advertising campaign, and pay the company based on certain performance criteria. For example, the product firm can pay the platform company based on the number of times the advertisements are displayed to mobile users, or based on the number of clicks generated from the advertising campaign, etc. The advertising platform company then develops specific campaign strategies, often in consultation with the product firm, to deliver the ads to the display areas in mobile applications, so that consumers using those mobile applications can view and click the advertisements. The advertising company is expected to choose the appropriate types of mobile applications and time periods in a day to deliver the ad impressions.

The advertising company then typically follows one of two scheduling mechanisms. In the first mechanism, once started, the advertising company will deliver an ad impression whenever possible, i.e. when an eligible mobile application is being used. Once the total number of ad impressions by contract has been delivered, though, no additional ad impressions will be delivered later on, even if the mobile application is still being used, so ad impressions can still be delivered. This mechanism is expected to lead to more ad impressions being delivered earlier in the day than in later hours. In the second mechanism, in contrast, the advertising company will seek to deliver ad impressions more evenly over time, such as by giving each hour a fixed quota depending on the total number of impressions to be delivered. This alleviates the concentration in earlier time of day, at the risk of not delivering enough impressions if consumers do not actively use the applications later. The advertising platform company considered all these factors in determining the delivery of ad impressions, although the particular strategies for delivering ad impressions are confidential and not known to researchers.

The advertising platform company runs an advertising engine on its server computers which handles the delivery of ad impressions. The technical aspect of delivering advertisements to mobile phones is illustrated in the bottom half of Figure TA.1-1. When a user opens a mobile application, the software program on the mobile phone sends a request to the advertising engine on server, seeking delivery of advertisements. As long as the application remains open, i.e. the consumer continues to use it, the program sends a new request every 20 seconds to refresh the advertisement. When a request is received, if there is a suitable advertisement to be delivered, the advertising engine on the server computer will “push” the advertisement to the mobile phone, which is then displayed to the mobile user. A same advertisement can be displayed repeatedly to the user, or can be intermingled with the displaying of other advertisements. Sometimes, the engine does not have an advertisement to send (for example, a firm may enter a contract with the advertising company to deliver its advertisement 10,000 times. After that, the advertising engine will stop sending the advertisement to mobile applications). In that situation, the in-app ad display area will be left empty.

**Figure TA.1-1: Mobile In-App Advertising**

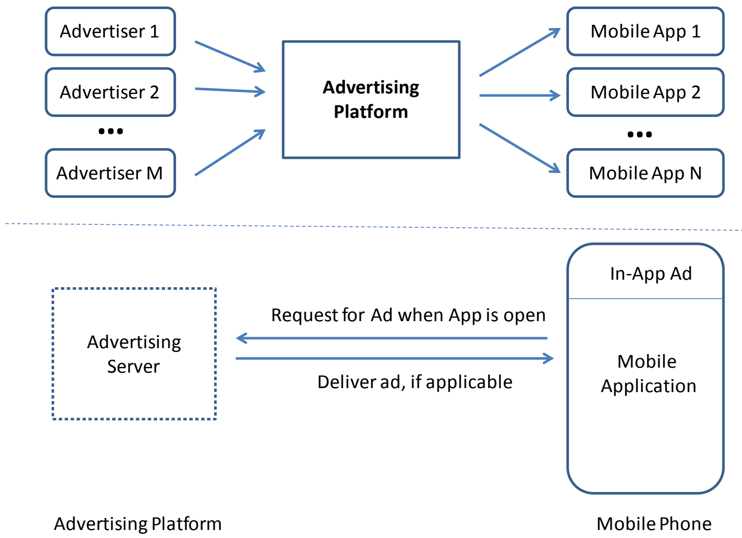


## Technical Appendix 2: Control for Ad Targeting

Marketing response modeling should account for the possibility that the marketing mix variables are not independent from response parameters (Manchanda et al. 2004). In our context of mobile in-app ads, it is possible that advertisers target their ad delivery based on certain knowledge of consumer responses. To the categories of applications or the time periods which are more likely to generate clicks, an advertiser may decide to deliver more ad impressions. To account for this, we model the delivery of ad impression as follows:

(TA.2-1) $$\ln\left(\mathrm{I}_{jkt}+1\right)=\omega_j+\omega^{\theta C}\theta^{C}_{kT(j)}+\omega^{\theta T}\theta^{T}_{t}+\omega^{A}\ln(A_{kt}+1)+\omega^{T}\mathrm{t}+\eta_{jkt}$$

In equation (TA.2-1), $\mathrm{I}_{jkt}$ is the total number of impressions of ad $j$ delivered at time $t$ in applications of category $k$. This is the dependent variable for ad delivery, for which we perform a log transformation as the numbers span across several orders of magnitude. The coefficient $\omega_j$ is an ad-specific fixed effect that reflects the size of the ad campaign – different ads have different total number of impressions, possibly determined by budget. $\theta^{C}_{kT(j)}$ is the same as in equation (8) in the paper, which represents the application's contextual effect on ad response, and $\omega^{\theta C}$ is the corresponding coefficient. A positive $\omega^{\theta C}$ would indicate that advertisers deliver more ad impressions to the application categories that are more conducive to consumer clicks.

Similarly, $\theta_t^T$ is the same as in equation (8) in the paper, which represents the time specific effect of ad response, and a positive $\omega^{\theta T}$, the corresponding coefficient, would indicate that advertisers deliver more impressions at more favorable time periods. In the fourth term $\omega^A \ln(A_{kt} + 1)$, $A_{kt}$ is the total amount of application usage of category $c$ at time $t$. Since the delivery of an ad impression is contingent upon a consumer using an application, the corresponding coefficient $\omega^A$ is expected to be positive. The next term $\omega^T \mathrm{t}$ captures the direct time effect. As discussed in the data section, advertisers often specify a certain number of impressions to be delivered, and the ad delivery would stop after that. This suggests that ad impression may decrease during the course of a day. Thus the coefficient $\omega^T$ is expected to be negative. Note that this is different from the term $\omega^{\theta T} \theta_t^T$ which, although also related to the time of day, actually captures how conducive the time is to ad clicks. This term $\omega^T \mathrm{t}$ instead captures the direct time trends arising from the mechanics of the delivery. Finally, $\eta_{jkt}$ is an independent error term that follows a normal distribution.

This ad delivery equation accounts for potential targeting of ad delivery by application category and by time. Based on the researchers' knowledge of the industry practice during the time covered by the dataset, there were targeting at application level and time level, but not at individual consumer level. This equation thus adequately controls for any potential endogeneity concerns arising from such targeting. Note that this model does not assume that advertisers make optimal ad delivery decisions. Instead, it merely posits that advertisers may have partial knowledge of the contextual and time effects of ad response, and their ad delivery decisions may be related to it. Both our understanding of the industry practice and the initial evidence from the dataset actually suggest that the current ad delivery practices leave much room for improvement.

Table TA.2-1 reports the parameter estimates for the ad delivery targeting equations. The estimates show that advertisers indeed delivered more ad impressions at time periods when consumers have higher interest in ads ($\omega^{\theta T}$ is 0.491 and statistically significant). In contrast, however, there is no evidence that advertisers delivered more ad impressions to applications which provide more favorable contexts ($\omega^{\theta C}$ is -0.070 and not statistically significant). As expected, more ad impressions were delivered when consumers used applications more intensively ($\omega^A$ is 0.485 and statistically significant), as a mobile in-app ad can be displayed only when the consumer is using the application. Meanwhile, there is a negative time trend for ad delivery ($\omega^T$ is -0.144 and statistically significant), suggesting that other things equal, more ads were delivered earlier in the day than later in the day. Taken together, these parameter estimates suggest both that it is necessary to control for potential endogeneity from advertisers' targeting practice when analyzing in-app ad response, as such targeting does exist to a certain extent, and that there is ample room for improvement for delivering ads to the right consumers at the right circumstances.

**Table TA.2-1: Parameter Estimate – Ad Delivery Targeting**

| Parameter | Mean | SD | 2.5% CI | 97.5% CI |
|---|---|---|---|---|
| $\omega^{\theta C}$ | -0.070 | 0.078 | -0.229 | 0.076 |
| $\omega^{\theta T}$ | 0.491 | 0.143 | 0.243 | 0.799 |
| $\omega^{A}$ | 0.485 | 0.032 | 0.426 | 0.553 |
| $\omega^{T}$ | -0.144 | 0.015 | -0.176 | -0.118 |

**Technical Appendix 3: Identification**

Identification of the model parameters rests on the temporal and cross-sectional variations of the application usage and advertising response data. Specifically, category-specific baseline involvement levels ($\beta_{ik}$) are identified by the individual-specific average usage of mobile applications of different categories. Time fixed-effects of involvement levels ($\delta_{tk}$) are identified from the overall change of application usage over the course of a day. Serial correlation coefficients of involvement levels ($\phi_{ik}$) are identified through the relationship between application usages of adjacent time periods, while the substitution or complementarity parameters for application usage ($\beta_{k1}$) are identified through the relationship between usage history of the day and the current usage amount. Advertising response parameters are generally identified through the click activities in response to the time and circumstance of the ad impression. Ad-specific quality parameters ($q_j$) are identified through the overall click-through rate of each in-app ad. The coefficients of underlying involvement on ad response ($\theta_k^{\beta}$) are identified through the relationship between the amount of application usage at the time and the click propensity. The coefficients for application context ($\theta^{C}_{kT(j)}$) are identified through the difference in click propensity when the same ad is shown in applications of different categories. The sequential effect parameters ($\theta^{S}_{aT(j)}$) are identified through the change in click propensity when the same ad is delivered repeatedly and intermingled with the deliveries of other ads. The time fixed effect parameters ($\theta^{T}_{t(ijn)}$) are identified through the overall click propensity at different time of day.

Certain normalizations are also needed for identification. The time fixed-ffects for application usage and the baseline application usage parameters cannot all be identified. Instead, we fix the time fixed effect for $t = 13$ (12pm to 1pm) to 0 – we choose to normalize a time period with more application usage than a time period with less application usage (e.g. $t = 1$) to avoid data sparseness issues. We also fix $\sigma^2_{\varepsilon} = 1$ (the variance term of the fluctuation in involvement levels) to address data sparseness and maintain estimation stability. Similarly, the ad quality parameters and time fixed-effect for ad response cannot all be identified, and we also fix the time fixed effect for ad response for $t = 13$ to 0. The ad quality parameters and the parameters for intrinsic interest in ad ($\theta_i$) cannot all be identified. Instead, we fix the population level mean for $\theta_i$ to be 1 (individual consumer level $\theta_i$ relative to population mean is still identified through the cross-sectional variation of click propensities). Finally, the parameters for application contextual effect on ad response ($\theta^{C}_{kT(j)}$) cannot all be identified as they are relative terms, and we fix $\theta^{C}_{kT(j)} = 0$ for the Entertainment category. Thus the coefficients for the other three categories should be interpreted as the contextual effect of those categories on ad clicks relative to the effect of the Entertainment category.

## References

Andrews, Michelle, Xueming Luo, Zheng Fang, and Anindya Ghose, (2015), "Mobile Ad Effectiveness: Hyper-Contextual Targeting with Crowdedness," *Marketing Science*, forthcoming

Bart, Yakov, Andrew T. Stephen, and Miklos Sarvary, (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* 51.3: 270-285.

Bhat, Chandra R, (2005), "A multiple discrete–continuous extreme value model: formulation and application to discretionary time-use decisions," *Transportation Research Part B: Methodological* 39.8: 679-707.

Braun, Michael, and Wendy W. Moe, (2013), "Online display advertising: Modeling the effects of multiple creatives and individual impression histories," *Marketing Science* 32.5: 753-767.

Cho, Chang-Hoan, (2003), "Factors influencing clicking of banner ads on the WWW," *CyberPsychology & Behavior* 6.2: 201-215.

Clancy, Kevin J., and David M. Kweskin, (1971) "TV commercial recall correlates," *Journal of Advertising Research* 11.2: 18-20.

Danziger, Shai, Jonathan Levav, and Liora Avnaim-Pesso, (2011), "Extraneous factors in judicial decisions," *Proceedings of the National Academy of Sciences* 108.17: 6889-6892.

eMarketer, 2015, "Mobile Ad Spend to Top $100 Billion Worldwide in 2016, 51% of Digital Market," url: http://www.emarketer.com/Article/Mobile-Ad-Spend-Top-100-Billion-Worldwide-2016-51-of-Digital-Market/1012299

Forbes, 2014(a), "Mobile App Usage By The Numbers," url: http://www.forbes.com/sites/niallmccarthy/2014/10/29/mobile-app-usage-by-the-numbers-infographic/

Forbes, 2014(b), "Study: Mobile Ads Actually Do Work - Especially In Apps," url: http://www.forbes.com/sites/roberthof/2014/08/27/study-mobile-ads-actually-do-work-especially-in-apps/

Ghose, Anindya, Avi Goldfarb, and Sang Pil Han, (2013), "How is the mobile Internet different? Search costs and local activities," *Information Systems Research* 24.3: 613-631.

Ghose, Anindya, and Sang Pil Han, (2011), "An empirical analysis of user content generation and usage behavior on the mobile Internet," *Management Science* 57.9: 1671-1691.

Goldberg, Marvin E., and Gerald J. Gorn, (1987), "Happy and sad TV programs: How they affect reactions to commercials," *Journal of Consumer Research*: 387-403.

Gupta, Sunil, (2013), "For Mobile Devices, Think Apps, Not Ads.(cover story)," *Harvard Business Review*, 91.3, 70-75.

Howard, Daniel J., and Thomas E. Barry, (1994), "The role of thematic congruence between a mood-inducing event and an advertised product in determining the effects of mood on brand attitudes," *Journal of Consumer Psychology* 3.1: 1-27.

Jacoby, Jacob, George J. Szybillo, and Carol Kohn Berning, (1976), "Time and consumer behavior: An interdisciplinary overview," *Journal of Consumer Research*: 320-339.

Kamakura, Wagner A, (2009), "American time-styles: A finite-mixture allocation model for time-use analysis," *Multivariate Behavioral Research* 44.3: 332-361.

Kamins, Michael A., Lawrence J. Marks, and Deborah Skinner, (1991), "Television commercial evaluation in the context of program induced mood: Congruency versus consistency effects," *Journal of Advertising* 20.2: 1-14.

Krugman, Herbert E, (1965), "The impact of television advertising: Learning without involvement," *Public Opinion Quarterly* 29.3: 349-356.

Krugman, Herbert E, (1983), "Television program interest and commercial interruption," *Journal of Advertising Research* (1983) 23, 21-23.

Lin, Chen, Sriram Venkataraman, and Sandy D. Jap, (2013), "Media multiplexing behavior: Implications for targeting and media planning," *Marketing Science* 32.2: 310-324.

Lord, Kenneth R., Robert E. Burnkrant, and H. Rao Unnava, (2001), "The effects of program-induced mood states on memory for commercial information," *Journal of Current Issues & Research in Advertising* 23.1: 1-15.

Luo, Lan, Brian T. Ratchford, and Botao Yang, (2013), "Why we do what we do: A model of activity consumption," *Journal of Marketing Research* 50.1: 24-43.

Luo, Xueming, Michelle Andrews, Zheng Fang, and Chee Wei Phang, (2013), "Mobile targeting," *Management Science* 60, 7: 1738-1756.

Manchanda, Puneet, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta, (2006), "The effect of banner advertising on internet purchasing," *Journal of Marketing Research* 43, no. 1: 98-108.

Manchanda, Puneet, Peter E. Rossi, and Pradeep K. Chintagunta, (2004), "Response modeling with nonrandom marketing-mix variables." Journal of Marketing Research 41.4: 467-478.

Moore, Robert S., Claire Allison Stammerjohan, and Robin A. Coulter, (2005), "Banner advertiser-web site context congruity and color effects on attention and attitudes," *Journal of Advertising* 34.2: 71-84.

Park, C. Whan, and Banwari Mittal, (1985), "A theory of involvement in consumer behavior: Problems and issues," *Research in Consumer Behavior* 1.3: 201-32.

Park, C. Whan, and S. Mark Young, (1986), "Consumer response to television commercials: The impact of involvement and background music on brand attitude formation," *Journal of Marketing Research*: 11-24.

Pavelchak, Mark A., John H. Antil, and James M. Munch, (1988), "The Super Bowl: An investigation into the relationship among program context, emotional experience, and ad recall," *Journal of Consumer Research*: 360-367.

Rosenstein, Ariel, (2015), "What App Usage Patterns Mean for Online Marketers: Case Study with 'Religious' Apps," SimilarWeb, url: http://www.similarweb.com/blog/app-usage-patterns-case-study-the-bible-the-quran

Salesforce, (2014), "2014 Mobile Behavior Report," url: https://www.exacttarget.com/sites/exacttarget/files/deliverables/etmc-2014mobilebehaviorreport.pdf

Shankar, V., Venkatesh, A., Hofacker, C., & Naik, P, (2010), "Mobile marketing in the retailing environment: current insights and future research avenues," *Journal of Interactive Marketing*, 24(2), 111-120.

Sharma, Andrew, (2000), "Recall of television commercials as a function of viewing context: The impact of program-commercial congruity on commercial messages," *Journal of General Psychology* 127.4: 383-396.

TechCrunch, 2014, "Mobile App Usage Increases In 2014, As Mobile Web Surfing Declines," url: http://techcrunch.com/2014/04/01/mobile-app-usage-increases-in-2014-as-mobile-web-surfing-declines/

VentureBeat, 2015, "10 mobile advertising trends for 2015 and beyond," url: http://venturebeat.com/2015/04/21/10-mobile-advertising-trends-for-2015-and-beyond/

Yoon, Carolyn, Michelle P. Lee, and Shai Danziger, (2007), "The effects of optimal time of day on persuasion processes in older adults," *Psychology & Marketing* 24.5: 475-495.

Zaichkowsky, Judith Lynne, (1985), "Measuring the involvement construct," *Journal of Consumer Research*: 341-352.

**Table 1: Application Usage Descriptive Statistics**

| | Daily Usage | | | |
|---|---|---|---|---|
| **Category** | **Mean** | **SD** | **Min** | **Max** |
| Entertainment | 113.60 | 598.98 | 0 | 33182 |
| Utility | 13.21 | 63.01 | 0 | 1784 |
| Information | 7.51 | 32.94 | 0 | 1029 |
| Social | 1.52 | 13.32 | 0 | 516 |
| Total Users | 3,988 | | | |
| Number of Days | 7 | | | |

**Table 2: Impressions and Clicks of Ad Campaigns**

| Ad Campaign | Ad Type | Impressions | Clicks | Click-Through Rate |
|---|---|---|---|---|
| Campaign 1 | Promotion | 11533 | 72 | 0.62% |
| Campaign 2 | Product Trial | 265451 | 6203 | 2.34% |
| Campaign 3 | Product Trial | 23838 | 106 | 0.44% |
| Campaign 4 | Promotion | 35024 | 272 | 0.78% |
| Campaign 5 | Promotion | 336744 | 3678 | 1.09% |
| Campaign 6 | Product Trial | 7642 | 43 | 0.56% |
| Campaign 7 | Product Trial | 17138 | 82 | 0.48% |
| Campaign 8 | Promotion | 17096 | 117 | 0.68% |
| Campaign 9 | Promotion | 10282 | 35 | 0.34% |
| Campaign 10 | Promotion | 12921 | 72 | 0.56% |
| Campaign 11 | Product Launch | 6062 | 46 | 0.76% |
| Campaign 12 | Product Launch | 6992 | 55 | 0.79% |
| Campaign 13 | Product Launch | 5454 | 43 | 0.79% |
| Campaign 14 | Promotion | 9161 | 39 | 0.43% |

**Table 3: CTR by Application Context**

| App Category | Ad CTR | Ad CTR Conditional on Above-Average Usage of App | | | |
|---|---|---|---|---|---|
| | | Entertainment | Utility | Information | Social |
| Entertainment | 1.72% | 1.56% | 0.81% | 1.77% | 5.78% |
| Utility | 1.14% | 0.52% | 0.93% | 1.40% | 1.92% |
| Information | 1.81% | 1.99% | 1.15% | 1.68% | 5.97% |
| Social | 1.96% | 1.18% | 0.00% | 0.00% | 1.96% |

**Table 4: Model Comparison**

| Model | In-Sample LMD | Out-of-Sample LL |
|---|---|---|
| Proposed Model | -1675628 | -362359 |
| Benchmark Model 1 | -1676136 | -365246 |
| Benchmark Model 2 | -1676170 | -365322 |
| Benchmark Model 3 | -2773499 | -387894 |

**Table 5: Parameter Estimate – Application Usage**

| Parameter | Mean | SD | 2.5% CI | 97.5% CI |
|---|---|---|---|---|
| *Baseline involvement level ($\bar{\beta}_k$)* | | | | |
| Entertainment | -3.985 | 0.110 | -4.201 | -3.778 |
| Utility | -7.946 | 0.155 | -8.248 | -7.656 |
| Information | -8.193 | 0.115 | -8.412 | -7.968 |
| Social | -11.769 | 0.086 | -11.923 | -11.598 |
| *Effect of prior usage ($\beta_{k1}$)* | | | | |
| Entertainment | -0.063 | 0.025 | -0.097 | -0.015 |
| Utility | -0.167 | 0.019 | -0.193 | -0.126 |
| Information | -0.123 | 0.020 | -0.149 | -0.081 |
| Social | 0.618 | 0.027 | 0.564 | 0.656 |
| *Persistence of involvement level ($\bar{\phi}_k$)* | | | | |
| Entertainment | 0.865 | 0.005 | 0.855 | 0.875 |
| Utility | 0.884 | 0.012 | 0.860 | 0.907 |
| Information | 0.768 | 0.021 | 0.726 | 0.808 |
| Social | 0.298 | 0.031 | 0.227 | 0.352 |

**Table 6: Parameter Estimate – Advertising Response**

| Parameter | Mean | SD | 2.5% CI | 97.5% CI |
|---|---|---|---|---|
| *Sequential Effect* ($\theta^{S}_{aT(j)}$) | | | | |
| Promotion Ad | | | | |
| Same Ad | 0.260 | 0.069 | 0.127 | 0.402 |
| Different Ad | -0.104 | 0.056 | -0.209 | -0.011 |
| Product Trial Ad | | | | |
| Same Ad | -0.176 | 0.024 | -0.224 | -0.131 |
| Different Ad | -0.009 | 0.019 | -0.045 | 0.027 |
| Product Launch Ad | | | | |
| Same Ad | 0.034 | 0.179 | -0.296 | 0.406 |
| Different Ad | 0.040 | 0.123 | -0.196 | 0.290 |
| *Effect of Involvement Levels* ($\theta^{\beta}_{k}$) | | | | |
| Entertainment | -0.174 | 0.016 | -0.203 | -0.141 |
| Utility | -0.107 | 0.008 | -0.122 | -0.092 |
| Information | -0.076 | 0.013 | -0.100 | -0.051 |
| Social | 0.363 | 0.012 | 0.342 | 0.387 |
| *Applications Contextual Effect* ($\theta^{C}_{kT(j)}$) | | | | |
| Promotion Ad | | | | |
| Utility | -0.045 | 0.282 | -0.617 | 0.475 |
| Information | -0.436 | 0.330 | -1.082 | 0.172 |
| Social | -0.486 | 0.794 | -2.151 | 1.074 |
| Product Trial Ad | | | | |
| Utility | 0.281 | 0.101 | 0.074 | 0.480 |
| Information | 0.622 | 0.109 | 0.404 | 0.828 |
| Social | -2.108 | 0.226 | -2.533 | -1.650 |
| Product Launch Ad | | | | |
| Utility | 0.058 | 0.524 | -1.019 | 1.037 |
| Information | -1.063 | 0.578 | -2.173 | 0.023 |
| Social | -0.181 | 0.737 | -1.676 | 1.342 |

Coefficient for the Entertainment category is normalized to 0 for the contextual effect

**Table 7: Simulation – Targeted Delivery of Ad Impressions**

| | Click-Through Rate (CTR) | | | |
|---|---|---|---|---|
| Target Number of Impressions | Even-Distribution (Benchmark) | Target By Time Only | Target By Population Level Estimate | Target By Individual Level Estimate |
| 5000 | 0.58% | 1.07% | 1.59% | 2.10% |
| 10000 | 0.52% | 0.82% | 1.45% | 1.91% |
| 20000 | 0.50% | 0.93% | 1.45% | 1.61% |
| 50000 | 0.55% | 0.87% | 1.09% | 1.28% |
| 100000 | 0.59% | 0.74% | 0.87% | 0.98% |

Target Number of Impressions: Number of impressions expected to serve per ad. Actual impressions differ slightly depending on consumers' actuall application usage amount

Cutoff thresholds for target strategies are chosen to deliver close to the target number of ad impressions.

## Figure 1: Mobile Application Usage by Hour and Category

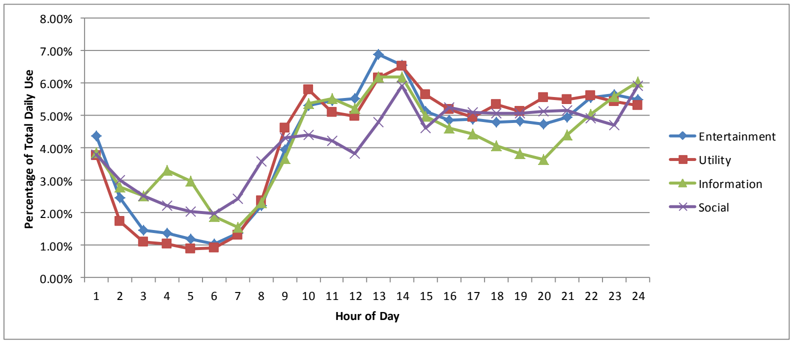

## Figure 2: Ad Delivery Statistics by Hour

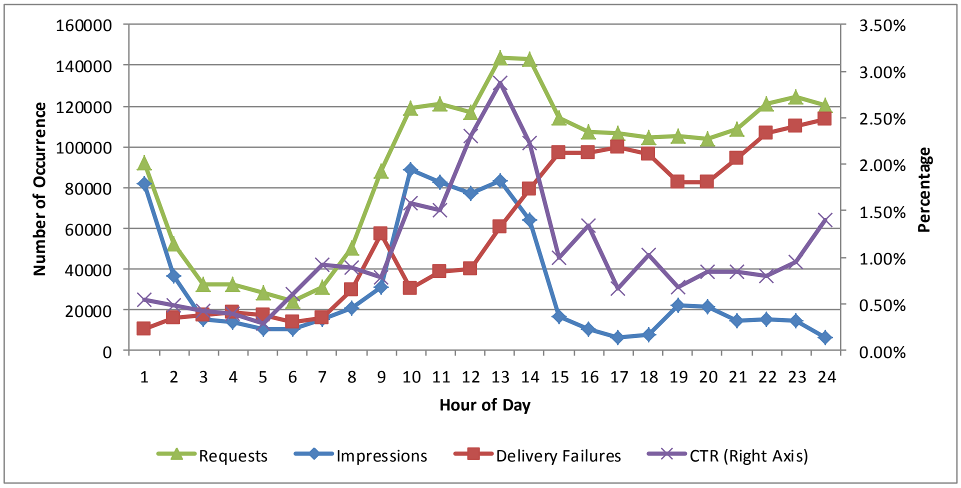

## Figure 3: Model Conceptual Framework

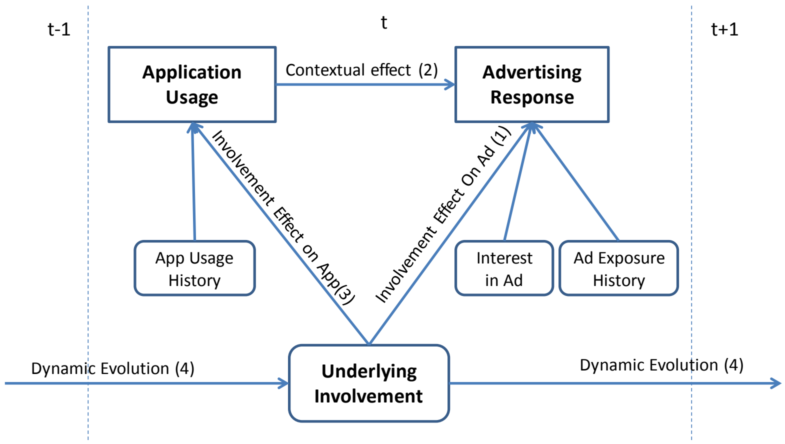


(1): Underlying involvement levels in different activities affect ad response
(2): Direct contextual effect of the mobile application on ad response
(3): Underlying involvement levels determine application usage
(4): Underlying involvement levels evolve over time

**Figure 4: Time Patterns of Involvement Levels**

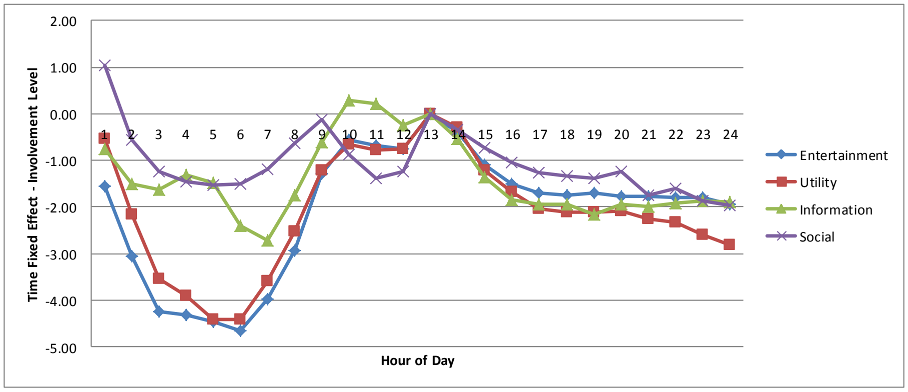

**Figure 5: Time Fixed Effect of Advertising Response**

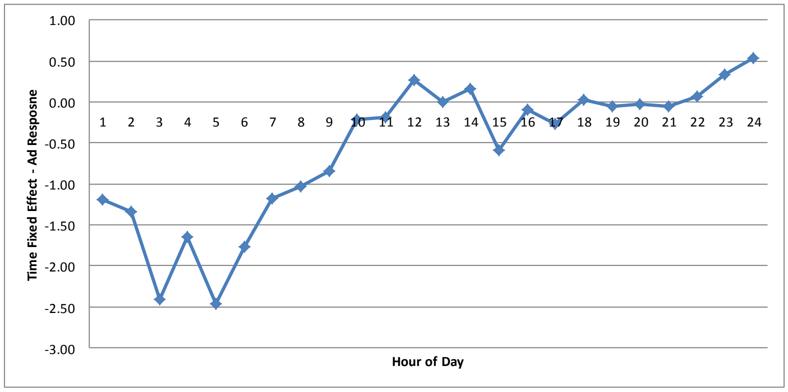

**Figure 6: Proportions of Ad Impressions by Hour of Day**