Player Analytics: Papers from the AIIDE Workshop
AAAI Technical Report WS-16-23

The Sensitivity of Retention to In-Game Advertisements: An Exploratory Analysis

Zachary Burns
University of San Francisco
San Francisco, California 94117
zcburns@usfca.edu

Isaac Roseboom
DeltaDNA
Edinburgh, Scotland
isaac.roseboom@deltadna.com

Nicholas Ross
University of San Francisco
San Francisco, California 94117
ncross@usfca.edu

Abstract

In-game advertisements have become a significant source of revenue for developers of freemium mobile games. Despite this, we are not aware of any published exploratory study looking into the effect of these advertisements on individual, user-level behavior. We consult existing literature on the psychology of advertising and interruption, and propose several metrics that could be used to measure user-level retention. Using a unique dataset representing 21 free-to-play mobile games, we create a simple model which attempts to determine the effect of advertisements on retention. We find only a weak relationship between advertisements and retention. More importantly we find that game-specific effects dominate all advertising effects, yielding important results for game development and design decisions. While the results presented have limitations, they provide a starting point for future research on individual level analysis.

Over the last decade, there has been a well-documented rise in "freemium" or free-to-play video games, where users are not charged for core gameplay and revenue is generated via selling advanced features or additional content (Kumar 2014; LeJacq 2012). In particular, it is estimated that over 80% of the $10 billion mobile video game market is using free-to-play (LeJacq 2012). Given that only a small percentage of players will spend money (Swrve 2014), many developers have turned toward in-app or in-game advertisements to bolster revenue.

There is little publicly available quantitative knowledge about the effect of these advertisements[1] though, anecdotally, game developers strongly believe that advertising increases revenue at the expense of retention (Raveh 2016):

> All too often, developers view serving in-game ads as a necessary evil and see a tradeoff between monetization and user retention.

This logic seems face-valid, as advertising is considered to be intrusive and contribute negatively to the gaming experience. Overall, there exists a strong negative sentiment to inclusion of advertising on the side of development (Munoz 2015; Raveh 2016)

The present paper attempts to quantify this aversion, using data at the level of an individual user. We first review existing literature in psychology to help us identify what retention might look like for an individual user. Then, using a novel data set covering over 20 mobile, free-to-play games, we attempt to determine how individualized behavior is affected by advertising.

Our primary result is perhaps surprising: using the metrics we propose, we find that across titles, in-game advertisements have a weak, insignificant effect on each measure of retention. We further find that the game itself significantly determines future retention metrics. This result has important implications for how games are designed and produced. Our analysis (which has a number of limitations described below) implies that the effect of advertising is marginal.

The organization of this paper is as follows. We begin with a technical description of in-app advertising, which presages the unavoidable limitations in our data, and review where this paper fits in existing literature. We next review literature which indicates what metrics might be used to measure individualized retention. We then describe our data, which leads to our model and conclusions. The final section contains a discussion with a focus on future research directions.

Primer on in-game advertising

Before continuing, we provide a brief description of the $100 billion market on in-game and in-app advertisements (eMarketer 2015). There are a number of different ways that advertisements can be served to users, though the most common is via an integrated SDK and by partnering with an advertising network ("ad network") or other partner (Allen 2014). In this situation, the content provider determines which subset of users and under which circumstances an advertisement is shown. To be explicit, consider the situation where a developer wishes to show an advertisement to users after they complete the tutorial of a game. Once the user finishes the tutorial, they game developer sends information to the ad network and the ad network responds with an advertisement (if any) to show that user. Importantly, both the developer and the ad network have private information re-

[1]The exception to this is from companies selling in-game advertising which purport no negative effects stemming from advertising, though, given their inherent bias it may be fair to assume selective reporting.

garding this transaction – the developer knows about the user while the ad network has information about the demand for showing ads.

A subtle part of this this process is that there are many different types of advertisements, such as video, banner and interstitial, and the developer has discretion over the type of advertisement shown (Foroughi 2015), though all might be considered intrusive by a user.

Given the immaturity of the field, the quickness with which it changes, and the difficulties of untangling the many threads above there are few academic studies directly addressing this topic. That said, this paper complements a number of different strands of literature, though it bears the strong similarity to studies looking at user-level data in video game analytics. For example, (Sifa et al. 2015; Weber et al. 2011; Mahlmann et al. 2010; Xie et al. 2015) attempt to predict events such as churn and monetization from user-level data in video games. In particular, the authors determine which factors of a user's game play predict managerially important metrics, though none focus on advertising.

There are also papers that have studied the effect of in-game advertising on user behavior (e.g. (Huang and Yang 2012; Glass 2007; Yang et al. 2006)). Researchers in this area tend to focus on the effect of advertising on user's preferences for the product advertised. Our paper approaches in-game advertising differently since it focuses squarely on user-level retention.

Measuring user-level retention

Retention is generally understood to be imperative in monetizing freemium games. While there are reasonably standard measures at the level of the game (e.g. comparing daily downloads to change in daily active users), analyzing retention with respect to advertising on the level of the game is difficult, as there is high churn and advertising is heterogeneous across users. By individuating users, we can use existing data to observe naturally varying rates of advertisement and how that affects retention. Because there is no agreed-upon measure of retention at the level of the user, we look to previous research to inform how we might create one.

Very broadly, a user's decision to continue playing considers the benefits relative to the costs of playing. The aversion to including in-game advertising during the design phase likely results from an intuition that it affects one of these: advertising might reduce the benefits of playing (by making the game less enjoyable) or might increase the costs of playing (by 'charging' the player more relative to other equally enjoyable pursuits). Marketing psychology literature has considered the effect of digital advertisements on individual's content consumption, and has found mixed evidence for these intuitions.

For example, (Goldstein, McAfee, and Suri 2013; Goldstein et al. 2014) evaluate how variously annoying advertising changes the length of time a person works on a paid task. They find that in the presence of more annoying (often animated) ads, people worked for a shorter amount of time. For a fixed benefit, increasing the cost results in less persistence. If advertising causes costly annoyance, the analyses should show that users shown more advertisements spend less total time playing.

(McCoy et al. 2004) studied the effect of pop-up and in-line ads on website visits, finding that self-reported behavioral intentions to return to sites with no ads were higher than those with ads. Put differently, advertising reduces the overall benefits, resulting in reduced interest in re-engagement. If advertising makes the gaming experience less enjoyable, the analyses should show that users shown more advertisements will wait longer to initiate new sessions, and initiate fewer total sessions.

In contrast, (McCoy et al. 2004) also contains evidence that might predict nearly opposite results when applied to games. Content from the target website was remembered equally well regardless of ad inclusion, indicating that the annoyance of ads did not interfere with engagement in a task for which there was no intrinsic motivation and enjoyment. If users' engagement in a game - a task selected by the user - is unaffected by ads, the rate of return to the game might not change. A further counterintuitive result is provided by (Nelson, Meyvis, and Galak 2009) demonstrating that people enjoyed television programs *more* when shown ads.[2] This is despite the fact they predicted they would enjoy the program less. Taken together, there is reason to predict that in-game advertising could have a neutral or positive impact on engagement and enjoyment (and hence retention). Furthermore, self-reported intentions or predictions may not be related to actual behavior.

Overall, the extent literature doesn't provide a conclusive answer to the effect of in-game advertising on retention. The literature does, however, give us insight into some observable behaviors (already present in most game datasets) we might use to measure retention on the individual user level.

Data

Our data comes from a sample of 21 free-to-play mobile games distributed via Google Play or Apple's App Store. The games in our sample use DeltaDNA as an analytics provider to track in-app advertising. DeltaDNA provides storage and analysis services for video game developers. DeltaDNA's system consists of an SDK, which developers integrate into their code, that sends information conditional on an event occurring to DeltaDNA's servers.

For our sample, we focus on users who installed within the two week period between June 29th and July 7th, 2016.[3] We focus on users who are new to the application during this period to avoid commingling other issues. Users are then tracked until two weeks after July 7th.

As part of our analysis we identify users who have only played a single session. Unfortunately, neither platform provides a consistent, user-level method of identifying if a user has uninstalled. We therefore assume that users who have

[2]The authors argue that people are likely to adapt to experiences, making them less enjoyable over time. (The first piece of candy corn is amazing, the n^{th} piece is disgusting). Periodic interruptions mitigate adaptation, and enhance enjoyment.

[3]Unless otherwise stated, all dates and times are in UTC.

Table 1: Summary Statistics. This table contains information about the characteristics of the 21 applications in the sample.

Statistic	Value
Installs Per Day, Per Game	338.84
Sessions Per User	9.50
Total Ads Shown	1,522,263
Average Users Per Game	4,743.81
Total Users	99,620

Table 2: Comparing Users with and without advertising

Statistic	With Ads	Without Ads
Number of Users	31,784	67,836
Lifetime Sessions	14.42	7.29
Total Minutes Played	131.01	41.42
Average Session Length	10.87	7.89
Ads Per Session	3.37	–

a single session and then fail to engage with the application again during this time period are considered single session users. Given the known retention rates in the industry, however, the likelihood of someone coming back after two weeks is negligible.

We restrict our attention to games that have more than 100 installs, 1,000 new users and 3,000 active users every day in the sample period. This is done to ensure that our sample contains only games which are active and have a sizable enough user base to be able to derive results. In order to make the size of the data more manageable, we randomly sampled a sixteenth of the users for each game.

Our sample is restricted to titles which show advertisements to their users. In order to identify these titles, we studied the events each game tracked for each user and identified events associated with advertisements. Our final preanalysis step removed users with obvious data errors, such as sessions spanning months or the existence of in-game events recorded before a user installs. Table 1 contains summary information on our sample, highlighting some important features of our sample.

Firstly, while we have taken a sample of users who installed these applications, our total sample size remains robust. This can be seen in the per-game statistics within our sample, which show almost 5,000 installs per game. Users in our sample have an average of 9.5 sessions, but this number is right skewed, as around a third of our dataset is users with only a single session.

Secondly, the total ads shown to users within our sample is around 1.5 million so that the average number of ads per user is 15. As Table 2 demonstrates this number is skewed toward a small group of users. This table shows user behavior as segmented between those who have and have not seen advertisements. Note that the segmentation in this table is forward-looking, as users who see an advertisement in a future session are classified as an ad viewer over their entire lifetime.

Table 2 demonstrates that ad viewers tend to be more engaged. The primary reason that this occurs is that ads only occur during gameplay. Users that install and quit shortly thereafter do not have the opportunity to be served an advertisement. As discussed in the introduction, users are subject to the advertising at a combination of the developer's and ad network's discretion. One commonly mentioned best practice is to only serve ads to engaged users. In other words, many game companies refrain from serving ads to users in their first few sessions.

Limitations of the data

We use a large dataset from an existing repository for the present analysis, which is exploratory in nature. In a world ideal for this paper, the authors would require developers to standardize conventions for both naming ad events and session measurement, as well as randomize both the timing of and pattern of advertising across users. The former requirements are impractical and the latter a probable impossible sell to studios (many of which may be adding advertising post-hoc, and not integrating it during the design process).

A caveat of our data, as hinted at in the previous discussion, is that we are only considering the real-world range of possible advertisements. In particular, our data is subject to what game developers are willing to do with advertisements. This has implications for our ability to measure the effect of in-game advertising on users and the interpretation of our results. For example, none of the games in the sample ever spam a user with more than one ad per minute over their lifetime.

A final limitation of our data deals with instrumentation between games. Because developers instrument their games differently and game mechanics vary, it is difficult to compare specific in-game behavior between titles. For example, one game may have a level-up event while another may eschew the concept of levels completely. In order to abstract from game specific measures which would needlessly decrease our sample, we choose to focus measures which are consistently defined between games. [4] As described in the next section, this leads us to three simple retention measures.

Model and Results

In this section we analyze how the presence of advertisements in the first session affect in-game user behavior, with a focus on retention. We chose three metrics to assess retention behavior in the games and consider two sets of data: our full dataset and the subset of users who have more than a single session. In other words, this second dataset excludes the roughly 1/3 of users who did not initiate a second session.

The first retention measure is the number of minutes between install and initiation of a second session. To reduce the impact of outliers, this metric, which captures an intuitive "thirst to replay," is log-transformed. Because this metric requires a second session, we only apply this measure to our second dataset, the subset of users with more than one session.

[4] Due to the nature of how DeltaDNA's SDK is implemented, the recording of session and install information is robust to developer specific implementations.

Table 3: Regression results describing the effect of advertising on multiple retention measures. Each of the columns (1) - (5) represents a linear regression on the dependent variable, listed as the column heading, and first session ad density, a constant and game specific factors (which can be found in Table 4). The first number is the estimated coefficient while the second, in parenthesis, is the standard error.

	log(Time to 2^{nd} Session)	log(Total Session Time)		log(Number Sessions)	
	(1)	(2)	(3)	(4)	(5)
Ad Density	0.333	−0.200	−0.865	0.236	−0.314
	(0.818)	(0.690)	(0.584)	(0.314)	(0.295)
Constant	8.368***	2.172***	3.976***	0.676***	1.391***
	(0.132)	(0.087)	(0.094)	(0.040)	(0.048)
Game specific effects in next table					
Observations	65,422	99,620	65,422	99,620	65,422
R^2	0.142	0.216	0.279	0.073	0.094
Adjusted R^2	0.142	0.216	0.279	0.072	0.094
Residual Std. Error	2.592	2.456	1.851	1.119	0.936
F Statistic	515.144***	1,309.938***	1,207.096***	371.498***	324.692***
Degrees of freedom	21, 65400	21, 99598	21, 65400	21, 99598	21, 65400

Note: *p<0.1; **p<0.05; ***p<0.01

The second measure is the total time spent playing the game, across all sessions. We analyze this metric both with and without users who only ever initiate a single session. This metric is likely to be the most important with respect to monetization of ads, as clearly, more time spent on the game allows for a greater number of ads to be served. We also log-transform this variable.

Our third and final retention measure is the number of sessions initiated. We analyze this metric both with and without single-session users, and log-transform it.

Each metric has its own advantages and disadvantages in terms of capturing retention. Due to our desire to capture behavior across as many games as possible, there is inherently great heterogeneity in both gameplay and event-capturing across games. For example, some games feature a real-time recharge; thus, highly engaged players interested in maximizing their sessions may choose (a) to wait longer to initiate a second session and (b) initiate fewer sessions across the sampling period. Thus, these metrics would incorrectly identify this user as being relatively weakly retained. Different games are likely to define sessions differently, and the number of sessions initiated, as well as total session time, might be artifacts of this definitional choice, not actual retention. Given that these attributes are unobservable in our data, and our desire to measure behavior over a broad range of games, we modeled all three metrics separately, the portfolio of which provides an overarching picture of retention behavior.

The pairwise correlations between the three retention metrics were quite small, none reaching more than 7% in absolute value. We chose the metrics to be simple, face-valid measures of retention, and find that they seem to be largely independent of one another.[5] This is perhaps unsurprising, given the preceding paragraph. It does, however, present a problem if the different metrics show different patterns. As reported below, this was not the case.

We used a single independent variable to predict each of the three retention behaviors: ads per minute shown in the first session (which we call 'first session ad density' for the remainder of the paper). As noted above, 1/3 of the sample didn't initiate a second session; for those users, this is the only session in which ads influence their decision not to return. For all users, this metric provides the cleanest slate for the downstream effects of advertising on behavior. By the second session, users have presumably incorporated ad frequency into their decision to continue playing the game.[6]

Continuing our discussion of the variables which feed into our model a few basic results emerge which foreshadow the main results reported below. One, first session ad density is independent of if a user initiates a second session ($t(91803) = 0.97$, *n.s.*). Secondly, across the entire sample, first session ad density is slightly negatively correlated with time spent in the first session ($r = -.007$, $p = .035$). These statistics are largely driven by the fact that users in the sample are rarely shown ads in their first session: fully 81,868 users were had no first-session ad events.[7]

Our main model uses linear regression to predict each of our retention metrics from first session ad density. We find no evidence that first session ad density impacts any of the three retention behaviors. We do, however, find large effects of individual games on the retention behaviors. Any impact of first session ad density is swamped by these individual game effects.

In Tables 3 and 4 the first column estimates the log seconds until initiation user's second session (if the user had

[5] Future research should focus on standardizing definitions and measurement of retention behavior. While our broad analyses don't necessitate this, finer-grained analyses are difficult to interpret without these standards.

[6] This is generally a strong assumption, but one that doesn't impact our results. Future research should begin to explore how changes in ad service over multiple sessions impact user continuation.

[7] The same analyses on only users who saw ads in their first session were not appreciably different. Ad density comparison: $t(8699.6) = 0.97$, *n.s.*. Correlation: $r = -.022$, $p = .003$.

Table 4: Regression results describing the effect of advertising on multiple retention measures. Each of the columns (1) - (5) represents a linear regression on the dependent variable, listed as the column heading, and first session ad density, a constant (both of which can be found in Table 3) and game specific factors. The first number is the estimated coefficient while the second, in parenthesis, is the standard error.

	log(Time to 2^{nd} Session)	log(Total Session Time)		log(Number Sessions)	
	(1)	(2)	(3)	(4)	(5)
Game B	3.210^{***} (0.237)	0.756^{***} (0.173)	-0.329^{*} (0.169)	0.424^{***} (0.079)	0.333^{***} (0.086)
Game C	3.561^{***} (0.161)	0.626^{***} (0.109)	-0.584^{***} (0.115)	0.056 (0.050)	-0.081 (0.058)
Game D	1.260^{***} (0.154)	1.287^{***} (0.107)	0.515^{***} (0.110)	0.695^{***} (0.049)	0.631^{***} (0.055)
Game E	0.921^{***} (0.140)	0.420^{***} (0.095)	-0.178^{*} (0.100)	0.645^{***} (0.043)	0.575^{***} (0.051)
Game F	-3.322^{***} (0.144)	-0.703^{***} (0.094)	-1.754^{***} (0.103)	-0.171^{***} (0.043)	-0.246^{***} (0.052)
Game G	-0.109 (0.177)	-11.377^{***} (0.118)	-12.135^{***} (0.127)	0.236^{***} (0.054)	0.392^{***} (0.064)
Game H	-0.235 (0.145)	-0.631^{***} (0.097)	-1.077^{***} (0.103)	0.137^{***} (0.044)	0.095^{*} (0.052)
Game I	-2.807^{***} (0.200)	1.129^{***} (0.163)	-0.428^{***} (0.143)	1.264^{***} (0.074)	0.685^{***} (0.072)
Game J	-0.101 (0.133)	0.625^{***} (0.088)	-0.267^{***} (0.095)	0.342^{***} (0.040)	0.181^{***} (0.048)
Game K	0.786^{***} (0.141)	0.972^{***} (0.096)	0.386^{***} (0.101)	1.011^{***} (0.044)	0.945^{***} (0.051)
Game L	0.048 (0.144)	-0.533^{***} (0.099)	-1.233^{***} (0.102)	0.690^{***} (0.045)	0.411^{***} (0.052)
Game M	-8.088^{***} (0.187)	-0.905^{***} (0.136)	-2.663^{***} (0.133)	1.131^{***} (0.062)	0.438^{***} (0.067)
Game N	0.025 (0.136)	0.828^{***} (0.091)	-0.322^{***} (0.097)	0.509^{***} (0.041)	0.275^{***} (0.049)
Game O	0.419^{***} (0.148)	-0.954^{***} (0.100)	-1.917^{***} (0.106)	0.288^{***} (0.046)	0.237^{***} (0.054)
Game P	0.230^{*} (0.137)	1.274^{***} (0.092)	0.202^{**} (0.098)	1.063^{***} (0.042)	0.901^{***} (0.050)
Game Q	1.292^{***} (0.187)	-0.316^{**} (0.124)	-0.747^{***} (0.134)	0.030 (0.056)	0.063 (0.068)
Game R	0.407^{***} (0.148)	1.069^{***} (0.103)	-0.102 (0.106)	0.791^{***} (0.047)	0.521^{***} (0.053)
Game S	-0.162 (0.142)	1.161^{***} (0.097)	0.104 (0.101)	0.957^{***} (0.044)	0.784^{***} (0.051)
Game T	0.530^{***} (0.142)	0.452^{***} (0.097)	-0.020 (0.101)	0.916^{***} (0.044)	0.847^{***} (0.051)
Game U	0.708^{***} (0.140)	0.463^{***} (0.094)	0.011 (0.100)	0.522^{***} (0.043)	0.551^{***} (0.051)

Note: *p<0.1; **p<0.05; ***p<0.01

more than one session). A lower number indicates a shorter time to initiate a second session, which we interpret as an increased retention signal. The constant therefore estimates the average time between the first and second session of a Game A user, which in the absence of ads is $e^{8.368} \approx 4307$ seconds or 1.2 hours.[8] This user, shown an additional ad per minute in the first session, will initiates a second session at 1.67 hours. By contrast, the average user of Game D initiates the second session 4.2 hours ($e^{8.368+1.260} \approx 15184$) after the first session. In other words, there is a large variance between the types of play experiences present in these games.

[8]Game is treated as a factor in our model and thus the constant term represents a use of the omitted game, which we define as Game A.

The second and third columns estimate the log total seconds spent by a user across all sessions. Column 2 estimates the full sample, while column 3 estimates only users who had multiple sessions (i.e. the same sample used in column 1). A higher number indicates more total time spent playing the game, which we interpret as an increased retention signal. The constant estimates that the Game A user spends 9 minutes in total on a game, or 53 minutes if they return for multiple sessions. This Game A user, shown additional ad per minute in the first session, now spends 7 minutes in total on a game, or 22 minutes if they return for multiple sessions. By contrast, the average user of Game D spends 31 minutes on the game, or 89 minutes if they return for multiple sessions.

The fourth and fifth columns estimate the log total sessions initiated by a user within the sample. Column 4 estimates the full sample, while column 5 estimates only users who had multiple sessions (i.e. the same sample used in column 1). A higher number indicates more total sessions initiated, which we interpret as an increased retention signal. The constant estimates that the Game A user initiates 2 sessions, or 4 sessions if they return for multiple sessions. This Game A user, shown additional ad per minute in the first session, now initiates 2.5 sessions, or 3 sessions if they return for multiple sessions. By contrast, the average user of Game P initiates 6 sessions, or 10 sessions if they return for multiple sessions.

The above examples are intended only to demonstrate the magnitude of difference between ad density and individual game effects. Ad density in the first session had a consistently non-significant effect on the three retention measures. By comparison, the effects of individual games were far more predictive of retention behavior.

The implication of this result is that the negativity surrounding the use of in-game advertisements should be tempered. While the reputation of in-game advertisements is negative, our results show that, as currently implemented in games, this result is not necessarily universal.

A strength of this analysis is the use of multiple titles with a consistent data generating process. While the process of by which an advertisement is shown is complex and subject to many decision points the analysis above abstracts above that, yielding an exploratory result which simply asks: What happens to a user's retention when they are exposed to an advertisement? Unlike many experimental results, including some listed in the literature review, our results do not attempt to control for all issues. Instead our analysis provides more of a snapshot of the relationship between ads, as they are currently served in-game, and retention.

In order to evaluate the results of this simple model we also reran the above including a number of other user level features including platform, country and device type. In each these situations, the broad strokes of our analysis remain unchanged: showing advertisements did not significantly influence retention. We also ran the analysis above with a 1/0 flag on if the user had seen ads, rather than the ad density and saw no changes. A final robustness check was completed by, instead of estimating each retention measure, we built a logistic regression model attempting to predict if a user would return, based on their advertising experience. As with the rest of the robustness checks, the results were consistent.

Discussion

The purpose of our paper was to consider the effect of ingame advertising on a user's behavior. Using a sample of 21 games that feature in-game advertising we found no systematic evidence of an effect of advertising on a user's retention. In other words, this paper stands in opposition to the common wisdom that in-game advertisements strongly and negatively affect in-game behavior. These results were robust to a number of different modeling choices, variables, and specifications, leading us to believe that they represent concrete evidence of the lack of an effect from in-game advertising.

We did find that the game itself was a significant predictor of retention, implying that game developers should focus their energy on creating satisfying in-game experiences in order to increase retention.

Our research is subject to a number of significant caveats. First, since we are using live data, our analysis is constrained by how developers are currently serving advertisements. Secondly, all of our retention measures were game-agnostic and did not focus on game-specific measures of engagement, such as levels played or points scored. Our final caveat is that showing an advertisement is the result of a complex (and unmodeled in our research) process involving multiple parties. Each of these caveats influences our results.

That being said, our analysis is exploratory and meant to address one of the more commonly held beliefs about advertisements and their effect on game play. To that end, this study paints a much more nuanced picture of the effect of in-game advertising on a user's behavior.

Finally, we believe that studying the effect of in-game advertisements could yield important information regarding player behavior that would be valuable for game designers and developers. Untangling some of the caveats above would yield significant changes to both how games are designed and understood.

References

Allen, S. 2014. Monetizing your mobile app: Ad networks. http://www.accella.net/knowledgebase/monetizing-your-mobile-app-ad-networks/.

eMarketer. 2015. Mobile ad spend to top $100 billion worldwide in 2016, 51% of digital market. http://www.emarketer.com/Article/Mobile-Ad-Spend-Top-100-Billion-Worldwide-2016-51-of-Digital-Market/1012299.

Foroughi, A. 2015. Which mobile ad formats are best? a down and dirty guide that spells it out. http://venturebeat.com/2015/03/10/which-mobile-ad-formats-are-best-a-down-and-dirty-guide/.

Glass, Z. 2007. The effectiveness of product placement in video games. *Journal of Interactive Advertising* 8(1):23–32.

Goldstein, D. G.; Suri, S.; McAfee, R. P.; Ekstrand-Abueg, M.; and Diaz, F. 2014. The economic and cognitive costs of annoying display advertisements. *Journal of Marketing Research* 51(6):742–752.

Goldstein, D. G.; McAfee, R. P.; and Suri, S. 2013. The cost of annoying ads. In *Proceedings of the 22nd international conference on World Wide Web*, 459–470. ACM.

Huang, J.-H., and Yang, T.-K. 2012. The effectiveness of ingame advertising: The impacts of ad type and game/ad relevance. *International Journal of Electronic Business Management* 10(1):61–72.

Kumar, V. 2014. Making 'freemium' work. *Harvard Business Review*.

LeJacq, Y. 2012. Something for nothing: How the videogame industry is adapting to a 'freemium' world. *International Business Times*.

Mahlmann, T.; Drachen, A.; Togelius, J.; Canossa, A.; and Yannakakis, G. N. 2010. Predicting Player Behavior in Tomb Raider: Underworld. In *Proc. of IEEE CIG*.

McCoy, S.; Everard, A.; Galletta, D.; and Polak, P. 2004. A study of the effects of online advertising: A focus on pop-up and in-line ads. *SIGHCI 2004 Proceedings* 11.

Munoz, M. 2015. The (real) effect of ads on mobile games retention. *Gamasutra*.

Nelson, L. D.; Meyvis, T.; and Galak, J. 2009. Enhancing the television-viewing experience through commercial interruptions. *Journal of Consumer Research* 36(2):160–172.

Raveh, J. 2016. In-game advertising in free-to-play mobile: Getting the , balance right. http://www.gamasutra.com/blogs/JonathanRaveh/20160317/268363/InGame_Advertising_in_FreeToPlay_Mobile_Getting_the_Balance_Right.php.

Sifa, R.; Hadiji, F.; Runge, J.; Drachen, A.; Kersting, K.; and Bauckhage, C. 2015. Predicting Purchase Decisions in Mobile Free-to-Play Games. In *Proc. of AAAI AIIDE*.

Swrve. 2014. Swrve Monetization Report 2014. http://landingpage.swrve.com/rs/swrve/images/swrve-monetization-report-0114.pdf.

Weber, B. G.; John, M.; Mateas, M.; and Jhala, A. 2011. Modeling Player Retention in Madden NFL 11. In *Proc. of IAAI*.

Xie, H.; Devlin, S.; Kudenko, D.; and Cowling, P. 2015. Predicting Player Disengagement and First Purchase with Event-frequency Based Data Representation. In *Proc. of CIG*.

Yang, M.; Roskos-Ewoldsen, D. R.; Dinu, L.; and Arpan, L. M. 2006. The effectiveness of" in-game" advertising: Comparing college students' explicit and implicit memory for brand names. *Journal of Advertising* 35(4):143–152.