



## Information Systems Research

Publication details, including instructions for authors and subscription information:  
<http://pubsonline.informs.org>

### When to Play Your Advertisement? Optimal Insertion Policy of Behavioral Advertisement

Subodha Kumar, Yinliang (Ricky) Tan, Lai Wei

To cite this article:

Subodha Kumar, Yinliang (Ricky) Tan, Lai Wei (2020) When to Play Your Advertisement? Optimal Insertion Policy of Behavioral Advertisement. Information Systems Research

Published online in Articles in Advance 09 Jun 2020

. <https://doi.org/10.1287/isre.2019.0904>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact [permissions@informs.org](mailto:permissions@informs.org).

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

# When to Play Your Advertisement? Optimal Insertion Policy of Behavioral Advertisement

 Subodha Kumar,<sup>a</sup> Yinliang (Ricky) Tan,<sup>b</sup> Lai Wei<sup>c</sup>
<sup>a</sup> Fox School of Business, Temple University, Philadelphia, Pennsylvania 19122; <sup>b</sup> A. B. Freeman School of Business, Tulane University, New Orleans, Louisiana 70118; <sup>c</sup> Antai College of Economics and Management, Shanghai Jiao Tong University, Shanghai 200030, China

**Contact:** subodha@temple.edu,  <https://orcid.org/0000-0002-4401-7950> (SK); ytan2@tulane.edu,

 <https://orcid.org/0000-0001-8971-818X> (Y(R)T); laiwei@sjtu.edu.cn,  <https://orcid.org/0000-0003-2083-1752> (LW)

**Received:** April 27, 2018

**Revised:** April 5, 2019; August 23, 2019

**Accepted:** October 2, 2019

**Published Online in Articles in Advance:** June 9, 2020

<https://doi.org/10.1287/isre.2019.0904>
**Copyright:** © 2020 INFORMS

**Abstract.** Digital advertisements offer a full spectrum of behavioral customization for timing and content capabilities. The existing research in display advertising has predominantly concentrated on the content of advertising; however, our focus is on optimizing the timing of display advertising. In practice, users are constantly adjusting their engagement with content as they process new information continuously. The recent development of emotional tracking and wearable technologies allows platforms to monitor the user’s engagement in real time. We model the user’s continuous engagement process through a Brownian motion. The proposed optimal policy regarding the timing of behavioral advertising is based on a threshold policy with a trigger threshold and target level. Specifically, the platform should insert the advertisement when the user’s engagement level reaches the trigger threshold, and the length of the advertisement should let the user’s engagement level drop to the target level. Analogous to the familiar idea of “price discrimination,” the methods we propose in this study allow the platforms to maximize their revenue by “discriminatory” customization of the timing and length of the advertisement based on the behavior of individual users. Finally, we quantify the benefits of the proposed policy by comparing it with the practically prevalent policies (i.e., preroll, midroll, and a mix of the two) through a simulation study. Our results reveal that, for a wide range of settings, the proposed policy not only significantly increases the platform’s profitability but also improves the completion rate at which consumers finish viewing the advertisement.

**History:** Xiaoquan (Michael) Zhang, Senior Editor; Zhengrui (Jeffrey) Jiang, Associate Editor.

**Funding:** Lai Wei received financial support from the National Science Foundation of China [Grants 71801151, 71531010, and 71831006] and Shanghai Pujiang Program [Grant 17PJJC065].

**Supplemental Material:** The e-companion is available at <https://doi.org/10.1287/isre.2019.0904>.

**Keywords:** behavioral advertisement • wearable technology • emotional analytics • optimal policies • Brownian motion

Strategy and timing are the Himalayas of marketing. Everything else is the Catskills. —Al Ries, cofounder and chairman of Ries & Ries (2005)

## 1. Introduction

According to a recent report in the *Wall Street Journal*, digital advertisement spending in the United States surged 22% from the previous year to record \$72.5 billion in 2016 (Shields 2017). More strikingly, in 2017, digital advertising in the United States surpassed television spending for the first time by a great margin, an increase fueled largely by mobile advertisement spending (Slefo 2017). In 2016, spending on mobile advertising surged 77% to \$33.6 billion (Shields 2017). Juniper Research estimates that global digital advertising spending across mobile, wearable, and online devices will exceed \$285 billion by 2020 (Bajpai 2016). The world is rapidly moving toward digitalized advertising, challenging older channels and introducing innovative strategies. However, finding

the right moment to reach consumers remains a crucial but open question for the digital advertising industry (Gupta 2015).

Compared with traditional television advertisements, for which the content provider has to play the same advertisement to the entire audience at the same time, the digital advertisement provider can choose not only what content to play but also when to play it. Thus, the digital advertisement offers a full spectrum of customization capabilities regarding the content and timing of the advertisement. Existing research and industry practice have predominantly considered the content dimension, in which the content provider can personalize the digital advertising content based on users’ browsing behavior, location, and search histories (Chen and Stallaert 2014, Guo et al. 2015). However, it is less clear what timing is optimal for the digital advertisement provider. For example, Facebook is poised to test preroll video advertising (i.e., a promotional message that plays before the

content) on its video-on-demand service “Facebook Watch” after a longtime ban of the preroll format (Sloane 2017). Around the same time, Google is limiting the use of preroll advertisements on YouTube (Reale 2017). It is unarguably true that timing is very important for the stakeholders involved, but there are no clear answers to questions such as when the platform should play the advertisement and how long each advertisement should be. Further, there is a lack of rigorous academic research on this crucial but overlooked topic. We aim to address this gap by investigating the optimal policy for the content provider regarding the *timing* of playing its advertisements on platforms.

With the rapid development of relevant technologies, monitoring the user’s engagement or feeling in real time is no longer science fiction but the hard reality. There are many successful stories showing how mobile devices, including wearable technologies, can track the user’s engagement activities. For example, Ginger.io uses continuous data to help people track their moods and help doctors track the health of their patients (Vivero 2018). Activity trackers such as Fitbit series and Spire can collect not only users’ biometric data (i.e., heart rate, distance walked or run, temperature, breath, quality of sleep) but also data on the user’s mood through hidden sensors (Reddy 2015). In most cases, these devices are connected and synced with the user’s mobile phone wirelessly, which provides a means to track behavior continuously. Recently, FOVE, a cutting-edge virtual reality (VR) manufacturer, introduced the first eye-tracking headset in the world. Eye tracking is a well-established research methodology to trace emotion and attention (Teixeira et al. 2012). Since marketing researchers began promoting this stream of research in 1978 (Russo 1978), eye tracking has been successfully applied in many fields of study (Wedel and Pieters 2008). Facial-tracking technology has also helped to identify people’s emotions. Both Apple and Nielsen bought startups specializing in identifying users’ feelings through facial-cue recognition (Foster 2016). In summary, rapid and disruptive technology innovations on emotional analytics allow the advertisement platform to track and analyze users’ emotions and engagement in real time.

Firms have also begun to use tracking information to improve advertisement decision. For example, during the Wimbledon tennis championships, Jaguar used wearable technology and court sensors to measure the mood and emotion of the crowd to facilitate a social media campaign (Faull 2015). Mars, Inc., a world-leading food company, used an emotion-detection app developed by Massachusetts Institute of Technology (MIT) to evaluate different advertisements (de Jesus 2018).

The key focus of our study is to address the optimal timing and length for digital advertising, especially in the context of displaying advertising. Against the backdrop of exponential growth of digital advertisement and emerging capabilities to track the user’s feelings, it is urgent and crucial for academic research to provide theoretical support to guide the practice. In this study, we formulate a dynamic advertising model in which, at each instant, the user’s engagement may increase or decrease from watching the content. The random update of user engagement behavior induces a stochastic process, which we capture through a drifted Brownian motion. Our study provides a framework to model and study the issue of timing in digital advertising. More importantly, our analysis provides critical insights to the practitioners by offering them actionable strategies regarding the timing of advertisement. Next, we highlight the key questions addressed in this study.

The first key question that we investigate is as follows: What should the optimal insertion policy be for digital advertisement when the platform can track the user’s engagement level? More specifically, what is the optimal timing and length of each advertisement? As we pointed out earlier, there is scant rigorous academic research to address the timing of advertisement, but the answer to this question is crucial for many platforms whose revenue mainly relies on advertisements. Currently, most platforms use the preroll (i.e., inserting the advertisement at the very beginning of the content), midroll (i.e., inserting the advertisement in the middle of the content), and mixed strategies. The answer to our research questions may shed light on this urgent issue. Intuitively, the structure of the optimal policy might be very complex, because users dynamically adjust their engagement according to content and random factors. Surprisingly, we find that a simple threshold policy is the optimal policy. When the user’s engagement level reaches a certain trigger threshold, an advertisement with the appropriate length is placed such that the user’s engagement level drops to the prespecified target level. This process will repeat every time that consumer’s engagement level reaches the trigger threshold until the consumer leaves the platform. From a conceptual point of view, our proposed policy is analogous to the familiar idea of “price discrimination,” which allows the platform or content provider to maximize its revenue by “discriminatory” customization of the timing and length of advertisements, based on the individual user’s behavior.

Next, to gain a deeper understanding of the optimal timing of advertisement under the finite horizon, we explore the following question: What is the structure of the optimal policy if the user can leave during the content period or the content length is restricted? One

may intuitively think that the policy should become dynamic, because there is an additional dimension of uncertainty involved. However, we find that this is not the case. We analytically show that a simple threshold policy can still be the optimal policy. Our analysis in this setting also reveals another interesting finding under the finite horizon case. Counter to our intuition that the optimal trigger threshold will always decrease as the time approaches the end of the finite horizon, we find that it is optimal for the platform to first increase then decrease the trigger threshold when approaching the end of the planning horizon, because the platform needs to achieve a balance between the insertion time and the length of each advertisement. Our results here further corroborate the results of the infinite horizon case, which showed that the proposed policy possesses the property of simple implementation (i.e., threshold policy), which is very important for practitioners.

Moreover, to examine the performance of our proposed policies, we raise the following research question: How do the proposed policies perform compared with the prevalent industry policies (e.g., preroll, midroll, and mix strategy) in terms of profitability and service level? The answer to this question is highly relevant to platform practitioners who strive to increase revenue from advertisements. One might intuitively expect that the proposed policy would affect the user's retention rate negatively, as the platform can precisely insert the advertisement at the climax of the user's engagement with the content. Surprisingly, contrary to this intuition, our results reveal that the proposed optimal policy can yield not only a much higher revenue but also a lower leaving rate compared with the current prevalent policies. The managerial implication here is that the timing and length of the advertisement can have a significant impact on the platform's profit and user retention rate. From a practical perspective, the platform should have a library of advertisements with different lengths to implement the proposed policies.

We have also conducted sensitivity analyses on different parameters and identified an efficient heuristic to the optimal policy under the finite horizon. One noteworthy finding is about the volatility of the consumer's engagement level. On the one hand, as volatility increases, the consumer may leave the platform early, which suggests that the platform should play a shorter advertisement. On the other hand, as volatility increases, it takes less time for the consumer's engagement level to rise, which suggests that the platform can insert a longer advertisement. It is not obvious how the optimal policy will change as volatility increases. Interestingly, we find that the advertisement length should increase as volatility increases. Further, we also identify that the optimal policies under the

infinite horizon can be an efficient heuristic to the optimal policy under the finite horizon. This result is practically important because the policy under the infinite horizon is computationally more efficient, as the platform only needs to calculate once, in contrast with the optimal policy under the finite horizon where the platform has to dynamically update the thresholds as time passes.

The rest of the paper is organized as follows. In the next section, we briefly review the relevant literature and highlight our contributions with respect to past studies. In Section 3, we describe the model setting and dynamics of the user's engagement level. In Section 4, we analyze the base model, outline the intuition behind the main results, and conduct a sensitivity analysis. In Section 5, we extend our base model in various ways to obtain new insights. In Section 6, we conduct a simulation study to quantify the benefits of the proposed policies. The paper concludes with managerial implications and possible avenues for future research in Section 7.

## 2. Literature Review

Our study has points of contact with the literature in (i) pricing and scheduling of digital advertising, (ii) emerging literature on behavioral targeting, and (iii) applications of Brownian motion in marketing and information systems (IS), but it also deviates from the existing literature in some essential aspects.

To begin with, our research is closely related to the pricing and scheduling of advertising literature in the IS area. Tan and Mookerjee (2005) examine the joint marketing–information technology (IT) spending allocation problem and propose the optimal spending allocation between advertising and IT for electronic retailers. Further, Fan et al. (2007) analyze the tradeoffs between pricing and advertising strategies for media providers. They find that when the content quality is high and online access cost is low, companies can sell commercial-free programs online, whereas it is more beneficial to offer free programs with advertisement when online access cost is relatively high. Kumar et al. (2007) propose a variety of efficient solutions that solve the advertisement scheduling problem to maximize website revenue. Kumar and Sethi (2009) use optimal control theory to characterize the policy for an online platform where the revenue comes from both subscription fees and advertisements. Liu et al. (2012) find that in the presence of information processing constraints, although advertising may still cause a user to switch, it may not result in a sale. Chen et al. (2016) compare two prevalent revenue models (brokerage and advertising) for a user-to-user platform. Specifically, they discuss how the chosen revenue model affects the revenue of a platform, buyers' payoffs, sellers' payoffs, and social

welfare. Mookerjee et al. (2017) develop a model whose objective is to maximize the advertising firm's revenue subject to a click-through-rate constraint. Hao et al. (2017) use a two-sided market model to analyze the agency pricing for app sales by incorporating both the app developer and the platform owner. Sun et al. (2017) propose the optimal sequencing policies for the fading advertisement in the mobile application environment. The existing studies in advertising have mainly focused on the selection and scheduling of advertisement, and different revenue models for the advertising. Unlike the aforementioned studies, we focus on the timing of the advertising in the digital space. Specifically, we develop a framework for modeling the user's engagement continuously to characterize the optimal policy facing the content provider regarding when to play its advertisement and how long the advertisement should be for each individual user.

Second, our study intersects the emerging stream of literature on behavioral targeting of online advertising. Because of rapid development in information technologies, firms can collect detailed user behavioral information (i.e., taste, location, habits, and browsing history) and target individual users with tailored advertisements. Beales (2011) uses data collected from online advertising networks and finds that prices and conversion rates for behaviorally targeted advertisements are more than twice as high as those for traditional advertising. Zhao and Xue (2012) explore advertising competition when competing firms are asymmetrically informed about the user value. Chen and Stallaert (2014) study the economic implications when an online publisher engages in behavioral targeting. Leveraging the data of mobile phone users, Andrews et al. (2016) find that commuters in crowded subway trains are about twice as likely to respond to a mobile offer by making a purchase compared with those in noncrowded trains. Lu et al. (2016) use a large, individual-level clickstream data set to examine the impact of behavioral and contextual targeting on users' click behavior and find that the combination of both targeting methods leads to better click-through rates. Chen et al. (2017) use a duopoly model to show that targeting users based on their real-time locations on mobile platforms can increase firm profit compared with traditional targeting and uniform pricing. Shen and Villas-Boas (2018) investigate the effects of behavior-based advertising when preferences for two products are correlated. Existing literature highlights the strategic use of behavioral targeting, which is mainly based on the content of the advertising. Our study complements this emerging stream of literature by investigating the optimal timing of the insertion policy of digital advertisement, which is not well understood in the literature.

Third, our study also contributes to the applications of Brownian motion. Brownian motion and variations of Brownian motion have been applied to model the stochastic process in both marketing and IS literature. In the marketing literature, scholars have used Brownian motion to model the dynamics of product information search and consumer preference. For example, Branco et al. (2012, 2016) develop a Brownian motion framework to study the consumer's optimal stopping rule for the product search process. Ke et al. (2016) extend the previous studies by considering the information search when facing more than one decision alternative. Villas-Boas (2018) studies firms' repositioning decision when the dynamics of consumer preference follow a Brownian motion. In the IS literature, Zhang and Zhang (2015) use Brownian motion to model online traders' order and discuss whether the internet stock trading affects financial market equilibrium. IS researchers have also applied real option pricing method, which is based on the geometric Brownian motion, to study the service-oriented architecture migration (Ghosh and Li 2013) and cybersecurity investment (Benaroch 2018). Our research not only provides a theoretical framework on how to model the user's continuous engagement with the platform but also makes a unique contribution to the application of Brownian motion by considering impulse control. Specifically, in most existing studies, a firm's action does not affect the underlying process; even for the very few papers that indicate a firm's action can affect the underlying process  $X_t$ , their control is  $dX_t$  rather than  $X_t$ . As a result, the underlying process  $X_t$  remains continuous after the control. In contrast to these studies, the control in our setting (i.e., insertion of advertisement) directly influences the dynamic of the process and makes the underlying process  $X_t$  become discontinuous. Consequently, our problem becomes an impulse control problem and is much more difficult to solve.

To the best of our knowledge, this is the first attempt to construct a theoretical model of user engagement to optimize advertisement insertion policy. Our study not only sheds light on the content provider of the digital advertisement but also provides practical guidance to the emerging wearable technology industry. This study also makes a theoretical contribution to the digital advertising literature by developing a general framework to model the user's engagement levels in continuous time.

### 3. The Model

In this section, we first introduce the basic model and assumptions. This is followed by a detailed discussion of how advertisement affects the user's engagement and decisions facing the online platform.

### 3.1. User Engagement Level Process

Let us consider a scenario in which a user watches a new episode of a television (TV) series on a video streaming platform, plays a popular mobile game on a smartphone, or reads a classic novel on a mobile app. In all of these cases, the user’s engagement with the content changes constantly. Specifically, users will update their engagement level continuously depending on how much they enjoy the content, how the story unfolds, or how the game progresses. Following the framework of Branco et al. (2012, 2016), we model the process of consumers’ continuous engagement through a Brownian motion.<sup>1</sup> There are three noteworthy differences between our modeling techniques and those of Branco et al. (2012, 2016). First, to reflect the fact that advertisement is intrusive to the consumer’s engagement, we introduce the stochastic impulse control such that the insertion of advertisement will affect the consumer’s engagement level discontinuously, whereas there is no sudden jump in Branco et al. (2012, 2016). Second, the platform is able to track the consumer’s engagement level instantaneously. As a result, the optimal decisions in our setting are no longer static as in previous literature. Third, we introduce a drift term,  $\mu$ , to characterize the overall (positive or negative) trend of the consumer’s engagement with the content; this term was set equal to zero in Branco et al. (2012, 2016). Next, we illustrate the details of our model.

To capture the user’s instantaneous interest level toward the content, we let  $X_t$  be the engagement level at time  $t$ . Consider a consumer who actively engages with the content on the platform. Apart from the fact that this consumer may enjoy or dislike the genre of the content, he or she will also update his or her engagement level continuously depending on how the story unfolds or the game progresses. That is, there are two factors contributing to the change of the user’s engagement level: deterministic trend and random component. Formally, to model the user’s continuous evolution of engagement, we let  $X_t = \mu t + Z_t$ , where the first and second terms capture the constant change and stochastic change, respectively.<sup>2</sup> Note that  $\mu$  is also known as the drift term, which captures the *trend* of the user’s overall engagement level over time. For example, when users watch a mystery movie, their general engagement level is likely to increase as the story approaches the finale. Another example is that for mobile games, on average, players are more engaged in the game as they spend more time on their smartphones. Both examples denote the situation where  $\mu > 0$ . The drift term  $\mu$  can also be negative, which reflects a case in which a user might initially have a high level of anticipation for the content but later find that the content does not match his or her expectations.

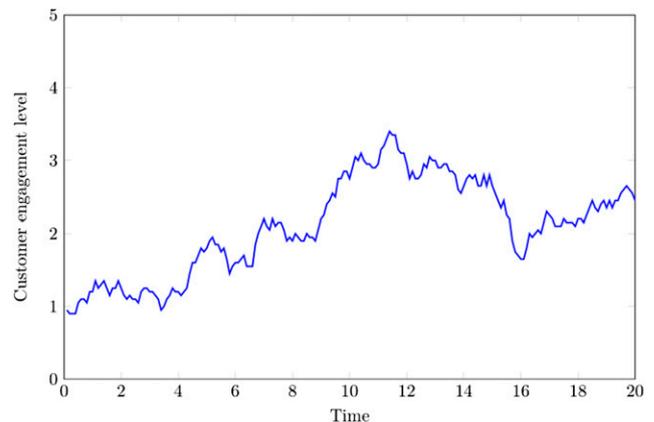
The drift term captures the deterministic component of how users update their engagement level, whereas the second term,  $Z_t$ , reflects the random change in the process (i.e., stochastic component). Note that the realization of random change can be either positive or negative, which reflects the fact that the content at time instant  $t$  can either fit the user’s preference or not. For example, game players may constantly change their engagement level due to the random outcomes of the game. If they receive their preferred in-game equipment, their interest level is likely to increase, and vice versa if they obtain undesired in-game equipment. To formally characterize this constantly changing randomness, we let  $Z_t = \sigma B_t$ , where  $\sigma > 0$  is known as the *volatility* reflecting the magnitude of the stochastic change, and  $B_t$  is the standard Brownian motion. For tractability, we focus on the case where  $\sigma$  is a constant across the time, but in general, the insights remain qualitatively the same as long as users do not know the exact value of  $\sigma$  prior to their engagement with the content.

Figure 1 provides an illustrative example of how user engagement level  $x$  changes along with time  $t$ . In summary, we model the user engagement level as  $X_t = \mu t + \sigma B_t$ , where the first component denotes the trend of the user’s engagement level over time, and the second component characterizes the stochastic change. Without loss of generality, we assume that the initial engagement level  $X_0$  is a nonnegative number; otherwise, the user would not choose to join the platform in the first place. After spending  $t$  time units, the current engagement level of the user becomes  $X_t$ . In the remainder of the study, we write  $x$  to replace  $X_t$  or ease of exposition.

### 3.2. Impact of Advertisement and Platform’s Decisions

In this subsection, we first discuss how an advertisement affects the user’s engagement level. This is followed by the platform’s decisions on the length

**Figure 1.** (Color online) A Sample Path of the User Engagement Level ( $\mu > 0$ )



and timing of the advertisement. It is widely accepted that advertising often annoys the user because it gets in the way of what viewers or players really want (Fan et al. 2007, Shen and Villas-Boas 2018).<sup>3</sup>

In this study, we assume that a provider's revenue is linear in the length of the advertisement, and there are two components of the advertisement's impact on the user's engagement level, which are fixed and variable parts. More specifically, if the content provider inserts an advertisement with the length of  $l$  units (i.e., seconds or bandwidth) during the content, the consequence can be described as follows. On the one hand, the provider can generate the revenue  $rl$ , where  $r$  denotes the revenue per unit.<sup>4</sup> This is known as time-based advertising in practice, which has gained popularity in recent years (Ryan 2016). On the other hand, the user's engagement level will drop from  $x$  to  $x - c_0 - cl$ , where  $c_0$  and  $c$  represent the fixed and variable parts of the user's dislike toward the advertisement, respectively.

Figure 2 illustrates a sample path of how a user's engagement level changes with respect to the advertisement inserted at time  $\tau$ . From a practical perspective, each individual user has a different value of  $c_0$  and  $c$  which can be estimated from the individual's existing browsing history or running experiments.

Thus, the provider's decisions are a sequence of insertion times  $(\tau_1, \tau_2, \dots)$  and corresponding lengths of advertisement  $(l_1, l_2, \dots)$  denoted as  $w = ((\tau_1, l_1), (\tau_2, l_2), \dots)$ , where  $\tau_k \leq \tau_{k+1}$ ,  $l_k \geq 0$ , and  $X_{\tau_k^-} - l_0 - cl_k \geq 0$ . Note that  $X_{\tau_k^-}$  denotes the user's engagement level immediately before watching the advertisement, and  $X_{\tau_k}$  represents the engagement level after the advertisement ends.

We denote the set of all such decisions as  $W$ . For a given strategy  $w$ , the dynamics of the user's engagement level are

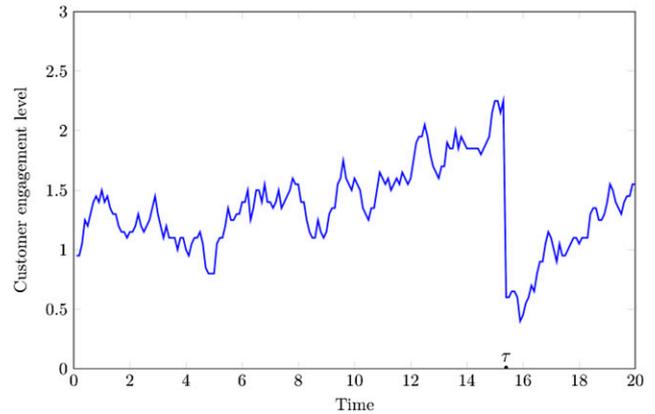
$$\begin{cases} dX_t = \mu dt + \sigma dB_t, & \tau_i \leq t < \tau_{i+1}, i \geq 0 \\ X_{\tau_i} = X_{\tau_i^-} - c_0 - cl_i & i \geq 1, \end{cases} \quad (1)$$

where the first line of the equation denotes the update of the user's engagement level over time, and the second line represents the change in user engagement level when viewing an advertisement. Users will leave the platform once their engagement drops below the level of 0. To capture the user's overall experience over the platform, we assume that the length of the advertisement is bounded such that the user's engagement level will not drop below  $e_0$  after viewing the advertisement.

For any given strategy  $w \in W$ , the provider's expected discounted profit at time  $t$  corresponding to the current user's engagement level at  $x$  is

$$J^w(t, x) = E_x \left[ \sum_{i=1}^{\infty} e^{-\beta \tau_i} r l_i \right],$$

**Figure 2.** (Color online) A Sample Path of the User Engagement Level with One Advertisement Inserted



where  $\beta$  is the discount factor. To maximize the total expected discounted profit function  $V(t, x)$ , the platform needs to find the optimal action strategy  $w \in W$  such that

$$V(t, x) = \max_{w \in W} J^w(t, x).$$

The platform faces two intertwined trade-offs in solving this optimization problem. First, regarding the length of advertisement, it is clear that users prefer advertisements with relatively short duration; nevertheless, the platform favors longer advertisements that can generate higher revenue.<sup>5</sup> The second trade-off concerns the timing of the advertisement. If the platform inserts the advertisement earlier, then users may leave the platform earlier or their engagement level may stay at a low level for an extended period of time, which prohibits the platform from inserting another advertisement. If the platform chooses to insert the advertisement at a later time, then the user may have already left the platform due to a low engagement level. Essentially, the platform needs to balance these trade-offs when determining the timing and length of the advertisement.

Note that traditional advertisements, such as TV advertisements, can be treated as a special case in our setting where the platform makes the decision based on the average engagement level rather than the individual dynamic engagement level. From a conceptual level, analogous to the familiar idea of price discrimination, the methods we propose in this study allow the platform or content provider to maximize revenue with discriminatory customization of the timing and length of the advertisement based on the behavior of individual users.

In the following analysis, we first assume that the length of the time horizon,  $T$ , is infinite. This assumption not only allows us to characterize the structural results but also reflects the idea that the game or TV series can be quite long. For example, users can view

an infinite number of images posted by another user on Instagram, or a game player can spend an extended period of time on a mobile game. Next, we consider an alternative setting where users may leave the platform at any time due to other factors, as well as the scenario in which the length of the content is bounded. Finally, we quantify the benefits of the optimal policy by comparing it with the prevalent industry policies (i.e., preroll, midroll, and mixed strategies) through a simulation study. For convenience, we summarize the notations used in this study in Table 1. Note that trigger threshold  $u$ , target level  $U$ , user's random departure time  $\theta$ , and user's random departure rate  $\lambda$  are formally defined in the later sections.

## 4. Optimal Solution and Managerial Insights

To characterize the structural results and closed-form solution, we first study the optimal advertisement insertion policy in an infinite horizon setting and generate the relevant managerial insights that can be applied in more general settings.

### 4.1. Optimal Insertion Policy and Insights

Before we give the solution, we first state some preliminary results to characterize the optimal policy. Let

$$g(x) := e^{\gamma_1 x} e^{-\gamma_2 x} \quad (2)$$

with  $\gamma_1 = [-\mu + \sqrt{\mu^2 + 2\beta\sigma^2}]/\sigma^2$  and  $\gamma_2 = [\mu + \sqrt{\mu^2 + 2\beta\sigma^2}]/\sigma^2$ . Define

$$p := r/c \quad (3)$$

$$\text{and } x_0 := \frac{2\ln\gamma_2 - 2\ln\gamma_1}{\gamma_1 + \gamma_2}. \quad (4)$$

**Table 1.** Parameters and Decision Variables

Symbol	Definition
$X_t$	User's engagement level at time $t$
$\mu$	Trend of user's overall engagement level along the time
$\sigma$	Volatility of the stochastic change
$\tau_i$	Insertion time of $i$ advertisement
$l_i$	Length of $i$ advertisement
$e_0$	Prespecified lower bound of user engagement level after the advertisement
$c_0$	Fixed part of user's dislike toward the advertisement
$c$	Variable part of user's dislike toward the advertisement
$r$	Unit revenue of the advertisement
$A$ and $A_1$	Constants in differential equations
$\beta$	Discount factor
$u$	Trigger threshold to insert the advertisement (decision variable)
$U$	Target level of the user engagement level after the advertisement (decision variable)
$\theta$	User's random departure time, following an exponential distribution with parameter ( $\lambda$ )
$\lambda$	User's random departure rate

**Lemma 1.** For  $g(x)$  and  $x_0$  defined in (2) and (4), for any  $p > 0$  and  $A \in (0, p/g'(x_0))$ , there exists a unique pair of  $(U(A), u(A))$  such that  $U(A) < x_0 < u(A)$ :

$$g'(x) \begin{cases} > \frac{p}{A} & \text{for } x < U(A) \\ = \frac{p}{A} & \text{for } x = U(A) \\ < \frac{p}{A} & \text{for } U(A) < x < u(A) \\ = \frac{p}{A} & \text{for } x = u(A) \\ > \frac{p}{A} & \text{for } x > u(A), \end{cases} \quad (5)$$

and

$$g''(x) \begin{cases} < 0 & \text{for } x < x_0 \\ = 0 & \text{for } x = x_0 \\ > 0 & \text{for } x > x_0, \end{cases} \quad (6)$$

$$U'(A) > 0. \quad (7)$$

All proofs are provided in the Online Appendix. Note that  $A$  is a constant in differential equations and can be uniquely determined by Equations (10), (11), and (12) in Lemma 2. We utilize the result of Lemma 1 to derive Lemma 2, which characterizes the parameters that are used to describe the optimal policy.

**Lemma 2.** For any  $p > 0$ ,  $c_0 > 0$ ,  $\gamma_1 = [-\mu + \sqrt{\mu^2 + 2\beta\sigma^2}]/\sigma^2$ ,  $\gamma_2 = [\mu + \sqrt{\mu^2 + 2\beta\sigma^2}]/\sigma^2$ , and  $x_0$  defined in (4), there exists unique  $(u, U, A)$  with

$$U < x_0 < u, \quad U < u - c_0 \quad (8)$$

and

$$A > 0, \quad (9)$$

such that

$$Ae^{\gamma_1 u} - Ae^{-\gamma_2 u} = Ae^{\gamma_1 U} - Ae^{-\gamma_2 U} + p(u - U - c_0), \quad (10)$$

$$A(\gamma_1 e^{\gamma_1 u} + \gamma_2 e^{-\gamma_2 u}) = p, \quad (11)$$

$$A(\gamma_1 e^{\gamma_1 U} + \gamma_2 e^{-\gamma_2 U}) = p. \quad (12)$$

In addition, for  $g(x)$  defined in (2), we have

$$g''(U) < 0 \text{ and } g''(u) > 0. \quad (13)$$

Note that  $u$  and  $U$  represent the trigger threshold to insert the advertisement and the target level of user engagement after the advertisement, respectively. Essentially, Lemma 2 not only characterizes the parameters  $(u, U)$  that will be used in Theorems 1 and 2 but also proves the existence and uniqueness of these parameters. Note that  $(u, U)$  determined in Lemma 2 are functions of  $(\mu, \sigma, \beta, p, c_0)$ . We analyze the optimization problems by separating them into two cases:  $U \geq e_0$  and  $U < e_0$ . In the first case, the (endogenous) target level  $U$  is greater than or equal to the prespecified

(exogenous) lower bound of the engagement level  $e_0$  (i.e., Theorem 1). In the second case, the target level  $U$  is less than the prespecified lower bound. As a result, instead of allowing the engagement level to drop to the target level, the platform will only agree with the advertisement length such that the engagement level decreases to exactly the prespecified engagement level,  $e_0$  (i.e., Theorem 2). The higher the  $e_0$  platform sets, the more the platform cares about the user's experience on the platform, and vice versa.

Let us define

$$\Theta := \{(\mu, \sigma, \beta, p, c_0, e_0) : U(\mu, \sigma, \beta, p, c_0) \geq e_0\}. \quad (14)$$

Note that when the parameters  $(\mu, \sigma, \beta, p, c_0, e_0) \in \Theta$ , we have  $U = U(\mu, \sigma, \beta, p, c_0) \geq e_0$ . This corresponds to the first case, and we provide the optimal solution in Theorem 1. When the parameters  $(\mu, \sigma, \beta, p, c_0, e_0) \notin \Theta$ , we have  $U(\mu, \sigma, \beta, p, c_0) < e_0$ , which corresponds to the second case. The optimal solution for this case is provided in Theorem 2.

**Theorem 1.** *When the target level  $U$  is greater than or equal to the prespecified lower bound of the engagement level  $e_0$  (i.e.,  $(\mu, \sigma, \beta, p, c_0, e_0) \in \Theta$ ), then the value function is*

$$V(t, x) = e^{-\beta t} \begin{cases} Ae^{\gamma_1 x} - Ae^{-\gamma_2 x} & \text{for } x < u \text{ (waiting region)} \\ Ae^{\gamma_1 U} - Ae^{-\gamma_2 U} + p(x - U - c_0) & \text{for } x \geq u \text{ (action region),} \end{cases} \quad (15)$$

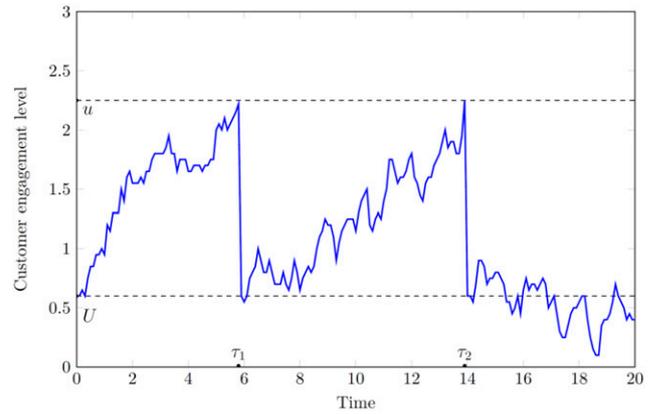
where  $p = r/c$ ,  $\gamma_1 = [-\mu + \sqrt{\mu^2 + 2\beta\sigma^2}]/\sigma^2$ ,  $\gamma_2 = [\mu + \sqrt{\mu^2 + 2\beta\sigma^2}]/\sigma^2$ , and  $uUA$  are defined in Lemma 2.

Theorem 1 characterizes the structure of the optimal policy. We begin by discussing the following property of the optimal policy.

**Proposition 1.** *When the target level  $U$  is greater than or equal to the prespecified lower bound of the engagement level  $e_0$ , then the optimal action strategy under the infinite horizon is a threshold policy that delays the insertion of the advertisement until the user's engagement level  $x$  reaches the trigger threshold  $u$  and then plays an advertisement with length  $(u - U - c_0)/c$ .*

Most importantly, we show that a threshold policy is an optimal policy among all classes of policies. This is an important result both theoretically and practically. Although one may think that a dynamic complicated policy should be optimal, here we analytically prove that a simple threshold policy can perform as well as a complex one, if it exists. Essentially, the platform only needs to insert the advertisement when the user's engagement level reaches the trigger threshold,  $u$ , and then the user will view an advertisement with length  $(u - U - c_0)/c$  such that the user's engagement level drops to the target level,  $U$ .

**Figure 3.** (Color online) Illustration of the Optimal Policy



Generally speaking, the parameters used to characterize the trigger threshold  $u$  and target level  $U$  (i.e.,  $\mu, \sigma, \beta, r, c_0, c$ ) are different for different users, as every user has his or her own distinct taste for the content and dislike toward the advertisement. The parameters heavily relying on individual users (i.e.,  $\mu$  and  $\sigma$ ) can be estimated through their browsing behaviors or by running experiments. For the parameters that are advertisement specific (i.e.,  $c_0$  and  $c$ ), the platform can estimate these through running a focus group of users. We provide a visual illustration of the implementation of this policy in Figure 3.

In the aforementioned illustrative example, the user has a positive attitude toward the content that is viewed or played, which is reflected by the upward trend of the user's engagement level. To maximize the revenue, the platform inserts the advertisement when the user's engagement level reaches the trigger threshold,  $u$ . Further, the length of advertisement is bounded such that the user's engagement level will not drop below the target level,  $U$ . Note that if the user's initial engagement level is sufficiently high (i.e.,  $x_0 > u$ ), our policy states that the platform should play the advertisement at the very beginning (i.e., preroll strategy), which is consistent with the practice of many platforms, such as YouTube. The difference here is that the platform should insert the advertisement only if the user's interest is high enough (i.e., high initial engagement level) under our policy, whereas the platform always inserts the advertisement at the very beginning regardless of the user's engagement levels under the current business practice. This may be suboptimal, driving away many valuable users.

One may be concerned that this proposed policy could affect the user's overall experience of the platform, as the platform precisely inserts the advertisement at the climax of the user's engagement with the content. Indeed, the insertion of the advertisement will mitigate the user's enjoyment of the content. However, we show that the leaving rate is actually

much lower in the proposed policy compared with the industry standard policies. Our policy mainly applies to platforms, websites, or apps whose revenue is solely driven by advertisement rather than sales, as the user's experience should be prioritized if sale of content is the key source of revenue. Further, with the surge of the subscription pricing model, our policy also contributes to the adoption of the user's subscription because our proposed policy may drive some users to pay for the service in order to avoid behavioral advertisements.

The optimal policy described here has many advantages, including easy implementation and improved profitability. Later, we demonstrate the benefits of this policy by comparing it with the existing popular advertisement policies, such as preroll, midroll, and a mix of the two. We provide the crux of the algorithm that can compute  $u$  and  $U$  efficiently in the proof of Proposition 1 in the Online Appendix. Next, we find that the target level  $U$  is less than the prespecified lower bound of the engagement level  $e_0$ . In this case, the platform should limit the advertisement length such that the consumer's engagement level does not drop below the level  $e_0$  after watching the advertisement.

**Theorem 2.** *When the target level  $U$  is less than the prespecified lower bound of the engagement level  $e_0$  (i.e.,  $(\mu, \sigma, \beta, p, c_0, e_0) \notin \Theta$ ), then the value function becomes*

$$V(t, x) = e^{-\beta t} \begin{cases} A_1 e^{\gamma_1 x} - A_1 e^{-\gamma_2 x} & \text{for } x < u_1 \text{ (waiting region)} \\ A_1 e^{\gamma_1 e_0} - A_1 e^{-\gamma_2 e_0} + p(x - e_0 - c_0) & \text{for } x \geq u_1 \text{ (action region),} \end{cases}$$

where  $p = r/c$ ,  $\gamma_1 = [-\mu + \sqrt{\mu^2 + 2\beta\sigma^2}]/\sigma^2$ ,  $\gamma_2 = [\mu + \sqrt{\mu^2 + 2\beta\sigma^2}]/\sigma^2$ ,  $u_1$  is decided by

$$e^{\gamma_1 u_1} - e^{-\gamma_2 u_1} = e^{\gamma_1 e_0} - e^{-\gamma_2 e_0} + (\gamma_1 e^{\gamma_1 u_1} + \gamma_2 e^{-\gamma_2 u_1})(u_1 - e_0 - c_0), \quad (16)$$

and

$$A_1 = \frac{p}{\gamma_1 e^{\gamma_1 u_1} + \gamma_2 e^{-\gamma_2 u_1}}. \quad (17)$$

Theorem 2, similar to the previous theorem, characterizes the structure of the optimal policy. In the following proposition, we highlight that the optimal length of advertisement does not always depend on the target level  $U$ .

**Proposition 2.** *When the target level  $U$  is less than the prespecified lower bound of the engagement level  $e_0$ , then the optimal action strategy under the infinite horizon is to wait until the user's interest level  $x$  reaches  $u_1$  and then add an advertisement with length  $(u_1 - e_0 - c_0)/c$ .*

Proposition 1 states that the advertisement length should be  $(u - U - c_0)/c$ , which is a function of the

target level  $U$ . One may intuitively think that the optimal advertisement length should always be determined by the target level, as it controls the consumer engagement level after watching the advertisement. However, Proposition 2 shows that the optimal policy does not always depend on the target level  $U$ . Although the structure of the optimal policy is similar to that under Proposition 1, the target level  $U$  is now replaced by the prespecified lower bound of the user engagement level  $e_0$ . The intuition of this result is that the optimal target level  $U$  is less than the prespecified lower bound  $e_0$ . To ensure the user's overall experience with the platform, the platform has to change the optimal length of the advertisement from  $(u - U - c_0)/c$  to  $(u_1 - e_0 - c_0)/c$  such that the engagement level drops to the prespecified engagement level  $e_0$  rather than the optimal target level  $U$ . From a practical perspective, this result suggests that the platform needs to make a trade-off between the length of the advertisement and the service level. If the platform puts high emphasis on the user's experience, it should play a relatively shorter advertisement.

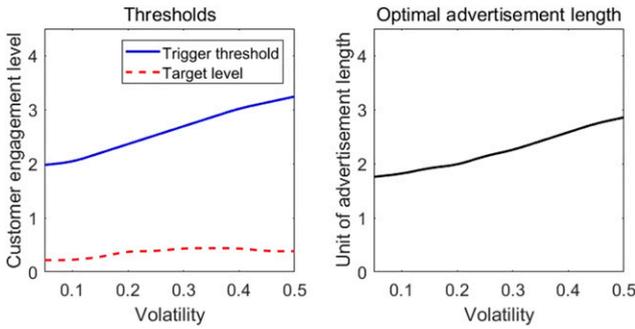
## 4.2. Additional Managerial Insights

In this subsection, we derive additional managerial insights by exploring how different parameters affect the optimal insertion policy. We are able to derive the impacts of a fixed decrease in engagement level (i.e.,  $c_0$ ) and the discount factor (i.e.,  $\beta$ ) analytically. These results are presented in Sections 4.2.3 and 4.2.4, respectively. However, because of the analytical intractability, the impacts of the volatility (i.e.,  $\sigma$ ) and the drift term (i.e.,  $\mu$ ) are analyzed numerically in Sections 4.2.1 and 4.2.2, respectively.

**4.2.1. The Volatility.** Recall that  $\sigma$  represents the magnitude of the volatility of the stochastic change in the consumer's engagement level. We reach the following observation from extensive numerical analysis.

**Observation 1.** As the volatility parameter  $\sigma$  increases, both the optimal trigger threshold and advertisement length will increase as well.

Figure 4 illustrates the impact of the volatility  $\sigma$  in the optimal insertion policy. If the volatility becomes higher, the consumer's engagement level will evolve more volatily. It is not obvious how the optimal policy will change in relation to an increase in volatility, as two opposite dynamics are possible. On the one hand, as volatility increases, the consumer's engagement level has a higher chance of reaching zero early, in which case the consumer will leave the platform. This seems to suggest that we should set a lower trigger threshold, such that the platform can generate revenue from these consumers before they leave. On the other hand, as volatility increases, it

**Figure 4.** (Color online) Impact of Volatility  $\sigma$ 

takes less time for the consumer's engagement level to reach the trigger threshold. Consequently, the platform should behave more patiently and set a higher trigger threshold to insert a longer advertisement. Interestingly, we find that the second scenario dominates the first one. This suggests that as the volatility level increases, the managers of the platform should set the optimal trigger threshold higher and the length of advertisement longer.

**4.2.2. The Drift Term.** Recall that the drift term  $\mu$  represents the average trend that the user is attracted by the content. As discussed earlier, we study the impact of this parameter on optimal insertion policy numerically due to analytical challenges.

**Observation 2.** As the drift term  $\mu$  increases, both the optimal trigger threshold and advertisement length will increase as well.

Figure 5 presents the impact of the trend of engagement level  $\mu$  in the optimal insertion policy. Note that we have conducted extensive numerical analyses and find that our results are robust to a variety of different parameter settings. From the numerical studies, we observe that both the trigger threshold and the advertisement length increase as the trend of engagement level increases. This is because, if the user has a high level of interest in the content, it is more likely that his or her engagement level will increase over time. Therefore, it is optimal for the platform to behave patiently (by setting a higher trigger threshold) and insert a longer advertisement. From a managerial perspective, this result suggests that the platform should have a longer version of the advertisement when it expects that the users will have higher engagement with the content.

**4.2.3. Fixed Decrease Level.** We are able to characterize analytically how the fixed decrease in level  $c_0$  influences the optimal insertion policy. The detailed results and proofs are presented in the Online Appendix. For simplicity, we only illustrate and discuss the results here. We find that as the fixed decrease level  $c_0$  increases, the trigger threshold  $u$  will increase,

the target level  $U$  will decrease, and the optimal advertisement length will increase.

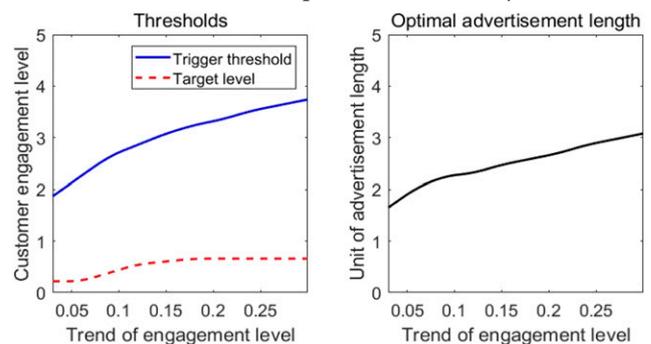
We illustrate the impact of the fixed decrease level  $c_0$  in Figure 6. The results here are in line with our expectations. Specifically, as the fixed decrease level  $c_0$  increases, we note that the platform behaves more patiently and inserts a longer advertisement when the opportunity arises. The intuition of this result is that, as the fixed decrease level  $c_0$  increases, the consumer becomes more annoyed each time the advertisement is played, so it becomes more costly for the platform to insert the advertisements, as the consumer is more likely to leave when the  $c_0$  level is high. Consequently, the optimal strategy facing the platform is to wait until the consumer becomes more engaged with the content, then insert a longer advertisement. The result here suggests that the platform should target users with varying thresholds and tolerance of advertisement lengths, because different users have different levels of tolerance for advertisements.

**4.2.4. Discount Factor.** The parameter  $\beta$  represents the discount factor.<sup>6</sup> We are able to analytically prove that as  $\beta$  increases, both the trigger threshold  $u$  and the target level  $U$  decrease, and the optimal advertisement length decreases. Similar to  $c_0$ , the detailed results and proofs are presented in the Online Appendix, and here we simply illustrate and discuss the results.

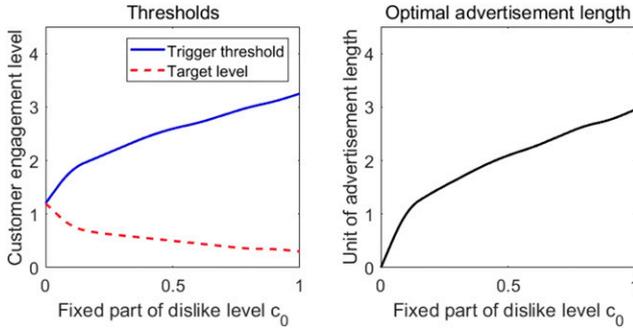
Figure 7 displays the impact of the discount factor  $\beta$  in the optimal insertion policy. As the discount rate  $\beta$  increases, we note that the platform behaves less patiently and inserts a shorter advertisement when the opportunity arises. This is because, as  $\beta$  increases, the platform places more value on the current profit as opposed to future profit. As a result, the platform prefers to generate revenue sooner rather than later. From a managerial perspective, when platforms have relatively long content, they should adjust the thresholds and advertisement lengths based on their discount factor.

## 5. Finite Horizon

In the previous section, we analyzed scenarios where the planning horizon could be regarded as infinite.

**Figure 5.** (Color online) Impact of Drift Term  $\mu$ 

**Figure 6.** (Color online) Impact of Fixed Decrease Level  $c_0$



In this section, we now focus on those scenarios for which the planning horizon can be regarded as finite by analyzing two alternative cases. In Section 5.1, we investigate the scenario in which users can leave the platform randomly during the content (i.e., stochastic horizon). In Section 5.2, we study the case in which the length of the planning horizon is finite and deterministic.

### 5.1. Stochastic Horizon

In this subsection, we consider the scenario in which the user can leave the platform unexpectedly. In addition to a low engagement level, the user may receive an important phone call, or the device's battery may be depleted. From the modeling perspective, we assume that, regardless of engagement level, the user may leave at any random time  $\theta$ , which follows an exponential distribution with parameter  $\lambda$ . Thus, the expected length of this horizon is  $E[\tau] = 1/\lambda$ , and the larger  $\lambda$  is, the smaller the expected horizon  $E[\tau]$  becomes. Note that, in this setting, the user can leave the platform because of either a low engagement level or a random occurrence. Similar to the base model, we solve the problem in two separate cases. Due to the similarity of these two cases, we first state the properties and then provide the discussion based on the two cases together.

**Theorem 3.** *When the target level  $U$  is greater than or equal to the prespecified lower bound of the engagement level  $e_0$  (i.e.,  $(\mu, \sigma, \beta + \lambda, p, c_0, e_0) \in \Theta$ ), the value function becomes*

$$V_d(t, x) = e^{-\beta t} \begin{cases} A_d e^{\gamma_{1,d} x} - A_d e^{-\gamma_{2,d} x} & \text{for } x < u_d \text{ (waiting region)} \\ A_d e^{\gamma_{1,d} U_d} - A_d e^{-\gamma_{2,d} U_d} + p(x - U_d - c_0) & \text{for } x \geq u_d \text{ (action region),} \end{cases}$$

where  $p = r/c$ ;  $\gamma_{1,d} = [-\mu + \sqrt{\mu^2 + 2(\beta + \lambda)\sigma^2}]/\sigma^2$ ;  $\gamma_{2,d} = [\mu + \sqrt{\mu^2 + 2(\beta + \lambda)\sigma^2}]/\sigma^2$ ; and  $u_d, U_d, A_d$  are defined in Lemma 2 with  $\gamma_{1,d}$  and  $\gamma_{2,d}$  replacing  $\gamma_1$  and  $\gamma_2$ .

Theorem 3 characterizes the structure of the optimal policy when the horizon is stochastic. We begin by discussing the following property of the optimal policy.

**Proposition 3.** *When the target level  $U$  is greater than or equal to the prespecified lower bound of the engagement level  $e_0$  (i.e.,  $(\mu, \sigma, \beta + \lambda, p, c_0, e_0) \in \Theta$ ), the optimal action strategy under the stochastic horizon is to delay the insertion of the advertisement until the user's engagement level  $x$  reaches  $u_d$  and then to insert an advertisement whose length equals  $(u_d - U_d - c_0)/c$ . In addition, we find that  $u_d < u$ , where  $u$  is the trigger threshold defined in Theorem 1.*

Similar to Proposition 1, we note that the optimal policy with the stochastic horizon is based on a threshold policy. The platform inserts the advertisement when the user's engagement level reaches the trigger threshold. When the target level is less than the prespecified lower bound, we have the following theorem.

**Theorem 4.** *When the target level  $U$  is less than the prespecified lower bound of the engagement level  $e_0$  (i.e.,  $(\mu, \sigma, \beta + \lambda, p, c_0, e_0) \notin \Theta$ ), the value function becomes  $V_d(t, x) = e^{-\beta t} \psi_d(x)$  with*

$$\psi_d(x) = \begin{cases} A_{1,d} e^{\gamma_{1,d} x} - A_{1,d} e^{-\gamma_{2,d} x} & \text{for } x < u_{1,d} \text{ (waiting region)} \\ A_{1,d} e^{\gamma_{1,d} e_0} - A_{1,d} e^{-\gamma_{2,d} e_0} + p(x - e_0 - c_0) & \text{for } x \geq u_{1,d} \text{ (action region),} \end{cases}$$

where  $p = r/c$ ,  $\gamma_{1,d} = [-\mu + \sqrt{\mu^2 + 2(\beta + \lambda)\sigma^2}]/\sigma^2$ ,  $\gamma_{2,d} = [\mu + \sqrt{\mu^2 + 2(\beta + \lambda)\sigma^2}]/\sigma^2$ , and  $u_{1,d}$  is determined by

$$e^{\gamma_{1,d} u_{1,d}} - e^{-\gamma_{2,d} u_{1,d}} = e^{\gamma_{1,d} e_0} - e^{-\gamma_{2,d} e_0} + (\gamma_{1,d} e^{\gamma_{1,d} u_{1,d}} + \gamma_{2,d} e^{-\gamma_{2,d} u_{1,d}})(u_{1,d} - e_0 - c_0) \quad (18)$$

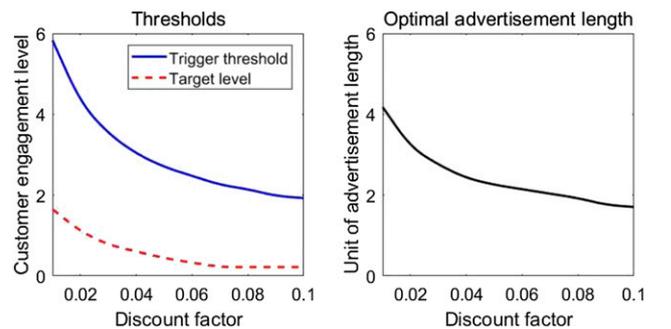
and

$$A_{1,d} = \frac{p}{\gamma_{1,d} e^{\gamma_{1,d} u_{1,d}} + \gamma_{2,d} e^{-\gamma_{2,d} u_{1,d}}}. \quad (19)$$

Theorem 4 characterizes the structure of the optimal policy. Once again, we note that the optimal length of advertisement does not need to depend on the target level  $U$ , which leads to the following property.

**Proposition 4.** *When the target level  $U$  is less than the prespecified lower bound of the engagement level  $e_0$  (i.e.,  $(\mu, \sigma, \beta + \lambda, p, c_0, e_0) \notin \Theta$ ), the optimal action strategy under*

**Figure 7.** (Color online) Impact of Discount Factor  $\beta$



the stochastic horizon is to delay the insertion until the user's interest level  $x$  reaches  $u_d$  and then to add an advertisement with length  $(u_{1,d} - e_0 - c_0)/c$ . In addition, we find that  $u_{1,d} < u_1$ , where  $u_1$  is the trigger threshold in Theorem 2.

The discussion here applies to both Propositions 3 and 4. When the horizon becomes stochastic, one may intuitively guess that the optimal policy should become more complicated, as there is additional uncertainty involved. However, we find that this is not the case here. We prove that, similar to Propositions 1 and 2, the optimal policies with the stochastic horizon are still based on a threshold policy, and they are optimal among the class of all policies. The key difference between the infinite horizon case and the stochastic horizon case is that we substitute  $\beta$  with  $\beta + \lambda$  in calculating the  $\gamma_{1,d}$  and  $\gamma_{2,d}$  (corresponding to  $\gamma_1$  and  $\gamma_2$  in Theorems 1 and 2).

Nevertheless, and interestingly, we find that the platform inserts the advertisement more urgently when the user can leave randomly. To understand the intuition of this result, recall that the larger the  $\lambda$  is, the smaller the expected random departure time  $E[\tau]$  becomes, which suggests that the user is more likely to leave randomly before the content ends. As a result, compared with the infinite horizon case, the platform becomes less patient in terms of the timing of the advertisement insertion (i.e.,  $u_d < u$  in Proposition 3 and  $u_{1,d} < u_1$  in Proposition 4).

## 5.2. Finite Planning Horizon

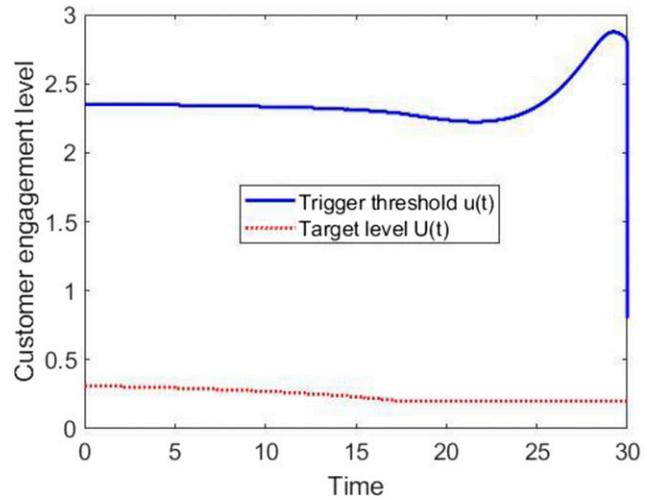
In this subsection, we discuss the optimal advertisement insertion problem with a finite and deterministic planning horizon. There are many practical situations that fit this setting; for example, the user can only watch one new episode of a popular TV series per week due to the release frequency.

It is worthwhile to point out that the key difference between the solutions to infinite and finite planning horizon cases lies in the fact that neither the trigger threshold nor the target level depends on time in the infinite horizon case, whereas either may depend on time in the finite horizon setting, thereby significantly increasing the complexity of the problem. Further, it is extremely difficult to find closed-form solutions for the stochastic impulse control problem in a finite planning horizon.<sup>7</sup> As a result, we resort to extensive numerical studies to analyze this case.

**Observation 3.** Under the finite horizon, the optimal trigger level  $u$  will first decrease, then increase, and finally decrease as the time approaches the end of the planning horizon. This is a result of the trade-off between the number of advertisements and advertisement length.

Figure 8 illustrates how the trigger threshold  $u$  and target level  $U$  change along with time.<sup>8</sup> Interestingly,

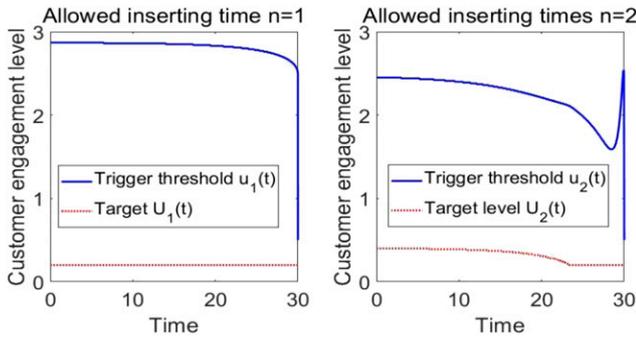
**Figure 8.** (Color online) Optimal Trigger Thresholds and Target Levels in Finite Planning Horizon



we observe that the optimal trigger level  $u$  first decreases, then increases, and finally decreases as the time approaches the end of the planning horizon. This result is robust with a variety of parameter settings. Intuitively, one may think that the platform might become less patient (i.e., decreased optimal trigger threshold  $u(t)$ ) to insert the advertisement when the time approaches the end of the planning horizon, as the platform would lose the opportunity to insert any further advertisement. However, we find that it is actually optimal for the platform to first increase and then decrease the trigger threshold when approaching the end of the planning horizon. The platform needs to reach a balance between the insertion time and the length of each advertisement. Specifically, when there is sufficient time left (i.e., in the beginning of the planning horizon), the platform initially plans to insert multiple episodes of the advertisement. However, if the user's initial engagement level stays low and the remaining planning horizon becomes shorter, the platform fears that it has fewer opportunities to insert the advertisement. As a result, the platform prefers to insert a longer advertisement instead of a shorter one, which is reflected by the increase in the optimal trigger threshold  $u(t)$ . Finally, if the planning horizon approaches the end, the platform has to reduce the trigger threshold because it has no time to wait toward the end of the planning horizon.

To confirm our intuition, we conduct further analysis. With everything else being equal, we now restrict the platform to inserting a maximum of  $n$  advertisements during the finite planning horizon. Figure 9 illustrates the optimal trigger threshold  $u(t)$  when  $n = 1$  and  $n = 2$ . When the platform is only allowed to insert one advertisement during the content (corresponding to the left panel), we observe that the optimal trigger

**Figure 9.** (Color online) Optimal Thresholds when There Is a Restriction on the Number of Insertions



threshold  $u(t)$  keeps decreasing. This result is in line with our conjecture that the platform becomes less patient to insert the advertisement when the time approaches the end of the planning horizon, and the platform loses the opportunity to insert any further advertisement. When the platform is allowed to insert more than one advertisement (i.e.,  $n = 2$  in the right panel), we find that the optimal trigger threshold  $u(t)$  exhibits the pattern that we observed before. That is, the optimal trigger threshold  $u(t)$  first decreases, then increases, and finally decreases again. This result has interesting implications for the platform manager regarding when to insert the advertisement. The optimal trigger threshold is a result of a balance between the advertisement length and the number of insertions.

## 6. Comparison with Prevalent Policies and Parameter Estimation

In this section, we first illustrate the outcomes of our proposed policies by comparing them with several practically prevalent advertisement insertion policies through a simulation study.<sup>9</sup> We then discuss how to estimate relevant parameters.

### 6.1. Comparison with Prevalent Policies

We focus on the case where the planning horizon is finite and deterministic, as most practical examples apply. The benchmark case is the proposed threshold policy. We compare this policy against the industry standard policies (i.e., preroll, midroll, and a mix of both placement strategies) in terms of profitability, leaving rate, and total number of advertisements placed. In practice, platforms typically adopt either preroll, midroll, postroll, or a mix of these strategies to place their advertisement. In the preroll placement strategy, the advertisement plays before the start of the content. Midroll placement allows the platform to play the advertisement during the content, which is similar to the schedule that viewers are accustomed to on broadcast television. In contrast, the postroll

placement simply means that an advertisement plays after the user views the content. We do not explicitly consider the postroll strategy in this study, because consumers are not likely to watch the advertisement after they have viewed the desired content. Further, postroll advertisement has the lowest completion rate in practice (Krishnan and Sitaraman 2013). Rather, we consider the postroll placement strategy indirectly. If the user views the content by using a playlist, a postroll advertisement can be considered as a special case of the midroll placement strategy. If users are interested in the upcoming content, they will stay through the advertisement to engage in what follows.

In the simulation study, we first consider the preroll placement strategy (i.e., preroll), where the advertisement is placed before the start of the content. Next, we consider the midroll strategy, in which the platform inserts one or more advertisements in the middle of the content with the same interval between them (i.e., midroll and midroll 2). Further, we also examine the mixed strategy of preroll and midroll, where the platform places advertisements at the very beginning as well as in the middle of the content (i.e., mix 2 and mix 3). We also consider a random strategy where the platform can insert the advertisement at any time during the content (i.e., random). Finally, we consider an alternative setting (i.e., sub-optimal case) where we set the trigger threshold  $u_t$  and target level  $U_t$  as fixed. We characterize the fixed levels by assuming that the planning horizon is infinite, which allows us to use the closed-form solution to impute both levels. This method is clearly sub-optimal compared with the case when trigger threshold  $u_t$  and target level  $U_t$  are dynamic in the finite planning horizon, but it is computationally appealing as the platform only needs to calculate once for each individual consumer, and the equations characterizing the thresholds in Proposition 1 can be readily solved. The objective here is to examine whether this heuristic method is good enough for practical use.

Without loss of generality, we allow the user's initial engagement level,  $X_0 \sim Uniform(0,5)$ , which reflects the fact that distinct consumers have different expectations for the content. We have also verified that our results are robust to a variety of parameter settings by conducting extensive analysis. For exposition, we next illustrate one representative case with the following parameter settings:  $\{\beta = 0, \mu = 0.1, \sigma = 0.2, r = 1, c_0 = 0.6, c = 1, e_0 = 0.2, \lambda = 0.05\}$ . In this experiment, each replication simulates 30 minutes of content length in which the consumer's engagement level updates every second. We summarize the results of profitability and leaving rate in Tables 2 and 3, respectively.

There are several important observations from the aforementioned simulation study.

**Table 2.** Profitability and Number of Advertisement of the Simulation Results

Ad length	Preroll	Midroll	Midroll 2	Mix 2	Mix 3	Random	Suboptimal	Optimal
1.0	13.43% (1)	9.63% (1)	19.66% (2)	20.56% (2)	27.83% (3)	10.03% (1)	88.59% (2.15)	100% (2.41)
1.5	18.96% (1)	14.59% (1)	29.86% (2)	28.09% (2)	37.36% (3)	15.22% (1)	88.97% (2.18)	100% (2.44)
2.0	23.40% (1)	19.47% (1)	39.01% (2)	33.83% (2)	43.94% (3)	19.69% (1)	88.57% (2.23)	100% (2.45)
2.5	27.86% (1)	24.94% (1)	48.10% (2)	38.21% (2)	46.70% (3)	24.78% (1)	88.37% (2.19)	100% (2.48)
3.0	30.00% (1)	30.09% (1)	56.40% (2)	38.86% (2)	44.14% (3)	29.13% (1)	88.79% (2.22)	100% (2.48)
3.5	33.93% (1)	34.50% (1)	62.32% (2)	41.23% (2)	44.12% (3)	33.22% (1)	88.23% (2.10)	100% (2.37)
4.0	34.32% (1)	38.83% (1)	66.54% (2)	38.52% (2)	38.72% (3)	36.76% (1)	89.20% (2.13)	100% (2.38)
4.5	34.43% (1)	44.89% (1)	67.75% (2)	34.99% (2)	35.01% (3)	39.74% (1)	88.89% (2.17)	100% (2.42)

*Notes.* Ad length only applies to preroll, midroll, midroll 2, mix 2, mix 3, and random cases, and denotes the difference between engagement level before and after each advertisement. It is proportional to the physical advertisement length. Midroll denotes the placement of advertisement at the middle of the content, whereas midroll 2 denotes the placement of advertisement at one-third and two-thirds of the content. Mix 2 represents the placement strategy where the platform inserts the advertisement at the very beginning as well as at the middle of the content, and mix 3 means the advertisement placement at the very beginning, at one-third of the content, and at two-thirds of the content. The percentage means the proportion of the revenue that the placement strategy can generate compared with the optimal strategy in the last column. The number in parentheses denotes the total number of advertisements played during the content.

**Observation 4.** Compared with the prevalent policies (i.e., preroll, midroll, and mixed strategy), the proposed policy in our study can generate a much higher profitability.

We note that the proposed optimal policy yields a much better outcome in terms of profitability compared with the industry standard policies. To illustrate, when the advertisement length is  $l = 3.0$ , the midroll placement strategy outperforms the other standard strategies, but it can only generate 56.40% of the revenue from the optimal policy. Further, for the same advertisement length, although the mix 3 strategy will play three advertisements during the course of the content, it only yields 44.14% of the revenue against the optimal policy of playing 2.48 advertisements during the content on average. As illustrated in Table 2, this result is robust to a wide variety of parameter settings. The timing and length of the advertisement make a significant impact on revenue, and our proposed optimal policy can fully operationalize along these two key dimensions and customize based on an individual consumer's behavior.

Recently, we have observed that many platforms are testing different advertisement insertion policies.

For example, Facebook is experimenting with preroll video advertising on its video-on-demand service, "Facebook Watch," after a longtime ban of the preroll format (Sloane 2017). Google is limiting the use of preroll advertisement by switching more toward the midroll placement strategy on YouTube (Reale 2017). One of the major reasons to test these alternative policies is because the platforms are unclear on which policy can yield higher profitability. Our analysis here sheds light on this issue. We illustrate that the proposed policies can yield a much higher profitability compared with the prevalent industry policies. Thus, one natural and intriguing question arises: Does this increased profitability come at the cost of lowering the user's experience on the platform? The answer to this question will be not only theoretically interesting but also critical to the managers in the relevant industries as they strive to deliver a better user experience. We summarize our finding in the following observation.

**Observation 5.** Compared with the prevalent policies, the proposed policy results in a *lower* leaving rate.

Note that there are three different factors contributing to the leaving rate. The first is a low engagement

**Table 3.** Leaving Rate of the Simulation Results

Ad length	Preroll	Midroll	Midroll 2	Mix 2	Mix 3	Random	Suboptimal	Optimal
1.0	30.10%	4.10%	4.50%	30.10%	30.30%	6.10%	5.50%	7.40%
1.5	40.10%	4.50%	5.10%	40.10%	41.40%	7.70%	5.60%	8.00%
2.0	48.10%	3.60%	5.50%	48.10%	51.20%	8.30%	4.40%	7.60%
2.5	60.00%	3.70%	8.50%	60.50%	66.80%	9.40%	4.50%	6.20%
3.0	71.70%	5.10%	17.80%	72.50%	81.40%	13.90%	5.50%	7.70%
3.5	79.30%	6.40%	27.00%	81.00%	89.80%	20.20%	4.60%	6.10%
4.0	88.80%	7.40%	38.30%	90.90%	96.30%	21.70%	5.40%	6.80%
4.5	98.70%	11.40%	52.20%	98.80%	99.70%	26.00%	4.30%	6.00%

*Notes.* The percentage means the proportion of the consumers who leave during the content due to either the low engagement level of stochastic change or low engagement level of watching the advertisement.

level due to stochastic change, where users may not enjoy the content. The second is a lengthy advertisement that drives away users. The third factor is the random departure of the user (i.e., stochastic horizon) due to some arbitrary causes that are irrelevant to the engagement level or advertisement length. We report the sum of the first two factors in Table 3, because the third factor applies equally across all policies, whereas the first two are directly influenced by placement policies. We observe that both the suboptimal and the optimal policy result in a leaving rate of less than 8% in a wide variety of parameter settings. Our result is also consistent with the empirical observation that midroll advertisements produce a higher completion rate (i.e., lower leaving rate) than preroll advertisements (Krishnan and Sitaraman 2013). Consumers are more engaged with the content after spending some time viewing it but are less engaged at the beginning of the content. Thus, both suboptimal and optimal policies lead to a lower leaving rate compared with the industry standard policies. More importantly, our results suggest that increased profitability does not come at the cost of the user’s experience with the platform.

**Observation 6.** The suboptimal policy generated from the infinite horizon can be an efficient heuristic for the optimal policy.

Recall that, in the suboptimal policy, we treat the planning horizon as infinite, which allows us to use the fixed trigger threshold  $u_t = u$  and target level  $U_t = U$  to determine when to place the advertisement. Compared with the optimal policy, this method is computationally more efficient, as the platform only needs to calculate once, in contrast with the optimal policy where the platform has to constantly update the trigger threshold  $u_t$  and target level  $U_t$  as time passes. We note that the suboptimal policy can outperform all of the industry standard policies and achieve nearly 90% of the revenue from the optimal policy. Thus, we argue that the suboptimal policy can serve as an efficient heuristic for the optimal policy in most practical situations.

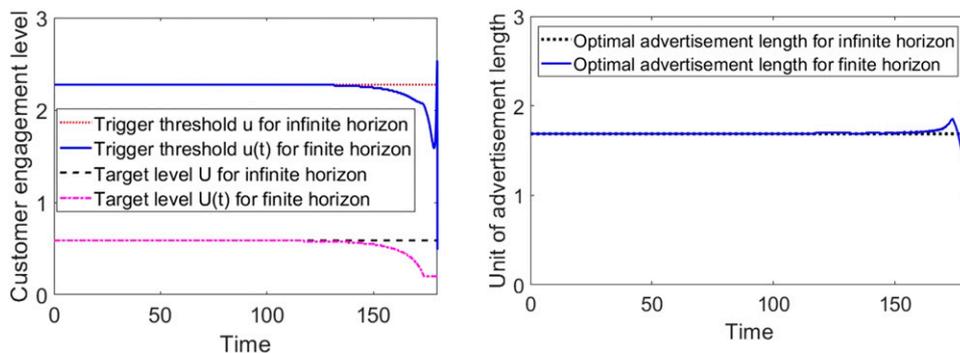
To further understand why and when the sub-optimal policy can be a viable substitute for the optimal policy, we consider the various content lengths, which leads to the following observation.

**Observation 7.** When the content length is long, the suboptimal policy can approximate the optimal policy well.

We have varied the content lengths and find that our observation is robust across different parameter settings. For exposition, we illustrate the thresholds and advertisement length when the content length is set equal to 180 minutes. From Figure 10, we observe that when the remaining time is sufficiently long, both the trigger threshold and the target level from the infinite horizon (i.e., suboptimal policy) are very close to those from the finite horizon (i.e., optimal policy), so the optimal advertisement length is very similar for both cases. This pattern continues until the content approaches the end of the horizon, at which point the trigger thresholds and advertisement length from the infinite horizon begin to differ significantly from those in the finite horizon. To illustrate, in a three-hour time window, the optimal trigger thresholds and advertisement length from the finite horizon are nearly the same as those from the infinite horizon for three-quarters of the time period. In more than half of the remaining one-quarter of the time period, the trigger thresholds and advertisement length from the infinite horizon are very similar to those from the finite horizon. Thus, the suboptimal policy adopting the thresholds from the infinite horizon can approximate the thresholds from the finite horizon very well, especially when the finite horizon is long.

In summary, the timing and length of the advertisement play a significant role in determining the platform’s revenue. The proposed policies in this study can not only generate a much higher revenue but also retain a higher proportion of consumers compared with the current industry standard policies. In addition, we find that the suboptimal policies can be efficient heuristics for the optimal policy, especially for longer content.

**Figure 10.** (Color online) Optimal Trigger Thresholds, Target Levels, and Advertisement Length for Both Infinite and Finite Horizons



## 6.2. Parameter Estimation

In order to implement the proposed policy, the platform needs to estimate the parameters used in this study. Generally speaking, the parameters used to characterize the trigger threshold  $u$  and target level  $U$  (i.e.,  $\mu, \sigma, \beta, \lambda, r, c_0, c$ ) are different for different users, as everyone has his or her own distinct taste for the content and dislike toward the advertisement. We next outline how the platform can estimate these individual's parameters.

To begin with,  $\mu$  is the trend term, which measures the user's overall engagement level of the content, and  $\sigma$  is volatility reflecting the magnitude of the stochastic change. After the platform tracks individual consumers' engagement level with a certain genre of the content (e.g., romance, action, and war), the platform can gather the data and discretize the continuous path by taking sufficient data points along the path. Then, the platform can run a simple linear regression analysis with engagement level being the dependent variable and time being the independent variable. The derived slope coefficient can be used to estimate the trend term,  $\mu$ , whereas the root mean squared error (RMSE) can be used to estimate the volatility term,  $\sigma$ . For each individual consumer, we need to estimate his or her parameters toward a certain genre; then, the platform can use these estimates to implement the customized advertisement based on different content.

The financial discount factor is denoted as  $\beta$ . When the content is relatively long or repeated (i.e., online games), the discount factor simply equals the compound interest rate, which is determined by the firm's financial capacity. Certainly, the length of most content is relatively short, and the consumer may leave the platform randomly (i.e., finite horizon with stochastic leave), which requires us to estimate another parameter,  $\lambda$ , in addition to the discount factor to characterize the thresholds. Recall that in the finite horizon with stochastic leave case, the user leaves the platform following an exponential distribution with parameter  $\lambda$ . That is, on an expectation, the user stays on the platform for a duration of  $E[\tau] = 1/\lambda$ , even when the user's engagement level is strictly positive. Thus, we only need to track the user's engagement with the platform several times and take the average of the time when the user leaves with his or her engagement still being positive to estimate the parameter  $\lambda$ .

Next, we discuss how the platform can estimate  $r, c_0$ , and  $c$ . Note that  $r$  is the revenue for the advertisement per unit, which is fully controlled by the platform or advertisement agency. The higher the value is, the more expensive the advertisement becomes. The platform can either charge a uniform price to all the advertisers or charge a different value to different advertisers based on consumer group.

The fixed and variable parts of a user's dislike toward the advertisement are represented by  $c_0$  and  $c$ , respectively. Both of the values can be estimated from the existing browsing history or by running experiments on a focus group of users. Specifically, the platform should track a user's engagement level before the insertion of the advertisement  $x_1$  and engagement level after the advertisement  $x_2$ . Then, the platform can run a simple linear regression with  $x_2 - x_1$  being the dependent variable and length of the advertisement being the independent variable. The derived slope coefficient is the estimate of  $c$ , whereas the constant coefficient can be the estimate of  $c_0$ .

## 7. Conclusions

In this paper, we study the optimal advertisement insertion policy when the user's engagement level can be tracked. We model the user's engagement level through a drifted Brownian motion. Each time the advertisement is played, the user's engagement level decreases. To characterize the structural results, we start with a model with an infinite horizon. Through the stochastic impulse control tools, we propose a threshold policy and analytically prove that it is optimal among all policies analytically. In addition, we analyze how the thresholds in the optimal policy change along with the discount factor, drift speed, volatility, and fixed decrease level of the user's engagement level from the advertisement insertion. Further, we investigate the case of a finite planning horizon. We show that the optimal policy is still based on a threshold policy, but the thresholds vary along with time. Finally, we conduct extensive simulation analyses to quantify the benefits of the proposed policies compared with the industry standard advertisement insertion policies. The results indicate that our proposed policies can not only generate significantly higher profits but also improve the user's retention rate.

Our research reveals several important and interesting theoretical findings. First, despite the complexity of the problem, we prove that a simple threshold policy is optimal among all class policies. This threshold policy has the advantage of easy implementation. Whenever the user's engagement level reaches a trigger threshold, it is optimal for the platform to insert an advertisement of appropriate length. Both the trigger threshold and optimal advertisement length can be solved rapidly. Based on the sensitivity analysis, our research shows that for the user who hates interruption during content (i.e., the fixed part of the user's dislike of the advertisement is large), is unlikely to leave due to an unexpected event (i.e., discount factor is small), has a high level of enjoyment of the material on the platform (i.e., the drift speed of the user's engagement level is high), or whose engagement level

varies dramatically (i.e., the volatility of the user's engagement level is large), it is optimal for the platform to set a higher trigger threshold and insert a longer advertisement. Moreover, for the finite planning horizon case, we find that the threshold policy is still optimal, but the optimal trigger threshold behaves very interestingly. Specifically, it first decreases, then increases, and finally decreases again as time approaches the end of the planning horizon. We show that this is a result of the balance between the advertisement length and the number of advertisements inserted by the content provider.

Our study also provides important and relevant insights for managers in related industries. First, it should be the platform rather than the content creator that decides the optimal timing of the advertisement. In the current practice, most platforms allow the content creator to choose the desired timing of the advertisement during the content to maximize revenue. If the platform can track users' engagement level, we show that using the optimal timing proposed by this study can yield a much higher profit than allowing the content creator to determine the timing of the advertisement (typically preroll, midroll, or a combination of the two). Further, our research suggests that the platform should prepare advertisements of differing lengths. This will help the platform to maximize its profit, as the optimal policy requires different advertisement lengths for different consumers. Finally, using the fixed thresholds derived from the infinite horizon case may produce similar results as those using the thresholds derived from the finite horizon case, without the heavy computation required for the latter.

We briefly note a few limitations of this study and provide some ideas for future research. First, we have briefly discussed how to measure the parameters to implement the customized advertising. In order to fully commercialize the idea of this study, more research needs to be conducted, which may require interdisciplinary efforts from neuropsychology to computer science. Second, this study is focused on the optimal timing of digital advertising. Future research can develop an algorithm that can jointly optimize the content and timing of digital advertising for each individual consumer. Third, we have compared the proposed policies with the prevalent policies through an extensive simulation study. Empirical researchers can further verify the results in a field experiment setting. Fourth, the platform requires advertisement with different lengths to implement the customized advertising. Usually, this is not an issue, as the production and editing costs are borne by the advertisers. However, if the platform bears the costs of advertisement production and editing, it might affect the

platform's revenue. Future research may study this issue explicitly and investigate the platform's performance in this scenario. Notwithstanding these limitations, the current study presents a first step in understanding the optimal timing of digital advertising and contributes to the emerging field of research on wearable technologies.

### Acknowledgments

The authors are listed alphabetically and contributed equally to this work. Lai Wei is the corresponding author of this article. We thank the senior editor Xiaoquan (Michael) Zhang, the associate editor Zhengrui (Jeffrey) Jiang, and three anonymous reviewers for their valuable and constructive suggestions.

### Endnotes

<sup>1</sup> The main reasons that we adopt Brownian motion (BM) rather than geometric Brownian motion (GBM) are as follows. To begin with, the GBM process always remains strictly positive. In contrast, BM can reach the zero level; this fits the scenario that consumers can leave the platform if their engagement level drops to zero, either because they watch the lengthy advertisement or because they are bored with the content. Further, the magnitude of the change in the GBM process intensifies as time proceeds. In our research context, the magnitude of increment or decrease of engagement level is relatively stable.

<sup>2</sup> We provide a more detailed description of why Brownian motion fits into our setting in the Online Appendix I.

<sup>3</sup> Our study can be applied to most online/digital content in which consumers cannot establish an expectation of when they will be shown with the digital advertisement. If the consumer has an expectation of timing of advertisement during the content (i.e., National Football League (NFL) Super Bowl), he or she may behave differently.

<sup>4</sup> We have also verified that our main results hold when the revenue function is convex in the advertisement length,  $rl^2$ . Some of the results are provided in Theorem 5 in the Online Appendix.

<sup>5</sup> Our model is flexible in terms of the specific advertisement being played. That is, the platform can choose which advertisement to insert depending on the consumer's preference and taste. The platform can gather the user's preference information through data mining and by analyzing clickstream data.

<sup>6</sup> When the planning horizon is infinite,  $\beta$  denotes the financial discount rate, which ensures that the revenue function is well behaved. Without this parameter, the revenue function approaches infinity. When the planning horizon is finite, we can actually disregard this parameter. Further,  $\beta$  will have a similar impact as the user's random departure rate  $\lambda$ . We show this in Section 5.1 when we introduce the stochastic horizon case.

<sup>7</sup> According to Björk (2009), in the context of a finite planning horizon, "Generally speaking, there is little hope of having an analytical solution of a free boundary value problem, so typically one has to resort to numerical schemes" (p. 111), and "there are no analytical formulas for the pricing function or the optimal boundary" (p. 345).

<sup>8</sup> In Figure 8, the parameter settings are as follows:  $\mu = 0.1$ ,  $\sigma = 0.2$ ,  $r = 0$ ,  $\beta = 0$ ,  $c_0 = 0.6$ , and  $e_0 = 0.2$ . We have verified that our results are robust to a wide variety of parameter settings. Results are available in the Online Appendix K.

<sup>9</sup> Codes used in this study are available on request. All models and algorithms in this section are implemented in MATLAB 2017b with 1,000 replications.

## References

- Andrews M, Luo X, Fang Z, Ghose A (2016) Mobile ad effectiveness: Hyper-contextual targeting with crowdedness. *Marketing Sci.* 35(2):218–233.
- Bajpai P (2016) The current and future trends of digital advertising. Accessed April 5, 2019, <https://www.nasdaq.com/articles/current-and-future-trends-digital-advertising-2016-08-23>.
- Beales H (2011) The value of behavioral targeting. Accessed April 5, 2019, [https://www.networkadvertising.org/pdfs/Beales\\_NAI\\_Study.pdf](https://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf).
- Benaroch M (2018) Real options models for proactive uncertainty-reducing mitigations and applications in cybersecurity investment decision making. *Inform. Systems Res.* 29(2):315–340.
- Björk T (2009) *Arbitrage Theory in Continuous Time*. 3rd ed. (Oxford University Press, Oxford, UK).
- Branco F, Sun M, Villas-Boas JM (2012) Optimal search for product information. *Management Sci.* 58(11):2037–2056.
- Branco F, Sun M, Villas-Boas JM (2016) Too much information? Information provision and search costs. *Marketing Sci.* 35(4):605–618.
- Chen J, Stallaert J (2014) An economic analysis of online advertising using behavioral targeting. *Management Inform. Systems Quart.* 38(2):429–449.
- Chen J, Fan M, Li M (2016) Advertising vs. brokerage model for online trading platforms. *Management Inform. Systems Quart.* 40(3):575–596.
- Chen Y, Li X, Sun M (2017) Competitive mobile geo targeting. *Marketing Sci.* 36(5):666–682.
- de Jesus A (2018) Artificial intelligence in video marketing – Emotion recognition, video generation, and more. Accessed April 5, 2019, <https://emerj.com/ai-sector-overviews/artificial-intelligence-for-video-marketing-emotion-recognition-video-generation-and-more/>.
- Fan M, Kumar S, Whinston AB (2007) Selling or advertising: Strategies for providing digital media online. *J. Management Inform. Systems* 24(3):143–166.
- Faull J (2015) Jaguar using wearable tech and ground sensors to measure crowd feeling at Wimbledon. Accessed April 5, 2019, <https://www.thedrum.com/news/2015/06/29/jaguar-using-wearable-tech-and-ground-sensors-measure-crowd-feeling-wimbledon>.
- Foster T (2016) Ready or not, companies will soon be tracking your emotions. Accessed April 5, 2019, <https://www.inc.com/magazine/201607/tom-foster/lightwave-monitor-customer-emotions.html>.
- Ghosh S, Li X (2013) A real options model for generalized meta-staged projects—Valuing the migration to SOA. *Inform. Systems Res.* 24(4):1011–1027.
- Guo H, Marston S, Chen Y (2015) Push or pull? Design of content delivery systems. *Decision Sci.* 46(5):937–960.
- Gupta S (2015) In mobile advertising, timing is everything. *Harvard Bus. Rev* (November 4), <https://hbr.org/2015/11/in-mobile-advertising-timing-is-everything>.
- Hao L, Guo H, Easley RF (2017) A mobile platform’s in-app advertising contract under agency pricing for app sales. *Production Oper. Management* 26(2):189–202.
- Ke TT, Shen ZJM, Villas-Boas JM (2016) Search for information on multiple products. *Management Sci.* 62(12):3576–3603.
- Krishnan SS, Sitaraman R (2013) Understanding the effectiveness of video ads: A measurement study. *Proc. ACM Internet Measurement Conf.* (ACM, New York), 149–162.
- Kumar S, Sethi SP (2009) Dynamic pricing and advertising for web content providers. *Eur. J. Oper. Res.* 197(3):924–944.
- Kumar S, Dawande M, Mookerjee VS (2007) Optimal scheduling and placement of internet banner advertisements. *IEEE Trans. Knowledge Data Engrg.* 19(11):1571–1584.
- Liu D, Kumar S, Mookerjee VS (2012) Advertising strategies in electronic retailing: A differential games approach. *Inform. Systems Res.* 23(3):903–917.
- Lu X, Zhao X, Xue L (2016) Is combining contextual and behavioral targeting strategies effective in online advertising? *ACM Trans. Management Inform. Systems* 7(1):1–20.
- Mookerjee R, Kumar S, Mookerjee VS (2017) Optimizing performance-based internet advertisement campaigns. *Oper. Res.* 65(1):38–54.
- Reale E (2017) Is this the end of pre-roll video ads? Accessed April 5, 2019, <https://www.marketingtechnews.net/news/2017/jun/15/end-pre-roll-video-ads/>.
- Reddy S (2015) Breathing for your better health: Controlling your breath is an easy way to improve mental and physical health. *Wall Street Journal* (January 26), <http://www.wsj.com/articles/breathing-for-your-better-health-1422311283>.
- Russo JE (1978) Eye fixations can save the world: A critical evaluation and a comparison between eye fixations and other information processing methodologies. *Adv. User Res.* 5(1):561–570.
- Ryan P (2016) Time-based advertising is taking off. Accessed April 5, 2019, <https://www.mediapost.com/publications/article/275913/time-based-advertising-is-taking-off.html>.
- Shen Q, Villas-Boas JM (2018) Behavior-based advertising. *Management Sci.* 64(5):2047–2064.
- Shields M (2017) More than half of digital advertising is mobile. *Wall Street Journal* (April 26), <https://www.wsj.com/articles/more-than-half-of-digital-advertising-is-mobile-1493218800>.
- Slefo G (2017) Desktop and mobile ad revenue surpasses TV for the first time. Accessed April 5, 2019, <http://adage.com/article/digital/digital-ad-revenue-surpasses-tv-desktop-iab/308808/>.
- Sloane G (2017) Facebook to lift longtime ban on pre-roll ads. Accessed April 5, 2019, <https://adage.com/article/digital/facebook-test-pre-roll-video-ads-ahead-watch-shows/311467>.
- Sun Z, Dawande M, Janakiraman G, Mookerjee V (2017) Not just a fad: Optimal sequencing in mobile in-app advertising. *Inform. Systems Res.* 28(3):511–528.
- Tan Y, Mookerjee V (2005) Allocating spending between advertising and information technology in electronic retailing. *Management Sci.* 51(8):1236–1249.
- Teixeira T, Wedel M, Pieters R (2012) Emotion-induced engagement in internet video advertisements. *J. Marketing Res.* 49(2):144–159.
- Villas-Boas JM (2018) A dynamic model of repositioning. *Marketing Sci.* 37(2):279–293.
- Vivero D (2018) The 3 healthcare benefits your millennial workforce wants, from the POV of a millennial CEO. *Forbes* (March 26), <https://www.forbes.com/sites/amino/2018/03/26/the-3-healthcare-benefits-your-millennial-workforce-wants-from-the-pov-of-a-millennial-ceo/#64f25fc5749d>.
- Wedel M, Pieters R (2008) A review of eye-tracking research in marketing. Malhotra NK, ed. *Review of Marketing Research*, vol. 4 (M. E. Sharpe, Armonk, NY).
- Zhang XM, Zhang L (2015) How does the internet affect the financial market? An equilibrium model of internet-facilitated feedback trading. *Management Inform. Systems Quart.* 39(1):17–37.
- Zhao X, Xue L (2012) Competitive target advertising and user data sharing. *J. Management Inform. Systems* 29(3):189–222.