

Split-second Decision-Making in the Field: Response Times in Mobile Advertising*

Khai Chiong
University of Texas at Dallas

Matthew Shum
California Institute of Technology

Ryan Webb
University of Toronto

Richard Chen
Happy Elements, Inc.

November 23, 2018

Preliminary and incomplete draft

Abstract

In this paper we take the class of *drift-diffusion* models from psychology and neuroeconomics, which were developed to jointly explain subjects' choices and response times in quick, split-second decision tasks in laboratory experiments, to a field setting – app users' response to advertisements on their mobile devices. We specify a two-stage drift-diffusion model to accommodate features of mobile advertisements. In most mobile advertising platforms, including our application, ads are “non-skippable”– that is, users are forced to watch the ad in its entirety. We use our estimates to simulate the counterfactual of “skippable” ads, and we find that permitting users to take an action while the ad is still playing would lead to lower click-through rates. While this finding rationalizes industry practice, the effects are very heterogeneous across users.

Keywords: Mobile advertising, Drift-diffusion model, Response times, Video advertisements, Skippable ads

JEL codes: L81, M37, D83, D87, C15, C22

*We thank Colin Camerer, Lawrence Jin, Xiaoxia Shi, Pengfei Sui and attendees at the Caltech Applied Micro Lunch, and the 2018 California Econometrics Conference for helpful comments.

1 Introduction

Large swaths of the empirical literature in economics is predicated on the notion of revealed preference, that agents' preferences can be recovered from the choices that they make. In many choice situations, however, *decision times* or *response times* – that is, the amount of time it takes decisionmakers to make a choice – can also be sharply informative about the decisionmakers' preferences. A swift food choice at a restaurant may indicate strong preferences for the chosen dish relative to the other available options; a legislator's delay and hesitation before announcing their vote may suggest that the policy options are very similar and difficult to distinguish in the legislator's mind. In this paper, we consider structural models of choice which incorporate not only the decisionmakers' eventual choices, but also their response times.

Specifically, taking a cue from psychology and neuroeconomics, we consider the *drift-diffusion* model, which has been postulated for modeling agents' choices in quick, "split-second" decisions, often in a matter of seconds. In the drift-diffusion framework, decisionmakers facing a choice problem are assumed to be accumulating information about the available options; the response time is modelled as the time it takes a decision-maker to accumulate enough information or evidence favoring a given choice. In this way, both the choice and response time are organically related as part of a dynamic choice process. While the drift-diffusion model has been well-tested in lab experimental environments, we apply this model to field data.

We apply the drift-diffusion model to a choice scenario arising in mobile advertising where we observe whether users respond favorably by "clicking through" on an interactive video advertisement, and the corresponding response time for users to reach that decision. Much advertising on mobile apps takes the form of short videos, which users of an app are shown and forced to watch before being able to continue using the app. After the ad video finishes playing, the app loads an 'end card', which is a small website which prompts users to either "click-through" to an app-store and download the advertised app, or "click-back" and return to the app which they were using. Decisions to click-through or click-back are typically made within no more than a few seconds, and this simple binary choice problem closely mimics the quick decision tasks which the drift-diffusion model has been used to calibrate in lab settings.

We use our estimates to address an important policy counterfactual for mobile advertising platforms: specifically, whether advertisers should make their ads "skippable" – that is, allow ad viewers to skip the ad and take a decision before the ad finishes. Most mobile advertising

platforms force viewers to watch the entire ad before making any decisions – that is, ads are “nonskippable” – and our paper studies an example from one such platform. Our results indicate that this practice is sensible, as making the ad completely skippable would lead to 44% lower clickthrough rates. But we also find that the clickthrough rates from forcing viewers to watch the entire 30-second ad are virtually the same as forcing them to watch only the first 10 seconds of the ad. This rationalizes the practice of some websites, notably YouTube, where users can skip an ad after some initial period, usually 5 or 10 seconds.

The results and methodologies in this paper may have applications to other markets and industries; this is useful as more and more “clickstream” datasets is becoming available to researchers in economics and marketing, in which timestamps are available for all the choices that agents are observed to make, and the methods used in this paper provide a way of integrating the timestamp data organically into the empirical analysis. More broadly, a take away from this paper is that incorporating response times may be important for accurate modelling of agents’ choice behavior in discrete-choice settings, as well as better estimation of agents’ preferences from available choice data.

As far as we are aware, this paper is one of the first attempting to fit a model in the drift-diffusion paradigm to field data, by estimating a structural econometric model. There is a sizeable literature in psychology and neuroeconomics testing implications of the DDM model in experimental data, but estimating the parameters of the DDM using experimental data has only been done in a handful of papers, including [Frydman & Nave \(2016\)](#) and [Clithero \(2018\)](#). Researchers in marketing have long had access to field datasets in which both consumers choices as well as response times were recorded, and there has been a small set of papers incorporating both of these outcomes in a structural choice models. The earliest of these papers is [Otter, Allenby & Van Zandt \(2008\)](#) while, more recently, [Seiler & Pinna \(2017\)](#) and [Ursu, Wang & Chintagunta \(2018\)](#) incorporate response time into a structural model of consumer search.

2 The Drift-Diffusion Model (DDM) Paradigm

Here we give a brief introduction to drift-diffusion models (DDM).¹ The drift-diffusion framework, comprising decision models in which an information accumulation process is represented by a stochastic diffusion process, originated in psychology, where it was used to model subjects' choices and reaction times in laboratory experiments involving quick, split-second tasks where decisions were typically made within a matter of seconds.² (A typical task is to display two numbers and ask subjects which number is larger.) The DDM has spawned a large literature; it has been successfully verified and calibrated in laboratory experiments, and mapped to neural activity in the brain, of both humans and animals.³

For the purposes of the mobile advertising example below, we consider a binary choice situation. Let $i \in \{0, 1\}$ denote the choices, and let time t be continuous. In a binary choice problem, the *difference* in the utilities of the two options determines an optimizing agent's choice. In a DDM framework, however, the utilities are assumed to be unknown to the agent, who only learns about them gradually over time. Specifically, the perceived utility difference between option 1 and option 0 is modelled mathematically as a Gaussian process Z_t with evolution given by the stochastic differential equation

$$dZ_t = (\mu + \gamma Z_t) dt + \sigma dW_t, \quad t > 0, \quad Z_0 = 0. \quad (1)$$

This stochastic process can be interpreted as a process of “evidence accumulation” favoring one or the other of the choices. In the above, μ is a drift term, σ is a standard deviation, and W_t is a standard Brownian motion (Wiener process). γ is known in the literature as the “leakage” parameter (Usher & McClelland (2001)), and allows users to systematically over- or underweight earlier information. When $\gamma = 0$, then the stochastic process (1) is essentially a continuous-time random walk with drift, while nonzero γ makes it an autocorrelated (Ornstein-Uhlenbeck) process, in which past and future pieces of evidence can be correlated.

Paired with this utility difference process is a decision rule. The most common decision rule is the following “first passage” rule: for some threshold or barrier $B > 0$, the agent chooses 1


¹We follow the presentation in Webb (2016). Book-length treatments of this modeling framework include Luce (1991) and Link (1992). Formal mathematical fundamentals of DDMs are presented in Smith (2000), who also provides closed-form expressions for the distributions of response times.

²See, eg. Luce (1991), Busemeyer & Townsend (1992), Ratcliff & McKoon (2008).

³Gold & Shadlen (2002), Krajbich, Lu, Camerer & Rangel (2012)

once $Z_t > B$ and chooses 0 once $Z_t < -B$. That is, the response time is the *first passage* time to either B or $-B$. See Figure 1.

Figure 1: Drift-Diffusion Model: an illustration



Sample path of Z_t in blue. In this example, DM chooses 1, with a response time of t^* .

For the purpose of this paper, we will take the decision rule as given. There has been a large literature studying optimal decision rules in this setting, starting with Wald's 1973 classic work in sequential analysis⁴ up to recent theoretical work by Fudenberg, Strack & Strzalecki (2017). For the purposes of specification and estimation, we will not assume that the decision rule is chosen according to an optimal learning theory, which potentially allows for testing of optimality.⁵ Furthermore, the DDM choice framework can be considered a "high-frequency" or "split-second" (Lu & Hutchinson (2017)) version of dynamic search or learning models which have been estimated in the empirical industrial organization and marketing literatures⁶, or duration models in statistics.⁷

⁴See also discussion in Lehmann (1959) and Bogacz, Brown, Moehlis, Holmes & Cohen (2006)

⁵As will be clear below, our estimation approach allows us to consider a wide variety of decision rules, including those where the threshold B may vary across time.

⁶See, eg., Erdem & Keane (1996), Crawford & Shum (2005), Hong & Shum (2006), Honka (2014)

⁷Cox & Oakes (1984)

3 Application: Video advertising on mobile platforms

Our empirical illustration comes from a dataset on mobile advertising. The number and variety of applications on mobile devices (“mobile apps”) has mushroomed in the last decade. A dominant revenue model for apps has emerged: the so-called “freemium” model in which users download the app for free, but are then subsequently monetized by periodically showing them video advertisements which interrupt their use of the app. As a result, the mobile advertising has recently become the most dominant segment of digital advertising. For instance, in the United States, businesses spending on mobile advertising is accounting for roughly 75% of all spending in digital ads, making mobile advertising larger than TV advertising for the first time in 2018.⁸ In what follows, we will refer to the originating app which the user was on as the “publisher”, and the creator of the advertisement as the “advertisers”. In most cases, the advertiser, like the publisher, is an app – an example could be users playing *Candy Crush* being interrupted by a video ad for *Clash of Clans*.⁹

An important, and perhaps most familiar, ad format in mobile advertising is the “non-skippable ad”: users on a particular publisher app are served a 30 seconds video trailer ad, typically for another app. They are not allowed to skip the ad. After the ad, a user has two choices. First, the user can close the ad (go back to the game). Second, the user can click on the ‘install’ button which takes the user to an App Store, where the user has the opportunity to find out more about the advertised app and install it. We will use “click-back” to denote the first action, and “click-through” to denote the second.

Our data come from a mobile ad network, which acts as an intermediary between advertisers and publishers in this two-sided platform. Publishers monetize by showing ads to their users, and advertisers are interested in acquiring new users. Both advertisers and publishers are mobile apps. Although we potentially have a large amount of data, we will focus in this paper on data from a *single* advertiser-publisher pair, that is, the same ad is shown to users within a single app (the publisher) – in industry parlance, this is a single ad campaign. While we cannot reveal the identity of the publisher or advertiser for confidentiality reasons, both are popular gaming apps, with daily average users in the tens of millions. Thus, even within this single advertiser-publisher pair, we observe $N = 194,510$ ad exposures, or “impressions”.

⁸<https://www.forbes.com/sites/johnkoetsier/2018/02/23/mobile-advertising-will-drive-75-of-all-digital-ad-spend-in-2018-heres-whats-changing/>

⁹See [Chen & Chiong \(2016\)](#) for a study of a mobile advertising auction market in which advertising rates are set.


Among these, 190,451 (97.91%) are clickbacks, and only 4,059 (2.09%) are clickthroughs. This matches the industry norm, where clickthrough rates are typically between 1.5-3%.

Table 1: Summary statistics.

Variables	Mean	St. Dev.
<i>Android Version</i>	2.82	2.71
<i>Device Volume</i>	0.541	0.314
<i>iOS Dummy</i>	0.477	0.500
<i>iOS Version</i>	4.69	4.91
<i>Samsung Dummy</i>	0.280	0.449
<i>Screen Resolution</i>	1.30	0.832
<i>WiFi</i>	0.959	0.198
<i>Clicks, y</i>	0.0209	0.143
<i>response time</i> $y = 1$	4.97	1.52
<i>response time</i> $y = 0$	4.14	1.73
# observations	194,510	

We define a binary outcome variable y_i describing user i 's choice: $y_i = 1$ if she clicks-through, and $y_i = 0$ if she clicks-back. We also observe the response time t_i^* , defined as the time (in seconds) elapsed from the end of the ad until a decision has been made. Across the observations, there is substantial observed heterogeneity across users, and we include a number of user-specific covariates X_i in our analysis. These are *Language*: the language used in the user's mobile device; *WiFi*: whether the device is connected to a WiFi network or not; *Device Volume*: a numeric value from $[0, 1]$ that describes the volume level of the user's device at the moment of ad serving; *Screen Resolution*: the number of pixels (per million) of the user's mobile device; *Android Version*: an integer-valued variable from 1 to 8 indicating the version number of the user's Android mobile operating system.

Figure 2: Empirical density of the actual observed response times



We plot smoothed kernel densities of response time in Figure 2. A clear feature which emerges is that users who decide to click-through take longer to make that decision. On average, users who click-through take 4.97 second (median=4.88) to make that decision. For those users clicking-back, however, the average response time is only 4.14 seconds (median=3.99), which is significantly lower.

The finding that the distribution of response times differs depending on the choice is noteworthy, as it is analogous to the so-called “slow errors” phenomenon commonly found in lab experimental settings involving attention tasks. For instance, in experiments where subjects were asked to judge which of two numbers is larger, mistakes tended to be rare events, but when subjects made them, the response time was longer. (See the discussion in Luce (1991).) Similarly, in our field setting, we observe that the “rare event” of clicking an advertised app tends to be chosen after a longer response time. While it would not be appropriate to interpret the choice of clicking as a “mistake”, one interpretation is that in both cases (making a mistake and choosing to click), more mental consideration and information gathering was involved – neither the mistakes nor the click decisions were hasty or lazy decisions made with little forethought – if anything, they involved *more* mental effort.

The standard DDM model – that is, one in which the utility difference is given by Eq. (1) with

$\gamma = 0$, and with a fixed threshold decision rule – cannot generate the “slow errors” phenomenon; instead, it implies that the distribution of response times is the same for both choices $y = 0$ or $y = 1$. Indeed, when the drift term $\mu \neq 0$, the two choices $y = 0$ and $y = 1$ will have different overall probabilities, but the average response time for both choices will be the same.¹⁰

3.1 A model of non-skippable mobile ads

The standard DDM model sketched in the previous section, while convenient and well-suited for modelling many binary choice tasks in laboratory experiments, cannot be applied to user choices in our mobile advertising platform, because it doesn’t take into account important features of the platform. Specifically, for the “non-skippable” ad which we analyze, users must watch the entire 30-second ad, with no opportunity to choose an action (neither clickback nor clickthrough) while the ad is playing. Only after the ad has finished playing can users make a decision.

This feature, that users’ choices are only permitted after an initial “inactive stage” while the ad is playing, requires adjusting the DDM model. Since users do learn about the advertised app from watching the ad, it would be wrong to assume that the DDM model begins only when the ad ends. At the same time, it also seem a stretch to assume that the same accumulation process in both the “inactive stage” (as the ad is playing) and the “active stage” (after the ad ends), as users may not be fully paying attention as the ad is playing.

For this reason, we develop here a “two-stage” DDM model. Users are not allowed to take any action (click or skip) during the initial inactive stage – that is, while the 30-second ad is playing. We model users’ information accumulation during the inactive stage as a diffusion process *without* barriers: specifically, for $0 \leq t \leq 30$, information accumulates as:

$$dZ_t = \mu_0 dt + \gamma_0 Z_t dt + \sigma_0 dW_t \quad (2)$$

Equation 2 is a re-parameterization of the Ornstein-Uhlenbeck process. Furthermore, we will normalize $\sigma_0 = 1$, as this parameter cannot be identified from the data at our disposal. Since

¹⁰Shadlen, Hanks, Churchland, Kiani & Yang (2006) provides a clear discussion of this. Essentially, this arises from the independence and symmetry of the increments in the basic DDM specification (1) with $\gamma = 0$. See also Woodford (2016).

there is no barrier during this stage, the probability density of the position of the accumulation process (2) at any point in time is characterized by a partial differential equation called the Fokker-Planck equation. The solution to this partial differential equation can be solved in closed-form as follows.

$$Z_t \sim \mathcal{N} \left(\frac{\mu_0}{\gamma_0} (e^{\gamma_0 t} - 1), \frac{\sigma_0^2}{2\gamma_0} (e^{2\gamma_0 t} - 1) \right), \text{ for } 0 \leq t \leq 30 \quad (3)$$

In the *active stage*, after the ad has finished playing and users have the opportunity to either clickback or clickthrough, the usual DDM runs. That is, the accumulation process is described by:

$$dZ_t = \mu_1 dt + \gamma_1 Z_t dt + \sigma_1 dW_t, \quad \text{for } t > 30 \quad (4)$$

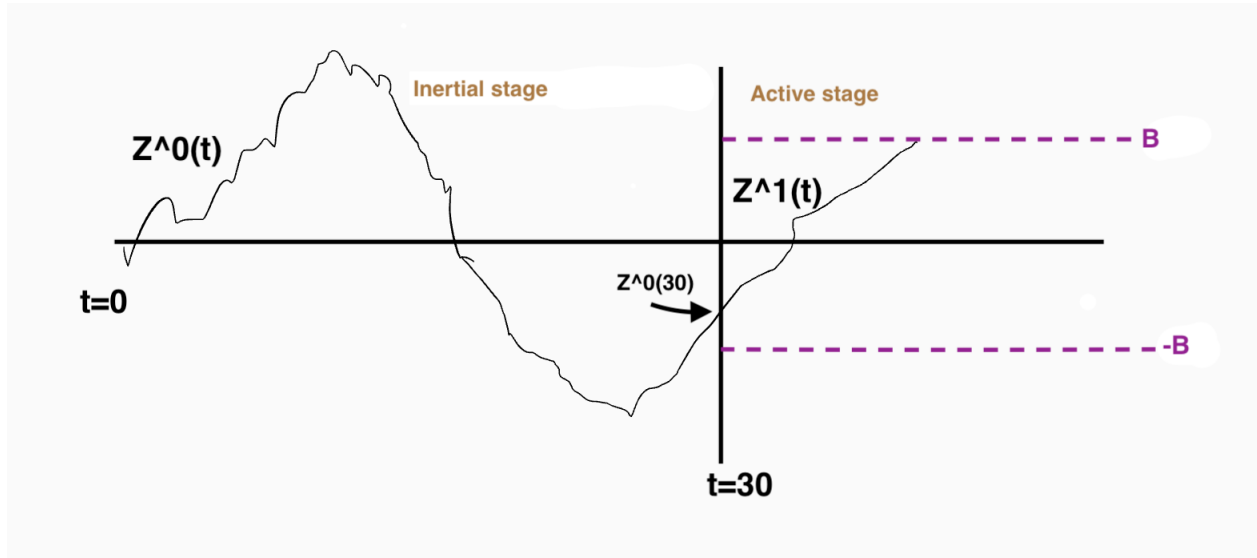
The initial condition for the DDM in Equation 4 (that is, Z_t at $t = 30$) is governed by Equation 3. Hence the initial starting point is itself random. The upper barrier is B and the lower barrier is $-B$. The stopping rule is as before: a user makes the choice $y = 1$ at the response time t^* if and only if the process first hits the barrier B at time t^* . That is, if t^* is the response time, then $Z_{t^*} \geq B$ but $Z_t < B$ for all $30 < t < t^*$. Similarly, we observe $y = 0$ at response time t^* if and only if $Z_{t^*} \leq -B$ but $Z_t > -B$ for all $30 < t < t^*$. See Figure 3.

3.1.1 Remarks

Essentially, and intuitively, the model here with an inactive stage followed by an active stage, is just the DDM model presented earlier in Eq. (1), but with a random starting value – that is, a starting value given by the normal distribution in Eq. (3) – instead of the usual starting value of zero.¹¹ This has important consequences for the “slow errors” issue discussed earlier, which is present in our data. Ratcliff & Rouder (1998) show that once users are assumed to have nonzero starting values, then the response time distribution will vary across the two choices – either “slow errors” or “fast errors” is possible. Thus, as we will see in our results below, our model here, with an inactive period preceding the active stage, not only models the

¹¹Frydman & Nave (2016) also consider a DDM model with random starting points across subjects.

Figure 3: Two-stage Drift-Diffusion Model



actual mobile advertising platform, but also accommodates the feature from the data that the distribution of response times varies across choices.¹²

We allow the accumulation processes to differ in the inactive and active stages, with the former characterized by the parameters μ_0 , σ_0 , and γ_0 , while the latter is characterized by μ_1 , σ_1 and γ_1 . These differences allow for the possibility that the information may have different precision between the two stages, or that users may be paying different levels of attention. As we will see below, our estimates suggest large differences in the accumulation processes in the two stages.

In addition, the observation that many subjects take up to 8 seconds to make a decision whether to click-through or click-back after viewing a video ad may be surprising or unexpected, relative perhaps to our self-introspection. But researchers in this setting have documented a “cold-start” phenomenon where users may be “starry eyed” and not fully aware that the ads have ended. In addition, users may not be paying attention the entire duration of the ads.

¹²In our empirical work, we will also allow extra flexibility by having the leakage parameter γ take values other than zero. Furthermore, the inclusion of the leakage parameter γ allows users to forget or de-emphasize earlier information, which distinguishes our model from the Poisson-based model of choice and reaction time estimated by Otter et al. (2008).

3.2 Parameter identification

Before proceeding, we provide some discussion regarding identification of the structural model parameters. The cleanest identification arguments are evident for the simplest single-stage Drift Diffusion Model without the leakage parameter, and so we start there. In this simple setting (corresponding to Eq. (1) with $\gamma = 0$), the parameters of the model are identified up to σ , using only two moments of the data: the mean response time,¹³ and the click-through probability. We have the following lemma (with proof in the Appendix).¹⁴

Lemma 1. *In the simple DDM, the drift μ and the bound B are identified using only two moments of the data: $\bar{t} \equiv \mathbb{E}[t^*]$, the mean response time, and $\bar{y} \equiv \mathbb{E}[y]$, the probability of click-through. The variance of the diffusion, σ^2 is normalized to 1. Specifically, the estimates $\hat{\mu}$ and \hat{B} are:*

$$\hat{\mu} = \begin{cases} -\sqrt{\frac{(2\bar{y}-1) \log\left(\frac{\bar{y}}{1-\bar{y}}\right)}{2\bar{t}}} & \text{if } \bar{y} < 0.5 \\ \sqrt{\frac{(2\bar{y}-1) \log\left(\frac{\bar{y}}{1-\bar{y}}\right)}{2\bar{t}}} & \text{if } \bar{y} \geq 0.5 \end{cases} \quad (5)$$

$$\hat{B} = \frac{1}{2} \sqrt{\frac{2\bar{t} \log\left(\frac{\bar{y}}{1-\bar{y}}\right)}{2\bar{y} - 1}} \quad (6)$$

The expression for the estimated drift in Equation 5 bears similarity to the expression for latent utilities in a Logit model. In particular, $\log\left(\frac{\bar{y}}{1-\bar{y}}\right)$ is the latent utility in terms of the choice probability \bar{y} . Hence, we may interpret the drift term μ as a latent utility. The important difference is that the magnitude of the estimated utility is affected by the response time – the magnitude of $\hat{\mu}$ is larger when the average response time \bar{t} is shorter. When a choice is made rather quickly, the preference of the individual must be more intense.

Moreover, the impact of response time on the estimated utilities is the greatest when the choice probability is further away from 0.5. That is, precisely when one of the choice is rare, that incorporating response time would greatly affect the estimated utilities. For instance, if the

¹³As discussed above, the simple DDM predicts the same mean response times *conditional* on choice $y = 0$ or $y = 1$.

¹⁴To our knowledge, such identification results are new in the primarily experimental literature where DDM is prevalent; probably this is because, in experimental settings, the drift parameter is usually assumed to be known.

probability is 0.02, the estimated utility would be $\hat{\mu} = -1.37$ when $\bar{t} = 1$, and $\hat{\mu} = -0.68$ when $\bar{t} = 4$.

Compared to this setting, our two-stage model contains the leakage parameter γ , as well as $(\mu_0, \sigma_0, \gamma_0)$, the parameters for the first-stage diffusion process. The identification discussion is less complete here. Given the above discussion, we have fixed both of the variances σ_0 and σ_1 to be equal to 1 in our estimation, but leave B as a free parameter to be estimated.

The identification of the remaining parameters is related to differences in the distribution of response times across click-backs vs. click-throughs. As we discussed earlier, the two-stage model is essentially a DDM with random starting values. The non-zero starting values, as well as non-zero leakage parameters, are both, by themselves, capable of generating differences between the response time distributions across choices. Moreover, the properties of diffusion processes imply that these starting values are normally distributed, and Eq. (3) shows how the structural parameters μ_0 and γ_0 enter the mean and variance of the starting value distribution.

4 Estimation

Our two-stage accumulation model of decision-making in the mobile advertising platform resembles an optimal-stopping model; therefore, it is not surprising that the econometric procedures for estimating this model share more similarities with methodologies for the estimation of structural dynamic models, than with static discrete choice models.¹⁵

For the simple DDM model, closed-form expressions are available for the optimal choice probabilities and associated response times (see Shadlen et al. (2006)), and software packages are available which exploit these closed forms for fitting the model to data.¹⁶ However, for our customized two-stage accumulation model, no such closed forms are available.¹⁷ Hence, we present a simulation-based Maximum Likelihood approach for estimation, which is based on

¹⁵Indeed, our simulation based approach resembles that in Pakes' (1986) optimal stopping model of patent renewal.

¹⁶See, for instance, Wiecki, Sofer & Frank (2013) or Jan Drugowitsch's code at <https://github.com/DrugowitschLab/dm>.

¹⁷There are no known closed-form solutions for the *first-passage time* of our two-stage process. Even for the standard Ornstein-Uhlenbeck process proposed in Equation 1, a closed-form density of the first-passage time is only available when there is a single absorbing barrier, as opposed to two absorbing barriers here (Alili, Patie & Pedersen (2005)).

discretizing time into small finite increments.¹⁸ Since there are no barriers during the first-stage, we only need to discretize the second-stage. Let $\tau = \{t_0, t_1, t_2, \dots\}$ denote a grid of time-point for the second active stage of decision-making process. Since the ad is 30 seconds long, and users are only allowed to make decisions after that, we set $t_0 = 30$. Denote the (constant) time increment by $\Delta t = t_{j+1} - t_j$ for all j . For our estimation, the time increment is a quarter of a second, that is $\Delta t = 0.25$.¹⁹ Based on this time-increment, we will introduce a discrete accumulation process as follows.

Let $Z_i(t_j)$ denote the value of the accumulation process for user i at time t_j . Let $\mu(X_i; \theta)$ denote the (time-independent) drift specific to a user with covariate X_i . For each candidate parameter vector θ , we can simulate the accumulation process $Z_i(t_j; \theta)$ at each $t_j \in \tau$ recursively as:

$$Z_i(t_j) - Z_i(t_{j-1}) = \mu(X_i; \theta)\Delta t + \gamma Z_i(t_{j-1})\Delta t + \sigma \Delta W_i(t_j) \quad (7)$$

$$Z_i(t_0) \sim \mathcal{N}\left(\frac{\mu_0(X_i; \theta)}{\gamma_0} (e^{30\gamma_0} - 1), \frac{1}{2\gamma_0} (e^{60\gamma_0} - 1)\right) \quad (8)$$

For each user i (and given parameters), we generate S number of independent sample paths described by Equations 7 and 8. The increments of the Brownian motion $\Delta W_i(t_j)$ is drawn i.i.d from $\mathcal{N}(0, \Delta t)$, for each user i and for each $t_j \in \tau$.²⁰ For each user i and each sample path $s = (1, \dots, S)$, we record $(y_{i,s}, t_{i,s}^*)$, where $t_{i,s}^*$ is the time the sample process s first hits the upper bound or the lower bound (whichever first) for user i . Similarly, $y_{i,s} = 1$ if the sample process hits the upper bound at $t_{i,s}^*$ and $y_{i,s} = 0$ otherwise. The number of sample paths S is set to 5,000.

¹⁸Although not called as such, the data-fitting exercise in Fudenberg et al. (2017) appears essentially to be a simulation-based approach based on simulating the underlying stochastic processes and stopping times.

¹⁹Section B in the appendix shows that this time-increment is good enough in the sense that the pdf of the discrete accumulation process closely approximates the pdf of the continuous accumulation process. We will show this in Section B.

²⁰Generating a large number of random numbers and subsequently building up the process according to Equation 7 is a highly-parallelizable problem. We develop a GPU-based simulated MLE that allows us very efficiently evaluate the simulated likelihood.

In the data, we observe (y_i, t_i^*, X_i) for each user i . Given our discretization, the response times t_i^* are discrete, implying that the likelihood of each observation $P(y_i, t_i^* | X_i, \theta)$ is multinomial, and can be approximated by the frequency of the observed choice and (discretized) response time among the simulation sequences of the accumulation processes. Thus we estimate θ by maximizing the full simulated log-likelihood function $\sum_{i=1}^n \log(\Pr(y_i, t_i^* | X_i, \theta))$.

For estimation, we sample $n = 10,000$ observations. Because the observations with $y = 1$ (clickthroughs) are much rarer, we oversampled the observations with $y = 1$ but correct for this oversampling in the estimation procedure.

4.1 Results

Table 2: Estimates of the single-stage DDMs.

	(1)	(2)
$B, Bound$	2.89	7.02
μ_1	-0.67	-1.36
σ_1	1 (fixed)	1 (fixed)
γ_1		0.163

The first column is the simple DDM – estimates are recovered using the Method of Moments estimators in Equations 5 and 6. The second column introduces the leakage parameter γ_1 . The estimates from this column are obtained via Simulated MLE.

4.1.1 Single-stage model results

As a preamble, we consider results from single-stage models, which ignores the inactive stage and assumes that all users begin the active stage with initial values of zero. These results are reported in Table 2. While this model is descriptively wrong for our mobile advertising platform, we will still discuss the results in some detail as their qualitative features persist in the subsequent estimates from the two-stage models below.

The drift parameter of the simple DDM is $\mu_1 = -0.67$, indicating that the process drifts towards the lower bound, corresponding to more frequently-made choice of $y = 0$. For com-

parison, we also fitted a simple binary logit model to the observed choice probabilities, and obtained an estimate of -3.84 for the net utility from click-through, which is six times the magnitude of the estimated μ_1 . From Eq. (5), such a large negative value for utility is consistent with a mean response time of only 0.1245 seconds, which is clearly only a fraction of the response times we observe in the data.

The estimate of the boundary $B = 7.02$. Compared to the magnitude of the variance of the active stage process, σ_1 , which is fixed at 1, the boundary is very wide. The drift term μ_1 is negative, and around the same magnitude (-1.36) as the fixed value for σ_1 . The negative drift in the process is expected, as it implies a bias towards the clickback ($y = 0$) choice, which is confirmed by the overwhelming frequency of click-back vs. click-through choices in our data. Finally, we estimate the leakage parameter γ_1 to be positive.

4.1.2 Two-stage model results

Next, in Table 3, we consider specifications of the two-stage model which, as described before, is essentially a DDM model augmented with a random non-zero starting value. Column (1) contains estimates of the simplest two-stage model, in which the parameters of the inertial and active stage inertial processes are assumed to be identical across users. We will refer to this as to “homogenous user” specification in what follows. The specification in column (2) allows μ_0 and μ_1 , respectively the drift terms in the inertial and active stage, to vary across users depending on observed covariates; we will refer to this as the “heterogeneous user” specification.

Users heterogeneity is substantively important, as the homogeneous and heterogeneous user specifications differ sharply in their implications for the drift terms μ_0 and μ_1 . In the homogenous user specification, $\tilde{\mu}_0$ is estimated to be negative (-1.172), which implies a negative drift in the inactive accumulation process (ie. $\mu_0 < 0$). This negative drift implies that users obtain overall pessimistic signals as the ad is playing, and hence become *less* inclined to click-through the longer they view the app.

On the other hand, in the heterogeneous specification, corresponding to the results in column (3) of Table 3, we find that while the constant in the specification of $\tilde{\mu}_0$ is negative (-1.12), the coefficients on a number of the covariates are positive and sizeable in magnitude, so that there are users with configurations of covariates for whom the drift during the inactive stage

Table 3: Two-stage model: Parameter estimates and bootstrapped standard errors

	(1)		(2)	
	Homogenous Users		Heterogeneous Users	
γ_1	0.195	(0.0165)	0.218	(0.0190)
B , Bound	9.215	(0.4342)	9.149	(0.4199)
μ_1 : Constant	-1.143	(0.0719)	-1.448	(0.0670)
μ_1 : Android Version			-0.0464	(0.0109)
μ_1 : Device Volume			0.319	(0.0182)
μ_1 : iOS Dummy			0.0947	(0.0113)
μ_1 : iOS Version			0.0282	(0.0116)
μ_1 : Samsung Dummy			-0.327	(0.0204)
μ_1 : Screen Resolution			0.101	(0.0102)
μ_1 : WiFi			-0.374	(0.0220)
σ_0 and σ_1	1	(fixed)	1	(fixed)
γ_0	-0.208	(0.0615)	-0.2909	(0.0169)
μ_0 : Constant	-0.207	(0.0564)	-0.0758	(0.0125)
μ_0 : Android Version			0.125	(0.0125)
μ_0 : Device Volume			0.0414	(0.0109)
μ_0 : iOS Dummy			0.0683	(0.0129)
μ_0 : iOS Version			-0.0578	(0.0105)
μ_0 : Samsung Dummy			0.0220	(0.0102)
μ_0 : Screen Resolution			-0.115	(0.0139)
μ_0 : WiFi			0.118	(0.0128)
Log-Likelihood	-100824.73		-96863.87	

is positive: that is, these users get *positive* signals while viewing the ad. From the coefficient estimates, these users are those with newer version of Android software (coefficient 0.125) and those using iOS (coefficient 0.0683). Those viewing the ad at higher volume (0.0414) are also more favorably inclined to the ad, which suggests that audial salience of the ad makes it more effective; however, the coefficient on screen resolution is negative (-0.115), so that visual salience has an opposite effect.

More formally, to summarize the heterogeneity in the drift terms μ_0 and μ_1 across users, we partitioned the users in our sample into ten clusters by the values of their covariates X using a k -means procedure; Table 4 reports the estimated values of μ_0 and μ_1 for the “medoids” of each cluster (the users with, roughly speaking, the median values of the covariates within each cluster). Clearly, there is substantial heterogeneity in the drift terms across users, with four of the ten medoids (#1-4) having $\mu_0 < 0$ and the others (#5-10) with $\mu_0 > 0$. Moreover, while $\mu_1 < 1$ for all medoids, there is a negative relationship between the two drift terms across medoids; the medoids with negative μ_0 have smaller magnitudes of μ_1 , while those with positive μ_0 have larger magnitudes of μ_1 . As we will see below, these big differences in the estimated drift terms for the homogeneous and heterogeneous user specification have important implications for the counterfactuals below.

For the remaining parameters, from Eq. 3, the variance of the accumulation process is $\frac{1}{2\gamma_0}(e^{2t\gamma_0} - 1)$. At the estimated value of $\gamma_0 = -0.227 < 0$, this variance is decreasing in t , and converges asymptotically to $-\frac{1}{2\gamma_0} = 0.220$. The variance term changes most rapidly between 0 and 4 seconds, but moves little beyond that. Clearly, information accumulates most quickly during the first few seconds of the ad. This may arise both from the intrinsic content and design of the video ad, or by decreasing user attention as the ad is playing.

While γ_0 plays a role in moderating the degree of divergence or convergence of the variance, the drift μ_0 during the inactive period determines the differences in the reaction times conditional on action, which we observe in the data. This is illustrated in Figure 4. Using nonparametric regression, we predicted, for each user, the difference between response times for clickthrough and clickback, conditional on the user’s covariates X_j . We then plotted these predicted differences in response times (which we call “slow errors”) against the user’s inertial stage drift parameter $\hat{\mu}_{i0}$, as implied by the estimates in Table 3, Column 2. From the figure, it is clear that there is a negative relationship – that is, the inertial drift μ_{i0} tends to take more

Table 4: Drift terms μ_0 and μ_1 for ten user medoids

Medoid	μ_0	μ_1
1	-0.763	-1.009
2	-0.516	-1.097
3	-0.477	-1.215
4	-0.570	-1.245
5	0.263	-1.532
6	0.589	-1.755
7	0.591	-1.799
8	0.713	-1.863
9	0.598	-2.049
10	0.624	-2.110

We partition users into ten clusters using k -means procedure. For the medoid of each cluster, we compute the drift parameters as $\mu_0^{(k)} = \mathbf{X}^{(k)} \boldsymbol{\beta}_0$ and $\mu_1^{(k)} = \mathbf{X}^{(k)} \boldsymbol{\beta}_1$, for $k = 1, \dots, 10$. The medoids are labeled in order from the largest μ_1 to the smallest μ_1 .


negative values for users for whom the predicted difference between the time to clickthrough vs. time to clickback is largest. This illustrates how the inertial drift parameter μ_{10} is identified from variation in the magnitude of “slow errors”. Generally, if we interpret the inertial drift as reflecting the “priming” effect of the ad on click decisions, when a user is observed to take a longer time to clickthrough, it must be that the ad provided a more negative prime in the inertial stage.

5 Counterfactuals: Nonskippable vs. skippable ads

In this section we use our estimation results to simulate behavior under alternative ad formats. As discussed before, the data used for estimation derive from a “nonskippable” ad presentation in which app users are constrained to view the entire ad, before deciding either to clickback and return to the app they were using before the ad, or clickthrough and be taken to the App Store to find out more (and potentially install) the advertised app. This nonskippable ad format is the most common in mobile advertising.

In these counterfactuals, we consider alternative ad presentations, in which users are per-

Figure 4: Slow error versus estimated first-stage drift.



The blue line is a 4th-degree polynomial fit.

mitted to skip all of part of the ad.²¹ We compute a range of simulations in which the duration of the inactive stage ranges between zero seconds, denoting a completely skippable ad, up to 30 seconds, which is the non-skippable ad setting as in our application. (As an example, YouTube allows their users to skip ads after 5 or 10 seconds.) In these simulations, we assume that the accumulation process in Equation 12 holds true during the inactive period, and the accumulation in Equation 4 governs users' behavior after the inactive period.

We evaluate the counterfactual change in the click-through probabilities, and report the results in Table 5. The column marked "Homogenous Users" present counterfactual results corresponding to the homogenous users model specification, in Column (1) of Table 3, a specification which ignores observed heterogeneity across users. We see that reducing the duration of the inactive stage (moving up the rows in the table) monotonically increases the clickthrough rates, from 2.7% in the benchmark non-skippable case (30 seconds) to 3.37% if the ad were made completely skippable, which is roughly a 20% increase in the clickthrough rate.

The columns marked "Heterogeneous Users" in Table 5 present counterfactual results using the estimates from the heterogeneous users specification, in Column (2) of Table 3. Since this specification of the model allows covariates to affect the parameters of the accumulation processes, we report the mean and median clickthrough probability for the 50 medoid users in our sample.²² Interestingly, in these results we see, contrary to the previous results, that the clickthrough rates *decrease* between the 30-second benchmark and the completely skippable (0 second) ad; this confirms the substantive importance of controlling for user heterogeneity in the empirical analysis. This decrease is nonmonotonic in the inactive duration for the reported average clickthrough rates, but monotonic for the median clickthrough rates.

The mean and median clickthrough rates for the heterogeneous users specification are aggregate statistics across different users. To see how *individual* users would hypothetically react to the counterfactual ad scenarios, we graph, in Figure 5, the counterfactual clickthrough rates across all treatments, for each of the ten medoid users introduced in Table 4 above. In this figure, the medoids are labeled such that Medoid 1 has the largest μ_1 and hence has the highest initial value during the active decision stage. Comparing this figure to the estimated drift terms μ_0 and μ_1 for each user medoid, as reported in Table 4, we see an interesting pattern that the sign of the inertial drift μ_0 , determines whether clickthrough probabilities increase or decrease

²¹Dukes, Liu & Shuai (2018) present a theoretical analysis of skippable ads.


²²That is, the median users in 50 clusters created by a k -means procedure.

Table 5: Counterfactual clickthrough probabilities.

<i>Treatments:</i> inactive duration (seconds)	Homogenous Users	Heterogeneous Users	
	Clickthrough probability	Clickthrough probability (mean)	Clickthrough probability (median)
0	0.0337	0.0307	0.0097
0.5	0.0347	0.0290	0.0137
1	0.0350	0.0282	0.0171
1.5	0.0352	0.0276	0.0198
2	0.0349	0.0276	0.0209
4	0.0335	0.0289	0.0271
10	0.0295	0.0331	0.0277
15	0.0281	0.0339	0.0285
20	0.0276	0.0340	0.0294
25	0.0275	0.0342	0.0299
30 (benchmark)	0.0272	0.0343	0.0299

In the case without heterogeneity, we use estimates from Column (1) of Table 3. In the case with heterogeneity, we use estimates from Column (2) of Table 3, and compute the counterfactual clickthrough probabilities of 50 medoids that represent the data. We then report the mean and median among the 50 medoids.

Figure 5: Counterfactual clickthrough probabilities for 10 medoid users



. These 10 medoid users correspond to the users in Table 4.

as the ad becomes less skippable: in other words, a positive μ_0 means that a longer inactive period will increase the clickthrough probabilities.

Authors have pointed out that internet advertising is oftentimes perceived as intrusive or “annoying” by users.²³ The results here show that, while some users, typified by Medoids 1-4, indeed are annoyed by ads, and are less inclined to clickthrough the longer they are exposed to the ad, a substantial number of users, typified by Medoids 5-10, are *not* annoyed by ads. In ongoing work, we are exploring how much ad annoyance (or not) depends on characteristics of the publisher or the advertiser.

These results also demonstrate that the substantial qualitative difference between the counterfactual results for the homogenous vs. heterogeneous users specification arises primarily from the difference in sign of μ_0 , the inertial stage drift, estimated from the two specifications.

²³See, for instance, Li, Edwards & Lee (2002) and McCoy, Everard, Polak & Galletta (2007). Both of these studies are primarily focused on early internet advertising (banner or pop-up ads), and do not consider video advertising on mobile devices.

μ_0 is large and negative in the homogeneous user specification, implying that users typically obtain negative signals from the ad; hence, allowing users to skip the ad can raise clickthrough rates, by essentially reducing the time that users are exposed to the pessimistic signals from the ad. In the heterogeneous users specification, as discussed before, the inertial drift μ_0 is actually *positive* for many users; these users receive positive signals while viewing the ad, and hence forcing them to view the ad as long as possible maximizes the clickthrough rates.

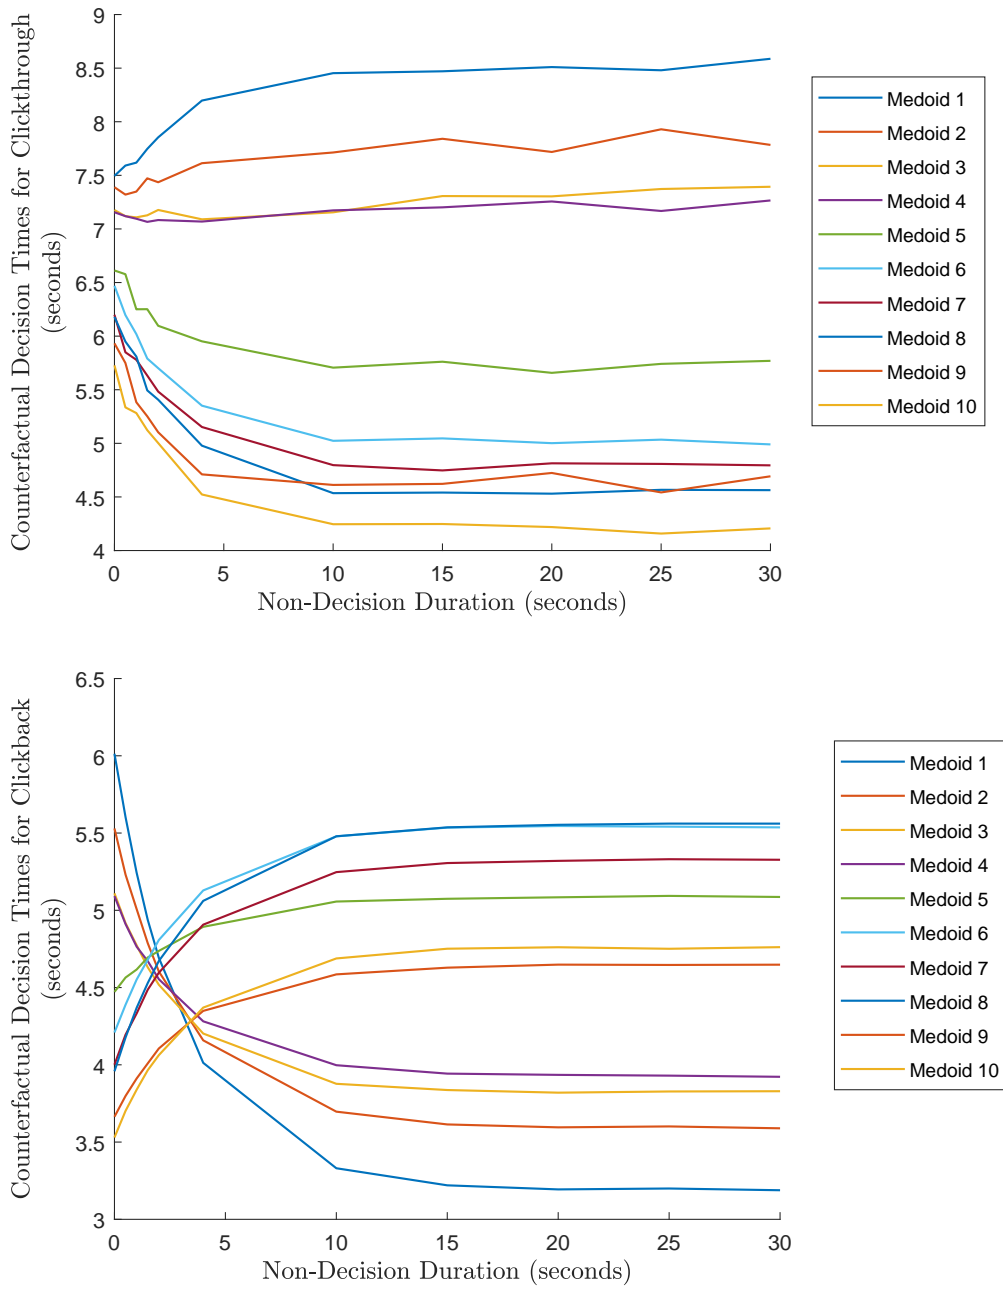
At the same time, what both sets of counterfactuals show is that the biggest changes in the clickthrough rate occur for inactive durations at 3 seconds or less – beyond this, clickthrough rates do not change much, even up to 30 seconds. That is, there is little difference between forcing viewers to watch only the first few seconds of a video ad, versus the entire 30-second ad, as the implied click-through probabilities do not change appreciably after 3 seconds. Mathematically, this arises from the convergent aspect of the variance of the accumulation during the inactive stage. In addition, it is consistent with an interpretation that users may only pay attention to the initial segment of an ad, and perhaps “switch off” after 3 seconds. In this respect, YouTube’s ad platform, which allows users to skip an ad after 5 seconds, seems reasonable.

Figure 6 presents the response times for the counterfactual scenarios, again for the 10 medoid users. For both clickthroughs and clickbacks, we see that most of the changes in response times occurs at small values for the inactive duration; for larger values, the response times are practically unaffected by the inactive duration. However, the relationship between response time and the sign of the inertial drift parameter μ_0 is different for clickthroughs vs. clickbacks. For clickthroughs, users with a negative μ_0 (such as medoid #1) have shorter response times for shorter inactive durations, while users with a positive inertial drift μ_0 (eg. medoid #10) have longer response times. These patterns are reversed for clickbacks: there, response times for users with negative μ_0 are decreasing in inactive durations, while response times for users with positive μ_0 are increasing in inactive duration.

5.1 Implications for mobile ad revenue

Previously, we see how varying the degree to which an ad is skippable would impact the ad’s clickthrough probability. Clearly, higher click-through probabilities directly benefit the publisher, and here we quantify the dollar magnitude of a higher clickthrough rate. In the mobile advertising market, publishers who run video ads on their app is paid whenever an ad eventually

Figure 6: Counterfactual Response Times: Clickthroughs (top) and Clickbacks (bottom)



leads a user of their app to install an advertised app. That is, the advertiser pays the publisher per user acquisition, unlike other advertising markets in which publishers are paid per click (as in banner ad markets) or per exposure (as in traditional media markets). In our dataset, we observe whether a click eventually leads to the user installing the app, and how much the advertiser pays to acquire that user, and here we use this information to compute estimates of counterfactual ad revenues from switching from a non-skippable to skippable ad format.

Across all users in our sample, the average install rate conditional on clickthroughs (“conversion rate”, in industry parlance) is 0.4789. The average revenue that the publisher earns for a given install is \$0.2794. Therefore for our sample size of $N = 194,510$, a one percentage point increase in clickthroughs leads to 1945.10 additional clicks; consequently, this generates $1945.10 \times 0.4789 = 931.51$ more installs, and $931.51 \times \$0.2794 = \260.26 additional revenue for the publisher. The heterogeneous user results in Table 5 imply that switching from a non-skippable to skippable format would increase click-through probabilities by around two percentage points, at the median, amounting to an extra \$520 (roughly) in ad revenue for this campaign.

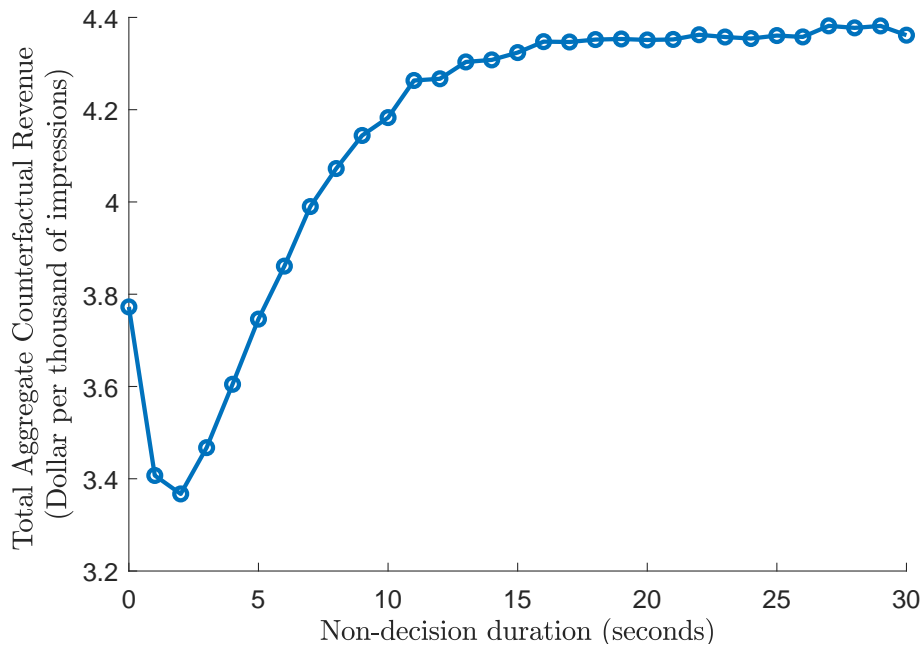
For a more in-depth revenue analysis, we consider the fact that users in our sample are heterogeneous, and advertisers pay different amounts to show their ad to different types of users. To incorporate this heterogeneity, we estimate the relationship between user’s characteristics and (i) the rate of installs conditional on clicks, (ii) the advertiser’s payment for that user conditional on installs, which is observed in our dataset. Given the large number of covariates and uncertain functional form of the relationship, we utilize non-parametric regressions with random forest in the estimation.

These regression estimates allow us to calculate the counterfactual revenue as a function of each user’s characteristics. We then calculate the total counterfactual revenues for each of the 50 medoid users, at different counterfactual points. Typically, revenue in the online ad industry is denominated in terms of CPM (dollars per thousand impressions). As such, we multiply the total counterfactual revenue by 20 to arrive at the final number.

We plot, in Figure 7, the aggregate counterfactual ad revenue at different values for the inactive duration, from 0 to 30 seconds. As the figure illustrates, the inactive duration has a non-monotonic effect on the publisher’s ad revenue. However, reflecting the counterfactual results reported in Table 5 for the heterogeneous users specification, ad revenue is maximized

for the non-skippable ad (30 seconds inactive duration), when the CPM is around \$4.40.²⁴

Figure 7: Total counterfactual revenue, aggregating the heterogeneity in users' characteristics.



6 Conclusions

We conclude with several remarks, mentioning several extensions to the analysis in this paper. First, we have assumed throughout that the choice threshold K is constant. Fudenberg et al. (2017) show that optimal decision rules in the DDM setting typically involve time-varying choice thresholds, and suggest reasonable functional forms for these thresholds which can be taken to the data.

Second, most clickstream datasets keep track of consumers' choices in multinomial choice settings, involving choice among more than two items. A common extension of the DDM model for multinomial choices is the "race" model, in which consumers face N competing accumulation processes, and the first process to hit a common "finish line" is chosen.²⁵ Formally, this model is

²⁴As a benchmark, we note that Instagram's average CPM for United States users was \$5.10 in 2016 Q1, according to Salesforce's Advertising Index Q1 2016 Report.

²⁵See Webb (2016) and Marley & Colonius (1992)

equivalent to the competing risk model from statistical duration analysis, and it will be interesting to explore these connections.

Third, given the results in this paper, one natural question is whether response time data can be useful for choice prediction. Such usefulness has already been demonstrated in choice experiments in the lab,²⁶ and it will be interesting to examine these benefits are also present in field data.

Finally, the use of structural estimation and modelling to address the effects of skippable vs. non-skippable ads is novel relative to existing methodologies for determining policy effects in online and mobile platforms, which typically involve randomized A/B testing.²⁷ These two approaches are complementary, as AB testing yields quick and convenient measures of causal effects while structural modelling allows us a glimpse into user behavior and preferences underlying these policy effects.

²⁶See Clithero (2018) and the cites therein.

²⁷Lewis & Rao (2015) discuss limitations of using randomized trials for evaluating online advertising campaigns.

References

- Allili, L., Patie, P., & Pedersen, J. L. (2005). Representations of the first hitting time density of an ornstein-uhlenbeck process. *Stochastic Models*, 21(4), 967–980.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, 113(4), 700.
- Busemeyer, J. R. & Townsend, J. T. (1992). Fundamental derivations from decision field theory. *Mathematical Social Sciences*, 23(3), 255–282.
- Chen, R. & Chiong, K. (2016). An empirical model of mobile advertising platforms. Working paper, UT Dallas.
- Clithero, J. A. (2018). Improving out-of-sample predictions using response times and a model of the decision process. *Journal of Economic Behavior & Organization*, 148, 344–375.
- Cox, D. & Oakes, D. (1984). *Analysis of Survival Data*, volume 21. CRC Press.
- Crawford, G. S. & Shum, M. (2005). Uncertainty and learning in pharmaceutical demand. *Econometrica*, 73(4), 1137–1173.
- Dukes, A. J., Liu, Q., & Shuai, J. (2018). Interactive advertising: The case of skippable ads. *SSRN*.
- Erdem, T. & Keane, M. P. (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing science*, 15(1), 1–20.
- Frydman, C. & Nave, G. (2016). Extrapolative beliefs in perceptual and economic decisions: Evidence of a common mechanism. *Management Science*, 63(7), 2340–2352.
- Fudenberg, D., Strack, P., & Strzalecki, T. (2017). Speed, accuracy, and the optimal timing of choices. *Available at SSRN*.
- Gold, J. I. & Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, 36(2), 299–308.
- Hong, H. & Shum, M. (2006). Using price distributions to estimate search costs. *The RAND Journal of Economics*, 37(2), 257–275.

- Honka, E. (2014). Quantifying search and switching costs in the us auto insurance industry. *The RAND Journal of Economics*, 45(4), 847–884.
- Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in psychology*, 3, 193.
- Lehmann, E. (1959). *Testing Statistical Hypotheses*. Wiley. first edition.
- Lewis, R. A. & Rao, J. M. (2015). The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics*, 130(4), 1941–1973.
- Li, H., Edwards, S. M., & Lee, J.-H. (2002). Measuring the intrusiveness of advertisements: Scale development and validation. *Journal of advertising*, 31(2), 37–47.
- Link, S. W. (1992). *The wave theory of difference and similarity*. Psychology Press.
- Lu, J. & Hutchinson, J. W. (2017). Split-second decisions during online information search: Static vs. dynamic decision thresholds for making the first click. Technical report, Wharton.
- Luce, R. D. (1991). *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford University Press.
- Marley, A. A. J. & Colonius, H. (1992). The “horse race” random utility model for choice probabilities and reaction times, and its compering risks interpretation. *Journal of Mathematical Psychology*, 36(1), 1–20.
- McCoy, S., Everard, A., Polak, P., & Galletta, D. F. (2007). The effects of online advertising. *Communications of the ACM*, 50(3), 84–88.
- Otter, T., Allenby, G. M., & Van Zandt, T. (2008). An integrated model of discrete choice and response time. *Journal of Marketing Research*, 45(5), 593–607.
- Pakes, A. (1986). Patents as options: Some estimates of the value of holding european patent stocks. *Econometrica*, 755–784.
- Ratcliff, R. & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, 20(4), 873–922.
- Ratcliff, R. & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9(5), 347–356.

- Seiler, S. & Pinna, F. (2017). Estimating search benefits from path-tracking data: measurement and determinants. *Marketing Science*, 36(4), 565–589.
- Shadlen, M. N., Hanks, T. D., Churchland, A. K., Kiani, R., & Yang, T. (2006). The speed and accuracy of a simple perceptual decision: a mathematical primer. In K. Doya, S. Ishii, A. Pouget, & R. Rao (Eds.), *Bayesian Brain Probabilistic Approaches to Neural Coding* (pp. 201–236). MIT Press.
- Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of mathematical psychology*, 44(3), 408–463.
- Ursu, R., Wang, Q., & Chintagunta, P. (2018). Search duration. Working paper, NYU Stern.
- Usher, M. & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological review*, 108(3), 550.
- Wald, A. (1973). *Sequential analysis*. John Wiley.
- Webb, R. (2016). The Dynamics of Stochastic Choice. *Management Science*, to appear.
- Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). Hddm: hierarchical bayesian estimation of the drift-diffusion model in python. *Frontiers in neuroinformatics*, 7, 14.
- Woodford, M. (2016). Optimal evidence accumulation and stochastic choice. Manuscript, Columbia University.

A Proof of Lemma 1

Proof. Derivation from [Shadlen et al. \(2006\)](#) shows that the simple DDM $dZ(t) = \mu dt + \sigma dB(t)$ implies:

$$\bar{y} = \mathbb{E}[y_i] = \frac{1}{e^{-2B\mu} + 1} \quad (9)$$

$$\bar{t} = \mathbb{E}[t_i^*] = \frac{B(e^{B\mu} - e^{-B\mu})}{\mu(e^{-B\mu} + e^{B\mu})} \quad (10)$$

When we solve for these equations in terms of μ and B , there are only two real solutions. Among the two solutions, one of the solutions is such that $\hat{\mu}$ is increasing in \bar{y} , and the other solution is such that $\hat{\mu}$ is decreasing in \bar{y} . Therefore the first solution is the valid one. \square

B How good is the discrete approximation?

Consider approximating the stochastic process $dZ_t = \mu dt + \gamma Z_t dt + \sigma dW_t$ via $Z_{t_j} - Z_{t_{j-1}} = \mu \Delta t + \gamma Z_{t_{j-1}} \Delta t + \sigma \Delta W_{t_j}$, where $t_j - t_{j-1} = \Delta t$ for all $j = 1, \dots$. Letting $J = \lfloor \frac{t}{\Delta t} \rfloor$ denote the number of discrete time intervals at time t . The discretized process at time t has the following distribution:

$$Z_t \sim \mathcal{N} \left(\mu \Delta t \sum_{i=0}^{J-1} (1 + \gamma \Delta t)^i, \Delta t \sum_{i=0}^{J-1} (1 + \gamma \Delta t)^{2i} \right) \quad (11)$$

or, equivalently:

$$Z_t \sim \mathcal{N} \left(\mu \frac{1 - (1 + \gamma \Delta t)^J}{-\gamma}, \Delta t \frac{1 - (1 + \gamma \Delta t)^{2J}}{1 - (1 + \gamma \Delta t)^2} \right). \quad (12)$$

As $\Delta t \rightarrow 0$, we check that $\lim_{\Delta t \rightarrow 0} \mu \frac{1 - (1 + \gamma \Delta t)^{\frac{t}{\Delta t}}}{-\gamma} = \frac{\mu}{\gamma} (e^{\gamma t} - 1)$ and $\lim_{\Delta t \rightarrow 0} \Delta t \frac{1 - (1 + \gamma \Delta t)^{\frac{2t}{\Delta t}}}{1 - (1 + \gamma \Delta t)^2} = \frac{1}{2\gamma} (e^{2\gamma t} - 1)$. Suppose that at $t = 30$, $\text{Var}(Z_{30}) = 2.0833$. This number is implied by the estimation result (Column 2 of [Table 3](#)). Using the time-increment $\Delta t = 0.25$, the resulting γ_0 is -0.247 . Now as $\Delta t \rightarrow 0$, we have $\gamma_0 \rightarrow -0.240$. Therefore, the time-increment of a quarter-second appears to do a reasonable job at approximating the continuous-time density.