

Online Display Advertising: Modeling the Effects of Multiple Creatives and Individual Impression Histories

Michael Braun

Cox School of Business, Southern Methodist University, Dallas, Texas 75275, braunm@smu.edu

Wendy W. Moe

Robert H. Smith School of Business, University of Maryland, College Park, Maryland 20742, wmoe@rhsmith.umd.edu

Online advertising campaigns often consist of multiple ads, each with different creative content. We consider how various creatives in a campaign differentially affect behavior given the targeted individual's ad impression history, as characterized by the timing and mix of previously seen ad creatives. Specifically, we examine the impact that each ad impression has on visiting and conversion behavior at the advertised brand's website. We accommodate both observed and unobserved individual heterogeneity and take into account correlations among the rates of ad impressions, website visits, and conversions. We also allow for the accumulation and decay of advertising effects, as well as ad wearout and restoration effects. Our results highlight the importance of accommodating both the existence of multiple ad creatives in an ad campaign and the impact of an individual's ad impression history. Simulation results suggest that online advertisers can increase the number of website visits and conversions by varying the creative content shown to an individual according to that person's history of previous ad impressions. For our data, we show a 12.7% increase in the expected number of visits and a 13.8% increase in the expected number of conversions.

Key words: online advertising; advertising response modeling; online visit and conversion rates; Bayesian models; targeting

History: Received: December 26, 2011; accepted: June 9, 2013; Preyas Desai served as the editor-in-chief and Puneet Manchanda served as associate editor for this article.

Introduction

During the past decade, online advertising budgets have grown steadily, often at the expense of off-line advertising budgets (Interactive Advertising Bureau 2011). Many marketers now prefer to advertise online because the interactive medium allows for the precise targeting of individual consumers. In this paper, we focus on online display advertising. Online targeting strategies have improved display ad performance (Interactive Advertising Bureau 2011); they have done so by using fairly straightforward targeting policies. For example, online ad networks examine individual clickstream histories to identify customers with an interest in a specific product category or brand. Interest is typically determined by whether the individual has previously visited Web pages related to that product or brand. At the next advertising opportunity, the ad network then exposes that individual to an advertisement that matches his or her interests (e.g., computer ads would be shown to individuals who have recently visited computer review websites, car ads would be shown to those who have recently visited car websites).

Current online targeting practices have achieved impressive results by leveraging the online data associated with an individual's history of page views. However, little consideration has been given to an individual's history of ad impressions. That is, online advertisers typically treat an individual who has repeatedly seen a given ad in the same way as an individual who has only occasionally seen it, but in all likelihood, the former will be less responsive to the *next* exposure than the latter would be (Chatterjee et al. 2003). This dynamic affects the decision of not only which product category to feature in the next ad exposure but also which of the creatives (versions of an ad) in the advertiser's advertising portfolio to serve.

In this paper, we explicitly model the effects of a given ad impression in the context of an individual's impression history. Additionally, because many online ad campaigns include multiple advertising creatives, we also allow the effect of each impression to vary depending on the creative content associated with that impression. We then examine the individual's response in terms of future visiting and conversion

behavior at the advertiser's website. This approach allows us to identify which creatives are most effective at increasing website visit rates and conversion probabilities. Although advertising agencies often pretest the effectiveness of different ad creatives, either in a controlled lab setting or in structured field experiments often referred to as A/B testing, once a campaign is launched, it can still be a challenge to attribute the gains from the advertising campaign to specific creatives within the campaign. Our research addresses this challenge.

To this end, we employ a hierarchical Bayesian approach that models an individual's ad impressions, visits to the advertiser's website (as opposed to the website on which the ad is placed), and conversion behavior at the website as observed outputs from three separate, but correlated, processes.¹ Similar to the long history of research on off-line advertising effectiveness, we incorporate an advertising goodwill construct that allows (1) advertising goodwill to accumulate and decay, (2) ad impressions to contribute differentially to goodwill according to differences in creative content, and (3) ad wearout and restoration effects when individuals are repeatedly exposed to an ad campaign.

The implications of our research for online advertisers are significant. First, we demonstrate the importance of accommodating differential ad effects across creatives in a given ad campaign. Second, we show how an individual's unique history of ad impressions can affect how he or she responds to subsequent ad exposures. In other words, our research highlights the opportunity for advertisers to further refine their ad targeting policies by considering both creative-specific effects and individual impression histories.

We demonstrate the managerial implications of our research by simulating a campaign in which the creative content that is shown to an individual is determined in part by that individual's ad impression history (note that the advertiser does not currently employ this practice). We compare the expected number of visits and conversions resulting from this simulated campaign to the outcome of the actual campaign observed in the data. In this context, we show how the advertiser, by customizing the creative content of ads based on individual impression histories, can achieve a 12.7% increase in the expected number of visits and a 13.8% increase in the expected number of conversions, compared to the policy that was actually employed.

The rest of this paper proceeds as follows. In the next section, we review advertising research, both

online and off-line, with a focus on the elements of previous research that should be incorporated into any model of online advertising response. From there, we present the specification of the model, a description of the data, and model results. After discussing our empirical results, we present a simulation that underscores the potential benefits of modeling both creative-specific effects and individual impression histories for online ad targeting.

Advertising Research Review

Off-line Advertising

Research on traditional "off-line" advertising campaigns has provided convincing empirical evidence that advertisements have both short-term and long-term effects on behavior. In a large-scale field experiment, Lodish et al. (1995) showed that a temporary increase in advertising expenditures can result in sustained sales benefits that extend beyond the advertising period; in some cases, elevated sales were observed two years into the future. In a meta-analysis, Tellis (2009) also documented long-term advertising effects. Specifically, his study concluded that advertising carryover effects were twice as large as any contemporaneous effects. These studies emphasize the importance of allowing for both short-term and long-term advertising effects by considering the impact of advertising on current and future behavior.

To accommodate both the long-term and short-term effects of advertising, Nerlove and Arrow (1962) proposed an advertising response model that incorporates a construct that they referred to as *goodwill*. In their model, goodwill accumulates with advertising expenditures but also decays from period to period. This dynamic of accumulation and decay has also been applied to the effect of advertising on awareness (Mahajan and Muller 1986). Specifically, aggregate brand awareness increases with advertising exposure but decreases as a result of consumers "forgetting" in the absence of advertising.

Other research (Mahajan and Muller 1986, Naik et al. 1998) has shown that the effect of any single advertisement is dynamic, and its impact is subject to both wearout and restoration effects.² *Wearout* refers to the decreased effectiveness of advertising copy over time. Naik et al. (1998) differentiated between repetition wearout and copy wearout. Repetition wearout results from repeated exposure to the ad, whereas copy wearout results from the passage of time. When there is a hiatus in advertising, restoration effects can reverse the repetition wearout

¹ Each process has a baseline rate, and the visit and conversion processes depend on a scaled advertising effect. The intercepts and coefficients are considered random effects whose distributions are modeled as priors in the Bayesian hierarchical model.

² *Wear-in* effects, distinct from wearout and restoration effects, refer to increasing ad effectiveness as the number of ad exposures increase. These effects have been shown to be negligible (Tellis 2009).

effects. During the hiatus, any degradation in ad effectiveness resulting from repetition wearout is restored gradually. Mahajan and Muller (1986) showed that these countering effects (wearout versus restoration) can have significant implications for ad scheduling decisions.

A variety of different approaches have been proposed to capture the aforementioned advertising effects (e.g., Little 1979; Feinberg 1992, 2001; Vakratsas et al. 2004). However, most studies, including those mentioned above, have centered on aggregate response models for which the dependent variable is total sales per time (and possibly per market). Although these aggregate response models can be valuable in planning off-line media schedules at a market level, they are unable to address the online advertising challenge of measuring advertising effects when each individual is exposed to a unique ad schedule. Additionally, most advertising models have focused on examining the response to total advertising expenditures, despite the fact that many ad campaigns consist of a variety of different creative copy (different “versions” of the ad within a campaign). To our knowledge, there are no extant models that examine the effectiveness, or associated dynamics, of a specific advertising creative within a campaign.³

Online Advertising

Despite the shared interest in advertising, online advertising research shares few similarities with its off-line counterpart. Whereas advertisers in the off-line environment tend to focus on long-term effects and brand building, online advertisers diligently (and almost exclusively) monitor performance metrics that can be observed immediately. Metrics such as click-through rates and conversion rates indicate an individual’s immediate reaction to an advertising impression but ignore any indirect effects that may result in changes in an individual’s future behavior. Highlighting the limitations of these metrics, Drèze and Hussherr (2003) showed that despite the fact that many individuals do not consciously attend to online ads and therefore do not click through on them, such ads still have a positive effect on brand measures, which, in theory, translates to increased future sales.

Manchanda et al. (2006) explored the delayed effects of banner ads on individual purchasing behavior. Rather than focusing on click-through or conversion rates, they showed that increased exposure to

banner ads shortens an individual’s repeat purchasing rate,⁴ providing empirical evidence of the indirect effects of online advertising.

Another complicating factor in the online advertising environment is the potential correlation between an individual’s online behavior and advertising exposure. For example, a highly active browser will probably see more ads simply because he visits more sites (including the advertiser’s site). Additionally, some online advertisers employ targeting policies that also lead to correlations between advertising exposure and behavior at the advertiser’s site. For example, advertisers may target specific websites on which to place their ads based on the types of visitors the site attracts (e.g., an automobile brand will target potential car purchasers by advertising on automobile-related sites). As a result, a car shopper is both more likely to be exposed to an ad and more likely to visit the advertiser’s website.⁵ In such cases, researchers must take care to separate the causal effect of ad impressions on behavior from simple correlated effects.

Model Development

To address these issues and challenges, we construct an individual-level advertising response model that allows for (1) creative-specific advertising effects, (2) wearout and restoration effects on ad copy, (3) non-immediate and sustained effects of advertising on behavior, and (4) correlation between ad impression rates and site visit and conversion behavior. We examine these effects with respect to both visiting and conversion behavior. Additionally, we incorporate a goodwill construct, similar to that used by Naik et al. (1998) and Nerlove and Arrow (1962), that allows for both the accumulation and decay of advertising effects over time. Rutz and Bucklin (2011) used a similar construct that they call “Ad Stock” in their model of paid search advertising.

Ad Stock and Ad Effect

First, we define E_{it} as the contemporaneous effect of all of the impressions of ads that were presented to individual i in week t . We define A_{it} as the accumulated Ad Stock for user i at time t , such that

$$A_{it} = \alpha A_{i,t-1} + E_{it}. \quad (1)$$

⁴ Specifically, Manchanda et al. (2006) modeled all repeat purchases irrespective of whether they resulted from a direct ad click-through or from manually entering the URL for the online retailer at a later time after being exposed to the ad.

⁵ Some online advertising networks will employ more sophisticated (and costly) “behavioral targeting” practices, in which an individual is exposed to ads based on more detailed browsing histories across multiple websites. This kind of targeting could also result in similar correlations. The advertiser represented in our data set did not employ these kinds of behavioral targeting techniques for this campaign.

³ We differentiate such attribution models that credit observed behaviors of interest to specific ad impressions from the practice of pretesting ads. In pretests, ads are evaluated in terms of how they affect attitudes, brand associations, etc., in a controlled testing environment. Our focus is on the advertisement’s effect on observable behaviors after the campaign has been launched in the market.

This formulation allows the overall advertising effect A_{it} to carry over into future periods subject to geometric decay with rate α . Ad Stock accumulates over time as individual i is exposed to additional advertising impressions. Modeling advertising effects through the Ad Stock specification above allows us to capture both direct and indirect effects of advertising. Note that if advertising has only a contemporaneous effect (i.e., there are no long-terms effects), then we would expect $\alpha = 0$.

Let E_{ijt} be the ad effect for an impression of creative j on individual i at time t , and let E_{it} be the aggregated effects across creatives. Thus E_{it} varies across individuals and evolves over time, based on which creatives are served to that individual as well as the individual’s previous impression history. We identify four types of effects that contribute to E_{it} :

1. An advertising campaign effect AD where exposure to any ad in the campaign contributes to Ad Stock
2. For each creative j , a marginal effect C_j that captures how some creatives may be more effective than others
3. A repetition wearout effect, denoted as δ , that allows for diminishing marginal effects of each exposure to a particular creative ($0 < \delta < 1$)
4. Restoration effects, denoted as R , that allow for the mitigation of wearout effects as time elapses since the last ad impression ($R > 1$)

We consider two types of wearout and restoration effects: (1) wearout/restoration associated with repeated exposures to any ad in the campaign and (2) wearout/restoration associated with exposures to a specific ad creative. Let x_{it} be the cumulative number of impressions of any ad in the campaign seen by individual i at time t , and let y_{ijt} be the cumulative number of impressions of creative j seen by individual i by time t . Our expression for the contemporaneous addition to Ad Stock is

$$E_{it} = AD[1 - (1 - \delta_1^{x_{it}}) + R_{1it}(1 - \delta_1^{x_{it}})] + \sum_j C_j[1 - (1 - \delta_2^{y_{ijt}}) + R_{2ijt}(1 - \delta_2^{y_{ijt}})]. \quad (2)$$

Overall, the effectiveness of each impression is determined by the baseline advertising effect, AD , less any campaign-level wearout ($1 - \delta_1^{x_{it}}$) plus any restoration of that wearout ($R_{1it}(1 - \delta_1^{x_{it}})$) and the effect of the creative itself, C_j , less any creative-specific wearout ($1 - \delta_2^{y_{ijt}}$) plus any restoration of that wearout ($R_{2ijt}(1 - \delta_2^{y_{ijt}})$).

Ad wearout effects are represented by the expressions $(1 - \delta_1^{x_{it}})$ and $(1 - \delta_2^{y_{ijt}})$, where the δ parameters represent the proportion of the ad’s effectiveness that is retained with each repeat impression.

Restoration effects are represented by the expressions $R_{1it}(1 - \delta_1^{x_{it}})$ and $R_{2ijt}(1 - \delta_2^{y_{ijt}})$, where the rates of restoration are captured by R_{1it} and R_{2ijt} . That is, R represents the percentage of ad wearout that is restored each week since the last ad impression. Since the restoration effect must be positive and cannot exceed the wearout, R_{1it} and R_{2ijt} must be between 0 and 1. Thus we specify R_{1it} and R_{2ijt} as follows:

$$R_{1it} = \frac{\rho_1 \tau_{1i}}{1 + \rho_1 \tau_{1i}} \quad \text{and} \quad R_{2ijt} = \frac{\rho_2 \tau_{2ij}}{1 + \rho_2 \tau_{2ij}}, \quad (3)$$

where ρ_1 and ρ_2 are nonnegative restoration parameters to be estimated, τ_{1i} is the time that has elapsed since person i was last exposed to any ad in the campaign, and τ_{2ij} is the time that has elapsed since person i was last exposed to ad creative j . (τ is equal to zero if there is no previous ad exposure or if time t is the time of the first exposure).

Impression, Visit, and Conversion Models

Our Bayesian hierarchical model is based on three separate but related processes for the arrival of ad impressions, visiting, and conversion behavior. Each individual has a baseline rate at which he is exposed to ads from the advertiser, a baseline rate at which he visits the advertiser’s website, and a baseline probability that a visitor “converts” (or engages in a target behavior). As the Ad Stock variable A_{it} evolves from period to period, it shifts the baseline visit rates and conversion probabilities.

We incorporate unobserved heterogeneity into our model by allowing the baseline rates and probabilities, as well as sensitivities to Ad Stock, to vary randomly across the population. The population-level mixing distribution is a joint one, so these heterogeneous parameters can be correlated across users. From our conversations with the advertiser who provided the data, we know that ads were placed on specific websites, but they were not targeted at specific individuals (we will discuss the details of the ad campaign later in the Data section). That is, individuals were not selectively exposed to ad impressions based on their potential interest in the advertiser or any other behavioral constructs. Of particular relevance for our analysis is that there was no matching at all of individuals to a specific creative, and thus we treat the specific schedule of creatives as exogenous.

Table 1 provides an overview of our model specification. We define m_{it} as the number of ad impressions

Table 1 Overview of Model Specification

Observed event	Symbol	Individual model
Impressions	m_{it}	Zero-inflated (r_m) Poisson (λ_i)
Visits	v_{it}	Zero-inflated (r_v) Poisson (μ_{it})
Conversions	s_{it}	Zero-inflated (r_s) binomial (p_{it}, v_{it})

seen by i in time t , v_{it} as the number of visits by i in time t to the advertiser’s website, and s_{it} as the number of “converted” visits. Many online marketers equate a converted visit with purchase, such as a purchase visit at an online bookstore. However, different companies or industries might define conversion in other ways, and for many product categories, purchasing simply does not occur online. Still, customers might engage in some kind of action that supports the purchasing process. For example, in many legal jurisdictions, new car purchases cannot be completed online. Instead, manufacturers’ websites encourage visitors to submit contact information so that a sales person can call or email them to facilitate a sale. From the firm’s point of view, these behaviors constitute successful website interactions, so they, and we, refer to them as conversions.

At the individual and week levels, we model the outcome parameters m_{it} , v_{it} , and s_{it} as realizations of two zero-inflated Poisson distributions (for impressions and visits) and a zero-inflated binomial distribution (for conversions). The zero-inflation allows for additional probability mass at zero, so we can accommodate a larger number of individuals who receive zero impressions, make zero visits the site, or convert zero times at the site than the Poisson or binomial would normally predict. When modeling purchasing behavior, previous researchers have interpreted this zero-inflation parameter as a hardcore never-buyer construct (Morrison and Schmittlein 1981). We could similarly interpret the complements of our zero-inflation parameters ($1 - r_m$, $1 - r_v$, and $1 - r_s$) as individuals who will never receive an ad impression, never visit the website, or never convert.

Specifically, if r_m is the probability that m_{it} could possibly be greater than zero for any t , and if λ_i is a heterogeneous rate parameter, then the conditional data likelihood is

$$f(m_{it}) = (1 - r_m)\mathbf{I}\left(\sum_{t=1}^T m_{it} = 0\right) + r_m \frac{e^{-\lambda_i} \lambda_i^{m_{it}}}{m_{it}!}. \quad (4)$$

Similarly, the data likelihood for visits is

$$f(v_{it}) = (1 - r_v)\mathbf{I}\left(\sum_{t=1}^T v_{it} = 0\right) + r_v \frac{e^{-\mu_{it}} \mu_{it}^{v_{it}}}{v_{it}!}, \quad (5)$$

where r_v is the probability that v_{it} could be greater than zero for any t and μ_{it} is a heterogeneous, time-varying visit rate. Finally, the probability of having the s_{it} successes out of the v_{it} visits is

$$f(s_{it}) = (1 - r_s)\mathbf{I}\left(\sum_{t=1}^T s_{it} = 0\right) + r_s \binom{v_{it}}{s_{it}} p_{it}^{s_{it}} (1 - p_{it})^{v_{it} - s_{it}}, \quad (6)$$

where r_s is the probability that $s_{it} > 0$ for all t and p_{it} is a heterogeneous, time-varying probability that a particular visit is a success.

One feature of most online data sets that are available to advertisers is that they include data from only those individuals who either received an ad or interacted with the company in some way. In other words, the data are *truncated*. To be included in our data set (and most other data sets available to online advertisers), an individual must be exposed to an ad, visit the website, or convert on the website. However, all conversions are associated with a website visit; thus an individual is included in the data if at least one of only two criteria are met: (1) the individual was exposed to an ad or (2) the individual visited the website. Put another way, if the individual meets at least one of these two criteria, the probability of being included in the database is 1, regardless of the number of conversions. We account for this data truncation by specifying a conditional likelihood (Gelman et al. 2003, Greene 2008, Winkelmann 2008).⁶ Specifically, the data likelihood contribution for a single user in our data set is the joint probability of the observed impression, visit, and conversion data across all observed time periods and conditional on model parameters, conditional on the probability that the individual is included in the data set:

$$l_i = \frac{\prod_t f(m_{it} | \lambda_{it}) f(v_{it} | \mu_{it}) f(s_{it} | v_{it}, p_{it})}{1 - \Pr(\sum_t m_{it} = 0, \sum_t v_{it} = 0 | \lambda_i, \mu_i)}. \quad (7)$$

To measure the effects of advertising on behavior, we model the visit rate μ_{it} and the conversion probability p_{it} as a function of accumulated Ad Stock, A_{it} . Ad Stock varies over time according to individual i ’s schedule of ad impressions up to time t . Additionally, we allow for time-varying effects in the frequency of ad impressions, site visits, and conversion behavior as follows:

$$\log \lambda_{it} = \log \lambda_{0i} + \gamma_\lambda X_t, \quad (8)$$

$$\log \mu_{it} = \log \mu_{0i} + \beta_{\mu i} A_{it} + \gamma_\mu X_t, \quad (9)$$

$$\text{logit } p_{it} = \text{logit } p_{0i} + \beta_{p i} A_{it} + \gamma_p X_t, \quad (10)$$

where X_t is a vector of weekly indicator variables to control for potential time-varying fixed effects and γ_λ , γ_μ , and γ_p are vectors of coefficients. These time-varying covariates allow us to control for any national off-line advertising or promotional campaigns or other unobserved variables that would affect all individuals in that week equally. Additionally, these weekly dummies also help control for any endogeneity resulting from time-varying unobserved variables

⁶ We assume that the truncated individuals have the same parameter distribution as the included individuals.

that affect all three processes.⁷ The weekly dummies can control for the effects of any national off-line advertising or promotional campaign, but they do not capture any potential interaction effects between off-line and online campaigns. After controlling for these covariates effects, any remaining variance can then be attributed to differences across individuals in terms of their online advertising exposures. The coefficient $\beta_{\mu i}$ represents the sensitivity of person i 's visit rate to advertising, and $\beta_{p i}$ captures the sensitivity of the conversion probability to advertising.

Moving up the hierarchical structure of the model, we let the baseline impression and visit rates, the baseline success probabilities, and the sensitivities to Ad Stock be heterogeneous and correlated. Specifically, we assume that all parameters (or appropriate transformations thereof) follow a multivariate normal distribution:

$$\log \lambda_i, \log \mu_{0i}, \text{logit } p_{0i}, \beta_{\mu i}, \beta_{p i} \sim \text{MVN}(\phi, \Sigma). \quad (11)$$

By specifying individual-level parameters in this way, we allow for unobserved heterogeneity. The elements of ϕ correspond to the mean log impression rate, log visit rate, logit conversion probability, marginal effect of Ad Stock on visit rate, and marginal effect of Ad Stock on conversion probability. Adding unobserved heterogeneity to the model allows for some individuals to be exposed to ads more frequently than others, for some to have a greater propensity to visit the site than others, for some to be more likely to convert than others, and for some to be more sensitive to the advertising effects than others.

By modeling the distribution of these latent parameters jointly, we control for correlation between an individual's latent rate of exposure to ads and his visit rates and conversion probabilities. For example, an individual who is an active online car shopper is both more likely to visit a car manufacturer's website and more likely to be exposed to a car ad (assuming that car advertisers place their ads on Web pages frequented by car shoppers). By explicitly capturing this potential correlation between ad impressions and visiting behavior through the multivariate normal distribution in Equation (11), we can then separate it from the effects of advertising on visits and conversions (Heckman 1979, Ying et al. 2006, Moe and Schweidel 2012).

The effects of advertising on individual behavior are revealed in posterior estimates of C_j , $\beta_{\mu i}$, and $\beta_{p i}$. For the visiting process, the element of the prior mean on ϕ that corresponds to $\beta_{\mu i}$ is constrained to be 1. Otherwise, multiplying all C_j by a constant would be equivalent to multiplying the mean of $\beta_{\mu i}$ by

that same constant. Thus, the creative-specific effects, C_j , represent the effects of impression histories on Ad Stock. For the conversion process, $\beta_{p i}$ represents the effect of advertising on the probability of conversion.

Note that there are two ways in which advertising can affect sales: increasing the rate at which individuals visit the advertiser's website and increasing the chance of conversion. It is not immediately clear what the signs of $\beta_{p i}$ should be. If $\beta_{\mu i} > 0$ and $\beta_{p i} = 0$, then advertising drives more customers to the site, but it has no effect on whether the customer converts. One could argue that individuals are less likely to convert on advertising-induced visits because those visitors were not intrinsically motivated to visit the site on their own. In this case, conversion rates would be lower for the advertising-induced visitors, and $\beta_{p i}$ would be less than 0. Alternatively, conversion rates could be greater for advertising-induced visits (compared with organic visits) if the ad copy effectively stimulated a purchasing need. A $\beta_{p i}$ greater than 0 would indicate such an effect. In fact, a possible extension to our model (with a richer data set) might be one that allows for a particular creative to have a different effect on the visit rate than it does on the conversion probability.

To complete the model specification, we place a weakly informative normal prior on a single vector that contains C_1, \dots, C_J , logit α , logit δ_1 , logit δ_2 , log ρ_1 , log ρ_2 , and the unconstrained elements of ϕ , as well as an inverse-Wishart prior on Σ . We then sample from the joint posterior distribution of the parameters of interest, conditional on the observed data and prior parameters, to collect an estimate of marginal posterior means and quantiles.

Data

Our data were provided by Organic, an online advertising agency that, as part of its services to clients, manages client websites and purchases online ad exposures. Our data set pertains to an advertising campaign run by a single automobile brand over a course of 10 weeks from June 15 to August 23, 2009. Specifically, ads featured either the parent brand or one of two models offered by the auto manufacturer. For our analysis, we will use 9 weeks to estimate the model and use the 10th week for holdout validation (and as a baseline to assess potential benefits of using the model).

Our data set describes the activities of 5,803 individuals who were randomly selected from a database collected and maintained by Organic. The complete database includes each and every individual who has seen an ad, visited the website, or converted at the website⁸ at least once during the observation period.

⁷We thank one of our reviewers for pointing this out.

⁸All conversions are associated with a visit. In other words, it is not possible to convert without visiting the website.

As such, it includes individuals who have never seen an ad but have visited the website (and perhaps converted) as well as those who were exposed to the ad but never visited the website (and thus could not convert).

In each week, we observe three types of data at the individual and weekly levels:

1. The number of impressions of each creative that were served
2. The number of browsing sessions that include at least one visit to the client website
3. The number of sessions that include a conversion behavior

One limitation of considering these events at the weekly level involves the temporal sequencing of ad impressions and visits within the week. That is, for ads to have a causal effect on visits, the ad must precede the visit. In our data, for weeks where an individual both sees an ad and visits the website, the ad precedes the visit in 82.4% of the observations. In the remaining cases, the ad may or may not precede future visits in subsequent weeks.

Table 2 describes the distribution of impressions, visits, and conversions across users in the data. Note that there are a number of individuals who were not exposed to any advertising. Although we do not explicitly designate these individuals as a “control” group, their presence in the data allows us to establish a baseline for visits and conversion behavior when estimating the effects of advertising.

During this campaign, the advertiser employed 15 unique banner ad creatives designed to promote awareness for the United States government’s “Cash for Clunkers” program. In the summer of 2009, the U.S. government established this program to provide consumers with rebates if they traded in their old vehicles for newer, more fuel-efficient vehicles. The advertiser in our data ran ads, both online and off-line, to promote the fact that their vehicles would qualify for the this rebate. Each ad creative shared the same strategic objective in that the ads included calls to action, but the specific content (e.g., actors, script,

setting) differed across each of the 15 creatives. Unfortunately, we do not have details describing each ad creative, and Organic was unable to provide them.

Table 3 summarizes the total number of impressions served of each creative and the number of unique users that each creative reached. Organic purchased banner ad space directly from websites, ranging from those devoted to cars (e.g., Edmunds.com) to general interest websites (e.g., MSN) to social networking sites (e.g., Facebook). Organic told us that they did not employ any behavioral targeting practices for this campaign. Empirical analysis of the data confirms that the distribution of ad creatives did not systematically vary across individuals depending on their frequency of visit or conversion behavior (see Figure 1). The process was similar to a traditional off-line “ad buy” in which ads are placed in specific newspapers, television programs, or other media outlets. Seven of the creatives in the campaign accounted for only 46 of the 23,205 total impressions in the data, so we pool them together and treat them as a single creative.

In our data, 3,327 individuals visited the advertiser’s website at least once, for a total of 4,631 visits. Of these visitors, 1,399 (42% of all visitors) visited the site even though they did not receive any ad impressions. We do not differentiate between visits generated directly through a click-through, search engine result, or direct URL entry. Although an ad click-through reflects the direct effect of advertising, there may also be indirect effects from when an individual who was previously exposed to an ad subsequently visits the website by manually entering a URL or by using a search engine. While this is not a direct click-through on the ad, the visit can still (at least in part) be attributed to the advertising, and restricting our focus solely on ad click-through would underestimate the ad effects. Thus, similar to Manchanda et al. (2006), we treat all visits equally regardless of whether they were initiated by an ad click-through or by manually entering a URL. In our data set, only 7.6% of all visits came from the consumer clicking through on an ad.


Of the 4,631 visits to the seller’s website, 1,828 ended in what the advertiser defines as a conversion

Table 2 Observed Impression, Visit, and Conversion Counts

Observed event	Impressions	Visits	Conversions
Number in data	23,205	4,631	1,828
Average per individual user	3.999	0.798	0.315
Number of individual users with . . .			
0 impressions/visits/conversions	1,399	2,476	4,497
1 impression/visit/conversion	1,612	2,744	1,059
2 impressions/visits/conversions	806	342	143
3 impressions/visits/conversions	450	109	46
4 impressions/visits/conversions	328	45	22
5 impressions/visits/conversions	199	29	15
> 5 impressions/visits/conversions	1,009	58	21

Table 3 Observed Creative Counts

Ad creative	No. of impressions	Unique users reached
Creative A	692	149
Creative B	13,078	3,595
Creative C	4,809	1,689
Creative D	3,667	1,382
Creative E	325	45
Creative F	135	29
Creative G	327	66
Creative H	126	85
Creatives I–O	46	24

Figure 1 Distribution of Ad Creatives by Individual Visits (Left) and Conversion Behaviors (Right)

behavior. For many product categories, a conversion is defined as a sale. In the automotive categories, final sales rarely occur online. Instead, new car sites focus on other conversion behaviors, such as building and pricing a car, getting a quote, finding a dealer, and searching inventory. In our analysis, we use the advertiser's definition of a conversion behavior. Because this comprises a fairly large set of activities, the overall conversion rate in our data (39%) is higher than the typical conversion rate at most retail websites. Typically, online retailers experience less than a 10% conversion rate. Whereas the conversion rate in our data is higher than the typical purchase conversion rate at retail websites (as a result of the liberal definition of conversion), it is simply a mean shift. We are less interested in the magnitude of the conversion rate and more interested in how advertising can increase or decrease visit and conversion behaviors. In our case, the overall conversion rate of 39% is noticeably greater than the 32% conversion rate among those visitors who were not exposed to any ad impressions. This difference suggests that there is an advertising effect that we capture with our model. As an additional point of comparison, the conversion rate for click-through visits is 33.7% compared with a conversion rate of 39.8% for all other visits, highlighting the importance of including non-click-through visits in our analysis.

Results

Model Fit and Benchmark Comparisons

Before we discuss our empirical results and applications of the model, it is important to determine

whether the model is a sufficiently good representation of the data and how model fit compares to a number of benchmarks.

To summarize, our proposed model allows for (1) advertising goodwill effects, (2) creative-specific ad effects, and (3) advertising wearout and restoration effects. Therefore, we compare our proposed model to benchmark models that vary along each of these dimensions. Table 4 provides an overview of the comparison models.

Model I is a baseline model that assumes no advertising effects and includes only weekly indicator variables to capture changes in behavior over time. Model II allows for advertising effects (through a goodwill construct with wearout and restoration) but assumes only campaign-level ad effects. That is, all ad impressions, regardless of creative content, are assumed to have the same impact on behavior. This model is the one that is most similar to existing advertising response models that do not accommodate creative-specific variation in ad effectiveness. Model III allows for creative-specific ad effects, but the effects do not persist from period to period (i.e., no goodwill Ad Stock is incorporated into the model). Finally, Model IV is the proposed model that includes all components. Comparing the fit of the proposed model against the benchmark models allows us to evaluate the contribution of each component of the model.

Table 5 compares the above models in terms of model fit. Specifically, we compare model predictions of the number of visits and conversions

Table 4 Overview of Comparison Models

Model	Campaign-level effects		Creative-specific effect	
	Contemporaneous ad effects	Goodwill carryover with wearout and restoration	Contemporaneous ad effects	Goodwill carryover with wearout and restoration
I	No	No	No	No
II	Yes	Yes	No	No
III	No	No	Yes	No
IV (proposed model)	Yes	Yes	Yes	Yes

against the observed number of visits and conversions. We evaluate the models in terms of mean absolute percentage error (MAPE) and compute MAPE for visits and conversions separately, where a small MAPE value represents a model that captures actual behavior with low error. The results presented in Table 5 confirm that the proposed model provides the best fit.

That Models II and III both show substantial improvement over Model I indicates the importance of incorporating advertising effects either through advertising goodwill, wearout, or restoration effects at the campaign level (Model II) or through creative-specific contemporaneous effects (Model III). The notable improvement of the proposed model over and above Models II and III shows the value of incorporating creative-specific effects and advertising goodwill with both campaign- and creative-level wearout and restoration together in an integrated model. Overall, MAPE decreases from 20.3% error for visits and 30.0% error for conversions in Model I to just 12.9% error for visits and 11.1% error for conversions in our proposed model (Model IV).

We also examine how the models fit for different subsets of the data to demonstrate the ability of the model to discriminate according to different impression patterns. For example, even if one model fits better than the others in aggregate, a model that includes restoration effects should fit better among those customers with at least one hiatus in their impression histories. Table 6 provides a description of the different subgroups we constructed, and Table 7 provides in-sample MAPEs for each subgroup. Across nearly all subgroups (with just a few exceptions), the model specification presented in this paper fits the behavior exhibited in our data better than the benchmark models do.

Table 5 Model Fit (MAPE) Comparisons

Model	Visits	Conversions
I	0.203	0.300
II	0.129	0.184
III	0.150	0.186
IV (proposed)	0.129	0.111

Table 6 Description of Subgroups Used for Posterior Predictive Checks

Subgroup label	Explanation
(a) No impressions	All users who received no impressions during the observation period (all of these users made at least one visit)
(b) One impression	All users who received exactly one impression during the observation period
(c) Two impressions	All users who received exactly two impressions during the observation period
(d) Three or more impressions	All users who were exposed to three or more impressions during the observation period
(e) One distinct creative	All users who were exposed to exactly one distinct creative (but possibly with multiple exposures of that creative)
(f) Two distinct creatives	All users who were exposed to at least two distinct creatives during the observation period
(g) Three or more distinct creatives	All users who were exposed to three or more distinct creatives during the observation period
(h) Restoration	All users who experienced some kind of restoration effect (e.g., there was at least one incidence of a skipped week between exposures to the same creative)
(i) Two impressions, same week	All users who received two impressions in the same week

Table 7 MAPE by Subgroup

Model	Subgroup									
	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)	
Panel A: Visits										
I	0.138	0.860	0.684	0.319	0.764	0.393	0.204	0.297	0.684	
II	0.138	0.726	0.546	0.168	0.624	0.254	0.046	0.136	0.546	
III	0.123	0.787	0.598	0.203	0.682	0.294	0.068	0.166	0.598	
IV	0.130	0.743	0.545	0.151	0.634	0.244	0.017	0.107	0.545	
Panel B: Conversions										
I	0.258	0.871	0.726	0.216	0.754	0.366	0.240	0.032	0.726	
II	0.250	0.690	0.530	0.026	0.564	0.179	0.134	0.167	0.530	
III	0.188	0.737	0.554	0.038	0.606	0.194	0.136	0.141	0.554	
IV	0.177	0.607	0.421	0.066	0.482	0.082	0.053	0.025	0.421	

Parameter Estimates

Baseline process estimates for each component model are provided in Table 8, panel A. The baseline rates of impressions, visits, or conversions are given in the first row of each section. The zero-inflation parameter is given in the second row. The coefficients (γ) for the weekly indicators are also provided in panel B of Table 8.

The zero-inflation parameter in the ad impressions model suggests that, in addition to what is predicted

Table 8 Baseline Process Estimates for Ad Exposures, Visits, and Conversion Behavior

	Quantile				
	2.5%	Median	97.5%		
Panel A: Estimates					
Ad impressions					
Baseline rate (λ_0)	0.087	0.097	0.094		
Zero-inflation (r_m)	0.387	0.406	0.423		
Week 1	—	—	—		
Week 2	0.426	0.526	0.636		
Week 3	0.999	1.086	1.174		
Week 4	1.105	1.183	1.260		
Week 5	1.510	1.593	1.670		
Week 6	1.864	1.953	2.021		
Week 7	1.756	1.841	1.917		
Week 8	1.741	1.822	1.904		
Week 9	1.595	1.688	1.761		
Visiting behavior					
Baseline rate (μ_0)	0.063	0.068	0.074		
Zero-inflation (r_v)	0.114	0.156	0.198		
Week 1	—	—	—		
Week 2	-0.310	-0.197	-0.068		
Week 3	-0.423	-0.296	-0.152		
Week 4	-0.657	-0.558	-0.456		
Week 5	-0.750	-0.592	-0.462		
Week 6	-0.740	-0.643	-0.514		
Week 7	-0.319	-0.187	-0.077		
Week 8	-0.172	-0.062	0.044		
Week 9	-0.417	-0.309	-0.183		
Conversion behavior					
Baseline rate (ρ_0)	0.524	0.582	0.639		
Zero-inflation (r_s)	0.316	0.358	0.394		
Week 1	—	—	—		
Week 2	-0.196	0.134	0.455		
Week 3	-0.439	-0.126	0.264		
Week 4	-0.600	-0.217	0.190		
Week 5	-0.491	-0.188	0.170		
Week 6	-0.584	-0.222	0.170		
Week 7	-0.185	0.098	0.424		
Week 8	-0.357	-0.028	0.253		
Week 9	-0.462	-0.135	0.155		
Panel B: Covariance matrix (Σ)					
	$\log \lambda_i$	$\log \mu_{0i}$	$\text{logit } \rho_{0i}$	$\beta_{\mu t}$	$\beta_{\rho i}$
$\log \lambda_i$	3.034				
$\log \mu_{0i}$	3.507	4.775			
$\text{logit } \rho_{0i}$	-0.129	1.561	98.316		
$\beta_{\mu i}$	-0.048	-1.447	1.027	3.988	
$\beta_{\rho i}$	3.386	4.051	19.353	0.962	9.058

Table 9 Advertising Impression Effects

	Quantile		
	2.5%	Median	97.5%
Baseline effect of any ad in campaign (AD)	0.585	0.747	0.871
Creative-specific effects (C_j)			
Creative A	0.342	0.543	0.751
Creative B	0.351	0.454	0.559
Creative C	0.373	0.507	0.602
Creative D	0.339	0.436	0.540
Creative E	0.259	0.511	0.772
Creative F	0.013	0.336	0.668
Creative G	0.471	0.789	0.960
Creative H	0.479	0.774	1.110
Creative I–O	-0.694	-0.103	0.540
$E[\beta_{\rho i}]$	0.016	0.113	0.195

by the Poisson distribution, another 40.6% of the consumers represented in our data are never exposed to an ad. Of the remaining users, the rate of exposure (λ_0) according to the Poisson is 0.087 ads per week. In terms of visiting behavior, the zero-inflation probability is much smaller (0.156). In the absence of any ad exposures, the average rate of visits (μ_0) for the remaining population is 0.068 visits per week. These numbers are consistent with the low ad click-through rates observed online. Finally, conversion model indicates a zero-inflation probability of 0.358. For the remainder of the population, the probability of conversion at each visit (p_0) is 0.582. Because the advertiser in our case classifies a broad range of activities to be conversion activities, our conversion estimate is higher than the typical conversion rate observed with online retailers, but it is not out of line with the advertiser’s expectations.

How Effective Is Each Ad Creative? Table 9 provides the estimated posterior medians and 95% highest posterior density intervals for each of the nine creative-specific effects (the C_j parameters). These parameters are indicators of the relative marginal contribution of each ad creative, before adjusting for wearout and restoration effects, on the contemporaneous ad effect. With the exception of creatives I–O, there is a high probability that all creatives contribute positively to Ad Stock. Based on median posterior estimates, we expect creatives G and H to be the most effective.⁹

The effect of advertising on conversion is captured by the same creative-specific and scheduling effects but is scaled by the coefficient $\beta_{\rho i}$ from Equation (10). The marginal posterior distribution of the mean of

⁹ Note that the range around the effect of creative H is quite large. However, even the lower range of the estimate indicates that it is an effective ad.

β_{pi} is in the last row of Table 9, from which we observe a significant positive mean effect of accumulated Ad Stock on conversion probabilities. These results suggest not only that advertising generates more visits but also that the advertising-induced visits are more likely to convert than organically generated visits. This metric can have significant managerial implications because it provides an important tool for advertisers to assess whether the increased number of visits they obtain from advertising provides valuable customer leads or just curious browsers.

How Does the Timing of Ad Impressions Influence Ad Stock Effects? The accumulation and decay of the advertising goodwill construct, Ad Stock, considers three separate effects: (1) the decay of Ad Stock over time, (2) wearout effects with repeated advertising exposures, and (3) restoration effects over time in the absence of repeated exposures. Table 10 presents the marginal posterior distributions for these effects.

The Ad Stock decay parameter captures the extent to which the effect of seeing ads in previous weeks carries over into subsequent weeks. We estimate that about 37.4% of the effects are lost from week to week.

The wearout effects, δ_1 and δ_2 , describe the extent to which advertising will retain its incremental effect after repeated viewing. We estimate this effect at both the campaign level (where repeated exposures to any ad in the campaign will create ad wearout) and the creative level (where repeated exposures of an ad with a given creative will result in wearout for that specific creative only). On average, each subsequent exposure to this ad campaign results in substantial wearout of all ad creatives ($1 - \delta_1 = 0.778$). Wearout from repeated exposures to the same ad creative is also substantial ($1 - \delta_2 = 0.403$) but substantially less than the campaign-level effects.

However, the wearout is gradually restored as time passes between exposures. At the campaign level, 2.7% ($0.028/(1 + 0.028)$) of the wearout is restored in each week for which there is a hiatus in advertising (i.e., the consumer is not exposed to any ad from the campaign). Creative-specific effects are restored at a

slightly faster rate, with 8.8% ($0.096/(1 + 0.096)$) of the effect that was previously “worn out” being restored each week since the previous exposure of the same ad creative.

Targeting Creatives Based on Impression Histories

Empirical Results

The results presented in the previous section clearly show how advertising impression histories can affect the response to subsequent ad exposures. To illustrate how impression histories could affect advertising targeting decisions, let us consider a stylized example of a campaign consisting of only two ad creatives: G (estimated to be one of the most effective) and A (estimated to be substantially less effective). Table 11 shows the potential effect that each creative can have on Ad Stock. Based on those effects, the final column indicates the ad creative that would generate the greatest response.

Panel A of Table 11 considers a situation in which the advertiser must be opportunistic. That is, we assume that advertising opportunities do not necessarily occur at regular intervals and that when an advertising opportunity presents itself, the advertiser must evaluate the potential effect of each creative before selecting it. In week 1, we assume that no ads from this campaign have been previously served, and thus the effects of G and A on Ad Stock are equivalent to the sum of the baseline ad effect and the creative-specific marginal effects presented in Table 9. If an advertising opportunity presents itself in week 1, the advertiser should serve the more effective ad (G). However, if a second advertising opportunity then presents itself in week 2 or 3, the advertiser should serve creative A, since the previous impression of G makes any subsequent impression subject to creative-specific wearout effects (in addition to campaign level wearout). However, if that second advertising opportunity were to come in week 4, the effectiveness of creative G will have been restored to a level that is comparable to that of A (after accounting for campaign-level and creative-specific wearout and restoration of both ad creatives).

Next, we consider an alternative scenario in which the advertiser has weekly advertising opportunities. In panel B of Table 11, we compare the effect of each creative on Ad Stock, assuming that an impression is served every week. Again, in week 1, the advertiser serves creative G. But once it is worn out, by week 2 creative A, which is still fresh, is more effective. In week 3, the effects of creative G are partially restored, but creative A is worn out, so G has a greater effect and therefore is the creative of choice for week 3.

Table 10 Impression History Effects

	Quantile		
	2.5%	Median	97.5%
Ad stock decay (α)	0.337	0.374	0.421
Campaign-level effects			
Wearout effects (δ_1)	0.130	0.222	0.380
Restoration effects (ρ_1)	0.011	0.028	0.059
Creative-level effects			
Wearout effects (δ_2)	0.507	0.597	0.693
Restoration effects (ρ_2)	0.016	0.096	0.463

Table 11 Creative Effects on Ad Stock

Week	G	A	Serve ad creative ...
Panel A: Assuming occasional advertising opportunities			
1 (first impression)	1.536	1.290	G
2 (possible second impression)	0.681	0.725	A
3 (possible second impression)	0.719	0.740	A
4 (possible second impression)	0.753	0.754	G or A
Panel B: Assuming weekly advertising opportunities			
1 (first impression)	1.536	1.290	G
2 (second impression)	0.681	0.725	A
3 (third impression)	0.578	0.399	G

The effects of different ad impression histories are then taken into consideration next. In this example, we construct ad exposure schedules for six hypothetical users whom we treat as if they were sampled from the same population of users that are included in our data set. We then simulated 5,000 posterior predictive data sets, conditioning on the data that are implicit in these profiles.

All six histories consist of at most two unique creatives (see Tables 12 and 13). The first two histories describe individuals who are exposed to one ad impression every week for four weeks. In the first history, all impressions are of ad G (the most effective creative in our sample), whereas in the second history, the impressions alternate between ads G and A. We also consider two additional impression histories for which impressions are concentrated in alternating

weeks. In history 3, we examine the case where all impressions are of ad G, and in history 4, we consider the case where both G and A are presented each week. The final two impression histories are scenarios for which the impressions are highly concentrated in the first week and significant time passes before the next impression opportunity in week 5. Again, we consider both a single creative history and a two-creative history for these highly concentrated impression schedules.

The effect of an additional ad in week 5 will depend on the individual's history of ad impressions and the ad presented in week 5. Therefore, using our model results, we simulate the expected number of visits and success conversions for various impression histories and compare the effects on visits and conversions that different creatives would have if shown in week 5 over the subsequent five weeks.¹⁰ In addition to ads G and A (which are present in the simulated impression histories), for comparison purposes we consider the effect of ad C, which has a baseline creative-specific effect comparable to that of A.

For each history, we compare the effects of each ad creative. In practice, the advertiser should favor the ad creative that generates the greatest number of visits and conversions, given the individual's history. We highlight these instances in bold in Tables 12 and 13.

When comparing histories in which only ad G was shown (histories 1, 3, and 5) to those where both G and A were shown (histories 2, 4, and 6), we see that histories with creative variety (histories 2, 4, and 6) result in more visits and conversions regardless of which ad creative is shown next. Additionally, after four repeated exposures, ad G is significantly worn out by the fifth week, whereas ads A and C are still fresh. As a result, in histories where only G is served (histories 1, 3, and 5), the advertiser would benefit similarly from showing ad A or ad C in the fifth week since they are comparable in baseline effectiveness

Table 12 Expected Visits for Various Ad Impression Histories

Impression history	Week				Ad		
	1	2	3	4	G	A	C
History 1	G	G	G	G	0.327	0.388	0.379
History 2	G	A	G	A	0.384	0.347	0.401
History 3	GG		GG		0.401	0.505	0.469
History 4	GA		GA		0.511	0.476	0.505
History 5	GGGG				0.403	0.489	0.458
History 6	GAGA				0.452	0.412	0.469

Notes. G is the most effective ad with $C_G = 0.789$. Ads A and C are comparable in effectiveness with $C_A = 0.543$ and $C_C = 0.507$.

Table 13 Expected Conversions for Various Ad Impression Histories

Impression history	Week				Ad		
	1	2	3	4	G	A	C
History 1	G	G	G	G	0.109	0.150	0.134
History 2	G	A	G	A	0.143	0.132	0.156
History 3	GG		GG		0.133	0.192	0.188
History 4	GA		GA		0.199	0.162	0.187
History 5	GGGG				0.155	0.185	0.167
History 6	GAGA				0.184	0.162	0.188

Notes. G is the most effective ad with $C_G = 0.789$. Ads A and C are comparable in effectiveness with $C_A = 0.543$ and $C_C = 0.507$.

¹⁰ We assume that no ads other than the ones described in Tables 12 and 13 are served.

Table 14 Number of Impressions by Ad Creative in Week 10

Creative	Strategy		
	Observed	Naïve	Model
Creative A	68	0	256
Creative B	906	0	57
Creative C	443	0	145
Creative D	437	0	59
Creative E	5	0	143
Creative F	11	0	47
Creative G	5	1,875	413
Creative H	0	0	746
Creatives I–O	0	0	9

($C_A = 0.543$, $C_C = 0.507$), and both would produce better results relative to a wornout ad G. In contrast, in histories where both ads G and A are subject to wearout effects (histories 2, 4, and 6), ad C would generate better results than ad A.

Furthermore, the highly concentrated impression schedules in histories 5 and 6 provide an interesting illustration of ad restoration effects. In both of those cases, the three-week hiatus in weeks 2–4 allowed the ads to gradually regain their effectiveness to a point where they are almost comparable to a fresh ad C.

Although the results presented in Tables 12 and 13 represent stylized examples of how different ad histories and ad creatives affect behavior, they begin to illustrate how advertisers can customize ad creatives for an individual’s impression history. To further test this application, we next present a larger simulation in which the individuals observed in our data are exposed to alternative ad content chosen by different advertising policies.

Simulation

Thus far, we have used weeks 1–9 of our data for model testing and estimation. We have held out a 10th week of data for use in this simulation. In week 10 of our data, 645 individuals (of a total of 5,803 in the entire data set) were exposed to an advertisement, resulting in a total of 1,875 impressions. The creative content served in these impressions is presented in the first column of Table 14.

As a benchmark scenario, we first consider a naïve advertising strategy in which the most effective ad

creative in the ad campaign is chosen to be the only ad in the campaign. In other words, all impressions are of a single ad creative (in our case, creative G). The second column of Table 14 describes this scenario.

We simulate, for each of the 645 individuals, alternative ad impressions in week 10 based each of their unique impression histories. Similar to the stylized examples presented above, we calculate the potential effect of each available creative if it were served in the next impression. After accounting for wearout and restoration effects (at both the campaign and creative levels), the effect of each creative will vary across individuals depending on his or her impression history. For each individual, our simulation will assume that the creative with the highest marginal effect will be served in the next impression opportunity. The last column of Table 14 summarizes the impressions that result from this process.

For the scenarios described in Table 14, we calculate Ad Stock for each individual based on his or her impression history. We assume that no other ad impressions were served after week 10. One could, in practice, relax that assumption, but for the purposes of this paper, this assumption allows us to more easily compare scenarios by limiting the number of factors varied across individuals in the simulation. Then, for each individual, we simulate the resulting visiting and conversion behavior. The results of these simulations are provided in Table 15.

The results in Table 15 highlight the value of using our proposed model to target ad creatives based on individual impression histories. The first four rows provide the expected number of visits and conversions per individual each week in weeks 10–13. These visits and conversions are based on the impression histories observed in weeks 1–9 and the impressions planned under each scenario in week 10. Because the available data end after week 10, we compare expected visits and conversions.

In week 10, the advertising impressions actually observed in the data are expected to generate 0.0545 visits per individual and 0.0298 conversions per individual. If, however, the creative content of the

Table 15 Simulated Behaviors Across Advertising Scenarios

Week	Expected visits per individual			Expected conversions per individual		
	Observed strategy	Naïve strategy	Model strategy	Observed strategy	Naïve strategy	Model strategy
Week 10	0.0545	0.0568	0.0614	0.0298	0.0311	0.0339
Week 11	0.0518	0.0524	0.0530	0.0282	0.0285	0.0289
Week 12	0.0509	0.0511	0.0513	0.0277	0.0278	0.0279
Week 13	0.0506	0.0507	0.0508	0.0276	0.0276	0.0276
Total for weeks 10–13	0.2078	0.2110	0.2165	0.1133	0.1150	0.1183

ads served in week 10 were chosen based on our proposed model, the expected number of visits would increase to 0.0614, and the number of conversions would increase to 0.0339. In other words, the number of visits would increase by 12.7%, and the number of conversions would increase by 13.8%. Note that, although less dramatic, the model-based ad policy also improves on a naïve policy in which the single most effective ad is used exclusively. Additionally, because ad effects carry over from week to week via the Ad Stock construct, increases in visiting and conversion behaviors are expected beyond week 10, even though our simulation assumes no additional ad impressions. Overall, the simulation presented above demonstrates the managerial value of considering ad impression histories when targeting ad creatives to individuals.

Conclusion

Although most online advertisers focus on the effects of a single ad impression, we examine the effect of an ad impression in the context of the individual's impression history. This has implications for online targeting practices. Specifically, when an individual visits a website on which an advertiser has purchased space, the advertiser has a choice of ads to present. Ideally, the advertiser would present the most effective ad. However, what is the "most effective" ad? Our research has shown that the most effective ad will depend on the individual's ad impression history, thus presenting an approach that advertisers can use to estimate the effect of each ad creative in the context of an ad impression history.

One facet of online advertising that we were not able to consider is whether there may be interactions in the order in which different creatives are presented. That is, some creatives might be more effective when presented before or after others. This would have further implications for the construction of impression schedules that would extend beyond just the consideration of wearout and restoration effects. Additionally, we did not consider the potential interaction effect between the ad and the website on which the ad appeared. We leave these modeling challenges to future research.

Additionally, we do not consider wear-in effects (Pechmann and Stewart 1990) or threshold effects (Dubé et al. 2005). Researchers have found that sales responses to advertising often follow an S-shaped curve (Little 1979). That is, when advertising is below a certain threshold, advertising has virtually no effect. Only when a critical or threshold level of advertising is reached do we observe any real response to advertising. In other words, some level of ad repetition is required for the ads to wear in and affect consumers.

We considered wearout effects in this paper, but we assumed that there is no wear-in period. To accommodate wear-in, the instantaneous effect of an ad impression (E_{it}) can be specified to allow ad effects to increase with repetition. We leave this extension for future researchers.

A natural extension of this research is to consider optimal ad schedules. For example, choosing an ad for the next period makes that ad less effective in subsequent periods, so in the presence of wearout, it might be better to save that ad for the future. In addition, the advertiser might want to select a creative whose expected effectiveness is low but for which the uncertainty of the effectiveness is high (e.g., a new creative). In this case, the advertiser might benefit by learning more about creatives that have been used less often, just in case they turn out to be more effective. Thus, there is a trade-off between "exploration" and "exploitation." One then might treat this problem as a Markov decision process that selects an impression schedule to maximize the total successes over time. The implementation of such an advertising policy resembles the classic "bandit problem" used to model reinforcement learning (Gittins et al. 2011) and extends well beyond the scope of this research. However, incorporating wearout and restoration effects into such a forward-looking decision process would be a welcome addition to any future research.

Acknowledgments

The authors thank the Wharton Customer Analytics Initiative and Organic, Inc., for providing the data used in this research and for their invaluable support and feedback throughout the research process.

References

- Chatterjee P, Hoffman DL, Novak TP (2003) Modeling the click-stream: Implications for Web-based advertising efforts. *Marketing Sci.* 22(4):520–541.
- Drèze X, Hussherr F-X (2003) Internet advertising: Is anybody watching? *J. Interactive Marketing* 17(4):8–23.
- Dubé J-P, Hitch GJ, Manchanda P (2005) An empirical model of advertising dynamics. *Quant. Marketing Econom.* 3(2):107–144.
- Feinberg FM (1992) Pulsing policies for aggregate advertising models. *Marketing Sci.* 11(3):221–234.
- Feinberg FM (2001) On continuous-time optimal advertising under S-shaped response. *Management Sci.* 47(11):1476–1487.
- Gelman A, Carlin JB, Stern HS, Rubin DB (2003) *Bayesian Data Analysis*, 2nd ed. (Chapman & Hall/CRC, Boca Raton, FL).
- Gittins J, Glazebrook K, Weber R (2011) *Multi-Armed Bandit Allocation Indices* (John Wiley & Sons, Chichester, UK).
- Heckman JJ (1979) Sample selection bias as a specific error. *Econometrics* 47(1):153–161.
- Interactive Advertising Bureau (2011) IAD Internet advertising revenue report. Report, IAB, New York. http://www.iab.net/media/file/IAB_Full_year_2010_0413_Final.pdf.
- Little JDC (1979) Aggregate advertising models—The state of the art. *Oper. Res.* 27(4):629–667.

- Lodish LM, Abraham MM, Livelsberger J, Lubetkin B, Richardson B, Stevens ME (1995) A summary of fifty-five in-market experimental estimates of the long-term effect of TV advertising. *Marketing Sci.* 14(3):G133–G140.
- Mahajan V, Muller E (1986) Advertising pulsing policies for generating awareness for new products. *Marketing Sci.* 5(2): 89–106.
- Manchanda P, Dubé JP, Goh KY, Chintagunta PK (2006) The effect of banner advertising on Internet purchasing. *J. Marketing Res.* 43(1):98–108.
- Moe WW, Schweidel DA (2012) Online product opinions: Incidence, evaluation, and evolution. *Marketing Sci.* 31(3):372–386.
- Morrison DG, Schmittlein DC (1981) Predicting future random events based on past performance. *Management Sci.* 27(9):1006–1023.
- Naik PA, Mantrala MK, Sawyer AG (1998) Planning media schedules in the presence of dynamic advertising quality. *Marketing Sci.* 17(3):214–235.
- Nerlove M, Arrow KJ (1962) Optimal advertising policy under dynamic conditions. *Economica* 29(114):129–142.
- Pechmann C, Stewart DW (1990) Advertising repetition: A critical review of wear-in and wear-out. MSI Report 90-106, Marketing Science Institute, Cambridge, MA.
- Rutz OJ, Bucklin RE (2011) From generic to branded: A model of spillover in paid search advertising. *J. Marketing Res.* 48(1):87–102.
- Tellis GJ (2009) Generalizations about advertising effectiveness in markets. *J. Advertising Res.* 49(2):240–245.
- Vakratsas D, Feinberg FM, Bass FM, Kalyanaram G (2004) The shape of advertising response functions revisited: A model of dynamic probabilistic thresholds. *Marketing Sci.* 23(1):109–119.
- Winkelmann R (2008) *Econometrics Analysis of Count Data* (Springer, Berlin).
- Ying Y, Feinberg F, Wedel M (2006) Leveraging missing ratings to improve online recommendation systems. *J. Marketing Res.* 43(3):355–365.