From Generic to Branded:

A Model of Spillover in Paid Search Advertising

Oliver J. Rutz

Randolph E. Bucklin*

May 2008

*Oliver J. Rutz is Assistant Professor of Marketing, Yale School of Management, 135 Prospect Street, New Haven, CT 06520, oliver.rutz@yale.edu. Randolph E. Bucklin is the Peter W. Mullin Professor, UCLA Anderson School, 110 Westwood Plaza, Los Angeles, CA 90095, rbucklin@anderson.ucla.edu. The authors wish to thank a collaborating firm for providing the data used in this study.

From Generic to Branded:

A Model of Spillover in Paid Search Advertising

Abstract

In Internet paid search advertising, many marketers pay for search engines to serve text ads in response to keyword searches that are generic (e.g., "Hotels") or branded (e.g., "Hilton Hotels"). While stand-alone metrics usually show that generic keywords have higher apparent costs to the advertiser than branded keywords, generic search may create a spillover effect on subsequent branded search. Building on the Nerlove-Arrow advertising framework, the authors propose a dynamic linear model to capture the potential spillover from generic to branded paid search. In the model, generic search ads expose users to information about the advertiser's brand, increasing its awareness level. This, in turn, affects future search activity for keywords which include the brand name. Using a Bayesian estimation approach, the authors apply the model to data from a paid search campaign for a major lodging chain. The results show that spillover is asymmetric. Generic search activity positively affects branded search activity via increased awareness but branded search does not affect generic search. Implications for improving metrics for paid search advertising are discussed.

Keywords: Internet, Advertising, Paid Search, Spillover, Awareness, Nerlove-Arrow Model, Bayesian Dynamic Linear Model (DLM)

INTRODUCTION

Paid search advertising is a service offered by Internet search engines in which the advertiser selects specific keywords and creates a text ad to appear when a user searches for those keywords. Paid search text ads are displayed in the sponsored section of the search results page and are separate from so-called organic search listings (see Figure 1). In 2007, paid search advertising accounted for roughly 44 percent of the $21.7 billion spent on Internet advertising, more than double that of Internet display advertising at 21 percent (PricewaterhouseCoopers 2007). Unlike paid search, where user queries lead to ad exposure, display ads typically appear as banners on web pages or in pop-ups.

In Internet paid search advertising, the serving of a text ad is called an impression. The position of the ad, which is based in part on an auction-like mechanism, describes the order in which the ad is listed in the sponsored sections, e.g., 1st or 3rd. Text ads are served for free but advertisers are charged when a user clicks on an ad. A click takes the user to the designated "landing page" of the sponsor's website. The selection of the keywords and the bid per keyword can be changed by the advertiser at any time. Advertisers can also limit the amount they are willing to pay for each click by setting a ceiling and a budget.

Most of the research in marketing on Internet advertising has, at least thus far, focused on understanding the effectiveness of banner ads. One long-standing metric for banner ads has been click-through rates (e.g., Novak and Hoffman 2000, Chatterjee, Hoffman and Novak 2003). However, click-through rates declined from about 7% in 1996 (Dreze and Hussherr 2003) to around 0.2% in 2007 (Business Week 2007) and have come under increasing scrutiny. Dreze and Hussherr (2003) found that consumers avoid looking at banner ads, implying a processing of such ads that is mostly pre-attentive. They conclude that traditional measures such as brand awareness and brand recall may be more appropriate to gauge the effect of online advertising.

Following up, Cho and Choen (2004) find that consumers' avoidance stems from perceived goal impediment. They recommend that advertisers should use highly customized context-congruent advertising messages to reduce perceived goal impediment.

Paid search advertising may provide a way to address some of the limitations of banner ads. Whereas traditional advertising, online and offline, is often seen as intrusive and disruptive, paid search advertising delivers a requested and highly context-congruent advertising message. Clearly, paid search advertising seems appealing from a theoretical perspective – and its fast growth (CAGR of 105% from 2000 to 2007, based on PricewaterhouseCoopers 2007) in the marketplace provides practical confirmation. On the other hand, little is known about how managers should assess the performance of a paid search campaign and what metrics are most appropriate to use.

At first glance, one might consider a standard approach to performance evaluation based upon calculating the marginal benefit of spending on each search keyword. Unfortunately, such an approach is infeasible with typical paid search data. Most search keywords – about 80 percent in our sample – have impressions and clicks (costs) associated with them on a daily basis. While a search always leads to an impression, it seldom leads to a click and, even more seldom, a sale. In the left panel of Figure 2, we illustrate this progression from the advertiser's perspective. Because conversion rates in paid search advertising are very low, on a daily, or even weekly, basis most keywords are not associated with any sales whatsoever. This precludes calculating the marginal benefit for most keywords, even over extended periods of time.

Generic versus Branded Keywords

In this study we focus on aggregate-level campaign performance metrics, due in part to the limitations just noted. At this level, one can examine metrics across all keywords included in an advertiser's campaign or, short of that, a large subset of those keywords. One of the clearest categorization schemes for keywords is to group them by generic versus branded. A generic keyword does not contain brand names (e.g., "Hotels LA") while a branded keyword does ("Hilton Hotels LA"). Using data from a paid search campaign for a lodging chain, we illustrate the striking differences in performance of branded versus generic keywords on data for the Google and Yahoo! search engines (see Table 1 and 2).[1]

We note first that the basic measures of consumer response in paid search are higher for branded keywords. On Google, click-through rates (e.g., 13.68% vs. 0.26%) and conversion rates (e.g., 6.03% vs. 1.05%) are substantially greater for branded versus generic keywords. (The Yahoo! data show a similar pattern.) What may account for this? We propose that a consumer who searches using a generic keyword may *not be aware* that a specific brand, e.g., "our" lodging company, is *relevant* for his current search. Conversely, a consumer using a branded keyword is likely to be *aware* that the brand is *relevant* to the search. This difference in *awareness of relevance* should then translate into differences in consumer response to both the text ad itself (click-through) and the likelihood of purchase, given click-through (in our case, a hotel reservation).

In addition to higher response rates, branded keywords are also less expensive on a cost-per-click basis (e.g., \$0.18 vs. \$0.55 for Google). In the lodging industry companies bid competitively for prime search keywords such as "Hotels LA" or "Lodging LA." These terms are of interest to many hospitality firms and to other travel companies such as Travelocity.com or Expedia.com. Branded keywords, on the other hand, are much less competitive. A given brand may have only a few direct competitors willing to bid on their branded keywords. For example,

---

[1] Data from other campaigns analyzed by the authors show similar patterns, but confidentiality restrictions do not allow us to report details. Conversations with numerous practitioners indicate that this pattern is also a widespread phenomenon.

"Hilton Hotels LA" might interest high-end hotel chains or general travel sites, but not others (e.g., Motel 6).

Advertisers can evaluate campaigns using stand-alone measures for performance (click-through or conversion rate) or financial return (the cost-per-click or cost-per-sale that can be attributed to the paid search advertising). Lower cost-per-click (CPC) combined with a higher rate of conversion then leads to a dramatically lower cost-per-conversion for branded keywords (e.g., $2.94 vs. $51.84 on Google). The purpose of this paper is to investigate whether these types of stand-alone metrics are likely to be valid when the paid search campaign includes large numbers of both generic and branded search terms.

Spillover Effects in Paid Search

One drawback of the simple metrics just described is that they do not account for the potential dynamic interaction that may occur between generic and branded search activity. We propose that generic search activity can create awareness that the brand is relevant for the search and consequently "spill over" to influence subsequent branded search activity. This awareness can then lead to future branded keyword searches in which the user seeks to research the brand in more detail. In this manner, we propose that generic search can create spillover to branded search via a latent construct for awareness.

We propose to model spillover in paid search advertising by building on the so-called "leaky bucket" approach to advertising (e.g., Nerlove and Arrow 1962, Naik et al. 2008). The basic premise of the model is that exposure to brand-related information ensuing from generic search increases awareness but that awareness also decays over time (i.e., leaks out of the bucket). Higher awareness, in turn, can then lead to an increase in branded search activity. To handle the dynamic nature of this process, we specify the model in a multivariate time-series framework. Specifically, we will use a Dynamic Linear Model (DLM) set up and estimated in a Bayesian Framework following the procedures from West and Harrison (1997).

Our dataset contains daily information on paid search advertising for a campaign conducted by a major lodging chain. The data set includes the number of daily impressions, clicks, and reservations associated with the generic and branded keyword categories. Though these data are quite detailed in many respects, they reflect the aggregate-level performance of the keywords. This means that we do not have information on the search activity of individuals. To the best of our knowledge, this type of information is currently not provided to advertisers by the search engines. Thus, the data we analyze have the same characteristics as the information managers currently use to evaluate their paid search campaigns.

Our intended contribution is to develop and test a modeling approach to assess the extent to which generic search activity creates spillover into branded search. Building on the leaky bucket conceptualization of advertising effects, we specify an aggregate-level model of paid search. We show that generic search spills over into branded search via a latent construct for awareness. Our findings imply the standalone metrics currently in widespread use could be misleading about the real performance of generic versus branded keywords and we discuss how these can be adjusted.

The paper is structured as follows. After a brief literature review, we present our model specification, describe our dataset in detail, and discuss the empirical results we obtained from estimating the model on the lodging data. Next, we discuss the implications of our findings for performance metrics in paid search. A concluding section summarizes, notes limitations of our approach, and discusses future research opportunities in search engine marketing.

LITERATURE

Most of the empirical research about *online advertising* in marketing has been focused on banner ads. As noted above, a key measure of banner advertising performance has been click-through rates (e.g., Novak and Hoffman 2000, Chatterjee, Hoffman and Novak 2003). Of course, click-through rates have declined dramatically since the 1990's. This has led some researchers to recommend that traditional measures of advertising effectiveness, such as awareness and recall, be used instead (e.g., Dreze and Hussherr 2003). Following up, Danaher and Mullarkey (2003) find that the duration of the exposure to a banner ad as well as user involvement (e.g., surf vs. goal-directed mode) increases the likelihood of recall. Cho and Choen (2004) find that consumers avoid looking at advertising on the Internet because of perceived goal impediment. To address this, they recommend using highly customized context-congruent advertising messages. Moore et al. (2005) also investigate the importance of congruity between the website and the ad and find that congruity has favorable effects on attitudes. Lastly, in an empirical modeling study, Manchanda et al. (2006) examine the relationship between banner advertising and purchase patterns. They find that banner ads can play a significant role in customer retention.

To the best of our knowledge, no empirical research on paid search advertising has yet appeared in the published marketing literature. Existing research on search engines has analyzed search engine visits (Telang et. al. 2004) and the effectiveness of search engines (Bradlow and Schmittlein 2000) in information retrieval. Recent theoretical papers investigate paid search auction mechanisms (Edelman and Ostrovsky 2007, Edelman et al. 2007) and paid search advertising as a product differentiation game (Chen and He 2008). In related work, Wilbur and Zhu (2008) investigate click fraud in paid search auctions from a theory perspective. Ghose and Yang (2008), in a recent working paper, investigate how bid choice and the search engine's position choice can be modeled simultaneously and investigate the potential for cross-selling by using paid search advertising. They do not address the potential spillover from generic to branded search.

MODELING APPROACH

Data on impressions, clicks, positions and costs are provided to the advertiser by the search engine/third party data sources on an aggregate level,[2] typically reported on a daily basis. The data are aggregated on the basis of a keywords. For each search keyword (e.g., rental cars LA) campaign managers have daily information on cost (in $), average position served (given by daily average placement rank, e.g., 2.3), number of impressions and clicks, and number of sales, or in our case, reservations. We build our model using this currently available paid search data.

Our modeling approach to capturing spillover effects is based, in part, on the notion of *awareness of relevance*. We propose that such awareness parsimoniously captures the effects of exposure to brand-related information. Consumers who conduct a generic search might not be aware of the brand at all or, even if they are, they might not be aware that the brand is *relevant* for the search. Generic search leads to brand-related exposures in the form of impressions, i.e., the text ads in the sponsored section of the search results page, and clicks, i.e., the searcher has clicked on the ad and has been taken to the advertiser's web site. These brand-related exposures may create and/or increase awareness[3] of the brand. We note that the impact of these two types of exposures might differ. An impression is a passive exposure to the brand's text ad, whereas a click is an active opt-in which leads to further information exposure at (and possibly after) the landing page. In our model, we investigate whether generic impressions and generic clicks have different effects in generating spillover to branded search.

[2] Note that Google as well as Yahoo! do not make searcher-level data (i.e., clickstream) available to advertisers.
[3] We use simply *awareness* in place of *awareness of relevance* for the remainder of the paper

Modeling the effects of advertising based on aggregate data has been done in a variety of ways (e.g., Tellis 2004). A popular class of models is the leaky-bucket type of approach which postulates that there is an "Ad Stock" that decays over time and is replenished by advertising activity. One of the most well-known leaky-bucket models is the Nerlove-Arrow model (N-A model, Nerlove and Arrow 1962). Recent examples of the use of the N-A model include Naik et al. (1998) and Naik et al. (2008). For example, Naik et al. (2008) investigate how a firm should optimally allocate resources to corporate versus product branding efforts for a multi-product firm. Other applications include Bass et al. (2007) in which the N-A model is linked with a demand model for telephone services. Building on this research stream, we propose a N-A type awareness model for the paid search spillover problem.

Model Specification

We begin our model specification by linking the change in awareness to generic search activity and a carryover effect of awareness. Following existing literature (e.g., Nerlove and Arrow 1962 or Naik et al. 1998) we specify the dynamic evolution of changes in awareness as

(1) $$\frac{dA_t}{dt} = \beta^{gen} Gen_t - \tilde{\alpha}^A A_t ,$$

where A_t is awareness at time t, Gen_t is a vector of generic search activity variables at time t, β^{gen} and $\tilde{\alpha}^A$ are parameters to be estimated. Note that in discrete time this model can be rewritten as

(1a) $$A_t = \beta^{gen} Gen_t + \alpha^A A_{t-1} ,$$

where $\alpha^A = (1 - \tilde{\alpha}^A)$ and α^A is the carry-over rate of awareness (e.g., Naik et al. 1998). [4]

Paid search data allow us to explore different measures for generic search activity, Gen_t. Following a traditional advertising approach, we can capture generic search activity by using the

[4] In discrete time $\Delta A = C - \tilde{\alpha}^A A_t$, where $\Delta A = A_t - A_{t-1}$ and C represents all other terms. Thus, $A_t = C + (1 - \tilde{\alpha}^A) A_{t-1}$. It follows that $A_t = C + \alpha^A A_{t-1}$ and α^A is called carry-over rate.

dollar amount spent on generic search by the advertiser on a daily basis. Because we have the number of daily impressions and clicks for generic keywords, we can model generic search activity with these variables. This allows us to investigate whether the actual exposure information provides a better measure than dollar spending. Possible saturation effects can also be easily studied by modeling, for example, the log of impressions and the log of clicks.

Equation (1) specifies the dynamics for how generic search activity affects awareness. Next, we need to specify the dynamics of how awareness affects branded search. We model the change in each branded search activity – impressions, clicks, and reservations -- as a function of exogenous variables such as seasonal effects, lagged branded search activity and latent awareness. The structure of this formulation closely follows the procedures developed by Naik et al. (1998, 2008) and Bass et al. (2007) and offers an appealing and parsimonious way to model the aggregate information available from paid search campaigns.

We define three dynamic models, one for changes in each branded search activity:

(2) $$\frac{dBr^{imp}}{dt} = \beta^{imp} I_t - \tilde{\alpha}^{imp} Br_t^{imp} + \gamma^{imp} A_t \ ,$$

(3) $$\frac{dBr^{cl}}{dt} = \beta^{cl} I_t - \tilde{\alpha}^{cl} Br_t^{cl} + \gamma^{cl} A_t \ ,$$

(4) $$\frac{dBr^{res}}{dt} = \beta^{res} I_t - \tilde{\alpha}^{res} Br_t^{res} + \gamma^{res} A_t \ ,$$

where Br_t^{imp} are branded impressions, Br_t^{cl} are branded clicks and Br_t^{res} are branded reservations, all at time *t.* I_t is a vector of indicator variables accounting for day-of-week and month (i.e., seasonality), and A_t is from equation (1). The coefficients $\tilde{\alpha}^{\cdots} = (1-\alpha^{\cdots})$ capture the carryover of branded search, where $\alpha^{\cdots}$ is the carryover rate. The coefficients γ^{imp}, γ^{cl}, and γ^{res} reflect the spillover effect from generic search as captured via the impact of awareness, $A_{t,}$. If there is no spillover from generic to branded, the γ coefficients will not be significant.

Bayesian Dynamic Linear Model

We use a Dynamic Linear Model (DLM) implemented in a Bayesian framework to integrate the models in equations (1) to (4). We estimate the model via Markov Chain Monte Carlo (MCMC) methods (West and Harrison 1997).[5] An appealing feature of the DLM is that it simultaneously captures the dynamic evolution of the branded search activities and latent awareness.

Because our data were collected at discrete points of time, i.e., daily, we can rewrite equations (1) through (4) as follows:

$$\text{(5)} \qquad \begin{bmatrix} Br_t^{imp} \\ Br_t^{cl} \\ Br_t^{res} \\ A_t \end{bmatrix} = \begin{bmatrix} \alpha^{imp} & & & \gamma^{imp} \\ & \alpha^{cl} & & \gamma^{cl} \\ & & \alpha^{res} & \gamma^{res} \\ & & & \alpha^{A} \end{bmatrix} \begin{bmatrix} Br_{t-1}^{imp} \\ Br_{t-1}^{cl} \\ Br_{t-1}^{res} \\ A_{t-1} \end{bmatrix} + \begin{bmatrix} d_t^{imp} \\ d_t^{cl} \\ d_t^{res} \\ d_t^{A} \end{bmatrix} + \begin{bmatrix} w_t^{imp} \\ w_t^{cl} \\ w_t^{res} \\ w_t^{A} \end{bmatrix}.$$

where the drift vector, d_t, is given by

$$\text{(6a)} \qquad d_t^{imp} = \sum_{days} I_t^{day} \beta_{day}^{imp} + \sum_{months} I_t^{month} \beta_{months}^{imp}$$

$$\text{(6b)} \qquad d_t^{cl} = \sum_{days} I_t^{day} \beta_{day}^{cl} + \sum_{months} I_t^{month} \beta_{months}^{cl}$$

$$\text{(6c)} \qquad d_t^{res} = \sum_{days} I_t^{day} \beta_{day}^{res} + \sum_{months} I_t^{month} \beta_{months}^{res}$$

$$\text{(6d)} \qquad d_t^{A} = Gen_t^{imp} \beta_{gen}^{imp} + Gen_t^{cl} \beta_{gen}^{cl}.$$

In equation (6a)-(6d), I_t^{day} is an indicator for weekday, e.g, Monday, I_t^{month} is an indicator for month, e.g., July, Gen_t^{imp} are the generic impressions at time t and Gen_t^{cl} are the generic clicks at time t. The potentially correlated error terms, w_t, capture the effect of other factors not included in the model.

[5] An alternative estimation procedure in a classical framework is given by the Kalman Filter procedure. Both methods solve the latent variable problem in a different statistical framework but lead to similar results. See Naik et al. (1998) and Naik et al. (2008) for Kalman Filter and Bass et al. (2007) for a Bayesian DLM application.

In the model given by equations (5) and (6), current generic activity does not influence current branded activity. We cannot distinguish whether a generic search occurred before or after a branded search on any given day. Note that the model specifies past generic search activity to influence current branded search activity through last period awareness. Current generic search activity influences current awareness which, in turn, influences branded activity in the next period.[6]

If awareness were known instead of latent, we could estimate the model in equations 5 and 6 using a traditional time-series approach (e.g., using a autoregressive distributed lag or ARDL model). Though some advertisers do track awareness levels alongside paid search campaigns, this is not routinely done nor are measures typically available on a daily basis. An advantage of the DLM approach is that it allows us to estimate the system of equations without data on awareness. Later, we will compare our DLM approach with a model that does not include latent awareness (i.e., it models the effect of generic search on branded search directly using a lag structure). We specify this alternate model in the section below.

To complete the set up of the DLM, we specify the observation equation linking the state variables to the observed branded search activities – branded impressions, clicks and reservations:

(7) $$\begin{bmatrix} Y_t^{imp} \\ Y_t^{cl} \\ Y_t^{res} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} Br_t^{imp} \\ Br_t^{cl} \\ Br_t^{res} \\ A_t \end{bmatrix} + \begin{bmatrix} v_t^{imp} \\ v_t^{cl} \\ v_t^{res} \end{bmatrix},$$

where the ν_t are potentially correlated error terms.

[6] Recently, search engines have begun to provide data over smaller time intervals (e.g., on an hourly basis). While we cannot investigate this here, with new data it may be possible to extend the model to capture intraday effects.

The model also can be expressed in the formal state space notation of West and Harrison (1997):

(8) $$Y_t = F\theta_t + v_t,$$

(9) $$\theta_t = T\theta_t + d_t + w_t,$$

where Y_t are the branded search activities at time t, i.e., impressions, clicks, and reservations, F is the 3x4 mapping matrix from equation (7), θ_t is the state vector at time t from equation (5), d_t is the drift vector at time t from equation (6), and T is the transition matrix from equation (5). Lastly, we assume that $v_t \sim N(0,\Sigma_v)$ and $w_t \sim N(0, \Sigma_w)$, while α, β, γ, Σ_v and Σ_w, are parameters (vectors and matrices) to be estimated. We estimate the DLM (equations 8 and 9) via sequential Gibbs sampling. (Please see the Appendix for a complete description of the estimation procedure).

Alternative Models

As noted above, one alternative modeling approach would be to use a traditional time series model to capture the dynamics of branded search activity as a function of prior generic search. In such an approach, there would be no latent construct for awareness entering the model. We test whether our DLM with awareness is superior in fit to a model in which generic activity enters as covariates in a traditional autoregressive distributed lag (ARDL) approach with (potentially) correlated errors. To do this, we specify a ARDL model with correlated errors in which branded search activity – i.e., impressions, clicks, and reservations -- is dependent on its own lagged values, seasonal effects, and lagged generic search activity (i.e., generic impressions and generic clicks). This model is given by

(10) $$\begin{bmatrix} Br_t^{imp} \\ Br_t^{cl} \\ Br_t^{res} \end{bmatrix} = \sum_{k=1}^{T_{Br}} \begin{bmatrix} \alpha_{t-k}^{imp} & & \\ & \alpha_{t-k}^{cl} & \\ & & \alpha_{t-k}^{res} \end{bmatrix} \begin{bmatrix} Br_{t-k}^{imp} \\ Br_{t-k}^{cl} \\ Br_{t-k}^{res} \end{bmatrix} + \sum_{k=1}^{T_{Gen}} \begin{bmatrix} \beta_{t-k}^{imp-imp} & \beta_{t-k}^{imp-cl} \\ \beta_{t-k}^{cl-imp} & \beta_{t-k}^{cl-cl} \\ \beta_{t-k}^{res-imp} & \beta_{t-k}^{res-cl} \end{bmatrix} \begin{bmatrix} Gen_{t-k}^{imp} \\ Gen_{t-k}^{cl} \end{bmatrix} + \begin{bmatrix} d_t^{imp} \\ d_t^{cl} \\ d_t^{res} \end{bmatrix} + \begin{bmatrix} \varepsilon_t^{imp} \\ \varepsilon_t^{cl} \\ \varepsilon_t^{res} \end{bmatrix},$$

where, as above, the *Br* terms reflect the various branded search activities and the *Gen* terms correspond to generic search activities. Again, d_t are drift terms, i.e., seasonal indicator variables as above, and α, β, T_{Br} and T_{Gen} are parameter (vectors) to be estimated.[7] As in the DLM model, only past generic search affects current branded search. We estimate the ARDL model via MCMC methods (to have comparable fit statistics across models, i.e., log-marginal density) and use Bayes Factors to determine the optimal lag structure, i.e., T_{Br}^* and T_{Gen}^*.

Test for Reverse Spillover

We have proposed a model to determine whether or not generic search "spills over" to affect branded search and, if so, to what extent. One might now wonder whether generic search is similarly affected by branded search. Since awareness is conceptualized to apply to a brand, not a generic entity, it is not clear how branded search activity would lead to greater generic search activity (via awareness). Nevertheless, we test for the effects of branded search on generic search with a traditional time series approach, i.e., a generic ARDL model with correlated errors analogous to the branded ARDL model described in Equation 10. We test whether past generic search activity and past branded search activity affect current generic search activity. Our generic ARDL model is given by:

$$(11)\quad \begin{bmatrix} Gen_t^{imp} \\ Gen_t^{cl} \\ Gen_t^{res} \end{bmatrix} = \sum_{k=1}^{T_{Gen}} \begin{bmatrix} \chi_{t-k}^{imp} & & \\ & \chi_{t-k}^{cl} & \\ & & \chi_{t-k}^{res} \end{bmatrix} \begin{bmatrix} Gen_{t-k}^{imp} \\ Gen_{t-k}^{cl} \\ Gen_{t-k}^{res} \end{bmatrix} + \sum_{k=1}^{T_{Br}} \begin{bmatrix} \delta_{t-k}^{imp-imp} & \delta_{t-k}^{imp-cl} \\ \delta_{t-k}^{cl-imp} & \delta_{t-k}^{cl-cl} \\ \delta_{t-k}^{res-imp} & \delta_{t-k}^{res-cl} \end{bmatrix} \begin{bmatrix} Br_{t-k}^{imp} \\ Br_{t-k}^{cl} \end{bmatrix} + \begin{bmatrix} d_t^{imp} \\ d_t^{cl} \\ d_t^{res} \end{bmatrix} + \begin{bmatrix} \xi_t^{imp} \\ \xi_t^{cl} \\ \xi_t^{res} \end{bmatrix},$$

where the *Gen* terms and the *Br* terms are the different generic and branded search activities, respectively, d_t are drift terms, i.e., seasonal indicator variables as discussed above, and χ, δ, T_{Br} and T_{Gen} are parameter (vectors) to be estimated.[8] We estimate the generic ARDL model via

---

[7] The notation $\beta^{imp\text{-}cl}$, for example, stands for the effect of generic clicks on branded impressions.

[8] As before, the notation $\delta^{imp\text{-}cl}$, for example, stand for the effect of branded clicks on generic impressions.

MCMC methods and use Bayes Factors for model selection, i.e., the optimal lag structure T_{Br}^{*} and T_{Gen}^{*}. For completeness we also estimate a generic DLM,[9] though we believe it is a poor fit to the situation and unlikely to perform well. Our objective in testing for reverse spillover is to investigate whether branded search activity influences subsequent generic search activity. In other words, is there spillover in both directions or just from generic to branded?

EMPIRICAL APPLICATION

Paid Search Lodging Data

Our data contain aggregate daily information for a paid search campaign for a major lodging chain which wishes to remain anonymous. For each search keyword (e.g., Hotels LA) we have category information (generic or branded), daily information on cost (in $), average position served (given by daily average placement rank, e.g., 2.3), and number of impressions and clicks. These data were provided to the company by the search engines. The company used a third party provider to assemble conversion data (reservations) linked to each keyword. The dataset includes campaign information from both Google and Yahoo!. The Google data run from March 1 to December 20, 2004 and the Yahoo! data are from May 6 to August 31, 2004. In both cases, the campaign included several hundred generic and branded keywords (the exact number is proprietary). Summary statistics are reported in Tables 1 and 2.

Compared to banner ads with click-through rates of 0.2% (Business Week 2007), paid search text ads can deliver high click-through rates (CTR) when the search is based on a branded keyword. On the other hand, if the search is based on a generic keyword, the CTR of 0.3% that we observe is in line with banner ads (Tables 1 and 2). Conversion rates (the percentage of clicks associated with a sale or, in our case, a reservation) also differ substantially. Again, branded search appears to be more effective for the firm. As Tables 1 and 2 show, generic clicks are

[9] Using the DLM specified in equations (4) - (7) we simply replace Br with Gen and vice versa.

more expensive than branded clicks ($0.55 versus $0.18 for the Google data). The difference between generic and branded becomes even more pronounced when looking at cost per reservation ($51.84 versus $2.94 at Google).

These standalone performance statistics could indicate to managers that generic search is a poor investment. In the case of the Google campaign, where 30 percent of the money was spent on generic search, a redistribution of funds towards branded search might seem warranted. If spillover is present and significant, however, it could be inadvisable for the firm to act on these standalone measures.

Before proceeding we note that impressions, clicks and reservations for both generic and branded search exhibit the same daily pattern (see Figure 3 for an illustrative snapshot from branded clicks). For the Google and Yahoo! data the point of highest activity is usually on Monday. Activity declines modestly up to Thursday. Starting on Friday, the weekend brings a steep drop in activity. This pattern is consistent with most online traffic coming from the workplace (Pauwels and Dans 2001). This suggests that indicator variables controlling for day of week effects will be significant components of the model.

Model Comparison – In-Sample

We estimate the DLM and the ARDL models in a Bayesian Framework and compare those using Bayes Factors. In Table 3, we report the results of the model comparisons for the Google data. (Results for the Yahoo! data are similar and available upon request.) We tested three different measures for generic activity: (1) absolute impressions and clicks, (2) log impressions and log clicks, and (3) dollars spent on generic search. We also tested several different lag structures for the ARDL model.

Within the ARDL models, the formulation with one lag term for both branded and generic activity provided the best fit. Note that excluding generic activity (0 Lag Gen formulations) reduced the fit of the ARDL models, regardless of the manner in which generic activity is specified. Next, we found that the models with absolute impressions and clicks provided better fits than the models with log or dollar specifications for generic activity. Thus, it appears superior to have data on actual exposure versus dollars spent. The better fit of the linear versus the log specification for impressions and clicks may suggest that the campaign is not yet operating in the range of diminishing returns. Finally, we also note that the lag selection results – one lag generic and branded fits best – also match-up with our expectations about the search process for a hotel/motel room, which we expect to be relatively short.

The most important comparison in Table 3 is between the best fitting ARDL model and the DLM model. We find that the best ARDL model is rejected in favor of our integrated DLM in all three cases. In particular, the Bayes Factor of 849.6 for the comparison between the DLM and one lag ARDL model (using absolute impressions and clicks) strongly favors the DLM.

Model Comparison – Out-of-Sample

We also assessed the out-of-sample forecast performance for the DLM and ARDL models, again on the Google data. Based on the previous comparisons, we use the absolute number of impressions and clicks as a measurement of generic activity. We estimate all models with two different sample cut-off points: t=100 and t=200 (the first sample has 100 data points, and the second 200). For both cases we generate an out-of-sample forecast for 10 time periods (Bass et al. 2007). In Table 4, we report the mean absolute percentage error (MAPE) across the 10 forecast periods for each branded search activity for each model. In all cases, the DLM outperforms the ARDL models, offering the lowest values for MAPE. The selected lag structure for the ARDL model approach is also confirmed by the holdout test (the best fitting ARDL model uses 1 Lag branded and 1 Lag generic).

In sum, our DLM model outperforms a more traditional ARDL model in-sample as well as out-of-sample. This suggests that incorporating the latent construct of awareness (DLM) offers a superior model versus one in which generic search directly enters the model through a specified lag structure (ARDL). Substantively, we find that measuring generic search activity on the basis of absolute impressions and clicks is preferred to doing so based on dollars spent.

Parameter Estimates

We begin our discussion of the parameter estimates based first on the results from the Google data set (Tables 5a and 5b). Google had a significantly higher level of daily activity than Yahoo! (at least for this search campaign) and also provided a somewhat longer time series. We will also briefly examine the Yahoo! results to corroborate the findings from the Google data (Tables 6a and 6b). In each case, the selection of the reported model was done by Bayes factors; parameters which did not produce an improvement were excluded from the final specification.

Indicator variables for day of week and month. All models include indicator variables to control for differences in search activity by day of week and month. The significant (and therefore retained) covariates are the same for Google and Yahoo!. As expected, the indicator variables capture the lower weekend search activity levels. For each branded search activity we find a negative effect for Saturday (or the start of the weekend) and a positive effect for Monday (or the start of the week). All other indicator variables were not significant and were omitted from the final model. We did not find any significant patterns of monthly seasonality.

Branded Impressions. We turn first to the results for the branded impressions equation reported in Table 5a. We find that lagged branded impressions have a coefficient of 0.75 (α^{imp}).

We find that the latent construct for awareness positively affects current branded impressions ($\gamma^{imp} = 0.0221$). This indicates that higher awareness leads to more searches for keywords that include the brand name. This, in turn, gives rise to more branded impressions. (We discuss the effect of generic search activity on awareness after findings for branded clicks and reservations.)[10]

Branded Clicks. The carryover from branded clicks from the previous period is strongly significant but slightly lower ($\alpha^{cl} = 0.7$) than for branded impressions. Like branded impressions, branded clicks are positively affect by awareness ($\gamma^{cl} = 0.0036$).[11]

Branded Reservations. The search activities which generate revenue from the company's standpoint are online reservations. We find that the carryover from past branded reservations is 0.72 (α^{res}), very similar to the effects discussed above. There is a significant positive effect of awareness on branded reservations ($\gamma^{res} = 0.0002$). While awareness is a dimensionless construct in our model, we do note that the varying coefficients across the three equations are related to mean level for the branded search activity (mean daily branded reservations are 16,753, mean branded clicks are 2,212 and mean branded reservations are 138).[12]

Awareness. In our "leaky-bucket" model formulation, changes in awareness are a function of the decay rate and increases associated with the brand-related exposure that comes from generic search activity. Table 5b reports the parameter estimates and coverage intervals for

[10] We also test for the effect of position. As expected, the position of the impression in the paid search listing is not significant (and therefore does not enter the model). This is because the position of the text ad revealed after the search which triggers the impression.
[11] A test of the effect of position is again not significant. Given that a higher position might be expected to lead to more clicks, the result is somewhat surprising. We believe this may be due to certain features of the campaign in our data set. The company has bid keywords into very similar positions over the course of the campaign. When taken together with the averaging of position over all branded keywords, this leads to very limited variation in the position variable.
[12] Once again, we did not find a significant effect of position, probably for much the same reason as for branded clicks.

this component of the model. We find that roughly 40 percent of current awareness is "carried over" to the next period (carryover rate $\alpha^A = 0.41$, see Table 5b). We do not believe that 60 percent of consumers who knew about a brand yesterday have forgotten that brand by today. In our framework this simply means that the brand is not relevant for the specific search anymore, i.e., awareness of relevance has decreased.

At the heart of our study is the question whether generic search activity "spills over" into branded search activity via awareness. Our modeling results indicate that this is the case. First, as discussed above, awareness has a positive impact on current branded search activity. Second, we find that generic clicks have a strong positive effect on awareness ($\beta_{gen}^{cl} = 39.28$). On the other hand, generic impressions do not have a significant effect on awareness ($\beta_{gen}^{imp} = -0.3$) as the 95 percent coverage interval [-0.50; 0.02] includes zero. This means that simply being exposed to the company's text-ad after a generic search does not "spill over" to increase branded search activity. However, if the user clicks on the ad and visits the company website, this leads – via awareness -- to an increase in branded search activity going forward. We can hypothesize that "inspecting" a brand after a generic click-through might lead the consumer to become aware of the relevance of it for current search goals. The user might search for the brand again, next time using a query that includes the brand name.

Results for the Yahoo! Data

We also estimated the DLM model on the Yahoo! dataset. This provides an interesting validation test across search engines. For example, Yahoo! used a different method to rank sponsored links on its site at the time our data set was collected. The two search engines also differed in site design and appearance and may attract different types of online users. (We have no direct evidence on user differences other than management's belief and industry white papers.) The parameter estimates for the Yahoo! data are reported in Tables 6a and 6b.

A comparison of Tables 5a and 5b (Google) with Tables 6a and 6b (Yahoo!) shows that all of the key findings are corroborated. Among the indicator variables, the effects for Saturday and Monday were found to be similar. As in the Google data, we found no significant seasonality when in monthly effects. Like the Google data, lagged branded activity has carryover coefficients in the 0.7 - 0.8 range (see Table 6a and the parameter estimates for α^{imp}, α^{cl} and α^{res},). In both cases, branded impressions have somewhat higher carryover coefficients than branded clicks or branded reservations. Awareness also significantly impacts branded impressions, clicks and reservations. (Because awareness is dimensionless, the coefficients are not directly comparable.)

The estimated awareness carryover rate for the Yahoo! data is similar in general magnitude (Yahoo! $\alpha^A = 0.32$ vs. Google $\alpha^A = 0.41$) to the one estimated for the Google data. The parameter estimates for generic impressions and clicks in the Yahoo! data parallel the results in the Google data. At both search engines, generic impressions had no significant effect on awareness but generic clicks did. Thus, the estimates from the Yahoo! data corroborate the positive spillover effects found in the Google data.

Testing for Reverse Spillover

We have found that generic search affects branded search positively through awareness. This finding could, however, be due to a general correlation between generic and branded search: on days with high search activity in the category, both generic and branded search might be similarly affected. Because of this alternative explanation, we also investigated whether branded search influences generic search. As we have noted above, the theory underlying the awareness model does not fit when generic activity is the dependent variable because the concept of awareness applies to a brand, not to a generic term. Thus, simply reversing the roles of branded and generic search in our awareness models does not seem appropriate.[13] Instead, we report results from the generic ARDL model in equation (11) in which we include daily and monthly indicator variables along with lagged generic and branded activity as covariates.

Analogous to the results shown above, we model three different dependent variables simultaneously (generic impressions, generic clicks and generic reservations) and estimate the model for both the Google and Yahoo! datasets. Following standard econometric procedures, we also tested for additional lag effects and autocorrelation. In all cases, no lag effects were significant beyond the initial one and autocorrelation was not material.

We find that the best representation of generic search activity is achieved when using one lag generic and no branded search activity. This indicates that branded search activity does not affect generic search activity. Even when lagged branded search activity is included in the model, we find that in no case are any of the coefficient estimates for branded search significant. In Table 7 we report a representative modeling result. The table shows parameter estimates for a model with one lag branded activity for the Google data. All six coefficient means for lagged branded activity lie within their 95 percent coverage intervals. In sum, we failed to find any significant impact of lagged branded search on generic search, indicating that spillover is asymmetric (generic affects branded but not vice versa) and consistent with the modeling framework that we have proposed.

---

[13] For completeness we also estimated the corresponding reverse form of the DLM model. As in the ARDL results in Table 7, in no case did branded search activity have a significant effect on awareness nor did awareness have an effect on generic search activity. This finding holds for both the Google and the Yahoo! datasets.

IMPLICATIONS FOR MANAGING SEARCH CAMPAIGNS

We investigate the implications of our findings based on the Google data. (A similar analysis is straightforward to perform on the Yahoo! data.) If the search campaign is evaluated purely on standalone measures (see Table 1), one might consider limiting spending on generic search to concentrate more on branded search. At average room rates in the range of $85-105 per night, the reservation cost of $51.82 associated with generic keywords might very well produce negative net margins. (Unfortunately, we lack exact cost or margin information for the lodging company in our data set.) However, if we take spillover from generic to branded search into account, spending on generic search should appear relatively more productive – assuming, as we find, that the spillover occurs only in that direction. We now explore how much of the branded search activity in our sample may be due to generic activity. This will permit us to propose a set of adjusted measures that reflect the spillover.

We use an impulse response approach to determine the value of generic search after accounting for spillover into branded search. Since generic impressions were not significant in the model, we will focus on generic clicks. It takes, on average, 95 generic clicks to "generate" one generic reservation (Table 1). We ask how much spillover, on average, these 95 generic clicks generate. (We could also investigate how much spillover one generic click generates – both methods will lead to the same results.) To do this, we shock the system in the first period with 95 generic clicks. Such a shock could come about, for example, if the company decided to "buy" additional generic keywords. (We assume that the company would pay, on average, the same per click as it does for the currently used generic keywords.) Based on the existing data and the model estimates, we calculate how branded search activity is affected by the generic shock in the first period. The results are reported in Table 8.

We find that the spillover effect immediately (i.e., in period 2) leads to an increase in branded activity, as we would expect based on our model. Brand-related exposures that are triggered by a generic search generate awareness for the brand which, in turn, will lead to more branded searches going forward. We see that the effect peaks in period 3 and decays afterwards. By period 7 (6 periods after the initial shock or, in our case, about one week) approximately 80 percent of the additional branded search activity is realized. By period 12 nearly 95 percent of the spillover has been realized and the effect of the initial shock has mostly left the system. This pattern of results fit well with the notion that the search for a hotel/motel room will be short. If our product was, say, a consumer durable like a new car, we might expect the timing of the spillover to differ, i.e., it might be more evenly spread out and not concentrated in the first couple of days.

The results in Table 8 show that the 95 additional generic clicks produced 5.5 reservations. One is a direct result of the generic search while 4.5 are "spillover" reservations from branded search. To get at the total cost of these 5.5 reservations we add the generic cost of \$51.82 to the branded cost of \$13.75 (76.41 branded clicks x \$0.18 per click, see Table 8) – resulting in a total cost of \$65.57. This implies that generic search has a cost per reservation of \$11.90 *after* accounting for spillover. This result is quite different from a cost-per-reservation of \$51.82 based on a stand-alone view.

We emphasize that our impulse response approach provides only an initial estimate of the adjustment needed to take into account the spillover from generic to branded search. Nevertheless, the striking changes in the economics for generic keywords highlight the potential value of using a model-based approach to gauge spillover effects in paid search.

CONCLUSION

Paid search is the fastest growing form of Internet advertising, with paid search campaigns now a key element of the marketing budget for many firms. Our objective is to help researchers and managers better understand and evaluate this form of advertising. In this paper, we examine two categories of paid search advertising – the text ads linked to generic keyword searches and those linked to branded keyword searches. We focus our analysis on the relationship between generic on branded search. More specific, does spillover occur from generic to branded search?

To study these questions, we developed a modeling framework based on the well-established theory that exposure to brand-related information can create awareness of the brand. Applying the framework to paid search, our approach holds that the exposure to brand-related information due to paid search (i.e., impressions from text ads, company web pages after click-through) can create awareness and spill-over into future branded search activity.

Following Naik et al. (1998, 2006) and Bass et al. (2007), we model the dynamic evolution of awareness using a "leaky-bucket" type model. However, in place of traditional (indirect) measures of brand-related exposures used (e.g., GRP or dollar expenditures), we use the direct measures of impressions and clicks that occur as a result of search activity. We combine the awareness model with a dynamic branded search activity model and estimate the two components together in a Bayesian framework.

We apply the modeling approach to daily data from paid search campaigns on Google and Yahoo! that were run by a major lodging chain in 2004. We model three measures of branded search activity as dependent variables simultaneously – impressions, clicks, and reservations. We find that generic search activity, specifically generic clicks, has a significant positive effect on awareness. In turn, we find that awareness significantly influences branded search activities (i.e., impressions, clicks and reservations). Thus, we conclude that generic search activity spills over to influence (and increase) branded search activity. We compare our DLM with a time-series model that does not use awareness and instead models the effect of generic search on branded search directly within a lag structure. We find that our DLM outperforms the alternate model in-sample as well as out-of-sample.

We also investigated whether branded search has a similar effect on generic search. In no instance did any form of branded search have a significant impact on generic search. This indicates that the spillover between generic and branded is asymmetric, moving from generic to branded but not vice versa.

The asymmetric spillover between generic and branded search points up the need to adjust currently used standalone financial performance measures in paid search advertising. Using the Google data set for illustration, we show how to adjust these measures for the asymmetric spillover by using an impulse response approach. The unadjusted measures available to management (e.g., cost per click and cost per reservation) show very poor economics for generic search when compared to branded. After adjusting the measures for the estimated spillover effects, the performance of generic versus branded keywords move much closer together.

While managers involved in paid search recognize that generic keywords may provide additional value, they lack model-based quantitative methods for assessing how much of an adjustment it is appropriate to make. We hope that our integrated DLM model, when applied to the type of aggregate daily data available to managers, will improve understanding and assessment of the effects of (and returns from) paid search advertising.

A key goal of this research is to develop and test a model based on the type of information that search campaign managers use on a day-to-day basis. This means that the approach would be able to aid managers without the need to obtain additional forms of data to estimate the model. A drawback of this approach, however, is that it does not track the search activity of individual users. Thus, we have refrained from attempting to develop a comprehensive theory of Internet search behavior as applied to paid search advertising. Instead, we propose the more general conceptualization of awareness of relevance for the search, distinguishing generic terms from branded by the notion that only a brand that has awareness of relevance will be searched using a branded term. Clearly, developing and testing a theory-driven model of individual level paid search should be a critical topic for future research.

Though we have data from both Google and Yahoo!, we lack information that would enable us to explore, in more detail, the differences between the search engines and their user bases. Given the innovation and competition currently taking place among search engines, we feel that this also represents an important topic for future research.

Links to the advertiser's web site may also appear in the organic search results that are returned along with the paid search ads. Unfortunately, we have no information on organic search in our data set. This means that we are unable to assess whether the absence of paid text ads would lead more users to click on organic links to the company's web site. To the extent that this occurs, the contribution of generic search to branded search will be overstated in our approach. We leave this as an important topic for future research. Additionally, one could imagine that an increase in awareness not only leads to an increase in branded search, but also to an increase in organic clicks as well as direct visits to the website, i.e., the visitors types in the

URL of the website directly. These effects could increase or decrease the net effectiveness of generic search. We believe that these will be important issues for future research.

In this paper we evaluate performance at the level of keyword categories (i.e., aggregations across keywords sharing certain characteristics). Managers also have the capability to make changes at the individual keyword level if they so choose. Developing models to aid managers in evaluating performance of individual keywords would also be worthwhile.

Lastly, we use aggregate data and a latent construct in our modeling approach. We have labeled this latent construct awareness and believe that, in so doing, it is consistent with both the overall approach and our empirical findings. An important next step would be to obtain actual measures for awareness, e.g., by using survey methods, to validate the construct in future modeling applications.

REFERENCES

Bass, Frank, Norris Bruce, Sumit Majundar and B.P.S. Murthi (2007), "A Dynamic Bayesian Model of the Advertising-Sales Relationship," *Marketing Science,* 26(2), 179-195.

Bradlow, Eric T. and David C. Schmittlein (2000), "The Little Engines that Could: Modeling the Performance of World Wide Web Search Engines," *Marketing Science,* 19 (1), 43-62.

Business Week (2007), "So many Ads, so Few Clicks", November 12th, 38.

Chatterjee, Patrali, Donna L. Hoffman, and Thomas P. Novak (2003), "Modeling the Clickstream," *Marketing Science,* 22 (4), 520-541.

Chen, Yongmin and Chuan He (2008), "Paid Placement: Advertising and Search on the Internet," *Working Paper,* University of Colorado Boulder.

Cho, Chang-Hoan and Hongsik J. Cheon (2004), "Why do People avoid Advertising on the Internet?," *Journal of Advertising,* 33 (4), 89-97.

Danaher, Peter and Guy Mullarkey (2003), "Factors Affecting Online Advertising Recall: A Study of Students," *Journal of Advertising Research,* 43 (3), 252-267.

Dreze, Xavier and Francois-Xavier Hussherr (2004), "Internet Advertising: Is Anybody Watching?," *Journal of Interactive Marketing*, 17 (4), 8-23.

Edelman, Benjamin and Michael Ostrovsky (2007), "Strategic Bidder Behavior in Sponsored Search Auctions," *Decision Support Systems,* 43 (1) 192-198.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz (2007), "Internet Advertising and the Generalized Second Price Auction," *American Economic Review,* 97 (1), 373-399.

Ghose, Anindya and Sha Yang (2008), "An Empirical Analysis of Search Engine Advertising: Sponsored Search and Cross-Selling in Electronic Markets," *Working Paper*, New York University.

Johnson, Eric J., Wendy Moe, Peter Fader, Steven Bellman, and Gerald L. Lohse (2004), "On the Depth and the Dynamics of Online Search Behavior," *Management Science*, 50 (3), 299-308.

Manchanda, Puneet, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta (2006), "The Effects of Banner Advertising on Internet Purchasing," *Journal of Marketing Research*, 43 (1), 98-108.

Moore, Robert, Claire Allison Stammerjohan, and Robin Coulter (2005), "Banner Advertiser-Web Site Congruity and Color Effects on Attention and Attitudes," *Journal of Advertising*, 34 (2), 71-84.

Naik, Prasad A., Murali K. Mantrala, and Alan G. Sawyer (1998), "Planning Media Schedules in the Presence of Dynamic Advertising Quality," *Marketing Science,* 17 (3), 214-235.

Naik, Prasad A., Kalyan Raman, and Shuba Srinivasan (2008), "Managing Corporate and Product Brands," *Working Paper,* University of California, Davis.

Nerlove, Marc and Kenneth J. Arrow (1962), "Optimal Advertising Policy under Dynamic Conditions," *Economica,* 29 (114), 129-142.

Novak, Thomas P. and Donna L. Hoffman (2000), "Advertising and Pricing Models for the Web," in *Internet Publishing and Beyond,* Cambridge: MIT Press.

PricewaterhouseCoopers (2007), IAB Internet Advertising Revenue Report, http://www.iab.net/media/file/resources_adrevenue_pdf_IAB_PwC_2007_Final.pdf

Telang, Rahul, Peter Boatwright, and Tridas Mukhopadhyay (2004), "A Mixture Model for Internet Search Engine Visits," *Journal of Marketing Research*, 41 (2), 206-214.

Tellis, Gerry (2004), *Effective Advertising,* Sage, Thousand Oaks, CA.

Willbur, Kenneth C. and Yi Zhu (2008), "Click Fraud," *Marketing Science*, forthcoming.

West, Mike and Jeff Harrison (1997). *Bayesian Forecasting and Dynamic Models*, Springer, New York, NY.

Table 1: Descriptive Statistics for Google Data

	Impressions	Clicks	Reservations	Cost
Generic	37,059,020	98,162	1,033	$ 53,549.52
Branded	4,925,351	673,971	40,671	$119,498.50
Total	41,984,371	772,133	41,704	$173,048.02

	Impressions/ Day	Clicks/ Day	Reservations/ Day	Cost/ Day	Average Position
Generic	126,051	334	4	$182.14	5.55
Branded	16,753	2,292	138	$406.46	1.55
Total	142,804	2,626	141	$588.60	

	Click-through Rate	Conversion Rate	Cost/ Click	Cost/ Reservation
Generic	0.26%	1.05%	$ 0.55	$ 51.84
Branded	13.68%	6.03%	$ 0.18	$ 2.94

Table 2: Descriptive Statistics for Yahoo! Data

	Impressions	Clicks	Reservations	Cost
Generic	2,118,555	5,608	108	$ 2,372.42
Branded	3,378,749	361,828	25,889	$118,024.09
Total	5,497,304	367,436	25,997	$120,396.51

	Impressions/ Day	Clicks/ Day	Reservations/ Day	Cost/ Day	Average Position
Generic	7,206	19	0.4	$ 8.07	4.92
Branded	11,492	1,231	88	$401.44	2.16
Total	18,698	1,250	88	$409.51	

	Click-through Rate	Conversion Rate	Cost/ Click	Cost/ Reservation
Generic	0.26%	1.93%	$ 0.42	$ 21.97
Branded	10.71%	7.16%	$ 0.33	$ 4.56

Table 3: Model Comparison – In-Sample Fit Measures

	Generic Activity measured as		Log Marginal Density		
			Impressions & Clicks	**Log(Imp.) & Log(Cl.)**	**Dollar Spending**
DLM Model			**-5,919.2**	-5,953.6	-5,997.5
ARDL Model	*1 Lag Br*	*0 Lag Gen*	-6,791.8	-6,791.8	-6,791.8
	1 Lag Br	*1 Lag Gen*	-6,768.8	-6,782.4	-6,802.3
	1 Lag Br	*2 Lags Gen*	-6,775.1	-6,787.0	-6,848.2
	2 Lags Br	*0 Lag Gen*	-7,458.0	-7,458.0	-7,458.0
	2 Lags Br	*1 Lag Gen*	-7,420.0	-7,438.2	-7,450.4
	2 Lags Br	*2 Lags Gen*	-7,427.1	-7,444.0	-7,467.2

	Generic Activity measured as		Bayes Factors*		
			Impressions & Clicks	**Log(Imp. & Log(Cl.)**	**Dollar Spending**
DLM Model			N/A	34.4	78.3
ARDL Model	*1 Lag Br*	*0 Lag Gen*	872.6	872.6	872.6
	1 Lag Br	*1 Lag Gen*	849.6	863.2	883.1
	1 Lag Br	*2 Lags Gen*	855.9	867.8	929.0
	2 Lags Br	*0 Lag Gen*	1,538.8	1,538.8	1,538.8
	2 Lags Br	*1 Lag Gen*	1,500.8	1,519.0	1,531.2
	2 Lags Br	*2 Lags Gen*	1,507.9	1,524.8	1,548.0

**In relation to the best Model, i.e., the full model that uses impressions and clicks as measures of generic activity.*

Table 4: Model Comparison – Forecast Performance

			MAPE (t=100)*		
			Impressions	**Clicks**	**Reser-vations**
DLM Model			**0.1113**	**0.0139**	**0.0026**
ARDL Model	*1 Lag Br*	*0 Lag Gen*	0.1933	0.0274	0.0037
	1 Lag Br	*1 Lag Gen*	0.1128	0.0165	0.0035
	1 Lag Br	*2 Lags Gen*	0.1130	0.0167	0.0037
	2 Lags Br	*0 Lag Gen*	0.1289	0.0447	0.0042
	2 Lags Br	*1 Lag Gen*	0.1195	0.0365	0.0038
	2 Lags Br	*2 Lags Gen*	0.1416	0.0409	0.0039

			MAPE (t=200)*		
			Impressions	**Clicks**	**Reser-vations**
DLM Model			**0.1341**	**0.0107**	**0.0020**
ARDL Model	*1 Lag Br*	*0 Lag Gen*	0.1729	0.0261	0.0030
	1 Lag Br	*1 Lag Gen*	0.1567	0.0242	0.0028
	1 Lag Br	*2 Lags Gen*	0.1734	0.0249	0.0029
	2 Lags Br	*0 Lag Gen*	0.1780	0.0436	0.0066
	2 Lags Br	*1 Lag Gen*	0.1581	0.0391	0.0056
	2 Lags Br	*2 Lags Gen*	0.1667	0.0441	0.0058

** We base forecast performance on 2 scenarios: using the first 100 data points (t=100) and using the first 200 data points (t=200). In each scenario we forecast the next 10 periods and compare models.MAPE is Mean Absolute Percent Error.*

Table 5a: DLM Parameter Estimates for the Google Data Set

			Dependent Variable		
			Branded Impressions*	**Branded Clicks***	**Branded Reser-vations***
Indicator Variables	*Saturday*	$\beta_{we}^{\cdots}$	-3,679.1	-391.8	-18.9
	Monday	$\beta_{wk}^{\cdots}$	4,907.8	694.8	35.1
Lagged Branded Activity	*Impressions*	α^{imp}	.75	---	---
	Clicks	α^{cl}	---	.70	---
	Reservations	α^{res}	---	---	.72
Awareness		γ^{imp}	0.0221	---	---
		γ^{cl}	---	0.0036	---
		γ^{res}	---	---	0.0002

** We report the estimate means and all reported estimates are within a 95% coverage interval that does not contain 0.*

We use Bayes factors for model selection. We only present the results for the best fitting model. All non-significant covariates, e.g., Tuesday, have been excluded from the model.

Table 5b: DLM Parameter Estimates for the Google Data Set

			Dependent Variable: Awareness		
			Mean	**Coverage Interval**	
Lagged Awareness		α^{A}	0.41	0.30	0.53
Generic Activity	*Impressions*	β_{gen}^{imp}	-0.30*	-0.50	0.02
	Clicks	β_{gen}^{cl}	39.28	26.90	52.05

** Estimate is <u>not</u> significant, i.e., the 95% coverage interval contains 0.*

Table 6a: DLM Parameter Estimates for the Yahoo! Data Set

			Dependent Variable		
			Branded Impressions*	**Branded Clicks***	**Branded Reser-vations***
Indicator Variables	*Saturday*	$\beta_{we}^{\cdots}$	-4,080.9	-422.6	-14.6
	Monday	$\beta_{wk}^{\cdots}$	5,529.8	552.8	26.8
Lagged Branded Activity	*Impressions*	α^{imp}	.81	---	---
	Clicks	α^{cl}	---	.71	---
	Reservations	α^{res}	---	---	.70
Awareness		γ^{imp}	0.184	---	---
		γ^{cl}	---	0.034	---
		γ^{res}	---	---	0.002

** We report the estimate means and all reported estimates are within a 95% coverage interval that does not contain 0.*

We use Bayes factors for model selection. We only present the results for the best fitting model. All non-significant covariates, e.g., Tuesday, have been excluded from the model.

Table 6b: DLM Parameter Estimates for the Yahoo! Data Set

			Dependent Variable: Awareness		
			Mean	**Coverage Interval**	
Lagged Awareness		α^{A}	0.32	0.20	0.44
Generic Activity	*Impressions*	β_{gen}^{imp}	0.20*	-0.12	0.71
	Clicks	β_{gen}^{cl}	21.70	11.88	31.46

** Estimate is <u>not</u> significant, i.e., the 95% coverage interval contains 0.*

Table 7: ARDL Model Parameter Estimates for the Google Data Set

			Dependent Variable		
			Generic Impressions	**Generic Clicks**	**Generic Reser-vations**
Indicator Variables	*Saturday*	$\beta_{we}^{\cdot\cdot}$	-643.7*	-37.9	0.2*
	Monday	$\beta_{wk}^{\cdot\cdot}$	814.9*	51.8	1.0
Lagged Generic	*Impressions*	χ^{imp}	.90	---	---
Activity	*Clicks*	χ^{cl}	---	.87	---
	Reservations	χ^{res}	---	---	.23
Lagged Branded	*Impressions*	δ^{imp}	-0.9253*	-0.0004*	0.0000*
Activity	*Clicks*	δ^{cl}	11.1508*	0.0179*	0.0009*

** Estimates are <u>not</u> significant, i.e., the 95% coverage interval contains 0.*

Table 8: Effect of a 95 Generic Click Shock for Google Data

		Time Period (t) since shock*					
		2*	**3**	**4**	**7****	**12*****	**294**
Effect	*Br. Impressions*	82.6	96.6	87.1	43.6	15.5	0.0
(Period t)	*Br. Clicks*	13.3	14.9	12.8	5.3	1.3	0.0
	Br. Reservations	0.7	0.8	0.7	0.3	0.1	0.0
Effect	*Br. Impressions*	82.6	179.2	266.3	438.1	542.3	574.9
(Cumulative)	*Br. Clicks*	13.3	28.2	41.0	63.5	73.3	76.4
	Br. Reservations	0.7	1.5	2.2	3.6	4.3	4.5
Effect	*Br. Impressions*	14.4%	31.2%	46.3%	76.2%	94.3%	100%
(% of total)	*Br. Clicks*	17.4%	36.9%	53.7%	83.1%	95.9%	100%
	Br. Reservations	15.6%	33.3%	48.9%	80.0%	95.6%	100%

** We pick 95 generic clicks as, on average, 95 generic clicks lead to one generic reservation. We shock the system in period 1 and investigate the effect of the shock through the data until t=294.*

*** and *** We pick two time periods for exemplary purposes: t=7 as 80% of the total effect has materialized in clicks and reservations and t=12 as 95% of the total effect has materialized on all variables.*

Figure 1: Search Results Page Example for Google

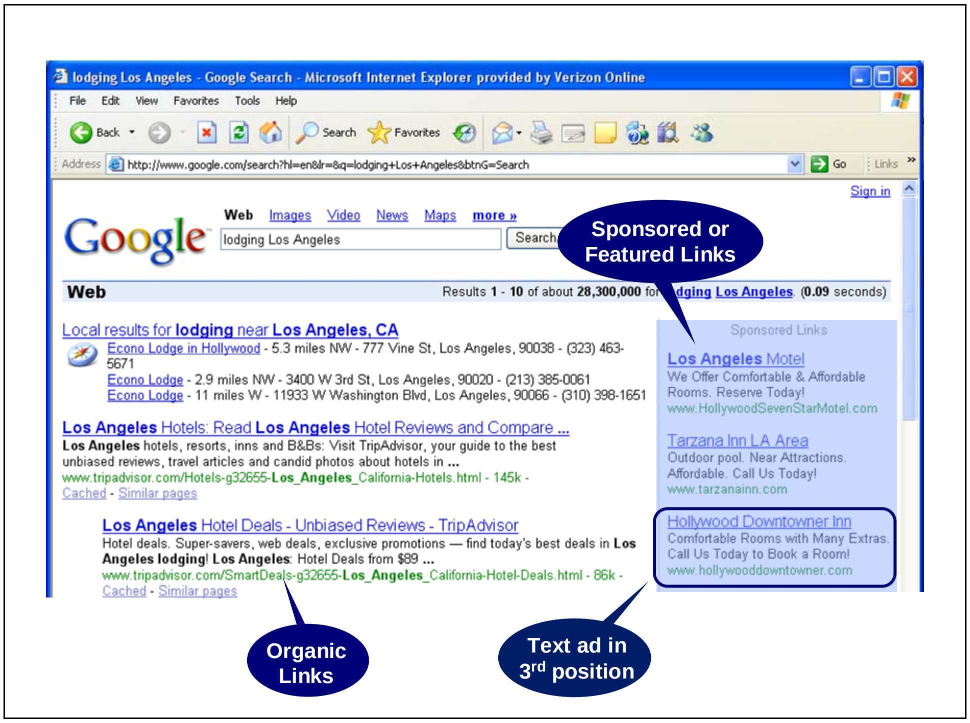

Figure 2: Illustration of the Search Process

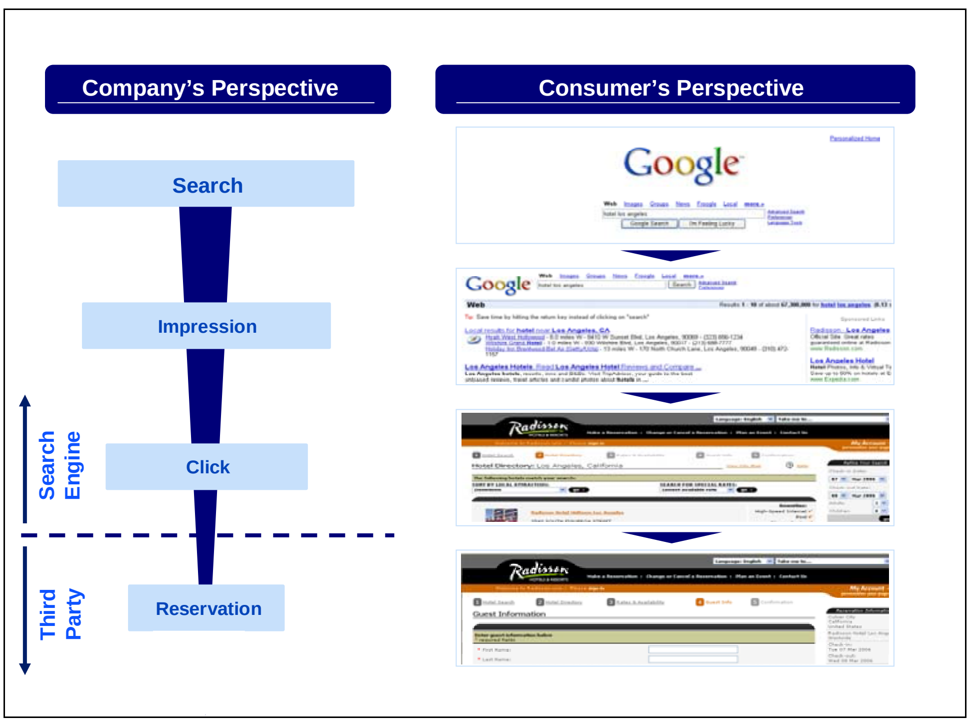

Figure 3: Daily Branded Click Counts for Google Data (Three-Week Period)

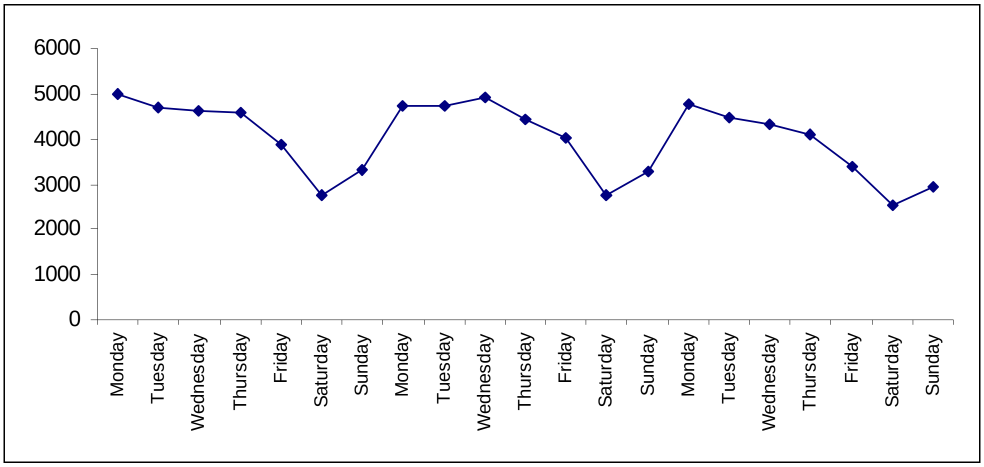

APPENDIX

Description of the Sampling Procedure

To estimate the model we assume prior distributions for all parameters – (1) a prior distribution for the parameters of the transition matrix T, α and γ, (2) prior distributions for the effects of drift parameters β, and (3) prior distributions for the observation and transition variances, Σ_v and Σ_w. In order to use Gibbs sampling, we employ non-informative conjugate priors. In addition, we need to specify an initial guess for the state vector θ_0. (For further details on the Bayesian DLM, please see West and Harrison 1997.)

We assume that the priors for Σ_v and Σ_w are independent and inverse Wishart. Given the independence assumption of the variances, the posteriors for Σ_v and Σ_w are also inverse Wishart distributions with the known properties of the conjugate model. For conjugacy to hold we pick independent Gaussian distributions as priors for α, β and γ, and the posteriors for α, β and γ, are Gaussian following standard theory.

We use Bayesian techniques to estimate our model. We follow the well established DLM methodology of West and Harrison (1997). We use a Forward-Filtering and Backward-Smoothing algorithm (e.g., Bass et al. 2007) to sample the latent state, which, in our case, is awareness. Conditional on Y, Σ_v, Σ_w, α, β and γ the distribution of Φ, $p(\Phi\,|\,\Sigma_v, \Sigma_w, \alpha, \beta, \gamma, \Psi)$, follows the standard normal DLM with a known covariance matrix, where $\Psi=\{Y(1),\ldots,Y(T)\}$ and $\Phi=\{\theta_1,\ldots,\theta_T\}$. We employ the Gibbs sampling procedure to estimate the remaining variables conditional on the latent state and the data. Our sampler has the following three steps:

Step 1: We sample $\theta(t)$ from the posterior distribution $p(\theta(t)|D(t))$ based on the forward filtering algorithm (see West and Harrison 1997) for all $t = 1,\ldots,T$, where D(t) is the set of all information available at time *t*.

Step 2: We sample $\theta(t\text{-}1)$ from the posterior distribution $p(\theta(t\text{-}1)|G(t),D(t))$ based on the backward smoothing algorithm (see West and Harrison 1997) for all $t = 1,\ldots,T$. The resulting draws from step 1 and step 2, $\Phi=\{\theta(0),\ldots,G(t)\}$, are draws from the full posterior distribution of the latent state.

Step 3: We sample Σ_v, Σ_w, α, β, and γ stepwise from the posterior distribution $p(\Sigma_v, \Sigma_w, \alpha, \beta, \alpha \mid \Psi, \Phi)$ where $\Psi=\{Y(1),\ldots,Y(T)\}$ and $\Phi=\{\theta(0),\ldots,\theta(T)\}$. We can do this as the error vectors of the observation equation (5) and the transformation equation (6) are assumed to be independent. Additionally, dependent on the states and the data, α, β and γ are also independent. We assume that both error vectors are normally distributed. Based on these assumptions we can employ the Gibbs sampler and sample Σ_v, Σ_w, α, β, and γ separately.

Details on drawing the coefficients of the Transition Matrix T (α and γ)

We allow for correlated errors in the state equation. Thus, we cannot simply sample the next transition matrix T, i.e., its elements α and γ. Instead, we employ conditional Normal Theory to generate the next transition matrix T. We show, exemplary for α^{imp}, how this is done.

We want to draw α^{imp} conditional on α^{cl}, α^{res}, β, γ, Σ_v, Σ_w, Ψ and Φ. We know that:

(1)
$$\theta_t = \begin{bmatrix} \alpha^{imp} & & & \gamma^{imp} \\ & \alpha^{cl} & & \gamma^{cl} \\ & & \alpha^{res} & \gamma^{res} \\ & & & \alpha^{A} \end{bmatrix} \theta_{t-1} + d_t + w_t \cdot$$

We can rewrite this as:

(1a)
$$\tilde{\theta}_t = \begin{bmatrix} \alpha^{imp} & & & 0 \\ & 0 & & 0 \\ & & 0 & 0 \\ & & & 0 \end{bmatrix} \theta_{t-1} + w_t \,,$$

where

$$\begin{bmatrix} \tilde{\theta}_t^1 \\ \tilde{\theta}_t^2 \\ \tilde{\theta}_t^3 \\ \tilde{\theta}_t^4 \end{bmatrix} = \begin{bmatrix} \theta_t^1 \\ \theta_t^2 \\ \theta_t^3 \\ \theta_t^4 \end{bmatrix} - \begin{bmatrix} 0 & & & \gamma^{imp} \\ & \alpha^{cl} & & \gamma^{cl} \\ & & \alpha^{res} & \gamma^{res} \\ & & & \alpha^{A} \end{bmatrix} \begin{bmatrix} \theta_{t-1}^1 \\ \theta_{t-1}^2 \\ \theta_{t-1}^3 \\ \theta_{t-1}^4 \end{bmatrix} - \begin{bmatrix} d_t^{imp} \\ d_t^{cl} \\ d_t^{res} \\ d_t^{A} \end{bmatrix}.$$

We partition Σ_w as follows:

$$\Sigma_w = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix},$$

where Σ_{11} is a (1x1), Σ_{12} is a (1x3), Σ_{21} is a (3x1), Σ_{22} is a (3x3) matrix.

Based on conditional Normal Theory it follows that:

(2) $$\tilde{\theta}_t^1 \mid \ldots \sim N(\mu_1, \Sigma_1),$$

where

$$\mu_1 = \alpha^{imp}\theta_{t-1}^1 + \Sigma_{12}\Sigma_{22}^{-1}\left(\begin{bmatrix} \tilde{\theta}_t^2 \\ \tilde{\theta}_t^3 \\ \tilde{\theta}_t^4 \end{bmatrix}\right), \text{ and}$$

$$\Sigma_1 = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}.$$

Picking a uninformative normal prior $p(\alpha^{imp})$, the posterior of $p(\alpha^{imp}|\ldots)$ is given by:

(3) $$p(\alpha^{imp} \mid \ldots) \sim p(\tilde{\theta}_t^1 \mid \ldots)\, p(\alpha^{imp}).$$

We sample from Equation 3 by Gibbs sampling. All other coefficients of the transition matrix T can be drawn using the same method.