

ASHISH AGARWAL, KARTIK HOSANAGAR, and MICHAEL D. SMITH*

The authors evaluate the impact of ad placement on revenues and profits generated from sponsored search. Their approach uses data generated through a field experiment for several keywords from an online retailer's ad campaign. Using a hierarchical Bayesian model, the authors measure the impact of ad placement on both click-through and conversion rates. They find that while click-through rate decreases with position, conversion rate increases with position and is even higher for more specific keywords. The net effect is that, contrary to the conventional wisdom in the industry, the topmost position is not necessarily the revenue- or profit-maximizing position. The authors' results inform the advertising strategies of firms participating in sponsored search auctions and provide insight into consumer behavior in these environments. Specifically, they help correct a significant misunderstanding among advertisers regarding the value of the top position. Furthermore, they reveal potential inefficiencies in current auction mechanisms that search engines use. The authors' results also reveal the information search strategies that consumers use in sponsored search and provide evidence of recency bias for immediate purchases.

Keywords: sponsored search, ad placement, hierarchical Bayesian estimation, online advertising, online auctions, search engine marketing

Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets

Internet advertising spending is currently growing faster than any other form of advertising and is expected to grow from \$23.4 billion in 2008 to \$34 billion in 2014 (Hallerman 2009). Of this ad spending, 40% occurs on sponsored searches, in which advertisers pay to appear alongside the regular search results of a search engine. Most search engines, including Google, Yahoo, and MSN, use auctions to sell their ad space inventory. In these auctions, advertisers submit bids on specific keywords based on their willingness to pay for a click from a consumer searching on that (or a closely related) keyword.

Search engines use a combination of the submitted bid and past click performance to rank order the advertisements. Sponsored search is unique compared with offline advertising and other forms of online advertising because it is presumed to occur close to a user's purchase decision and is matched on the basis of the user's stated information need (Hosanagar and Cherapanov 2008). As a result, many advertisers spend a greater share of their advertising budgets on search engine marketing and are often engaged in intense bidding wars to win the top slots in the list of sponsored results (Goodman 2006; Steel 2007). The following quote posted on a search engine forum succinctly summarizes this thinking:

I believe that people who think it's better to be anything other than #1 are just fooling themselves The fact [is] that you'll get 3 1/2 times more traffic being #1 as opposed to #2, and the numbers keep sliding from there.¹

*Ashish Agarwal is Assistant Professor of Information Management, McCombs School of Business, University of Texas at Austin (e-mail: ashish.agarwal@mcombs.utexas.edu). Kartik Hosanagar is Associate Professor of Internet Commerce, The Wharton School, University of Pennsylvania (e-mail: kartikh@wharton.upenn.edu). Michael D. Smith is Professor of Information Systems and Marketing, Heinz College, Carnegie Mellon University (e-mail: mds@cmu.edu). The authors thank the two anonymous *JMR* reviewers for their valuable comments and suggestions. Fred Feinberg served as associate editor for this article.

¹The quote is based on analysis of click-through rates (CTRs) observed in the top ten algorithmic search positions in AOL's data set. However, similar thinking is prevalent for sponsored search as well.

The rationale behind these bidding wars for top positions is that click-through rate (CTR) typically decreases exponentially with ad position, and thus the top few positions receive the majority of clicks (see Feng, Bhargava, and Pennock 2007).² Because of these CTR observations, most advertisers aggressively seek the topmost positions in their bidding, and search engine marketing firms that offer bidding services to advertisers often provide guarantees to clients of securing the top positions. However, there have been few formal studies of the impact of ad position on CTRs, conversion rates (i.e., the likelihood that a consumer will buy a product after clicking an advertisement), and advertising costs.³

Moreover, there is conflicting theory regarding the impact of ad position on consumers' postclick behavior. Some studies show that prescreening information is irrelevant in subsequent search behavior, which in turn suggests that conversion rate may be independent of ad position (Chakravarti, Janiszewski, and Ulkumen 2006). Other studies show that primacy affects brand and product recall, which in turn suggests that conversion rates may actually decrease with position. Finally, some studies show that in sequential choice environments, consumers are disproportionately influenced by the attractiveness of the most recently evaluated product. This implies that lower positions may have higher conversion rates when consumers click them. Empirical evidence in sponsored search environments is limited. Brooks (2004) finds that conversion rate increases with position for several keywords, but he only reports average values and does not control for keyword attributes and other sources of heterogeneity. In contrast, Ghose and Yang (2009) find that conversion rate decreases with position. However, their results are aggregated over a large number of product categories, and they do not study differences across the top few positions, which is the focus of the current work. Furthermore, to the best of our knowledge, there have been no studies that compare an ad position's impact on clicks and conversions and on advertising costs.

In this article, we address this question by analyzing how ad position in sponsored search affects an advertiser's revenues and overall profits. We use a field experiment to generate a unique panel data set of daily clicks, orders, and cost for multiple keywords in the sponsored search ad campaign of an online retailer and then use a hierarchical Bayesian model to analyze the probabilities of clicking and ordering in this environment. Our findings suggest that, contrary to conventional wisdom, the topmost positions for keywords in our data set are associated with lower profits than are lower (and less expensive) positions. Our results confirm that ad CTR decreases with position. However, we find that the conversion rate increases with position. We also find

that revenue increases with ad position for keywords associated with more specific search. For nonspecific keywords, the revenue decreases with position. However, the costs are disproportionately higher in the top positions, resulting in higher profits at lower positions.

Our research makes several contributions. First, we provide key managerial insights for advertisers. A common assumption in the industry is that the value of a click from a sponsored search campaign is independent of the position of the advertisement. Our results indicate that this is not true. Rather, a click from an advertisement at the top position may have lower expected revenue than a click from the same advertisement placed lower in the list of advertisements. As a result, we find that the top ad positions do not necessarily maximize advertiser revenues or profits, and thus advertisers should revisit the assumptions driving current bidding wars for the top ad positions.

Second, our results highlight potential inefficiencies in the rules commonly used in sponsored search auctions and suggest the need for further investigation of these pricing mechanisms. If advertisers with the best combination of bid and CTR are assigned the top position but lower positions would generate higher revenues for those advertisers, current auction rules may be doing them a disservice. Our results suggest that using CTR and other click-oriented measures alone to determine ad ordering may not be sufficient and that more conversion-oriented metrics for ad ordering and pricing can help increase market efficiency.

Finally, we provide insight into consumer behavior in sponsored search environments. We find that CTRs decrease with ad position, which is consistent with prior findings of decay in user attention (Ansari and Mela 2003; Brooks 2004; Hoque and Lohse 1999). We find that buying consumers are more likely to visit lower positions, resulting in an increase in conversion rate with position. We also find that buying consumers with specific queries (high specificity) may prefer to buy from the same advertisement when it is ranked in a lower position. This suggests that buying consumers with specific search queries may have a recency bias in their search behavior.

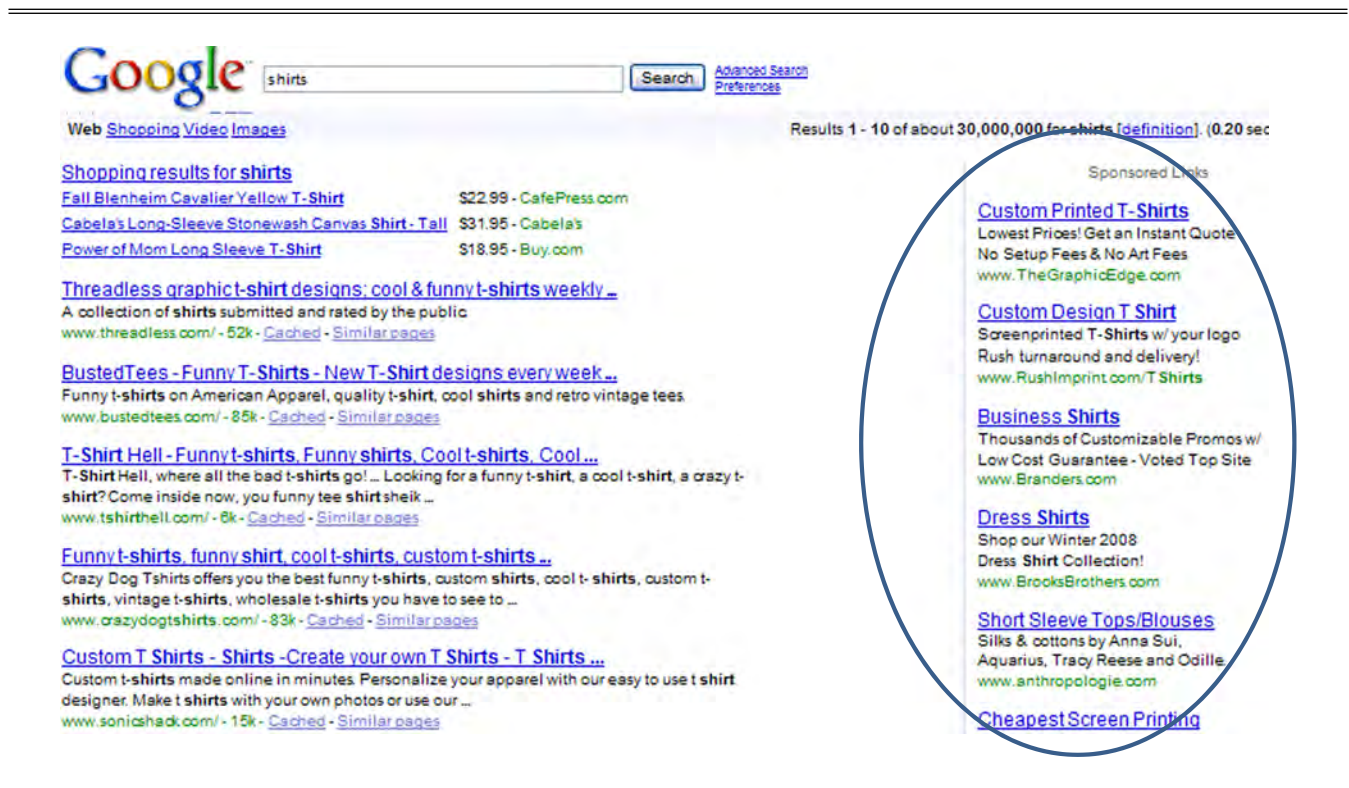
It is important to note that we conducted our study from the perspective of transactional revenues and profit. We do not consider nontransactional benefits, such as increased product or brand awareness, in our analysis. Thus, although we believe our results are applicable to a wide range of industries, they will be less applicable in industries in which the goal of the advertisement is primarily to increase exposure, awareness, or branding.⁴ In additional robustness tests, we consider spillover effects—namely, scenarios in which advertisements corresponding to some keywords create awareness about the retailer such that consumers return to the retailer's website at a later point and ultimately make a purchase. However, we find that the spillover value for such advertisements is not affected by ad position, suggesting that our results are robust to spillover effects.

²Similarly, a recent study (Ruby 2010) found that the top position receives as many clicks as the next four positions and many more clicks than Positions 6–20.

³The term "conversion" is more commonly used in the industry than "order" because the definition of successful customer acquisition varies by firm. For example, for some firms, such as a free e-mail service provider, user registration is referred to as a conversion.

⁴Surveys of online advertisers indicate that 99% of advertisers use search engine advertising to drive direct transactional benefits such as immediate sales or profits (Kitts et al. 2005).

Figure 1
SEARCH RESULTS



SPONSORED SEARCH BACKGROUND

When a consumer enters a search query—for example, “shirts”—the search engine displays algorithmic (i.e., regular) and sponsored search results, as we show in Figure 1. The algorithmic results are determined according to their relevance to the query. The sponsored results are ranked according to continuous real-time auctions run by the search engines. Advertisers bid on sponsored search keywords of relevance to them. Upon receiving a query, the search engine identifies the advertisers bidding on closely related keywords and uses data on bids and past click performance of advertisements to rank order the advertisements that appear in the list of sponsored results.

An advertiser pays the search engine only when the consumer clicks on the advertiser’s advertisement. The cost per click (CPC) is determined using a generalized second price auction mechanism; that is, whenever a user clicks on an advertisement in position k , the advertiser pays an amount equal to the minimum bid needed to secure that position (Lahaie and Pennock 2007). After clicking on the advertisement, the consumer is redirected to the advertiser’s website and then chooses whether to purchase a product or register for a service (which we define as a conversion).

The search engines provide daily reports to advertisers on the status of their campaigns. These reports provide statistics on the number of impressions and clicks and the average position for each keyword in the advertiser’s portfolio. The continuous nature of the auction enables an advertiser to change the portfolio of keywords as well as the bids, ad copy, and landing page for each keyword in real

time. The advertiser’s submitted bid implicitly determines the target position for the advertisement. These decisions ultimately drive the advertiser’s return on ad spending, a key metric used to evaluate return on investments in advertising. In our study, we focus on the impact of the advertisement’s position in the list of sponsored search results on revenues and profitability for a given set of keywords. The ad copy and landing pages associated with these keywords do not change over time for the advertiser under consideration.

LITERATURE REVIEW

Consumers’ Online Search Behavior

The literature most relevant to our study includes prior research on consumers’ online search behavior, with a special emphasis on the impact of message order on consumer choice, and the research focused on advertisers’ performance in sponsored search markets. An important consideration in evaluating the performance of sponsored search advertisements is consumer response to the ad position in terms of both clicks and conversions. Prior work in traditional media has demonstrated that message ordering influences ad persuasion (Brunel and Nelson 2003; Rhodes et al. 1979), and similar results have been shown in online environments. Hoque and Lohse (1999) find that consumers are more likely to choose advertisements near the beginning of an online directory than they are when using paper directories. Ansari and Mela (2003) find that the higher position of links in an e-mail campaign can lead to higher probability of clicking. Johnson et al. (2004) find that consumers

searched fewer than two stores during a typical search session. Similarly, Brynjolfsson, Dick, and Smith (2009) find that only 9% of shopbot users select offers beyond the first page. In general, because of the cognitive costs associated with evaluating alternatives, consumers often focus on a small set of results (Montgomery et al. 2004).

Search engines can also be viewed as tools that aid consumer decision making. Häubl and Trifts (2000) find that the use of decision aids reduces the size of consumers' consideration sets but improves the quality of both their consideration sets and the ultimate purchase decision in an online shopping environment. This again suggests that consumers are likely to evaluate only a few sponsored search results because they might expect that the results are in decreasing order of relevance. In this regard, Feng, Bhargava, and Pennock (2007) find evidence of an exponential decrease in the number of clicks for an advertisement with its rank and attribute this to decay in user attention.

Click behavior also depends on the type of consumer. Online consumers include both buying consumers and information seekers (Moe 2003; Moe and Fader 2004; Montgomery et al. 2004). Moe (2003) shows that consumers with high purchase intent tend to be focused in their search, targeting a few products and categories than consumers with low purchase intent, who tend to have broad search patterns. Using path analysis, Montgomery et al. (2004) show that consumers with directed search have a higher probability of purchase. Urbany, Dicksoti, and Wilkie (1989) show that consumers with greater uncertainty about information for alternatives are likely to search less, while consumers with uncertainty about choice search more. Brucks (1985) and Srinivasan and Ratchford (1991) show that product knowledge increases search. Similarly, Moorthy, Ratchford, and Talukdar (1997) show that consumers with low expertise are likely to search less. A similar pattern is likely in sponsored search: Consumers may be heterogeneous in their purchase intent, expertise, and resulting search behavior. The keywords consumers use can potentially reflect their underlying purchase intent and expertise. For example, a common belief in the industry is that the use of more specific keywords may reflect a higher proportion of buyers than the use of less specific keywords. White and Morris (2007) and White, Dumais, and Teevan (2009) confirm that advanced users submit longer, more specific queries and click further down the search results. Therefore, we expect that more specific keywords are associated with deeper search (i.e., generate clicks at lower positions) than are less specific keywords.

Advertiser revenues depend on both clicks and conversion probability. One possibility is that position has no impact on purchase probability conditional on clicking. Some studies show that consumers tend to deemphasize prescreening information in their search process (Chakravarti, Janiszewski, and Ulkumen 2006; Diel, Kornish, and Lynch 2003). This suggests that the criteria used for selecting an advertisement may not affect the final order as much as information obtained after visiting the associated website. Thus, if a consumer discounts all prescreening information and buys from the website that maximizes his or her utility, conversion rate may be independent of ad position.

Alternatively, consumers may form expectations that advertisements are arranged in decreasing order of relevance

or quality and may therefore prefer to buy from a website in the top position. This is likely in sponsored search because consumers are accustomed to relevance-based ordering of nonsponsored results. Sequential evaluation can also influence conversion performance. In sponsored search, consumers most likely cannot perfectly recall product information from all visited websites, as they may need to view several pages across each website to get to the product of interest. Traditional advertising studies have demonstrated primacy effects in the recall of brand and product information (e.g., Pieters and Bijmolt 1997). Both these factors (expectation of relevance ordering and primacy in evaluation) can result in a conversion rate that decreases with position.

Finally, Wyer and Srull (1986) show recency effects under conditions of high information load. Wedel and Pieters (2000) also find a recency effect in the recall of advertisements in a print magazine. Häubl, Benedict, and Bas (2010) show that in the context of sequential choice, consumers are disproportionately influenced by the attractiveness of the most recently evaluated product. This suggests that the consumers who are likely to buy are more likely to do so from the website they evaluate last rather than the website they evaluate early in their sequential search. In addition, if top positions draw clicks from both information seekers and buyers and lower positions draw clicks primarily from buyers, the lower positions may be associated with higher conversion rates. Brooks's (2004) industry report suggests that click-to-order probability increases with position for low-volume search advertisements and decreases for high-volume search advertisements and that low-volume search advertisements are associated with very specific keywords. In summary, there is conflicting prior work regarding whether conversion rate is independent of, decreases with, or increases with position.

Sponsored Search Markets

Existing work in sponsored search has focused on auction design, consumer behavior, and advertiser strategy. In terms of work on auction design, Edelman, Ostrovsky, and Schwarz (2007) and Varian (2007) compute the equilibriums of the generalized second price sponsored search auction and demonstrate that the auction, unlike the Vickrey-Clarke-Groves mechanism, is not incentive compatible. Thus, advertisers will bid strategically in these auctions. Edelman and Ostrovsky (2007) examine paid search auction data and find evidence of strategic bidder behavior. Feng, Bhargava, and Pennock (2007) and Weber and Zhang (2007) compare the performance of various ad-ranking mechanisms, finding that a yield-optimized auction, with rankings based on a combination of the submitted bid and ad relevance, maximizes revenue to the search engine.

Recent empirical studies have modeled consumer choice in sponsored search. Rutz and Bucklin (2010) show that there is a spillover effect from generic keywords to branded keywords, in which searches using generic keywords may not generate immediate sales but may instead drive future searches and sales using branded keywords. The authors focus only on conversions associated with branded keywords. In our sample, there are more direct conversions than conversions due to a sequence of keywords. Furthermore, Rutz and Bucklin refer to branded keywords as those that take the consumer to the site of the brand. For a

retailer, these are typically keywords with the retailer brand name. There is typically no (or very limited) competition for such keywords, and thus ad position is irrelevant for such keywords. Consequently, although the study analyzes consumer click and conversion behavior in sponsored search, Rutz and Bucklin do not study the effect of ad position on revenues and profitability, and their study does not help answer our research question. Ghose and Yang (2009) model clicks and orders and show that profits can be non-monotonic in position due to poor bid efficiency of keywords at higher positions. They find that both CTRs and conversion rates decrease with position. They also study the effect of keyword characteristics such as the presence of national brand information. Their study focuses on a wide range of positions, and Ghose and Yang do not separately evaluate differences between the top few positions. This, in turn, drives several differences in our results, as we explain subsequently.

In terms of work on advertiser strategies, several recent studies have focused on optimal bidding strategies for advertisers (Cary et al. 2007; Feldman et al. 2007; Hosanagar and Cherapanov 2008). These studies present optimization models to compute the bids for all keywords in an advertiser's portfolio to maximize advertiser profits subject to a budget constraint. The models capture the notion that a high bid ensures a top position and therefore generates a large volume of clicks but correspondingly incurs a high CPC. However, none of these studies explicitly investigate the impact of ad position on advertiser revenues and profitability. Ganchev et al. (2007) evaluate bids submitted in sponsored search auctions and find that bids decay exponentially with position. Arbatskaya (2007) studies equilibrium pricing in markets with ordered search for homogenous goods and shows that equilibrium prices and profits decrease in the order of search. The results provide a theoretical explanation for the decay in bids with position based on the observation that revenues decay with position.

In summary, prior research reveals two themes. First, the literature on consumer search behavior suggests that ad position is likely to influence consumer response, but the overall impact on advertiser revenues is ambiguous because the impact of ad position on conversion rates is not clear. Second, recent work in sponsored search has emphasized the need for advertisers to track the bid efficiency of various ad positions. These studies confirm that CTRs decay with ad position but indicate that costs also decay. In other words, the net effect on profits is not easily predictable. Thus, the net impact of position on revenues and profitability is an open and managerially significant research question, which we address with the current study.

FIELD EXPERIMENT AND DATA

We generated our main data set through a field experiment for a sponsored search ad campaign on Google for an online retailer of pet products. The data were generated by submitting randomized bids for 68 keywords and measuring consumer response in terms of clicks and orders at different ad positions for the keywords. We chose the keywords randomly from a set of keywords in the campaign related to food products that had generated orders in the past for the retailer. Google allows advertisers to use "broad," "exact,"

or "phrase" matches for their keywords (for more information, see <http://adwords.google.com>). An exact or a phrase match ensures that the search query contains the keywords in the same order and thus indicates a better match with the consumer's search intention. Therefore, we considered only keywords with phrase and exact matches.

Our objective was to rotate the advertisements associated with the experimental keywords across multiple positions on the results page and measure consumer response in terms of clicks and orders as a function of position. Search engines use the advertiser's bid and the advertisement's quality score to determine the position of an advertisement. To influence the ranking, we randomly varied the bids for these keywords. The bid range was wide enough for each keyword to ensure that corresponding advertisements could be placed in various slots available on the first page.⁵ In addition, we ensured that the advertisement for a keyword appears in a particular position for several days. We retained the advertiser's original ad copy for each keyword and the associated landing page and did not change them for the duration of the field experiment. We used no other performance criteria to determine the bids. In our data set, there are fewer than 100 observations for Positions 7 and lower; therefore, we restrict our analysis to the first seven positions.⁶

The data set consists of 3187 observations of daily impressions, clicks, and orders for 68 keywords over a 45-day period from June 2009 to July 2009. Table 1 provides summary statistics. Note that the observations represent daily aggregate data for advertisements corresponding to the sample keywords for our advertiser, and the data set is typical of the information received by advertisers in sponsored search. We do not have information on the performance of competing advertisements or detailed information on how an individual consumer makes a choice during a search session. In addition, the position reported for any keyword is the average position on a given day. The position of an advertisement can vary within a day because the set of advertisers may be different for different queries associated with a keyword. For example, the advertisement for the keyword

⁵We used Google's keyword tool to determine the range of bids for each keyword.

⁶Several studies (e.g., Brynjolfsson, Dick, and Smith 2009; Johnson et al. 2004) have shown that consumer search depth is limited. For example, Sherman's (2006) survey shows that the vast majority of consumers do not search beyond the first page of search results.

Table 1
KEYWORD PERFORMANCE SUMMARY STATISTICS

Variable	<i>M</i>	<i>SD</i>	<i>Minimum</i>	<i>Maximum</i>
Impressions	73	132	1	1,092
Clicks	1.14	2.1	0	24
Orders	.02	.16	0	3
Average position	3.5	1.22	1	6
Bid	.52	.3	.05	2
Size	2.33	.9	1	4
Brand	.75	.44	0	1
Specificity	.81	.68	0	2
LQScore	7.9	.39	6	10

“red dress” may be shown if the consumer types “red party dress” or “red dress.” The competitors and their bids may be different for these two queries, causing the position to vary. By eliminating keywords with a “broad” match, we eliminate a major source of such intraday variation in position. Moreover, our results hold for the subset of keywords with an “exact” match. An advertisement for a keyword with an exact match is shown only if the query is exactly the same as that keyword, and thus the competitors are fixed for such keywords. Another reason the position may vary is that competitors may change their bids several times within a day. Although firms change their bids periodically, typically weekly and sometimes even daily, we do not find significant intraday variation in ad position for keywords with exact and phrase match.⁷

SIMULTANEOUS MODEL OF CLICKS, CONVERSIONS, AND AD POSITION

Consider an advertiser placing bids for a keyword to ensure that its advertisements are visible in the list of sponsored results for a query related to that keyword. The search engine uses this bid and expected ad performance to determine the ad position in the list of sponsored search results. Consumers see the advertisements, decide to click on particular ones, and subsequently choose whether to make a purchase. We simultaneously model consumers’ click-through and conversion behavior as well as the search engine’s keyword ranking decision for the advertisement.

The advertiser receives aggregated information on a daily basis from the search engine regarding the number of impressions the advertisement for the keyword received, the average position of the advertisement, and the number of times it was clicked. The advertiser is also able to record daily orders that were generated for each search engine keyword. Competitive information is not available to the advertiser. The advertiser’s expected profit from a keyword as a function of ad position n is as follows:

$$(1) \quad \pi(n) = I \times [CTR(n) \times CONV(n) \times RPO - CTR(n) \times CPC(n)],$$

where I is the expected number of ad impressions, $CTR(n)$ is the click-through rate, or the fraction of ad impressions that generate clicks; $CONV(n)$ is the conversion rate per click, or the fraction of clicks that generate orders (given n); RPO is the revenue per order; and $CPC(n)$ is the average cost per click charged to an advertiser assigned to position n . We assume that the number of impressions is independent of the position of the advertisements. We make a few assumptions in Equation 1. First, the equation assumes that the expected number of impressions, I , is independent of ad position. This assumption is reasonable for the top few positions that appear on the first page of search results. However, for ad positions in subsequent pages, the number of impressions is clearly lower because consumers rarely evaluate advertisements beyond the first page. Our main

analysis focuses on the top ad positions that appear in the first page, and impressions do not seem to depend heavily on position for these top positions.

A more critical assumption is that the revenues from an order are independent of ad position. Specifically, the advertiser’s prices and costs associated with selling the products are independent of ad position. In our main analysis based on random assignment of ad positions, we kept prices constant for all the products. Thus, our analysis ensures that the customers are not influenced by difference in product prices for different positions. In addition, for our analysis of a secondary data set from the field, we verified that product prices were independent of ad positions. However, it is possible that advertisers endogenize product prices on the basis of different ad positions. For example, if consumers search sequentially and end their search early because of high search costs, advertisers may be able to charge a higher price for their product at the top position. The changes in prices may also affect consumers’ click and purchase behavior. This would affect the firm’s revenue and profit as a function of position; however, this is outside the scope of our analysis. We now discuss how we model each of the components of the profit equation.

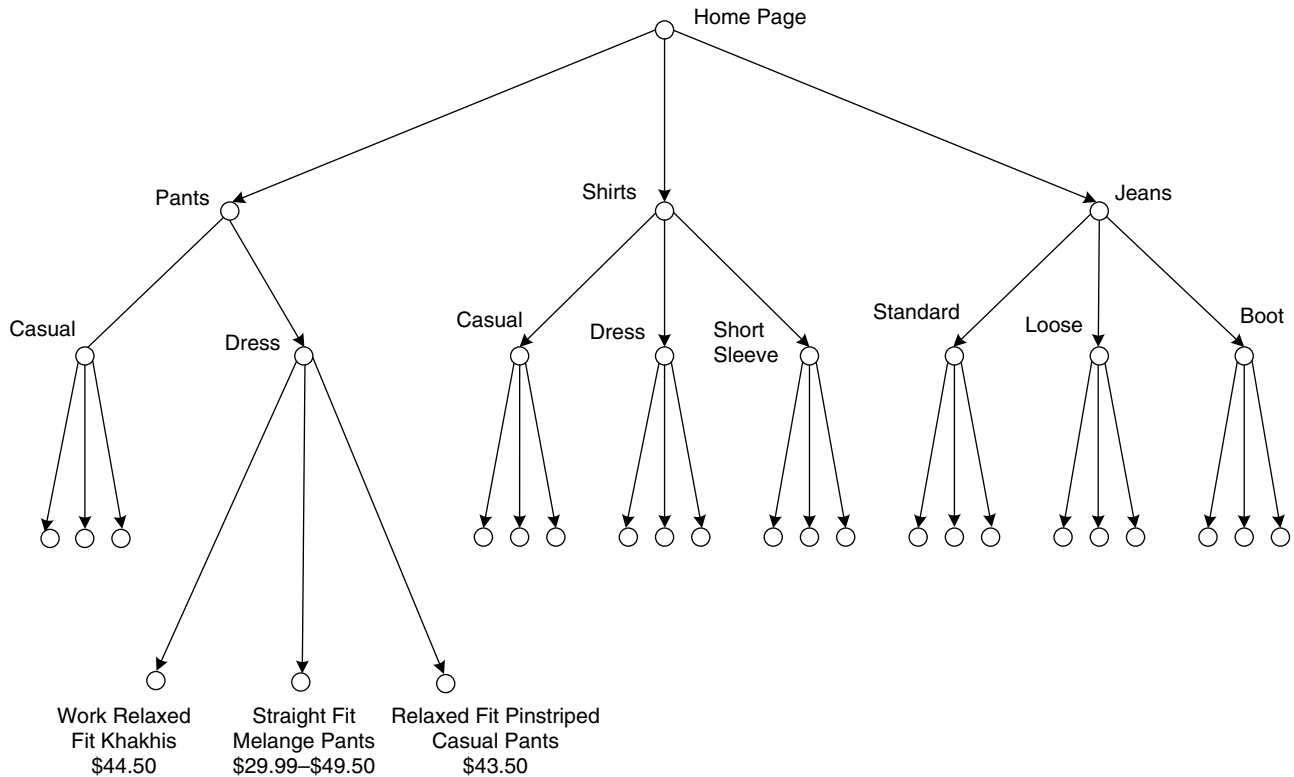
CTR per Impression

Our unit of analysis is a keyword because the search engine auction is keyword specific. Keyword characteristics are an indication of the underlying search behavior, which varies across consumers. For example, the keyword “shirt” is less specific and indicates an initial stage of information search, whereas more specific keywords such as “Levi’s shirt” and “formal blue shirt” indicate a more advanced and directed stage of information search. To account for these differences across keywords, we capture how specific a keyword is using two measures: specificity and brand. The specificity of a keyword is based on the nearness of its landing page to a product. Advertisers organize their websites hierarchically to accommodate the search intent of users and to reduce their search cost. Various levels in the hierarchy represent product categories, subcategories, and products. For example, Figure 2 shows the hierarchy for men’s clothing in a representative website. When consumers are routed through a search engine, the landing page coincides with a level in the website hierarchy chosen according to the search intent of the consumer as reflected in the keyword. We define specificity as the level in the product hierarchy of the advertiser. For example, a top level such as “men’s clothing” would have a specificity value of 0, a second level such as “shirts” would have a specificity value of 1, and so on.

A keyword can also represent the national brand preference of the consumer. For example, the keyword “Levi’s jeans” would indicate that the consumer has a preference for the Levi’s brand and is further along in his or her search. We use a dummy variable to represent the presence of national brand information in the keyword. In addition to brand and specificity, there can be other variables that capture keyword characteristics. For example, the presence of retailer information captures preference for the retailer. However, in our data sets, there is no competition for such keywords, which results in a single ad position when these

⁷We separately verified this for our sample keywords by monitoring the relative ad positions across multiple queries in a day. For the keywords with a phrase match, we used a large set of queries that had been associated with these keywords in the past several months.

Figure 2
HIERARCHY FOR MEN'S CLOTHING WEBSITE



keywords are used. Although this may enable us to measure performance of such keywords compared with other keywords, it is not within the scope of this research study. Another variable that has been used in prior studies is the size of the keyword, which indicates the number of words in the keyword. This can capture additional aspects of consumer preference. However, this variable is redundant for our main data set because the information is already captured in specificity and brand variables. We use this variable in our secondary data set for another retailer.

Each keyword has its own set of competing advertisers. Consumers form an expectation of advertiser quality on the basis of their familiarity with the advertiser and the quality of the advertisement compared with other advertisers. We proxy for this perceived quality using a measure of relative quality, called the “listed quality score,” or LQS-score, maintained by the search engine and available to the advertiser (see <http://adwords.google.com/support/aw/bin/answer.py?hl=en&answer=6111>). This measure of quality represents the click propensity of an advertiser and is based on several metrics such as the relative click performance of the advertiser for the keyword, the relative overall click performance of the advertiser, the relative quality of the advertisement, and the relevance of the advertisement to the keyword. There are other unobservable characteristics associated with a keyword that can influence consumer choice; for example, the regular search results are different for different keywords.

We use a hierarchical model to capture the effect of keyword characteristics. This provides a flexible random component specification, enabling us to incorporate observable and unobservable keyword-specific heterogeneity given the limited observations for each keyword. Hierarchical models are commonly used to draw inferences on individual-level characteristics (Rossi and Allenby 2003). Hierarchical Bayesian models have also been used to study sponsored search data with keyword as a unit of analysis (Ghose and Yang 2009; Yang and Ghose 2010).

We assume an i.i.d. extreme value distribution of the error term for individual choices and use a logit model to represent the click probability for a keyword k at time t as follows:

$$(2) \quad \Lambda_{k,t}^{CTR} = \frac{\exp(U_{kt}^{CTR})}{1 + \exp(U_{kt}^{CTR})},$$

where U_{kt}^{CTR} is the latent utility of clicking. This depends on the position of the advertisement and the expected ad quality. For a keyword k at time t , this latent utility for a can be expressed as follows:

$$(3) \quad U_{kt}^{CTR} = \theta_0^k + \theta_1^k \text{Pos}_{kt} + \theta_2 \text{AdQuality}_{kt} + \sum_d \delta_{kt}^d \theta_{DOW_d} + \theta_{\text{Time}} \text{Time}_{kt} + \varepsilon_{kt}^\theta, \\ \theta^k = \Delta^\theta z_k + u_k^\theta, u_k^\theta \sim N(0, v^\theta), \quad \text{where } \theta^k = [\theta_0^k, \theta_1^k],$$

where Pos represents the position of the advertisement in sponsored search results and AdQuality is the expected

quality of the advertisement (proxied by LQScore, as noted previously); z_k represents keyword-specific characteristics: brand and specificity; Δ^θ is a matrix capturing the relationship between the keyword characteristics and the mean values of coefficients; u_k^θ represents the unobservable heterogeneity for the random coefficients, which we assume are normally distributed with a mean 0 and covariance matrix V^β ; and ε_{kt}^θ represents the time-varying unobserved keyword attributes that are common for all consumers. Buying patterns can change during the week, and we use day-of-the-week dummies to control for this: δ_{kt}^d . We also control for the time dynamics of the auction using a time variable Time_{kt} .

Conversion Rate per Click (CONV)

Next we discuss the model for conversion rate, $\text{CONV}(n)$, another key input into the advertiser’s profit function presented in Equation 1. Assuming an i.i.d. extreme value distribution of the error term for individual choices, we can express the conversion probability as follows:

$$(4) \quad \Lambda_{kt}^{\text{CONV}} = \frac{\exp(U_{kt}^{\text{CONV}})}{1 + \exp(U_{kt}^{\text{CONV}})},$$

where U_{kt}^{CONV} is the latent utility of conversion, which may depend on the position of the advertisement. For a keyword k at time t , we can express this latent utility as follows:

$$(5) \quad U_{kt}^{\text{CONV}} = \beta_0^k + \beta_1^k \text{Pos}_{kt} + \sum_d \delta_{kt}^d \beta_{\text{DOW}_d} + \beta_{\text{Time}} \text{Time}_{kt} + \varepsilon_{kt}^\beta,$$

$$\beta^k = \Delta^\beta z_k + u_k^\beta, \quad u_k^\beta \sim N(0, V^\beta), \quad \text{where } \beta^k = [\beta_0^k, \beta_1^k].$$

We have controls for time and a constant term similar to the CTR model and similar to that of Ghose and Yang (2009) and Yang and Ghose (2010).

Ad Position

The search engine determines the position of an advertisement for a keyword according to the product of the current bid and the advertisement’s quality compared with competing advertisements. As we mentioned previously, this relative quality measure is the quality score and is available to the advertisers as the LQScore (see <http://adwords.google.com/support/aw/bin/answer.py?hl=en&answer=6111>, <https://adwords.google.com/support/aw/bin/answer.py?hl=en&answer=100305>, and <https://adwords.google.com/support/aw/bin/answer.py?hl=en&answer=115967>). The dependence of ad position on bid and past performance introduces two sources of endogeneity related to the advertiser’s decision and the search engine’s decision. Advertisers can influence the position by changing their bids. In particular, advertisers might choose bids to obtain positions that yield the best performance for them. As a consequence, position is endogenously determined. Furthermore, search engines might assign advertisers to specific positions that yield the search engine the highest revenues.

To correct for the resulting bias, we must account for the advertiser’s bid choices as well as the position the search engine assigns. In our setup, we randomized bids for the sample keywords. Thus, the advertiser did not control the bids during the field experiment, taking away any strategic effect of our advertiser. Using a wide range of random

bids also ensures that even if other advertisers are bidding using their own objective functions, the advertisements in our experiment are exposed to consumers over a wide range of positions.

Position can also be endogenous because search engines use ad performance data to compute an advertisement’s position. To account for this, we explicitly model the search engine’s decision. We express the ad position for a keyword k at time t as follows:

$$(6) \quad \text{Pos}_{kt} \propto \sigma_0^k (\text{Bid}_{k,t})^{\sigma_1^k} (\text{LQScore}_{k,t})^{\sigma_2^k}.$$

Note that the position of the advertisement is the daily average position and is a continuous variable. The functional form ensures that the bid and the LQScore are required to determine the rank and explicitly incorporates the provision that the ad position is not randomized even if advertiser bids are random. To account for the effect of competition, we also use the maximum competitive bid, CompBid , for each keyword, which can be obtained from Google’s keyword tool (see <https://adwords.google.com/select/KeywordToolExternal>). Substituting, taking the log, and using controls for day and time, we obtain the following:

$$(7) \quad \ln(\text{Pos}_{kt}) = \alpha_0^k + \alpha_1^k \ln(\text{Bid}_{k,t}) + \alpha_2 \ln(\text{LQScore}_{k,t}) + \alpha_3 \text{CompBid}_{kt} + \sum_d \delta_{kt}^d \alpha_{\text{DOW}_d} + \alpha_{\text{Time}} \text{Time}_{kt} + \varepsilon_{kt}^\alpha,$$

with $\alpha^k = \Delta^\alpha z_k + u_k^\alpha$ and $u_k^\alpha \sim N(0, V^\alpha)$.

Finally, because the position of the advertisement depends on the search engine’s decision and is endogenous, the unobservable time-varying keyword attributes for the equations representing consumer decisions will be correlated with the error term for the equation representing the search engine decision. As such, we use the following distribution to account for correlation between the error terms for CTR, conversion rate, and position equations:

$$(8) \quad \begin{bmatrix} \varepsilon_{kt}^\theta \\ \varepsilon_{kt}^\beta \\ \varepsilon_{kt}^\alpha \end{bmatrix} \sim N(0, \Omega), \quad \text{where } \Omega = \begin{bmatrix} \Omega_{11} & \Omega_{12} & \Omega_{13} \\ \Omega_{21} & \Omega_{22} & \Omega_{23} \\ \Omega_{31} & \Omega_{32} & \Omega_{33} \end{bmatrix}.$$

Identification

The preceding set of simultaneous equations represents a triangular system that has been addressed in classical econometrics studies (Greene 1999; Hausman 1975; Lahiri and Schmidt 1978) and in Bayesian econometrics (Zellner 1962). We can represent it as follows:

$$U_{kt}^{\text{CTR}} = f(\text{Position}, X1, \varepsilon_{kt}^\theta),$$

$$U_{kt}^{\text{CONV}} = f(\text{Position}, X2, \varepsilon_{kt}^\beta), \quad \text{and}$$

$$\text{Position} = f(X3, \varepsilon_{kt}^\alpha).$$

In this construction, position is endogenous, and variables $X1$ – $X3$ are exogenous. Identification is possible because rank is completely determined by the exogenous variables bid and LQScore. In our setting, bid for each keyword is randomized by the experiment. LQScore is a value that the search engine internally calculates for each keyword, and

it remains stable for the short period unless the advertisers change their advertisements or landing pages to influence the quality score. Rank, in turn, influences click and conversion performance. Thus, the rank and order conditions are satisfied for identification purposes (Greene 1999).

Lahiri and Schmidt (1978) show that the parameter estimates for a triangular system can be fully identified using generalized least squares. Hausman (1975) shows that the likelihood function for a triangular system is the same as for seemingly unrelated regressions. Zellner (1962) addresses triangular systems from a Bayesian point of view and shows that the posterior probability distribution function is the same as in a seemingly unrelated regressions setting. Triangular systems have been estimated using the classical approach (Godes and Mayzlin 2004) and, more recently, in sponsored search using the Bayesian approach (Ghose and Yang 2009; Yang and Ghose 2010).

We estimated the model using a Bayesian approach, applying Markov chain Monte Carlo (MCMC) sampling because of the nonlinear characteristics of our model (Rossi and Allenby 2005). (We discuss the priors and conditional posteriors of this model in the Appendix.) For the hierarchical Bayesian models, we ran the MCMC simulation for 80,000 draws, discarding the first 40,000 as burn-in. To ensure that our parameter estimates were accurate, we simulated the clicks, orders, bids, and positions using our estimates. By repeating the estimation with this simulated data set, we were able to recover our parameter estimates, indicating that our parameters are fully identified.

RESULTS

CTR

Table 2 provides the mean values for the posterior distribution of the Δ^0 matrix and the estimates for V^0 from Equation 3. The coefficient for Pos is negative and significant, indicating that click performance decays with position. In addition, the brand and the keyword specificity have no significant impact on the click performance of keywords or its rate of decay with position. Although more specific keywords are, on average, associated with a higher CTR


Table 2
ESTIMATES FOR THE CTR

	Intercept	Brand	Specificity
Const	-3.87 (.25)**	-.45 (.31)	.02 (.21)
Pos	-.39 (.06)**	.04 (.13)	.04 (.09)
LQScore	.14 (.03)**		
Day 1	.01 (.06)		
Day 2	.001 (.06)		
Day 3	-.09 (.07)		
Day 4	-.13 (.07)		
Day 5	-.18 (.05)*		
Day 6	-.1 (.09)		
Time	-.01 (.001)**		
V^0	Const	Pos	
Const	.76 (.15)**	-.18 (.05)**	
Pos		.16 (.03)**	

*Statistically significant at 5%.
**Statistically significant at 1%.

Figure 3

CTR AS A FUNCTION OF POSITION FOR SAMPLE KEYWORDS



and a lower rate of decay with position, these coefficients are not significant. Figure 3 shows the mean CTR for the first five positions for a few sample keywords with different combinations of brand and specificity values. We used the posterior distribution of the keyword parameter estimates to calculate the CTR for these keywords.

Conversion Rate

Table 3 provides the mean values for the posterior distribution of the Δ^B matrix and the estimates for V^B from Equation 5. Not accounting for brand and specificity, the coefficient for Pos is positive and significant, indicating that, on average, conversion rate increases with position. Thus, we conclude that ad position has an impact on conversion rate, suggesting that the serious buyers are visiting the lower positions more than information seekers and are buying from these positions. Brand information does not seem to have a significant impact on either conversion rate or the rate at which conversion rate decays with position. In contrast, the specificity of the keyword seems to have

Table 3
ESTIMATES FOR THE CONV

	Intercept	Brand	Specificity
Const	-2.51 (.22)***	.16 (.36)	-.54 (.25)**
Pos	.39 (.06)***	.2 (.14)	.2 (.1)**
Day 1	-.46 (.06)***		
Day 2	-.4 (.09)***		
Day 3	-.28 (.06)***		
Day 4	-.1 (.07)		
Day 5	-.13 (.07)*		
Day 6	.08 (.07)		
Time	.01 (.001)***		
V^B	Const	Pos	
Const	.93 (.2)***	-.04 (.06)	
Pos		.15 (.03)***	

*Statistically significant at 10%.
**Statistically significant at 5%.
***Statistically significant at 1%.

a significant impact. The increase in conversion rate with position seems to be far more prominent for more specific keywords. Notably, Brooks's (2004, p. 4) industry report observes that "conversion rates for low volume keywords may very well increase as rank falls." Given that more specific keywords in our data set are associated with lower search volume, our study is the first to formally verify their observation.

Our results suggest that buying consumers visit lower positions more than information-seeking consumers do. This finding appears to contradict the work of Ghose and Yang (2009), whose results suggest that conversion rate decreases with position. However, Ghose and Yang evaluate their results with a large range of positions (1–131) compared with our study, which investigates only the top seven positions. Very low conversion rates at the low positions they study might drive these differences. Although our primary data set does not have sufficient data for such low positions, we conducted an additional analysis with our secondary data set with a wider range of positions. We find that the strong increasing trend disappears, and overall conversion rate is independent of ad position. This suggests that although serious buyers are likely to visit the top few positions, they do not visit very low positions and that the effect on conversion rate over a wider range of positions may be nonmonotonic. It also shows that a clear and significant trend for the top positions can be masked by using a large number of positions.

An additional driver of the differences may be the differences in the type of retailers studied. Ghose and Yang (2009) study a large *Fortune* 500 retailer with several hundred retail stores, whereas the retailers in our study are pure online retailers. It is possible that the actual conversion behavior is not fully captured for the retailer in Ghose and Yang's study because consumers are making purchases in the physical store after an online search.

Figure 4 shows the mean conversion rate for the top positions of a few sample keywords with different combinations of brand and specificity values. We use the posterior distribution of the keyword parameter estimates to calculate the

Figure 4
CONV AS A FUNCTION OF POSITION FOR SAMPLE
KEYWORDS

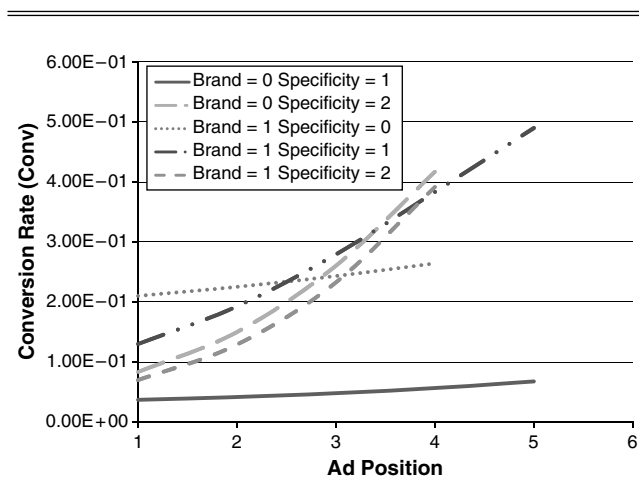


Table 4
ESTIMATES FOR THE AD POSITION

	<i>Intercept</i>	<i>Brand</i>	<i>Specificity</i>
Const	.53 (.26)*	-.04 (.12)	.09 (.02)
Bid	-.38 (.05)**	.07 (.1)	-.04 (.02)
LQScore	-.27 (.09)**		
Comp_Bid	.1 (.04)**		
Day 1	.02 (.02)		
Day 2	.001 (.02)		
Day 3	.03 (.02)		
Day 4	.02 (.02)		
Day 5	.05 (.02)**		
Day 6	.001 (.02)		
Time	.001 (.001)		
V^α	<i>Const</i>	<i>Bid</i>	
Const	.33 (.06)**	.001 (.03)	
Bid		.15 (.03)**	

*Statistically significant at 5%.

**Statistically significant at 1%.

conversion rate for these keywords. The figure illustrates how conversion rate increases with position, and the effect is more pronounced for specific keywords.

Ad Position

Table 4 provides the mean values for the posterior distribution of the Δ^α matrix and the estimates for V^α from Equation 7. In these results, higher bids lead to a higher current position. Similarly, a higher LQScore leads to higher current position. This is reasonable because both bid and LQScore are the primary inputs used to compute ad rank, and higher values should move the ad higher in the list of results. A higher maximum competitive bid also leads to a lower current position, which is reasonable because it indicates that higher competing bids would lower the advertiser's rank.

Finally, Table 5 shows covariance between unobservables for CTR, CONV, and ad positions from Equation 8. Covariance between the unobservables for CONV and CTR is significant. This indicates that the unknown factors influencing consumer clicks also influence the subsequent conversion behavior. The covariance between the unobservables for CONV and position is statistically significant. Similarly, covariance between the unobservables for CTR and position is statistically significant. This correlation between the error terms for CONV and CTR with the error term for ad position shows that position is endogenous, and the proposed simultaneous equation model helps capture the endogeneity effect.

Table 5
ESTIMATES FOR THE COVARIANCE MATRIX Ω

	<i>CONV</i>	<i>CTR</i>	<i>Pos</i>
CONV	.4 (.04)**	-.19 (.02)**	-.02 (.01)*
CTR	-.19 (.02)**	.27 (.02)**	.013 (.006)*
Pos	-.02 (.01)*	.013 (.006)*	.08 (.001)**

*Statistically significant at 5%.

**Statistically significant at 1%.

Revenue as a Function of Position

To determine the overall effect of position on revenue, we calculated revenue using the posterior distribution of the CTR and CONV coefficients for each keyword (assuming a \$10 average value of the associated products). We found that revenue increases with position for keywords with high specificity. Note that conversion rate may be increasing with position if the top positions are drawing relatively more information seekers than serious buyers. However, this cannot by itself cause revenues to increase with position. A potential explanation is that serious buyers, who are using more specific keywords, show a recency bias and ultimately buy from lower positions. That is, these consumers may be evaluating multiple advertisements and, because of a recency bias, are more likely to buy from our advertiser when they evaluate its advertisement at a lower position. Figure 5 shows revenue as a function of position for the top positions for a few sample keywords with different combinations of brand and specificity values.


Profitability as a Function of Position

We expect CPC to decay with position. Therefore, it follows that the top position is not the most profitable for keywords with higher revenue for lower positions. To assess the impact of ad position on profitability for other keywords, we need to know the impact of position on the cost of the keywords. Therefore, we use the relationship between the search engine rank and the advertiser's bid (Equation 7) to determine the cost. For a given bid and position j , we assume that the actual CPC is the bid for position $j + 1$.⁸ We use the posterior distribution of the parameter estimates for CONV, CTR, and the position equations to compute profits for each keyword for the top

⁸The actual cost for position j is determined using a generalized second price auction and depends on the bid and the position adjusted CTR of the advertiser in position $j + 1$. Our calculation assumes that the position adjusted CTR of the advertiser in position $j + 1$ is same as that for our advertiser. This is an approximation for convenience.

Figure 5

REVENUE PER IMPRESSION AS A FUNCTION OF POSITION FOR SAMPLE KEYWORDS



positions using Equation 1. Figure 6 shows the mean values for some sample keywords with various combinations of brand and specificity. We find that for our advertiser, lower ad positions generate higher profit for almost all the keywords, because cost decays at a faster rate than revenue. There has been some evidence (Kitts and LeBlanc 2004) that bid efficiency is not the highest for the top position. Recently, Ghose and Yang (2009) have established similar results for bid efficiency. Our results also indicate that the top positions may not always be profitable and, moreover, may not maximize revenue in the first place.

Robustness of Results

In this subsection, we outline several steps we took to evaluate the robustness of these results. Because of space constraints, we present only the high-level description for alternative model specifications and the spillover effect. (For detailed descriptions of the related models along with the estimates, see the Web Appendix at <http://www.marketingpower.com/jmrdec11>.)

Model without keyword heterogeneity. We evaluated an alternate model without keyword heterogeneity and compared it with our original model using Bayes factors. We used harmonic mean (Newton and Raftery 1994) to calculate the log-marginal density based on the MCMC output. We report log-marginal densities and the Bayes factors in Table 6. Using Bayes factors, we find strong evidence supporting our model with keyword heterogeneity (Bayes factor = 86).

Holdout sample analysis. We attempted to verify the prediction accuracy of our results using a holdout sample. To do this, we used data for the first four weeks as the estimation sample and data for the remaining two weeks as the holdout sample. We used mean absolute percentage error (MAPE) for daily CTR and CONV values at the aggregate and the keyword levels. Table 7 reports error values, which indicate that the model prediction accuracy is similar for both the estimation and the holdout samples. This suggests that our model estimates are robust.

Figure 6

PROFIT PER IMPRESSION AS A FUNCTION OF POSITION FOR SAMPLE KEYWORDS




Table 6
PREDICTION ACCURACY AND FIT FOR DIFFERENT MODELS

Models	Log-Marginal Density	Log-Bayes Factor Main Model Versus	CTR Fit (MAPE)		CONV Fit (MAPE)	
			Aggregate	Keyword	Aggregate	Keyword
Main model	-17,851		.45	.43	.28	.27
Model with rank dummies	-17,868	17	.44	.43	.25	.25
Model without keyword heterogeneity	-17,937	86	.49	.49	.36	.35

Notes: MAPE = mean absolute percentage error. Aggregate MAPE is the average MAPE across all data points. Keyword MAPE is the average of the average MAPE for different keywords.

Alternative model specification. Because our position values for a keyword advertisement are averaged on a daily basis, it is possible that the actual clicks and conversions could have occurred at a higher position in a manner that is not reflected in the daily averages. If so, this would imply that the value of higher positions in driving clicks and conversions are underestimated. However, this is unlikely because the actual positions demonstrate limited intraday variation in our data set. To further test for the impact of a bias, we used an alternative approach in which we rounded the average position values to the nearest lower integer. We used position dummies in this model to reflect these integer values. We found that the results hold for this alternative model specification. Table 6 compares the model fit for this model with the original mode. (for detailed descriptions along with the estimates, see the Web Appendix at <http://www.marketingpower.com/jmrdec11>).

Spillover effect. It is possible that a keyword can generate clicks and consumer awareness on the initial search and the consumer can return to the website later using different keywords to purchase the product. This is commonly referred to as a spillover effect (Rutz and Bucklin 2010). If clicks associated with higher positions have greater spillover value than clicks from lower positions, this can potentially confound our results. In additional analysis, we found that there is limited spillover between keywords in our sample. We found that there are keywords in our sample that assist the order generation process for other keywords not in our sample (e.g., keywords with the retailer's name) but that position has no impact on the ability of the keywords to assist other keywords in generating orders; that is, although keywords have spillover value, it is not affected by ad position.

Analysis for another retailer. We ran a similar analysis for another retailer as an additional robustness test for our results. In this case, we use archival data for a random

sample of 225 keywords from the advertising campaign of a specialty women's apparel online retailer. The data set consists of daily impressions, clicks, and orders for the sample keywords over a 90-day period from April 2007 to June 2007. In selecting the sample, we considered only advertisements for high-specificity keywords that appeared in the top five positions during this period. In addition, because the retailer sells only its own brand, we do not have a brand attribute for the keywords. Instead we use keyword size (i.e., number of words in the keyword) to capture additional details beyond specificity. For example, "trendy urban men's clothing" and "men's clothing" both have specificity 0, but the longer keyword conveys more information about consumer preference. Table 8 provides summary statistics for this data set.

An important difference compared with the analysis of the main data set is that bids in this data set are not randomized. Rather, the advertiser selects the bids, and this decision is typically based on the past performance and future expected performance for each keyword. Thus, bid choice is endogenous, and we model the advertiser's bid for each keyword on a day as a function of past position and past CPC for different positions. Ghose and Yang (2009) and Yang and Ghose (2010) adopt a similar approach to account for the advertiser's bidding decision. We measure past performance in terms of the performance over the previous seven days.⁹ We use the following reduced form equation to represent the bid for a keyword in the current period:

$$(9) \ln(\text{Bid}_{k,t}) = \gamma_0^k + \gamma_1 \text{CPC}_{k,t-1} + \gamma_2 \text{Pos}_{k,t-1} + \sum_d \delta_{k,t}^d \gamma_{\text{DOW}_d} \\ + \gamma_{\text{Time}} \text{Time}_{kt} + \varepsilon_{kt}^\gamma, \\ \gamma^k = \Delta^\gamma z_k + u_k^\gamma u_k^\gamma \sim N(0, V^\gamma),$$

where $\text{Pos}_{k,t-1}$ is the average position for keyword k for the past seven days, and $\text{CPC}_{k,t-1}$ is the average CPC for keyword k for the past seven days.

We use a log-normal representation because the bids are nonnegative. Another difference compared with our primary data set is that the LQScore was not available for this advertiser. Instead, we used average CTR for the keyword for the past seven days as a proxy for LQScore. We

Table 7
PREDICTION ACCURACY FOR ESTIMATION AND HOLDOUT SAMPLES

Models	CTR Fit (MAPE)		CONV Fit (MAPE)	
	Aggregate	Keyword	Aggregate	Keyword
Estimation sample	.45	.43	.27	.26
Holdout sample	.44	.45	.29	.28

Notes: MAPE = mean absolute percentage error. Aggregate MAPE is the average MAPE across all data points. Keyword MAPE is the average of the average MAPE for different keywords.

⁹The decision to use past seven days of data as a measure of past performance was based on a bidding strategy described to us by the search engine marketing firm that bids on behalf of the advertiser.

Table 8
KEYWORD PERFORMANCE SUMMARY STATISTICS FOR ANOTHER RETAILER

Variable	M	SD	Minimum	Maximum
Impressions	87	191	1	5,481
Clicks	2.29	4.7	0	127
Orders	.016	.16	0	5
Average position	2.85	.95	1	5
AdQuality (LQScore)	.05	.05	.001	.5
Bid	.39	.14	.05	.85
Size	3.1	.83	1	5
Specificity	2.4	.68	2	3
Average position _{t-1}	3.19	1.17	1	7.9
CPC _{t-1}	.2	.13	0	.86
Comp_Bid	.39	.12	.1	1

used the following distribution to account for the correlation between the error terms for CTR, CONV, bid, and position:

$$(10) \begin{bmatrix} \epsilon_{kt}^\beta \\ \epsilon_{kt}^\theta \\ \epsilon_{kt}^\gamma \\ \epsilon_{kt}^\alpha \end{bmatrix} \sim N(0, \Omega), \text{ where } \Omega = \begin{bmatrix} \Omega_{11} & \Omega_{12} & \Omega_{13} & \Omega_{14} \\ \Omega_{21} & \Omega_{22} & \Omega_{23} & \Omega_{24} \\ \Omega_{31} & \Omega_{32} & \Omega_{33} & \Omega_{34} \\ \Omega_{41} & \Omega_{42} & \Omega_{43} & \Omega_{44} \end{bmatrix}.$$

Tables 9–13 show our parameter estimates. On average, the coefficient for Pos is negative for CTR (Table 9) and positive for CONV (Table 10). This suggests that buying consumers using these keywords visit lower positions more than the information seekers do. In addition, the interaction terms between position and size as well as position and specificity are significant and positive for CONV (Table 10). This suggests that conversion rate increases with position for more specific keywords and for longer keywords. Using the posterior distribution for CTR and CONV, we find that the revenue for these keywords increases with position for the range of positions in our sample. This suggests that for specific keywords and longer keywords, consumers are more likely to buy immediately when the advertisement appears in a lower position than when it is at the topmost position.

Table 9
PARAMETER ESTIMATES FOR CTR FOR ANOTHER RETAILER

	Intercept	Size	Specificity
Const	-3.534 (.717)**	.302 (.078)**	-.346 (.124)**
Pos	-.125 (.039)**	-.031 (.021)	.029 (.029)
CTR	4.249 (.402)**		
Day 1	.025 (.021)		
Day 2	-.008 (.028)		
Day 3	-.009 (.023)		
Day 4	-.048 (.026)*		
Day 5	-.024 (.03)		
Day 6	-.013 (.031)		
Time	-.004 (.00)**		
V ^θ	Const	Pos	
Const	1.295 (.165)**	-.168 (.028)**	
Pos		.05 (.008)**	

*Statistically significant at 10%.
**Statistically significant at 1%.

Table 10
PARAMETER ESTIMATES FOR CONV FOR THE ADDITIONAL RETAILER

	Intercept	Size	Specificity
Const	-2.963 (.263)**	-.159 (.068)*	.134 (.09)
Pos	.123 (.032)**	.036 (.015)*	.049 (.021)*
Day 1	.094 (.065)		
Day 2	-.043 (.037)		
Day 3	.094 (.035)**		
Day 4	.212 (.026)**		
Day 5	.125 (.035)**		
Day 6	.219 (.076)**		
Time	-.001 (.001)		
V ^β	Const	Pos	
Const	1.018 (.163)**	-.055 (.017)**	
Pos		.038 (.006)**	

*Statistically significant at 5%.
**Statistically significant at 1%.

Table 11
PARAMETER ESTIMATES FOR BID FOR ANOTHER RETAILER

	Intercept	Size	Specificity	Unobservable Heterogeneity
Const	-1.08 (.03)***	.06 (.03)**	-.08 (.05)	.13 (.02)***
Pos _{t-1}	.022 (.002)***			
CPC _{t-1}	.058 (.015)***			
Day 1	.004 (.004)			
Day 2	.005 (.004)			
Day 3	.005 (.004)			
Day 4	.001 (.004)			
Day 5	.007 (.004)*			
Day 6	.004 (.004)			
Time	.001 (.001)			

*Statistically significant at 10%.
**Statistically significant at 5%.
***Statistically significant at 1%.

Table 12
PARAMETER ESTIMATES FOR POSITION FOR ANOTHER RETAILER

	Intercept	Size	Specificity
Const	2.765 (.185)***	-.08 (.061)	-.156 (.106)
Bid	.994 (.108)***	-.036 (.058)	-.112 (.106)
CTR	-.63 (.121)***		
Comp_Bid	-1.961 (.267)***		
Day 1	-.006 (.01)		
Day 2	-.02 (.01)**		
Day 3	.005 (.01)		
Day 4	-.004 (.01)		
Day 5	-.02 (.01)**		
Day 6	-.017 (.01)*		
Time	.001 (.00)***		
V ^α	Const	Bid	
Const	.243 (.04)***	.131 (.033)***	
Bid		.202 (.034)***	

*Statistically significant at 10%.
**Statistically significant at 5%.
***Statistically significant at 1%.

Table 13
ESTIMATES FOR THE COVARIANCE MATRIX Ω FOR ANOTHER
RETAILER

	<i>CONV</i>	<i>CTR</i>	<i>Bid</i>	<i>Rank</i>
CONV	.196 (.007)	-.099 (.004)	.005 (.001)	-.024 (.008)
CTR	-.099 (.004)	.157 (.009)	-.005 (.001)	.021 (.005)
Bid	.005 (.001)	-.005 (.001)	.012 (.001)	-.012 (.001)
Rank	-.024 (.008)	.021 (.005)	-.012 (.001)	.089 (.002)

Notes: All values are statistically significant at 1%.

DISCUSSION AND CONCLUSION

In this article, we analyze the impact of position on the revenues and profitability of sponsored search advertisements that appear alongside regular algorithmic search results in search engines. A widely held belief in the industry is that the higher the ad placement, the better is the performance. Most of these statements are based primarily on an observed exponential decay in the CTR of the advertisements as a function of their position rather than on a careful analysis of resultant orders and revenues.

We analyze the impact of position on ad profitability using a unique data set generated from a field experiment of an online retailer's ad campaign on Google. This data set documents the daily impressions, clicks, orders, and costs for a select sample of keywords in the ad campaign for different positions for the corresponding advertisements. We also validate our results, using an archival data set for the ad campaign performance of another online retailer. Consistent with the prior literature, our study confirms that CTR decreases rapidly with the rank of the advertisement. However, for advertisers that are interested in maximizing revenues or profit (rather than exposure benefit), this tells only part of the story. Our results show that an advertisement's conversion rate increases with position, and revenue increases with position for more specific keywords. Because the ranking mechanism the search engines use does not account for conversion rate, advertisements placed in the top position do not always maximize revenues. We also show that even for keywords for which the revenue decreases with ad position, the top position may not maximize profit, because costs rapidly decrease with position.

These findings are important to the industry because advertisers are currently engaged in intense bidding wars to secure the top positions in sponsored search results. Our results suggest that these bidding strategies may be based on faulty assumptions about the relationship of CTR, CPC, and conversion probability as a function of position. Specifically, our results suggest that, at least at present, advertisers seeking to maximize transactional benefits are often better off in the short run placing less weight on obtaining top positions. Note that this is not an equilibrium argument, and the strategy will not work in the long run if all advertisers follow the same approach. However, it emphasizes the importance of tracking orders when measuring the effectiveness of sponsored search campaigns.

Our study also points to potential inefficiencies in the auction mechanisms that popular search engines use. If advertisers with the best combination of bid and CTR are assigned the top position and lower positions generate higher revenues for certain keywords, this may be doing

them a disservice. An alternative approach available to search engines is to invest in technologies to track postclick consumer action and to charge advertisers per order (also known as pay-per-action auctions). To this end, we note that several search engines are currently testing pay-per-action auction strategies (see, e.g., Claburn 2007).

Finally, our study sheds light on consumer behavior in sponsored search environments. While CTR decreases with position, conversion rate increases with position. This suggests consumers with greater purchase intent visit lower positions relatively more frequently than those with lower purchase intent. In addition, they are more likely to buy from the same advertisement in a lower position, suggesting a recency bias. Although revisiting the product pages of a previously clicked advertiser requires only a few additional clicks, there is increasing evidence that consumers often associate a relatively high cost with making a few clicks (e.g., Hann and Terwiesch 2003).¹⁰ If this is the case, placing advertisements at lower positions may be an effective way to reach buying consumers without paying more for the top positions.

As with any empirical analysis, there are several limitations of our study. We evaluate the impact of only the top seven positions because of the nature of our data set. However, the top positions garner 80% of the traffic, according to AOL data (see Hearne 2006). Furthermore, although our results explain some information search behavior of consumers at an aggregate level, the aggregate nature of our data limits our ability to account for the actions of individual consumers. This calls for further research using clickstream data to empirically evaluate the behavior of different types of consumers in sponsored search. In addition, we were forced to use only advertiser-specific information to determine the rate at which cost decays with position because sponsored search auctions are now implemented as closed auctions and the true cost of securing other positions is not known. Access to bid data from other advertisers would help increase the accuracy of our findings. However, we do not expect the direction of findings to reverse with such analysis.

An additional limitation is that our analysis of orders is based on measurements taken by a search engine marketing firm that tracked consumer action during the entire search session. This is potentially problematic because consumers may click on an advertisement and visit the advertiser's landing page without converting but return on a later day (even using a different search engine query) to buy the product. In such cases, the future purchases are not properly attributed to the original keyword. In this case, our results for position hold if the subsequent actions are all initiated using refined search engine queries. For example, if a consumer queries "shirt" to shortlist the advertisers and then uses "blue dress shirt" to finally buy the product, our results show that the consumer is still likely to buy when the advertisement corresponding to the second query appears in a lower rather than higher position. Although we determined that this is not true for most of the orders in our data set, our model does not explicitly account for this

¹⁰For example, Hann and Terwiesch (2003) find that the cost of rebidding on an ascending auction is on the order of \$5.00. In this setting, all that is required to rebid is a series of clicks.

possibility. This calls for developing consumer-level models to account for the entire search process to gain further insights.

Finally, our analysis focuses only on transactional benefits from advertising. We believe that this is a reasonable approach in our data setting. However in other settings, nontransactional benefits such as branding and awareness may be more important to advertisers. We find that keywords that help increase awareness and generate orders at a later point do not have a higher chance of success when the corresponding advertisements are placed in higher position. Further research should investigate different strategies consumers use in the buying process and how advertisers should evaluate the performance of the related keywords.

APPENDIX: MCMC ALGORITHMS

We can write the model in the following hierarchical form:

$$\begin{aligned} &U_{kt}^{\text{CONV}} \mid \beta^k, \beta, X_{kt}^{\beta^k}, X_{kt}^{\beta}, \Omega \\ &U_{kt}^{\text{CTR}} \mid \theta^k, \theta, X_{kt}^{\theta^k}, X_{kt}^{\theta}, \Omega \\ &b_k \mid \{U_{kt}^{\text{CONV}}\}, \{U_{kt}^{\text{CTR}}\}, X_{kt}^{\beta^k}, X_{kt}^{\beta}, b, z, \Delta^b, V^b, \Omega \\ &b \mid \{U_{kt}^{\text{CONV}}\}, \{U_{kt}^{\text{CTR}}\}, \{b_k\}, X_{kt}^{\beta^k}, X_{kt}^{\beta}, V^b, \Omega, \bar{b} \\ &\Omega \mid \{U_{kt}^{\text{CONV}}\}, \{U_{kt}^{\text{CTR}}\}, \{b_k\}, \{b\}, X_{kt}^{\beta^k}, X_{kt}^{\beta}, v_\Omega, S_\Omega \\ &V^b \mid \{b_k\}, z, \Delta^b, v, S \\ &\Delta^b \mid \{b_k\}, z, A_0, \bar{\Delta}^b, \end{aligned}$$

where $b_k = [\beta^k \ \theta^k \ \gamma^k \ \alpha^k]$, $b = [\beta \ \theta \ \gamma \ \alpha]$, $V^b = [V^\beta \ V^\theta \ V^\gamma \ V^\alpha]$, $\Delta^b = [\Delta^\beta \ \Delta^\theta \ \Delta^\gamma \ \Delta^\alpha]$, $\bar{b} = [\bar{\beta} \ \bar{\theta} \ \bar{\gamma} \ \bar{\alpha}]$, and $\bar{\Delta}^b = [\bar{\Delta}^\beta \ \bar{\Delta}^\theta \ \bar{\Delta}^\gamma \ \bar{\Delta}^\alpha]$ and where $X_{kt}^{\beta^k}$ are independent variables with keyword-specific coefficients and X_{kt}^{β} are independent variables with common coefficients in Equations 2, 4, 7, and 9.

Note that we provided the approach for estimating a model in which the advertiser is also making a bidding decision (Equation 9). However, for estimating the model parameters for our main data set, we do not include this equation in our analysis. We used 0 as the initial value for elements of b_k , b , and Δ^b and an identity matrix as an initial value for elements of V . Next, we describe the MCMC algorithm.

Step 1

Draw U_{kt}^{CONV} and U_{kt}^{CTR} . We use a data augmentation approach and a random walk Metropolis–Hastings algorithm for sampling:

$$\begin{aligned} U_{kt} &= (U_{kt}^{\text{CONV}}, U_{kt}^{\text{CTR}}) \text{ (Rossi and Allenby 2005);} \\ U_{kt}^{\text{CTR}^{\text{new}}} &= U_{kt}^{\text{CTR}^{\text{old}}} + \delta^{\text{CTR}}, \text{ where } \delta^{\text{CTR}} \sim N(0, .021); \text{ and} \\ U_{kt}^{\text{CONV}^{\text{new}}} &= U_{kt}^{\text{CONV}^{\text{old}}} + \delta^{\text{CONV}}, \text{ where } \delta^{\text{CONV}} \sim N(0, .021). \end{aligned}$$

The draws are accepted with a probability α , where

$$\alpha = \min \left\{ \frac{\exp[-1/2(U_{kt}^{\text{new}} - Y_{kt} - E_{kt})' A (U_{kt}^{\text{new}} - Y_{kt} - E_{kt})] 1(U_{kt}^{\text{new}})}{\exp[-1/2(U_{kt}^{\text{old}} - Y_{kt} - E_{kt})' A (U_{kt}^{\text{old}} - Y_{kt} - E_{kt})] 1(U_{kt}^{\text{old}})}, 1 \right\}$$

and where $1(U_{kt})$ is the likelihood of orders and clicks

$$\begin{aligned} 1(U_{kt}) &= \prod_{k=1}^K \prod_{t=1}^T [(A)_{kt}^{\text{CONV}} \times \Lambda_{kt}^{\text{CTR}}]^{\text{Orders}_{kt}} \\ &\quad \times [(1 - \Lambda_{kt}^{\text{CONV}}) \times \Lambda_{kt}^{\text{CTR}}]^{\text{Clicks}_{kt} - \text{Orders}_{kt}} \\ &\quad \times (1 - \Lambda_{kt}^{\text{CTR}})^{\text{Impressions}_{kt} - \text{Clicks}_{kt}}, \\ e_{kt} &= \begin{bmatrix} e_{kt}^1 \\ e_{kt}^2 \end{bmatrix}, \text{ where } e_{kt}^1 = \ln(\text{bid}_{kt}) - \gamma^k X_{kt}^{\gamma^k} - \gamma X_{kt}^{\gamma} \text{ and} \\ e_{kt}^2 &= \ln(\text{pos}_{kt}) - \alpha^k X_{kt}^{\alpha^k} - \alpha X_{kt}^{\alpha}, \\ E_{kt} &= W_{12} W_{22}^{-1} e_{kt}, \\ A^{-1} &= W_{11} - W_{12} W_{22}^{-1} W_{21}, \\ W_{11} &= \begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{bmatrix}, \quad W_{22} = \begin{bmatrix} \Omega_{33} & \Omega_{34} \\ \Omega_{43} & \Omega_{44} \end{bmatrix}, \\ W_{12} = W_{21} &= \begin{bmatrix} \Omega_{13} & \Omega_{14} \\ \Omega_{23} & \Omega_{24} \end{bmatrix}. \end{aligned}$$

Step 2

Draw $b_k = [\beta^k \ \theta^k \ \gamma^k \ \alpha^k]$. We define

$$\begin{aligned} x_k &= \begin{bmatrix} X_k^{\beta'} & 0 & 0 & 0 \\ 0 & X_k^{\theta'} & 0 & 0 \\ 0 & 0 & X_k^{\gamma'} & 0 \\ 0 & 0 & 0 & X_k^{\alpha'} \end{bmatrix}, \quad y_k = \begin{bmatrix} U_{kt}^{\text{CONV}} - \beta X_{kt}^{\beta} \\ U_{kt}^{\text{CTR}} - \theta X_{kt}^{\theta} \\ \ln(\text{bid}_{kt}) - \gamma X_{kt}^{\gamma} \\ \ln(\text{pos}_{kt}) - \alpha X_{kt}^{\alpha} \end{bmatrix}, \\ V &= \begin{bmatrix} V^\beta & 0 & 0 & 0 \\ 0 & V^\theta & 0 & 0 \\ 0 & 0 & V^\gamma & 0 \\ 0 & 0 & 0 & V^\alpha \end{bmatrix}, \quad \bar{b}_k = \begin{bmatrix} \Delta^\beta Z_k \\ \Delta^\theta Z_k \\ \Delta^\gamma Z_k \\ \Delta^\alpha Z_k \end{bmatrix}, \end{aligned}$$

$$Q_k = [(x_k' \Omega x_k)^{-1} + V^{-1}]^{-1}, \text{ and } \tilde{b}_k = Q_k [x_k' \Omega^{-1} y_k + V^{-1} \bar{b}_k].$$

Then $b_k \sim N(\tilde{b}_k, Q_k)$.

Step 3

Draw $b = [\beta \ \theta \ \gamma \ \alpha]$. We define

$$x = \begin{bmatrix} X^{\beta'} & 0 & 0 & 0 \\ 0 & X^{\theta'} & 0 & 0 \\ 0 & 0 & X^{\gamma'} & 0 \\ 0 & 0 & 0 & X^{\alpha'} \end{bmatrix}, \quad y = \begin{bmatrix} U_{kt}^{\text{CONV}} - \beta^k X_{kt}^{\beta^k} \\ U_{kt}^{\text{CTR}} - \theta^k X_{kt}^{\theta^k} \\ \ln(\text{bid}_{kt}) - \gamma^k X_{kt}^{\gamma^k} \\ \ln(\text{pos}_{kt}) - \alpha^k X_{kt}^{\alpha^k} \end{bmatrix},$$

$$\bar{V} = 100I, \quad \bar{b} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix},$$

$$Q = [(x' \Omega x)^{-1} + \bar{V}^{-1}]^{-1}, \text{ and } \tilde{b} = Q [x' \Omega^{-1} y + \bar{V}^{-1} \bar{b}].$$

Then $b \sim N(\tilde{b}, Q)$.

Step 4

Draw Ω .

$$\Omega \sim \text{IW} \left(v_\Omega + N, \sum_{k=1}^K \sum_{t=1}^T Y_{kt}' Y_{kt} + S_\Omega \right), \text{ where}$$

$$Y_{kt} = \begin{bmatrix} U_{kt}^{CONV} - \beta^k X_{kt}^{\beta k} - \beta X_{kt}^{\beta} \\ U_{kt}^{CTR} - \theta^k X_{kt}^{\theta k} - \theta X_{kt}^{\theta} \\ \ln(\text{bid}_{kt}) - \gamma^k X_{kt}^{\gamma k} - \gamma X_{kt}^{\gamma} \\ \ln(\text{pos}_{kt}) - \alpha^k X_{kt}^{\alpha k} - \alpha X_{kt}^{\alpha} \end{bmatrix},$$

N = number of observations, and $v_{\Omega} = 10, S_{\Omega} = 10I$.

Step 5

Draw $V^{\beta} V^{\theta} V^{\gamma} V^{\alpha}$.

$$V^{\beta} \sim IW \left[\left(v + N, \sum_{k=1}^K (\beta^k - \Delta^{\beta} z_k)' (\beta^k - \Delta^{\beta} z_k) + S \right) \right],$$

where N = number of keywords and $v = 10, S = 10I$;

$$V^{\theta} \sim IW \left[\left(v + N, \sum_{k=1}^K (\theta^k - \Delta^{\theta} z_k)' (\theta^k - \Delta^{\theta} z_k) + S \right) \right],$$

where N = number of keywords and $v = 10, S = 10I$;

$$V^{\gamma} \sim IW \left[\left(v + N, \sum_{k=1}^K (\gamma^k - \Delta^{\gamma} z_k)' (\gamma^k - \Delta^{\gamma} z_k) + S \right) \right],$$

where N = number of keywords and $v = 10, S = 10I$; and

$$V^{\alpha} \sim IW \left[\left(v + N, \sum_{k=1}^K (\alpha^k - \Delta^{\alpha} z_k)' (\alpha^k - \Delta^{\alpha} z_k) + S \right) \right],$$

where N = number of keywords and $v = 10, S = 10I$.

Step 6

Draw $\Delta^{\beta} \Delta^{\theta} \Delta^{\gamma} \Delta^{\alpha}$. Then,

$$\Delta^{\beta} \sim N(\tilde{\Delta}^{\beta}, q_{\beta}),$$

where $q_{\beta} = [(z'_k z_k)^{-1} + A_0]^{-1}$ and $\tilde{\Delta}^{\beta} = q_{\beta} (z'_k \beta^k + A_0 \overline{\Delta}^{\beta})$, with $\overline{\Delta}^{\beta} = 0, A_0 = .01I$;

$$\Delta^{\theta} \sim N(\tilde{\Delta}^{\theta}, q_{\theta}),$$

where $q_{\theta} = [(z'_k z_k)^{-1} + A_0]^{-1}$ and $\tilde{\Delta}^{\theta} = q_{\theta} (z'_k \theta^k + A_0 \overline{\Delta}^{\theta})$, with $\overline{\Delta}^{\theta} = 0, A_0 = .01I$;

$$\Delta^{\gamma} \sim N(\tilde{\Delta}^{\gamma}, q_{\gamma}),$$

where $q_{\gamma} = [(z'_k z_k)^{-1} + A_0]^{-1}$ and $\tilde{\Delta}^{\gamma} = q_{\gamma} (z'_k \gamma^k + A_0 \overline{\Delta}^{\gamma})$, with $\overline{\Delta}^{\gamma} = 0, A_0 = .01I$; and

$$\Delta^{\alpha} \sim N(\tilde{\Delta}^{\alpha}, q_{\alpha}),$$

where $q_{\alpha} = [(z'_k z_k)^{-1} + A_0]^{-1}$ and $\tilde{\Delta}^{\alpha} = q_{\alpha} (z'_k \alpha^k + A_0 \overline{\Delta}^{\alpha})$, with $\overline{\Delta}^{\alpha} = 0, A_0 = .01I$.

REFERENCES

Ansari, Asim and Carl Mela (2003), "E-Customization," *Journal of Marketing Research*, 40 (May), 131-45.
 Arbatskaya, Maria (2007), "Ordered Search," *RAND Journal of Economics*, 38 (1), 119-26.
 Brooks, Niko (2004), "The Atlas Rank Report II: How Search Engine Rank Impacts Conversions," (accessed September 13, 2011), [available at <http://www.atlassolutions.com/pdf/Rank-ReportPart2.pdf>].
 Brucks, Marie (1985), "The Effects of Product Class Knowledge on Information Search Behavior," *Journal of Consumer Research*, 12 (1), 1-16.

Brunel, Frederick F. and Michelle R. Nelson (2003), "Message Order Effects and Gender Differences in Advertising Persuasion," *Journal of Advertising Research*, 43 (3), 330-41.
 Brynjolfsson, Eric, Astrid A. Dick, and Michael D. Smith (2009), "A Nearly Perfect Market? Differentiation vs. Price in Consumer Choice," *Quantitative Marketing and Economics*, 8 (1), 1-33.
 Cary, Matthew, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin et al. (2007), "Greedy Bidding Strategies for Keyword Auctions," in *Proceedings of the 8th ACM Conference on Electronic Commerce*. New York: Association for Computing Machinery.
 Chakravarti, Amitav, Chris Janiszewski, and Gulden Ulkumen (2006), "The Neglect of Prescreening Information," *Journal of Marketing Research*, 43 (November), 642-53.
 Claburn, Thomas (2007), "Google Launches Test of Pay-Per-Action Ads," *InformationWeek*, (March 24), (accessed September 19, 2011), [available at <http://www.informationweek.com/news/198500474>].
 Diehl, Kristin, Laura J. Kornish, and John G. Lynch (2003), "Smart Agents: When Lower Search Cost for Quality Information Increases Price Sensitivity," *Journal of Consumer Research*, 30 (6), 56-71.
 Edelman, Ben and Michael Ostrovsky (2007), "Strategic Bidder Behavior in Sponsored Search Auctions," *Decision Support Systems*, 43 (1), 192-98.
 ———, ———, and Michael Schwarz (2007). "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97 (1), 242-59.
 Elberse, Anita and Jehoshua Eliashberg (2003), "Demand and Supply Dynamics for Sequentially Released Products in International Markets: The Case of Motion Pictures," *Marketing Science*, 22 (3), 329-54.
 Feldman, Jon, S. Muthukrishnan, Martin Pal, and Cliff Stein (2007), "Budget Optimization in Search-Based Advertising Auctions," in *Proceedings of the 8th ACM Conference on Electronic Commerce*. New York: Association for Computing Machinery.
 Feng, Juan, Hemant Bhargava, and David Pennock (2007), "Implementing Sponsored Search in Web Search Engines: Computational Evaluation of Alternative Mechanisms," *INFORMS Journal on Computing*, 19 (1), 137-48.
 Ganchev, Kuzman, Alex Kulesza, Jinsong Tan, Ryan Gabbard, Qian Liu, and Michael Kearns (2007), "Empirical Price Modeling for Sponsored Search," paper presented at WWW2007, Banff, AB, Canada (May 8-12).
 Ghose, Anindya and Sha Yang (2009), "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets," *Management Science*, 55 (10), 1605-1622.
 Godes, David and Dina Mayzlin (2004), "Using Online Conversations to Study Word of Mouth Communication," *Marketing Science*, 23 (4), 545-60.
 Goodman, Andrew (2006), "Bid Fight," *Target Marketing*, (February), (accessed September 13, 2011), [available at <http://www.targetmarketingmag.com/article/understand-search-environment-then-plan-your-keyword-strategy-optimize-your-search-budget-33406/1>].
 Greene, William (1999), *Econometric Analysis*, 4th ed. Upper Saddle River, NJ: Prentice Hall.
 Hallerman, David (2009), "US Online Ad Spend Turns the Corner," *eMarketer*, (December 11), (accessed September 13, 2011), [available at <http://www.emarketer.tv/Article.aspx?R=1007415>].

- Hann, Il-Horn and Christian Terwiesch (2003), "Measuring the Frictional Costs of Online Transactions: The Case of a Name-Your-Own-Price Channel," *Management Science*, 49 (11), 1563–79.
- Häubl, Gerald, Benedict G.C. Dellaert, and Bas Donkers (2010), "Tunnel Vision: Local Behavioral Influences on Consumer Decisions in Product Search," *Marketing Science*, 29 (3), 438–55.
- and Valerie Trifts (2000), "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids," *Marketing Science*, 19 (Winter), 4–21.
- Hausman, Jerry A. (1975), "An Instrumental Variable Approach to Full Information Estimators for Linear and Certain Nonlinear Econometric Models," *Econometrica*, 43 (4), 727–38.
- Hearne, Richard (2006), "SERP Click Through Rate of Google Search Results—AOL-data.tgz—Want to Know How Many Clicks The #1 Google Position Gets?" Red Cardinal blog, (August 12), (accessed September 19, 2011), [available at <http://www.redcardinal.ie/google/12-08-2006/clickthrough-analysis-of-aol-datatz/>].
- Hoque, Abeer Y. and Gerald L. Lohse (1999), "An Information Search Cost Perspective for Designing Interfaces for Electronic Commerce," *Journal of Marketing Research*, 36 (August), 387–94.
- Hosanagar, Kartik and Vadim Cherapanov (2008), "Optimal Bidding in Stochastic Budget Constrained Slot Auctions," in *Proceedings of ACM Conference on Electronic Commerce*. New York: Association for Computing Machinery.
- Johnson, Eric J., Wendy W. Moe, Peter S. Fader, Steven Bellman, and Gerald L. Lohse (2004), "On the Depth and Dynamics of Online Search Behavior," *Management Science*, 50 (3), 299–308.
- Kitts, Brendan, Parameshvya Laxminarayan, Benjamin LeBlanc, and R. Meech (2005), "A Formal Analysis of Search Auctions Including Predictions on Click Fraud and Bidding Tactics," in *Proceedings of the 7th ACM Conference on Electronic Commerce*. New York: Association for Computing Machinery.
- and Benjamin LeBlanc (2004), "Optimal Bidding on Keyword Auctions," *Electronic Markets*, 14 (3), 186–201.
- Lahaie, Sebastian and David M. Pennock (2007), "Revenue Analysis of a Family of Ranking Rules for Keyword Auctions," in *Proceedings of the 8th ACM Conference on Electronic Commerce*. New York: Association for Computing Machinery.
- Lahiri, Kajal and Peter Schmidt (1978), "On the Estimation of Triangular Structural Systems," *Econometrica*, 46 (5), 1217–21.
- Moe, Wendy W. (2003), "Buying, Searching, or Browsing: Differentiating Between Online Shoppers Using In-Store Navigational Clickstream," *Journal of Consumer Psychology*, 13 (1–2), 29–40.
- and Peter S. Fader (2004), "Dynamic Conversion Behavior at e-Commerce Sites," *Management Science*, 50 (3), 326–35.
- Montgomery, Alan L., Kartik Hosanagar, Rammaya Krishnan, and Karen B. Clay (2004), "Designing a Better Shopbot," *Management Science*, 50 (2), 189–206.
- , Shibo Li, Kannan Srinivasan, and John C. Liechty (2004), "Modeling Online Browsing and Path Analysis Using Clickstream Data," *Marketing Science*, 23 (4), 579–95.
- Moorthy, Sridhar, Brian T. Ratchford, and Debabrata Talukdar (1997), "Consumer Information Search Revisited: Theory and Empirical Analysis," *Journal of Consumer Research*, 23 (4), 263–77.
- Newton, Michael A. and Adrian E. Raftery (1994), "Approximate Bayesian Inference with the Weighted Likelihood Bootstrap," *Journal of the Royal Statistical Society, Series B*, 56 (1), 3–48.
- Pieters, Rik G.M. and Tammo H.A. Bijmolt (1997), "Consumer Memory for Television Advertising: A Field Study of Duration, Serial Position, and Competition Effects," *Journal of Consumer Research*, 23 (March), 362–73.
- Rhodes, Edward W., Norman B. Teferman, Elizabeth Cook, and David Schwartz (1979), "T-Scope Tests of Yellow Pages Advertising," *Journal of Advertising Research*, 19, 49–52.
- Rossi, Peter E. and Greg M. Allenby (2003), "Bayesian Statistics and Marketing," *Marketing Science*, 22 (3), 304–328.
- and ——— (2005), *Bayesian Statistics and Marketing*. West Sussex, UK: John Wiley & Sons.
- Ruby, Daniel (2010), "The Value of Google Result Positioning," Chitika Insights, (May 25), (accessed September 15, 2011), [available at <http://insights.chitika.com/2010/the-value-of-google-result-positioning/>].
- Rutz, Oliver and Randolph Bucklin (2010), "From Generic to Branded: A Model of Spillovers in Paid Search Advertising," *Journal of Marketing Research*, 48 (February), 87–102.
- Sherman, Chris (2006), "Searcher Behaviour Research Update," *Search Engine Watch*, (April 10), (accessed December 2009), [available at <http://searchenginewatch.com/article/2049435/Searcher-Behavior-Research-Update/>].
- Srinivasan, Narasimhan and Brian T. Ratchford (1991), "An Empirical Test of a Model of External Search for Automobiles," *Journal of Consumer Research*, 18 (2), 233–42.
- Steel, Emily (2007), "Keywords: a Growing Cost for News Sites; Media Firms Place Bids to Secure Top Positions with Search Engines," *The Wall Street Journal*, (April 30), (accessed September 13, 2011), [available at <http://online.wsj.com/article/SB117788946503386423.html>].
- Urbany, Joel E., Peter R. Dicksoti, and William L. Wilkie (1989), "Buyer Uncertainty and Information Search," *Journal of Consumer Research*, 16 (September), 208–215.
- Varian, Hal R. (2007), "Position Auctions," *International Journal of Industrial Organization*, 25 (6), 1163–78.
- Weber, Thomas A. and Zhiqiang Zheng (2007), "A Model of Search Intermediaries and Paid Referrals," *Information Systems Research*, 18 (4), 414–36.
- Wedel, Michel and Reek Pieters (2000), "Eye Fixations on Advertisements and Memory for Brands: A Model and Findings," *Marketing Science*, 19 (4), 297–312.
- White, Ryan W., Susan T. Dumais, and Jaime Teevan (2009), "Characterizing the Influence of Domain Expertise on Web Search Behavior," paper presented at ACM International Conference on Web Search and Data Mining, Barcelona, Spain (February 9–12).
- and Dan Morris (2007), "Investigating the Querying and Browsing Behavior of Advanced Search Engine Users," in *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York: Association for Computing Machinery, 255–62.
- Wyer, Robert S. and Thomas K. Srull (1986), "Human Cognition in its Social Context," *Psychological Review*, 93 (3), 322–59.
- Yang, Sha and Anindya Ghose (2010), "Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence?" *Marketing Science*, 29 (4), 602–623.
- Zellner, Arnold (1962), "An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests for Aggregation Bias," *Journal of the American Statistical Association*, 57 (298), 348–68.