

Analyzing Search Engine Advertising: Firm Behavior and Cross-Selling in Electronic Markets

Anindya Ghose
Stern School of Business
New York University
New York, NY-10012
aghose@stern.nyu.edu

Sha Yang
Stern School of Business
New York University
New York, NY-10012
syang0@stern.nyu.edu

ABSTRACT

The phenomenon of sponsored search advertising is gaining ground as the largest source of revenues for search engines. Firms across different industries have are beginning to adopt this as the primary form of online advertising. This process works on an auction mechanism in which advertisers bid for different keywords, and final rank for a given keyword is allocated by the search engine. But how different are firm's actual bids from their optimal bids? Moreover, what are other ways in which firms can potentially benefit from sponsored search advertising? Based on the model and estimates from prior work [10], we conduct a number of policy simulations in order to investigate to what extent an advertiser can benefit from bidding optimally for its keywords. Further, we build a Hierarchical Bayesian modeling framework to explore the potential for cross-selling or spillovers effects from a given keyword advertisement across multiple product categories, and estimate the model using Markov Chain Monte Carlo (MCMC) methods. Our analysis suggests that advertisers are not bidding optimally with respect to maximizing profits. We conduct a detailed analysis with product level variables to explore the extent of cross-selling opportunities across different categories from a given keyword advertisement. We find that there exists significant potential for cross-selling through search keyword advertisements in that consumers often end up buying products from other categories in addition to the product they were searching for. Latency (the time it takes for consumer to place a purchase order after clicking on the advertisement) and the presence of a brand name in the keyword are associated with consumer spending on product categories that are different from the one they were originally searching for on the Internet.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms: Performance, Measurement, Economics.

Keywords: *Online advertising, Search engines, Web 2.0, Hierarchical Bayesian modeling, Paid search advertising, Electronic commerce.*

1. INTRODUCTION

Search engines like Google, Yahoo and MSN have discovered that as intermediaries between users and firms, they are in a

unique position to try new forms of advertisements without annoying consumers. In this regard, the advent of sponsored search advertisements – the delivery of relevant, targeted text advertisements as part of the search experience, makes it increasingly possible for firms to attract consumers to their websites. These keyword advertisements are based on customers' own queries and are thus considered far less intrusive than online banner advertisements or pop-ups. In many ways, one could imagine that this enabled a shift in advertising from 'mass' advertising to more 'targeted' advertising. By allotting a specific value to each keyword, an advertiser only pays the assigned price for the people who click on their listing to visit its website. Because listings appear when a keyword is searched for, an advertiser can reach a more targeted audience on a much lower budget. Hence, it is now considered to be among the most effective marketing vehicles available in the online world.

Despite the growth of search advertising, we have little understanding of how consumers respond to sponsored search advertising on the Internet. In this paper, we focus on two previously unexplored questions: (i) For a given set of keywords, what is the spread between the optimal bid prices and the actual cost-per-click incurred by the advertiser in the aftermath of an auction? (ii) Can firms benefit from cross-selling or spillovers in paid search advertising? While an emerging stream of theoretical literature in sponsored search has looked at issues such as mechanism design in keyword auctions, no prior work has empirically analyzed these questions. We adopt the model and estimates from [10] to explore this divergence between actual cost per clicks and the bid price that would maximize advertiser profits. Further, using a panel dataset of several hundred keywords collected from a large nationwide retailer that advertises on Google, we empirically estimate the impact of keyword attributes (such as the presence of *retailer information*, *brand information* and the *length of the keyword*) on consumer purchase propensities across different categories after clicking on a specific keyword. This enables us to evaluate the cross-selling potential of sponsored search by tracking the cross-category spillover effects of a click-through on a given keyword advertisement.

2. DATA

Our data is the same as [10] and contains weekly information on paid search advertising from a large nationwide retail chain,

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2008, April 21–25, 2008, Beijing, China.

ACM 978-1-60558-085-2/08/04.

which advertises on Google.¹ The data span *all keyword advertisements* by the company during a period of three months in the first quarter of 2007, specifically for the 13 calendar weeks from January 1 to March 31. Unlike most datasets used to investigate on-line environments which usually comprise of browsing behavior only, our data are unique in that we have individual level stimulus (advertising) and response (purchase incidence).

Each keyword in our data has a unique advertisement ID. Once the advertiser gets a rank allotted (based on the bid price) to display its textual ad, these sponsored ads show up on the top left, right and bottom of the computer screen in response to a query that a consumer types on the search engine. The serving of a text ad in response to a query for a certain keyword is denoted as an impression. If the consumer clicks on the ad, he is led to the landing page of the advertiser's website. This is recorded as a click, and advertisers usually pay on a per click basis. In the event that the consumer ends up purchasing a product from the advertiser, this is recorded as a conversion. The time between a click and an actual purchase is known as latency. This is usually measured in days. In the majority of cases the value of this variable is 0, denoting that the consumer placed an order at the same time as when they landed on a firm's website.

Our data consists of the number of impressions, number of clicks, the average cost-per-click (CPC), the rank of the keyword, the number of conversions, the total revenues from a click (revenues from conversion) and the average order value for a given keyword for a given week. Given that these are second price auctions, the CPC is likely to be highly correlated with the actual bid price in the case of a successful bid. While a search can lead to an impression, and often to a click, it may not lead to an actual purchase (defined as a conversion). The product of CPC and number of clicks gives the total costs to the firm for sponsoring a particular advertisement. Thus the difference in revenues and costs gives the profits accruing to the retailer from advertising a given keyword in a given week. Our dataset includes 5147 observations from a total of 1799 unique keywords that had at least one positive impression.

2.1 Keyword Characteristics

As described in [10], there are three important keyword specific characteristics for a firm (the advertiser) when it advertises on a search engine. This includes whether the keyword should have (i) retailer-specific information, (ii) brand-specific information, (iii) and the length (number of words) of the keyword. A consumer seeking to purchase a digital camera is as likely to search for a popular brand name such as NIKON, CANON or KODAK on a search engine as searching for the generic phrase "digital camera" on the same search engine. Similarly, the same consumer may search directly for a retailer such as "BEST BUY" or "CIRCUIT CITY" on the search engine. In recognition of these electronic marketplace realities, search engines do not merely sell generic identifiers such as "digital cameras" as keywords, but also well-known brand names that can be purchased by any third-party advertiser in order to attract consumers to its Web site. The length of the keyword is also an important determinant of search and

purchase behavior but anecdotal evidence on this varies across trade press reports. Some studies have shown that the percentage of searchers who use a combination of keywords is 1.6 times the percentage of those who use single-keyword queries [19]. To investigate the impact of the length of a keyword, we constructed a variable that indicates the number of words in a keyword that a user queried for on the search engine (in response to which the paid advertisement was displayed to the user).

The dataset was enhanced by introducing some keyword-specific characteristics such as Brand, Retailer and Length. For each keyword, we constructed two dummy variables, based on whether they were (i) branded or unbranded keywords and (ii) retailer-specific or non-retailer specific keywords. To be precise, for creating the variable in (i) we looked for the presence of a brand name (either a product-specific or a company specific) in the keyword, and labeled the dummy as 1 or 0, with 1 indicating the presence of a brand name. For (ii), we looked for the presence of the advertising retailer's name in the keyword, and then labeled the dummy as 1 or 0, with 1 indicating the presence of the retailer's name. There were no keywords that contained both retailer name and brand name information. This enabled a clean classification in our data. This classification is similar in notion to [2, 13] who classify user queries in search engines as *navigational* (searching for a specific firm or retailer), *transactional* (searching for a specific product) or *informational* (longer keywords).

Table 1: Summary Statistics (Keyword level)

Variable	Mean	Std. Dev.	Min	Max
Impressions	383.37	2082.08	1	97424
Clicks	32.915	519.555	0	33330
Orders	0.483	8.212	0	527
Click-through Rate (CTR)	0.008	0.059	0	1
Conversion Rate	0.013	0.073	0	1
Cost-per-Click (CPC)	0.294	0.173	0.005	1.410
Lag Rank	4.851	6.394	1	64
Log (Lag Profit)	0.106	1.748	-5.160	10.710
Rank	5.179	7.112	1	64
Lag CTR	0.007	0.053	0	1
Retailer	0.057	0.232	0	1
Brand	0.398	0.490	0	1

3. POLICY SIMULATIONS

A primary goal of research is to evaluate and recommend optimal policies for marketing actions. One way of doing this is to assess current decision-making behavior and compare them with optimal decision-making behavior. Towards this objective, we estimate the optimal bid price for each keyword and assess how much the advertiser's decision (actual bid price) deviates from the optimal bid price based on our model estimates. We adopt the model from [10] and estimate it using Markov Chain Monte Carlo (MCMC) methods (see [16] for a detailed review of such models). We use the Metropolis-Hastings algorithm with a random walk chain to generate draws. ([4]). Rather than describe the entire Hierarchical

¹ The firm is a Fortune-500 firm but due to the nature of the data sharing agreement between the firm and us, we are unable to reveal the name of the firm.

Bayesian model, we refer interested readers to [10] and provide a summary description of the theoretical framework below.

Assume for search keyword i at week j , there are n_{ij} click-throughs among N_{ij} impressions (the number of times an advertisement is displayed by the retailer), where $n_{ij} \leq N_{ij}$. Suppose that among the n_{ij} click-throughs, there are m_{ij} click-throughs that lead to purchases, where $m_{ij} \leq n_{ij}$. Let us further assume that the probability of having a click-through is p_{ij} and the probability of having a purchase is q_{ij} . In our model, a consumer faces decisions at two levels – one, when she sees a keyword advertisement, she makes decision whether or not to click it; two, if she clicks on the advertisement, she can take any one of the following two actions – make a purchase or not make a purchase.

Thus, there are three types of observations. First, a person clicked through and made a purchase. The probability of such an event is $p_{ij}q_{ij}$. Second, a person clicked through but did not make a purchase. The probability of such an event is $p_{ij}(1 - q_{ij})$. Third, an impression did not lead to a click-through or purchase. The probability of such an event is $1 - p_{ij}$. Then, the probability of observing (n_{ij}, m_{ij}) is given by:

$$f(n_{ij}, m_{ij}, p_{ij}, q_{ij}) = \frac{N_{ij}!}{m_{ij}!(n_{ij} - m_{ij})!(N_{ij} - n_{ij})!} \times \{p_{ij}q_{ij}\}^{m_{ij}} \{p_{ij}(1 - q_{ij})\}^{n_{ij} - m_{ij}} \{1 - p_{ij}\}^{N_{ij} - n_{ij}} \quad (1)$$

Thereafter, click-through rates, conversion rates, cost-per-click, and keyword ranks are analyzed by jointly modeling the consumers' search and purchase behavior, the advertiser's bid pricing behavior, and the search engine's keyword rank allocating behavior. The decision of whether to click and purchase in a given week is modeled as a function of the probability of advertising exposure (for example, through the *Rank* of the keyword) and individual differences, both observed and unobserved heterogeneity (for example, through the keyword-specific attributes like *Retailer*, *Brand* and *Length*). The advertiser's bid price decision is modeled as a function of keyword attributes and other variables such as lagged values of *Rank* and *Profit*. The search engine's ranking decision is modeled as a function of keyword attributes and other factors such as *CPC* and lagged values of *CTR* in accordance with institutional practices. The Metropolis-Hastings algorithm with a random walk chain is adopted to generate draws in the MCMC methods ([4]).

Using the parameter estimates from the click-through, conversion and rank models from [10] and the data on click-through rates, conversion rates, revenues and actual *CPC* of each advertisement, we estimate the expected profit of the firm. We assume the advertiser determines the optimal bid price for each keyword to maximize the expected profit (Π) from each consumer impression of the advertisement:

$$\Pi_{ij} = p_{ij}(q_{ij}r_{ij} - CPC_{ij}) \quad (2)$$

In equation (2), p_{ij} is the expected click-through rate for keyword i at week j , q_{ij} is the expected conversion rate conditional on a click through, r_{ij} is the expected revenue from a conversion that is

observed from our data, and CPC_{ij} is the actual cost per click paid by the advertiser to the search engine for each keyword. p_{ij} , q_{ij} and $Rank_{ij}$ are predicted based on equations (4.2), (4.6) and (4.15) respectively in [10], using the estimates obtained from the proposed model. We conduct the optimization routine to maximize the expected profit from each consumer impression of the advertisement for each keyword at each week using the grid search.

Our simulation results highlight that there is a considerable amount of spread in the optimal bid prices and the actual cost per click for a given keyword, with the average deviation being 23.3 cents per bid. Given that this is a second price auction and the firm actually pays the bid price of the next highest bidder plus a small increment of 1 cent, we find that a vast majority of the keyword CPCs actually highlight that the firm is overbidding relative to the optimal bid price. Specifically, 6% of the CPCs are below the optimal bid prices with the average difference being 67 cents, while the remaining 94% of the CPCs (and thus the firm's bid price) are above the optimal bid price with the average difference being 28.7 cents.

We also examined the deviation from the optimal bid prices based on whether the keyword advertisement had retailer or brand information. On an average, the firm was underbidding by 11.2 cents for each ad that had retailer information in it and was overbidding by 16.4 cents for each ad that had brand information in it. For those keywords that did not have retailer or brand information in them the firm was generally overbidding with the range going from 25.4 cents to 27.7 cents. These results are very intuitive: the lack of competition for retailer-specific keywords is likely to be driving the underbidding behavior while the presence of intense competition in branded or generic keywords would be driving the overbidding behavior.

Consequently, there is a significant amount of divergence between optimal expected profits and actual profits accruing to the firm from their current bid prices, with the average difference being 1.14 times the expected profits with actual bid prices. Next we examined the sample based on overbidding or underbidding behavior. We found that the average difference in profits is 1.15 times the expected profits with actual bid prices when the firm is overbidding. When the firm is underbidding, the ratio is 1.05. When the firm is underbidding, the ratio is 1.05. Figure 1a and 1b highlight the differences from the optimal and actual bid prices.

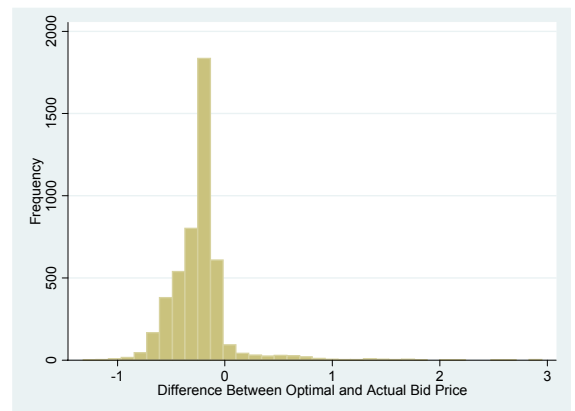


Figure 1a: Distribution of the Difference between Optimal and Actual Bids

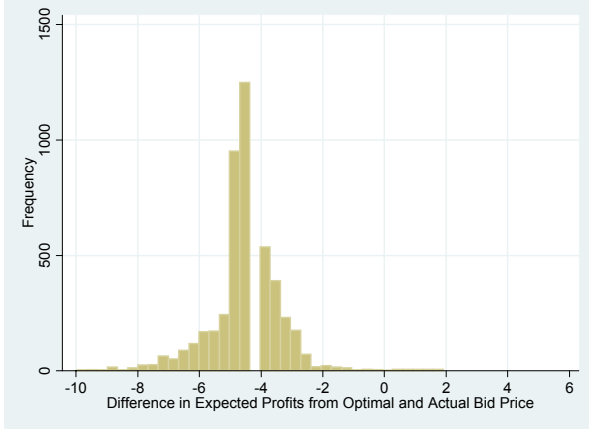


Figure 1b. Distribution of the Log of Difference in Expected Profits using Optimal and Actual Bids

In order to investigate how the three keyword level covariates are associated with optimal bid prices, we ran OLS regressions with keyword-level random effects. The dependent variable was the optimal bid price. Our analysis reveals that the presence of retailer-specific information (*Retailer*) or brand-specific (*Brand*) information leads to an increase in the optimal bid price, while longer keywords (*Length*) is associated with a lower optimal bid price. Specifically, the presence of retailer and brand information should lead to an increase in the optimal bid prices by 21.5% and 3.9%, respectively while an increase in the length of the keyword by one word should lead to a decrease in the bid price by 2.3%. Note that these results are in contrast to the results from [10] wherein using actual bid prices we found that the firm is actually incurring a lower *CPC* when it has either retailer or brand information in the keywords, and incurs a higher *CPC* for longer keywords.

To summarize, while the firm is exhibiting some learning behavior over time in terms of deciding on bid prices based on its rank and profit in the previous period, our simulations suggest that it can improve its profits dramatically by bidding optimally. Further, it would be better off by placing higher bids on keyword advertisement that either have retailer or brand information in them, and lower bids as keywords become longer. Moreover, we also find that expected profits from retailer-specific keywords are likely to be much higher than those from brand-specific keywords.

4. ECONOMETRIC MODEL: IMPACT OF SPONSORED SEARCH ON CROSS-SELLING

In this section, we investigate the impact of sponsored search advertising in a given category on consumer’s propensity to buy products across other categories. Our dataset has detailed information on the various categories of products that were eventually purchased by consumers after they had clicked on any given paid advertisement. There are six product categories in our data: bath, bedding, electrical appliances, home décor, kitchen and dining. Due to the confidentiality agreement with the firm that gave us the data, we are not able to reveal any more details

about the individual products within these categories. Since, our analysis is about the cross-selling potential of a given product-based advertisement, we exclude advertisements that only have the *retailer* information in them but no product information. Hence, we focus on the 166 keywords that have some product or product category information imbedded in them. Table 2 reports the summary statistics of the data. As shown, the average spending is 79 dollars on the searched product category, and 21.8 dollars on the non-searched product category. The average latency is about one day. These statistics provide some evidence suggesting that keyword advertising can lead to purchases on a non-searched product category, and consumers may wait for a while after starting the search to complete an order.

Table 2: Summary Statistics of the Cross-Selling Data

Variable	Mean	Std. Dev.	Min	Max
Order Value– Own	79.007	100.812	0	930
Order Value– Cross	21.805	78.534	0	1249
Latency	1.062	3.527	0	29
Rank	1.257	1.999	1	40.25
Brand	0.883	0.322	0	1
Length	2.410	0.956	0	5

We cast our model in a hierarchical Bayesian framework and estimate it using Markov chain Monte Carlo methods (see [19] for a detailed review of such models). We use the Metropolis-Hastings algorithm with a random walk chain to generate draws. ([4]). Each order can lead to a purchase from the searched product category and/or from any of the other five non-searched product categories. We model the consumer purchase behavior as a two-stage decision process. In the first stage, the consumer decides on how much to spend on the searched product category. We adopt the Tobit model specification to account for a large number of zeros in consumer spending on either the searched product category or non-searched product categories. Let’s denote y_{ij}^{own} as the money spent on the searched product category in order *j* for the searched keyword *i*. We assume there is latent spending intention (z_{ij}^{own}) that determines how much to spend on the searched product category, that is,

$$y_{ij}^{own} = z_{ij}^{own} \quad \text{if } z_{ij}^{own} > 0 \tag{3.1}$$

$$y_{ij}^{own} = 0 \quad \text{if } z_{ij}^{own} \leq 0 \tag{3.2}$$

We model the latent buying intention of the searched category as:

$$z_{ij}^{own} = \alpha_i^{own} + \sum_{k=1}^{K-1} \gamma_k^{own} Search_{ik} + \beta_1^{own} Latency_j + \beta_2^{own} Rank_j + \beta_3^{own} Brand + \beta_4^{own} Length + \varepsilon_{ij}^{own} \tag{3.3}$$

where $Search_{ik} = 1$ if the searched category is the k^{th} product category for keyword *i*, and $Search_{ik} = 0$ if the searched category is not the k^{th} product category for keyword *i*. $Latency_{ij}$

is the time duration in number of days between the search and the order j for keyword i . $Rank_{ij}$ is the average rank of keyword i for order j . $Brand_i$ is a dummy variable indicating whether a brand name is included in the search keyword i . $Length_i$ is the number of words included in the search keywords i . We have a total of 6 product categories, that is, $K=6$ and without loss of generality, we use category 6 as the baseline. To complete the model specification, we assume the following distributions regarding the error term and intercept term:

$$\varepsilon_{ij}^{own} \sim N(0, \sigma_{own}^2) \tag{3.4}$$

$$\alpha_i^{own} \sim N(\overline{\alpha^{own}}, \tau_{own}^2) \tag{3.5}$$

In the second stage, the consumer decides on how much to spend on the non-searched product categories in total conditional on the spending on the searched product category. Let's denote y_{ij}^{cross} as the money spent on the non-searched product category in order j for the searched keyword i . We assume there is latent spending intention (z_{ij}^{cross}) that determines how much to spend on the non-searched product category, that is,

$$y_{ij}^{cross} = z_{ij}^{cross} \quad \text{if } z_{ij}^{cross} > 0 \tag{3.6}$$

$$y_{ij}^{cross} = 0 \quad \text{if } z_{ij}^{cross} \leq 0 \tag{3.7}$$

We model the latent buying intention of the non-searched category as follows:

$$z_{ij}^{cross} = \alpha_i^{cross} + \sum_{k=1}^{K-1} \gamma_k^{cross} Search_{ik} + \beta_1^{cross} Latency_{ij} + \beta_2^{cross} Rank_{ij} + \beta_3^{cross} Brand_i + \beta_4^{cross} Length_i + \beta_5^{cross} y_{ij}^{own} + \varepsilon_{ij}^{cross} \tag{3.8}$$

To complete the model specification, we assume the following distributions regarding the error term and intercept term:

$$\varepsilon_{ij}^{cross} \sim N(0, \sigma_{cross}^2) \tag{3.9}$$

$$\alpha_i^{cross} \sim N(\overline{\alpha^{cross}}, \tau_{cross}^2) \tag{3.10}$$

Equations (3.1) – (3.3), and (3.6) – (3.8) lead to a non-linear fully non-recursive simultaneous equations model. Note that γ_k^{own} , γ_k^{cross} as well as $\beta_1^{own} - \beta_5^{own}$ are modeled as fixed effects due to the empirical identification with our data.

4.1 Results and Analysis

We next discuss the findings from our analysis. In table 3a, the coefficient, γ_1^{own} is negative and significant suggesting that consumer average spending on the searched category is lower in category 1 than category 6. On the other hand, the coefficient, γ_2^{own} is positive and significant suggesting that the consumer average spending on the searched category is higher in category 2

than category 6. The coefficients, γ_3^{own} , γ_4^{own} , and γ_5^{own} are statistically insignificant suggesting that on an average, and consumers spend the same amount in each of these categories (3, 4 and 5) as they do in category 6 when they search for a product in each of these categories.

What are the main factors that affect this kind of consumer behavior? Based on the estimates in Table 3a and 3b, we find that *Latency* tends to decrease consumer spending on the searched category, but increase their average spending on the non-searched category. Recall that latency is the time between when consumers click on an advertisement and when they actual purchase the product from the website. Intuitively, this result suggests that if consumers delay the final purchase of the product after the initial click on the ad, they are likely to digress from their original spending intention in the searched category and increasing their purchase of products in other non-searched categories. Note also that the coefficient of y^{own} is negative suggesting that if a consumer has already spent a lot on the category that they had originally searched for, then they are likely to spend less on the other categories.

Table 3a: Estimates on Consumer Spending on the Searched Product Category

Intercept	Latency	Rank	Brand	Length
$\overline{\alpha^{own}}$	β_1^{own}	β_2^{own}	β_3^{own}	β_4^{own}
8.349	-0.410	0.024	-1.756	-1.061
(2.974)	(0.079)	(0.145)	(1.496)	(0.900)
Search ₁	Search ₂	Search ₃	Search ₄	Search ₅
γ_1^{own}	γ_2^{own}	γ_3^{own}	γ_4^{own}	γ_5^{own}
-17.845	6.569	4.619	-0.252	-4.739
(4.255)	(2.250)	(2.658)	(2.263)	(3.100)
σ_{own}^2	τ_{own}^2			
114.361	12.167			
(6.910)	(4.740)			

Table 3b: Estimates on Consumer Spending on Non-Searched Product Category

Intercep t	Latency	Rank	Brand	Length	y^{own}
$\overline{\alpha^{cross}}$	β_1^{cross}	β_2^{cross}	β_3^{cross}	β_4^{cross}	β_5^{cross}
-9.973	0.583	-0.311	7.256	1.770	-0.086
(4.926)	(0.131)	(0.327)	(2.345)	(1.486)	(0.016)
Search ₁	Search ₂	Search ₃	Search ₄	Search ₅	
γ_1^{cross}	γ_2^{cross}	γ_3^{cross}	γ_4^{cross}	γ_5^{cross}	
12.718	-11.600	-17.056	-3.576	-2.714	
(4.767)	(3.478)	(4.486)	(3.319)	(4.128)	
σ_{cross}^2	τ_{cross}^2				
260.199	7.779				
(27.040)	(3.236)				

Interestingly, we find that the presence of *Brand* information in the search keyword advertisement does not affect the amount that consumers spend on the category that they originally searched for on the search engine. However, note from Table 3b that it does significantly increase consumers' spending in the other categories. This implies that the presence of a brand name in a keyword advertisement can have a strong switching effect on consumer's purchasing propensities. It has a similar flavor to the bait and switch strategies used by retailers, when they attract consumers to their stores based advertisements in one category and then induce them to buy a product in a different category addition to the original product, perhaps through some marketing promotion. Thus, our analysis indicates a strong cross-selling potential of a sponsored search advertisement that contains a brand name in it. The statistically significant estimates of γ_1^{cross} , γ_2^{cross} , and γ_3^{cross} in Table 3b indicate that there are complementary demands for three product categories at each purchase incidence. In particular, we see in Table 3b that categories 1, 2, and 3 (bath, bedding and electrical appliances) exhibit the strongest opportunities for cross-selling.

We find that neither *Rank* nor the *Length* has any impact on consumers' spending either on the searched category or the non-searched category. This is not too surprising. Both these attributes are likely to influence consumer click-through behavior but are unlikely to affect their latent spending intention once they have already landed on the retailer's web page. As a robustness check, we also fit a model that controls for the potential endogeneity in *Rank*. We found similar results on the coefficient estimates. We also included dummies for different categories of landing pages such as search page, shop, home page, information page, product page and category page. This did not affect the qualitative nature of the results, and moreover the estimates on the dummies were not statistically significant.

5. RELATED WORK

Our paper is related to several streams of research. First, it contributes to recent research in online advertising in economics and marketing by providing the first known empirical analysis of sponsored search keyword advertising. Much of the existing academic (e.g., [5], [6], [7]) on advertising in online world has focused on measuring changes in brand awareness, brand attitudes, and purchase intentions as a function of exposure. This is usually done via field surveys or laboratory experiments using individual (or cookie) level data. In contrast to other studies which measure (individual) exposure to advertising via aggregate advertising dollars ([12]), we use data on individual search keyword advertising exposure. [17] looks at online banner advertising. Because banner ads have been perceived by many consumers as being annoying, traditionally they have had a negative connotation associated with it. Moreover, it was argued that since there is considerably evidence that only a small proportion of visits translate into final purchase ([3], [6], [18]), click-through rates may be too imprecise for measuring the effectiveness of banners served to the mass market. Interestingly however, [14] found that banner advertising actually increases purchasing behavior, in contrast to conventional wisdom. These studies therefore highlight the importance of investigating the impact of other kinds of online advertising such as search keyword advertising on actual purchase behavior, since the success of keyword advertising is also based on consumer click-through rates.

A large literature in economics sees advertising as necessary to signal some form of quality ([11]). There is also an emerging theoretical stream of literature exemplified by [9, 10] that examines auction price and mechanism design in keyword auctions. Despite the emerging theory work, very little empirical work exists in online search advertising. The handful of empirical studies that exist in search engine advertising have mainly analyzed publicly available data from search engines. [1] looks at the presence of quality uncertainty and adverse selection in paid search advertising. [13] classifies queries as *informational*, *navigational*, and *transactional* based on the expected type of content destination desired and analyze click through patterns of each. In a paper related to our work, [20] studied the conversion rates of hotel marketing keywords to analyze the profitability of different campaign management strategies. Finally, in our prior work [10], we only analyzed the impact of keyword attributes on consumer and firm behavior. This paper goes well beyond it by conducting policy simulations to examine the optimality of advertiser strategies. We also substantially extend that work in this paper by examining the potential for cross-selling products through sponsored search advertisement and quantifying the actual impact of specific variables like brand and latency.

Our paper is also related to the stream of work in cross-selling. Amongst the first papers that formally model sequential ordering and the cross-selling opportunities is [14]. Their research applies latent trait analysis to position financial services and investors along a common continuum. [15] present next product-to-purchase models that can be used to predict what is to be purchased next and when. [16] model consumers' sequential acquisition decisions for multiple products and services, a behavior that is common in service and consumer technology industries. We thus contribute to the literature by demonstrating the cross-selling potential of paid search advertising in an online context, thereby supplementing the existing stream of work on cross-selling.

6. CONCLUSIONS AND FUTURE WORK

The phenomenon of sponsored search advertising is gaining ground as the largest source of revenues for search engines. In this research, we aim to analyze how advertiser's actual cost per click may differ from optimal bid prices. In addition to this, our second objective is to enhance our understanding of how sponsored search advertising affects consumer purchasing patterns on the Internet by analyzing its cross-selling potential. Using a unique panel dataset of several hundred keywords collected from a nationwide retailer that advertises on Google, we empirically model the relationship between different keyword attributes and consumer search and purchase behavior across multiple product categories. We use a Hierarchical Bayesian modeling framework and estimate the model using Markov Chain Monte Carlo (MCMC) methods.

We conduct simulations to assess the relative profit impact from changes in bid prices, and find that despite some learning, the advertiser is not bidding optimally. What are some of the implications? Retailer-name searches are navigational searches, and are analogous to a customer finding the retailer's phone number or address in the White Pages. These searches are driven by brand awareness generated by catalog mailings, TV ads, etc, and are likely to have come from more 'loyal' consumers. Even

though the referral to the retailer's website came through a search engine, the search engine had very little to do with generating the demand in the first place. On the other hand, searches on product or manufacturer specific brand names are analogous to consumers going to the Yellow Pages—they know they need a product or service, but don't yet know where to buy it [10]. These are likely to be “competitive” searches. Even for loyal buyers, a “branded” search means the searcher is surveying the market and is vulnerable to competition. If the advertiser wins the click and the order, that implies they have taken market share away from a competitor. Thus, retailer-specific keywords are likely to be searched and clicked by 'loyal' consumers who are inclined towards buying from that retailer whereas brand-specific keywords are likely to be searched and clicked by the 'shoppers or searchers' who can easily switch to competition. Our policy simulations suggest that the average profitability from conversions generated by 'retailer' keywords is much higher than that from 'brand' keywords. Our results thus provide some managerial insights for an advertiser of sponsoring such retail store keywords (retailer-specific keywords) with national-brand keywords (brand-specific keywords).

We have shown some evidence that although the average click-through and conversion rates are typically low in sponsored search, there are other potential benefits from such advertising. Specifically, retailers can not only refine their keyword purchases on search engines, but also set up relevant cross-selling opportunities on their own websites by advertising 'brand-specific' keywords. The strategy is that when a consumer searches for a specific product and lands deep within the retailer's website by clicking on its keyword advertisement, the retailer can pair that product with other products that sell well with that keyword and prominently feature them on its website. This provides a retailer with an opportunity to not only convert someone on the product they had searched for, but also get other opportunities for cross-selling. From the retailer's perspective, there could be synergies in promoting both categories simultaneously rather than separately. Indeed anecdotal evidence suggests that retailers are engaging in the practice of looking up the most-searched and the top-converting keywords on their websites, and bidding for them on search engines. They are taking cross-selling reports from other marketing mix campaigns and putting up the top cross-selling product for the searched product on the same page.

Interestingly, we find that latency in purchases is not necessarily detrimental for a firm that is sponsoring the keyword advertisement. While it is in general associated with a reduction in product purchases in the category that the consumer was originally searching for, it increases consumers' spending in other product categories. In a way, it has an impact similar to a bait and switch strategy. This effect is particularly strong in keywords that have a brand name in it, since consumers who click on branded keywords typically tend to spend more on other categories than the one they were originally searching for. Thus, online advertisers can focus on investing more often in such keywords relative to the generic keywords, especially if the cannibalization effect of drawing out consumers from one category is smaller relative to revenue expansion effect. From the point of view of the manufacturer, such dependencies across categories may be exploited by running cooperative promotions within brands but across categories. Of course, such decisions would need a detailed profitability analysis based not only on the potential from cross-

selling in other product categories but also the performance of the keyword in its own category.

We are cognizant of the limitations of our paper. These limitations arise primarily from the lack of information in our data. For example, we do not have data on competition. That is, we do not know the keyword auction ranks or other performance metrics such as click-through rates and conversion rates of the keyword advertisements of the competitors of the firm whose data we have used in this paper. Future research can use data on competition and highlight some more insights on how firms should manage a paid search campaign by running more detailed policy simulations that incorporate competitive bid prices. Further, we do not have any knowledge of the other marketing variables such as any promotions during consumers' search and purchase visits. Future work can investigate the value to firms from participating in such sponsored search advertising by comparing the performance of sponsored searches with natural searches using a common pool of keywords during the same time period. By collecting information on the ranks and page numbers of the natural search listings for the same keywords as those in paid search, one can study the impact of natural search listings on paid search advertisements and vice-versa. We hope that this study will generate further interest in exploring this important emerging area in web search.

7. ACKNOWLEDGMENTS

The authors would like to thank the anonymous company that provided data for this study. This work was partially supported by a grant from the NET Institute and Marketing Science Institute. Anindya Ghose also acknowledges the generous financial support of the National Science Foundation through CAREER Award IIS-0643847. The usual disclaimer applies.

8. REFERENCES

- [1] Animesh A., Ramachandran, V., and Viswanathan, S. Quality uncertainty and adverse selection in sponsored search markets. *Proceedings of the 2006 NET Institute Conference*.
- [2] Broder, A. Taxonomy of web search, *SIGIR Forum*, vol. 36, 2002, 3-10.
- [3] Chatterjee, P., Hoffman, D., and Novak, T. Modeling the clickstream: implications for web-based advertising efforts. *Marketing Science*: 22(4), 2003, 520-541.
- [4] Chib, S., and Greenberg, E. Understanding the Metropolitan-Hastings algorithm. *The American Statistician*, 49, 1995, 327-335.
- [5] Cho, C., Lee, J., and Tharp, M. Different forced-exposure levels to banner advertisements, *Journal of Advertising Research*, 41(4), 2001, 45-56.
- [6] Dahlen, M. Banner advertisements through a new lens, *Journal of Advertising Research*, 41(4), 2001, 23-30.
- [7] Danaher, P., and Mullarkey, G. Factors affecting online advertising recall: A Study of Students, *Journal of Advertising Research*, September, 2003, 252-267.
- [8] Edelman, B., Ostrovsky, M., and Schwarz, M. Internet advertising and the generalized second-price auction: Selling

- billions of dollars worth of keywords. *American Economic Review*, 97(1), 2007, 242-259.
- [9] Feng, J, Bhargava, H., and Pennock, D. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *Inform Journal on Computing*, 19(1), 2007, 137-148.
- [10] Ghose, A., and Yang, S. An Empirical Analysis of Sponsored Search Performance in Search Engine Advertising. *Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM)*, 2008.
- [11] Grossman, G., and Shapiro, C. Informative advertising with differentiated products. *Review of Economic Studies*. 51(1), 1984, 63-81.
- [12] Ilfeld, J., and Winer, R. Generating website traffic. *Journal of Advertising Research*, 42, 2002, 49-61.
- [13] Jansen, B., and Spink, A. The effect on click-through of combining sponsored and non-sponsored search engine results in a single listing, *Proceedings of the 2007 Workshop on Sponsored Search Auctions, WWW Conference, 2007*.
- [14] Kamakura W., S. Ramaswami, and Srivastava, R. 1991. Applying latent trait analysis in the evaluation of prospects for cross-selling of financial services. *Int. Journal. of Research in Marketing*, 8: 329-349.
- [15] Knott A., A. Hayes, and Neslin, S. Next-Product-to-Buy models for cross-selling applications. *Journal of Interactive Marketing*, 16 (3), 2002, 59-75.
- [16] Li, S., Sun, B., and Wilcox, R. Cross-selling sequentially ordered products: An application to consumer banking services. *Journal of Marketing Research*, XLII, May, 2005, 233-239.
- [17] Manchanda, P., Dubé, J., Goh, K., and Chintagunta, P. The effect of banner advertising on internet purchasing. *Journal of Marketing Research*, 43(1), 2006, 98-108.
- [18] Moe, W., and Fader, P. Dynamic conversion behavior at e-commerce sites. *Management Science*, 50(3), 2003, 326-335.
- [19] Rossi, P. E., and Allenby, G. Bayesian statistics and marketing. *Marketing Science*, 22, 2003, 304-329.
- [20] Rutz, O., and R. E. Bucklin 2007. A model of individual keyword performance in paid search advertising, Unpublished Mimeo, UCLA.