

A Dynamic Model of Sponsored Search Advertising

Song Yao

Carl F. Mela¹

September 15, 2010

¹Song Yao (email: s-yao@kellogg.northwestern.edu, phone: 847-467-2767) is an Assistant Professor of Marketing at the Kellogg School of Management, Northwestern University, Evanston, Illinois, 60208. Carl F. Mela (email: mela@duke.edu, phone: 919-660-7767, fax: 919-681-6245) is a Professor of Marketing, The Fuqua School of Business, Duke University, Durham, North Carolina, 27708. The authors would like to thank seminar participants at Cornell University, Dartmouth College, Duke University, Emory University, Erasmus University, Georgia Institute of Technology, Georgia State University, Harvard Business School, INSEAD, London Business School, New York University, Northwestern University, Ohio State University, Rice University, Stanford University, University of British Columbia, University of California at Berkeley, University of California at Davis, University of California at Riverside, University of Chicago, University of Maryland, University of Rochester, University of Southern California, University of Texas, University of Tilburg, University of Wisconsin, Yale University, and the 2008 Marketing Science Conference, NET Institute Conference 2009, NBER Summer Institute 2009, AMA Summer Conference 2009 as well as J.P. Dubé, Wes Hartmann, Günter Hitsch, Han Hong, Wagner Kamakura, Anja Lambrecht, Andrés Musalem, Harikesh Nair, Peter Rossi and Rick Staelin for their feedback. We gratefully acknowledge financial support from the NET Institute (www.netinst.org) and the Kauffman Foundation.

Abstract: A Dynamic Model of Sponsored Search Advertising

Sponsored search advertising is ascendant – Jupiter Research reports expenditures rose 28% in 2007 to \$8.9B and will continue to rise at a 26% CAGR, approaching 1/2 the level of television advertising and making it one of the major advertising trends to affect the marketing landscape. Yet little empirical research exists to explore how the interaction of various agents (*searchers*, *advertisers*, and the *search engine*) in keyword markets affects consumer welfare and firm profits. The dynamic structural model we propose serves as a foundation to explore these outcomes. We fit this model to a proprietary data set provided by an anonymous search engine. These data include consumer search and clicking behavior, advertiser bidding behavior, and search engine information such as keyword pricing and website design.

With respect to *advertisers*, we find evidence of dynamic bidding behavior. Advertiser value for clicks on their links averages about 26 cents. Given the typical \$22 retail price of the software products advertised on the considered search engine, this implies a conversion rate (sales per click) of about 1.2%, well within common estimates of 1-2% (gamedaily.com). With respect to *consumers*, we find that frequent clickers place a greater emphasis on the position of the sponsored advertising link. We further find that about 10% of consumers do 90% of the clicks.

We then conduct several policy simulations to illustrate the effects of changes in *search engine* policy. First, we find the search engine obtains revenue gains of 1% by sharing individual level information with advertisers and enabling them to vary their bids by consumer segment. This also improves advertiser revenue by 6% and consumer welfare by 1.6%. Second, we find that a switch from a first to second price auction results in truth telling (advertiser bids rise to advertiser valuations). However, the second price auction has little impact on search engine profits. Third, consumer search tools lead to a platform revenue increase of 2.9% and an increase of consumer welfare by 3.8%. However, these tools, by reducing advertising exposures, lower advertiser profits by 2.1%.

Keywords: Sponsored Search Advertising, Two-sided Market, Dynamic Game, Structural Models, Empirical IO, Customization, Auctions

1 Introduction

Sponsored search is one of the largest and fastest growing advertising channels. In January of 2010 alone, Internet users conducted 15.2B searches using the top 5 American search engines compared to 13.5B in the previous January, indicating a robust 13% year over year increase.¹ In the United States, annual advertising expenditures on sponsored search is forecast to grow to \$25B by 2012.² By contrast, overall 2007 television advertising spending in the United States is estimated to be \$62B, an increase of only 0.7% from the preceding year.³ Hence, search engine marketing is becoming a central component of the promotional mix in many organizations.

Given the increasing ubiquity of sponsored search advertising, the topic has seen substantially increased attention in marketing as of late (Ghose and Yang, 2009; Rutz and Bucklin, 2007; Rutz and Bucklin, 2008; Goldfarb and Tucker, 2008). To date, empirical research on keyword search has been largely silent on the perspective of the search engine, the competition between advertisers, and the behavior of the searcher. Given that the search engine interacts with advertisers and searchers to determine the price and consumer welfare of the advertising medium (and hence its efficacy), our objective is to broaden this stream of research to incorporate the role of all three agents: the search engine, the advertisers, and the searchers. This exercise enables us to determine the role of search engine marketing strategy on the behavior of advertisers and consumers as well as the attendant implications for search engine revenues. Our key contributions include:

1. From a theoretical perspective, we conceptualize and develop an integrated model of web searcher, advertiser and search engine behavior. Much like Yao and Mela (2008), we construct a model of a two-sided network in an auction context. One side of the two-sided network includes the searchers who generate revenue for the advertiser. On the other side of the network are advertisers whose bidding behavior determines the revenue of the search en-

¹“January 2010 U.S. Search Engine Rankings,” comScore, Inc. (<http://ir.comscore.com/releasedetail.cfm?ReleaseID=444505>). “January 2009 U.S. Search Engine Rankings,” comScore, Inc. (<http://ir.comscore.com/releasedetail.cfm?ReleaseID=366442>).

²“US Interactive Marketing Forecast, 2007 to 2012,” Forrester Research, 2007.

³“Insider’s Report,” 2007, McCann WorldGroup, Inc.; <http://www.tns-mi.com/news/01082007.htm>

gine. In the middle lies the search engine. The goal of the search engine is to price consumer information, set auction mechanisms, and design webpages to elucidate product information so as to maximize its profits.

2. From a substantive point of view, we offer concrete marketing policy recommendations to the search engine. In particular, the two-sided network model of keyword search we consider allows us to address the effect of the following policy simulations (and would enable us to address many others) on auction house and advertiser profits as well as consumer welfare:

- **Search Tools.** Many search engines, especially specialized ones such as `Shopping.com`, provide users options to sort/filter search results using certain criteria such as product prices. On one hand, the search tools may mitigate the desirability of bidding for advertisements because these tools can remove less relevant advertisements. This would lower search engine revenues. On the other hand, these tools can also attract more users to the site, leading to a potential increase in advertising exposures and searchers. This would increase revenues. Our analysis indicates that positive consumer effects on search engine profits (5.5%) outweigh the corresponding negative advertiser effects on search engine profits (-2.6%) and that overall the sort/filter options enhance platform profits by 2.9%. Consistent with this result, there is a corresponding increase in consumer welfare of 3.8% and an attendant loss in advertiser profits of 2.1%.
- **Segmentation and Targeting.** Most search engines auction keywords across all market segments. However, it is possible to auction keywords by segment. This targeting tends to reduce competition between advertisers within segments as markets are sliced more narrowly, leading to lower bids and hence lower potential revenues for the search engine. Yet targeting also enhances the efficiency of advertising, which tends to increase advertiser bids. Overall, we find that the latter effect dominates (2.1%) the former effect (-1.1%) and that search engine revenue increases 1% by purveying keywords by consumer market segments. Moreover, we find advertiser profits improve by 6% (from

reduced competition in bidding and more efficient advertising) and consumer welfare (as measured by utility) increases 1.9%. Hence, this change leads to considerable welfare gains across all agents.

- **Mechanism Design.** The wide array of search pricing mechanisms raises the question of which auction mechanism is the best in the sense of incenting advertisers to bid more aggressively thereby yielding maximum returns for the search engine. We consider two common mechanisms: a first price auction (as used by the considered firm in our analysis) and a second price auction (wherein a firm pays the bid of the next highest bidder). Virtually no revenue gains accrue to the platform from a second price auction (0.02%). However, advertiser bids under second price auction are close to bidders' true values (bids average 98% of valuations), while bids under the first price auction are much lower (70%). This finding is consistent with theory that suggests first price auctions lead to bid shading and second price auctions lead to truth telling (Edelman et al., 2007). Hence, we lend empirical validation to the theoretical literature on auction mechanisms in keyword search.

3. From a methodological view, we develop a dynamic structural model of keyword advertising. This dynamic is induced by the search engine's use of past advertising performance when ranking current advertising bids. The dynamic aspect of the problem requires the use of some recent innovations pertaining to the estimation of dynamic games in economics (e.g., Bajari et al., 2007). Overall, we find that there is a substantial improvement in model fit when the advertiser's strategic bidding behavior is considered, consistent with the view that their bidding behavior is dynamic. One key finding from this model is that advertisers in our application have an average value per click of \$0.26. Given that the average price of software products advertised on the site in our data is about \$22, this implies these advertisers expect about 1.2% (i.e., $\$0.26/\22) of clicks will lead to a purchase. This is consistent with the industry average of 1-2% reported by `GameDaily.com`, suggesting good external validity for our model.

Though we cast our model in the context of sponsored search, we note that the problem, and hence the conceptualization, is even more general. Any interactive, addressable media format (e.g., DVR, satellite digital radio) can be utilized to implement similar auctions for advertising. For example, with the convergence in media between computers and television in DVRs, simple channel or show queries can be accompanied by sponsored search, and this medium may help to offset advertising losses arising from ads skipping by DVR users (Bronnenberg et al., 2009; Kempe and Wilbur, 2009). In such a notion, the research literature on sponsored search auctions generalizes to a much broader context, and our model serves as a basis for exploring search based advertising.

The remainder of this paper proceeds as follows. First we overview the relevant literature to differentiate our analysis from previous research. Given the relatively novel research context, we then describe the data to help make the problem more concrete. Next, we outline the details of our model, beginning with the clicking behavior of consumers and concluding with the advertiser bidding behavior. Subsequently, we turn to estimation and present our results. We then explore the role of targeted bidding, advertising pricing, and webpage design by developing policy simulations that alter the search engine marketing strategies. We conclude with some future directions.

2 Recent Literature

Research on sponsored search, commensurate with the topic it seeks to address, is nascent and growing. Heretofore this literature can be characterized along two distinct dimensions: theoretical and empirical. The theoretical literature details how agents (e.g., advertisers) are likely to react to different pricing mechanisms. In contrast, the empirical literature measures the effect of advertising on consumer response in a given market but not the reaction of these agents to changes in the platform environment (e.g., advertising pricing, information state or the webpage design of the platform). By integrating the theoretical and empirical research streams, we develop a complete representation of the role of pricing and information in the context of keyword search.

Foundational theoretical analyses of sponsored search include Edelman et al. (2007), Varian (2007), Chen and He (2006), Athey and Ellison (2008), Katona and Sarvary (2008), Iyengar and

Kumar (2006), and Feng (2008). Summarizing the key insights from this stream of work, we note that i) there are three types of agents interacting in the sponsored search context, Internet users who engage in keyword search, advertisers that bid for keywords, and the search platform, ii) searchers affect advertisers bidding behavior by reacting to the search engine's web page design and hence advertisers payoffs, iii) bidders affect searcher behavior by the placement of their advertisements on the page, and iv) changes in advertiser and consumer behavior are incumbent upon the strategies of the platform.

In spite of these insights, several limits remain. First, because equilibrium outcomes are incumbent upon the parameters of the system, it is hard to characterize precisely how agents will behave. This implies it would be desirable to estimate a model of keyword search in order to measure these behaviors. Second, a static advertiser game over bidding periods is typically assumed, which is inconsistent with the pricing practices used by search engines. Search engines commonly use the preceding period's click-throughs together with current bids to determine advertising placement, making this an inherently dynamic game. Third, this research typically assumes no asymmetry in information states between the advertiser and the search engine even though the search engine knows individual level clicking behaviors and the advertiser does not. We redress these issues in this paper.

Empirical research on sponsored search advertising is also proliferating (including Rutz and Bucklin, 2007; Rutz and Bucklin, 2008; Ghose and Yang, 2009). Though extant empirical research on sponsored search establishes a firm link between advertising, slot position, and revenues – and indicates that these effects can differ across advertisers, some limitations of this stream of work remain. First, it emphasizes a single agent (one advertiser), making it difficult to predict how advertisers in an oligopolistic setting might react to a change in the policy of the search engine. Further, an advertiser's value to the search engine pertains not only to its direct payment to the search engine but also to the indirect effect that advertiser has on the intensity of competition during bidding. Second, the advertisers' actions affect search engine users and vice-versa. For example, with alternative advertisers being placed at premium slots on a search result page, it is likely that

users' browsing behaviors will be different. As advertisers make decisions with the consideration of users' reactions, any variations of users' behaviors provide feedback on advertisers' actions and thus will ultimately affect the search engine revenue.

Integrating these two research streams suggests it is desirable to both model and estimate the equilibrium behaviors of all the agents in a network setting. In this regard, sponsored search advertising can be characterized as a two-sided market wherein searchers and advertisers interact on the platform of the search engine (Rochet and Tirole, 2006). This enables us to generalize a structural modeling approach advanced by Yao and Mela (2008) to study two-sided markets. However, additional complexities exist in the keyword search setting including: i) the aforementioned information asymmetry between advertisers and the search engine and ii) the substantially more complex auction pricing mechanism used by search engines relative to the fixed fee auction house pricing considered in Yao and Mela (2008). Moreover, unlike the pricing problem addressed in Yao and Mela (2008), sponsored search bidding is inherently dynamic owing to the use of lagged advertising click rates to determine current period advertising placements. Hence we incorporate the growing literature of two-step dynamic game estimation (e.g., Hotz and Miller, 1993; Bajari et al., 2007; Bajari et al., 2008; Pesendorfer and Schmidt-Dengler, 2008). Instead of explicitly solving for the equilibrium dynamic bidding strategies, the two-step estimation approach assumes that observed bids are generated by equilibrium play and then use the distribution of bids to infer underlying primitive variables of bidders (e.g., the advertiser's expectation about the return from advertising). A similar method is also used in an auction context in Jofre-Bonet and Pesendorfer (2003). Equipped with these advertiser primitives, we solve the dynamic game played by the advertiser to ascertain how changes in search engine policy affect equilibrium bidding behavior.

3 Empirical Context

The data underpinning our analysis is drawn from a major search engine for high technology consumer products. Within this broad search domain, we consider search for music management software because the category is relatively isolated in the sense that searches for this product do

not compete with others on the site.⁴ The category is a sizable one for this search engine as well. Along with the increasing popularity of MP3 players, the use of music management PC software is increasing exponentially, making this an important source of revenue. The goal of the search engine is to enable consumers to identify and then download trial versions of these software products before their final purchase.⁵ It is important to note that the approach we develop can readily generalize to other contexts and that we consider this particular instantiation to be an illustration of a more general approach.

3.1 Data Description

The data are comprised of three files, including:

- Bidding file. Bidding is logged into a file containing the bidding history of all active bidders from January 2005 to August 2007. It records the exact bids submitted, the time of each bid submission, and the resulting monthly allocation of slots. Hence, the unit of analysis is vendor-bid event. These data form the cornerstone of our bidding model.
- Product file. Product attributes are kept in a file that records, for each software firm in each month, the characteristics of the software they purvey. This file also indicates the download history of each product in each month.
- Consumer file. Consumer log files record each visit to the site and are used to infer whether downloads occur as well as browsing histories. A separate but related file includes registration information and detailed demographics for those site visitors that are registered. These data are central to the bidding model in the context of complete information.

⁴The search engine defines music management broadly enough that an array of different search terms (e.g., MP3, iTunes, iPod, lyric, etc.) yield the same search results for the software products in this category. Hence we consider the consumer decision of whether to search for music software on the site and whether to download given a search. This search algorithm allows us to abstract away from issues pertaining to consumer search and advertisers bidding across multiple keywords. Recognizing the importance of these issues, we call for future research on these dimensions.

⁵A “click” and a “download” are essentially the same from the perspectives of the advertiser, consumer, and search engine. In the “click” case, a consumer makes several clicks to investigate and compare products offered by different vendors and then makes a final purchase. In the “download” case, a consumer downloads several products and makes the comparison before final purchase. Hence there is no difference for a “click” and a “download” in the current context. We use “click” and “download” interchangeably throughout the paper.

We detail each of these files in turn.

3.1.1 Bidding File

Most search engines yield “organic” search results that are often displayed as a list of links sorted by their relevance to the search query (Bradlow and Schmittlein, 2000). Sponsored search involves advertisements placed above or along side the organic search results. Given that users are inclined to view the topmost slots in the page (Ansari and Mela, 2003), advertisers are willing to pay a premium for these more prominent slots (Goldfarb and Tucker, 2008).

To capitalize on this premium, advertising slots are auctioned off by search engines. Advertisers specify bids on a per-click basis for a search term. While there is considerable variation in the nature of the auctions they use, the most widely adopted approach is the one developed by Google. Google’s algorithm factors in not only the level of the bid, but the expected click-through rate of the advertiser. This enhances search engine revenue because these revenues depend not only on the per-click bid, but also the number of clicks a link receives. Winning advertisers pay the next bidder’s bid (adjusted for click-through rates).⁶

The mechanism used by the firm we consider is similar to that of Google except that the considered search engine uses a first price auction in place of a second price auction (we intend to compare the efficacy of this mechanism to that of Google in our policy experiments). Winning bids are denoted as *sponsored search results* and the site flags these as sponsored links. The site we consider affords up to five premium slots which is far less than the 400 or so products that would appear at the search engine. Losing bidders and non-bidders are listed beneath the top slots on the page and like previous literature we denote these listings as *organic search results*.

The search engine collects bidding and demographic data on all advertisers (products attributes, products download history, and bids from active bidders). Table 1 reports summary statistics for the bidding files. At this search engine, bids were submitted on a monthly basis. Over the 32 months

⁶With a simplified setting, Edelman et al. (2007) show that the Google practice may result in an equilibrium with bidders’ payoffs equivalent to the Vickrey-Clarke-Groves (VCG) auction, whereas VCG auction has been proved to maximize total payoffs to bidders. Iyengar and Kumar (2006) further show that under some conditions the Google practice induces VCG auction’s dominant “truth-telling” bidding strategy, i.e., bidders will bid their own valuations.

from January 2005 to August 2007, 322 bids (including zeros) were submitted by 21 software companies.⁷ As indicated in Table 1, bidders on average submitted about 22 positive bids in this interval (slightly less than once per month). The average bid amount (conditioned on bidding) was \$0.20 with a large variance across bidders and time.

Table 1: Bids Summary Statistics

	Mean	Std. Dev.	Minimum	Maximum
Non-zero Bids (ϕ)	19.55	8.32	15	55
Non-zero Bids/Bidder	6.40	10.46	1	30
All Bids (ϕ)	8.14	11.04	0	55
Bids/Bidder	23.13	9.68	1	32

3.1.2 Product File

Searching for a keyword on this site results in a list of relevant software products and their respective attributes. Attribute information is stored in a product file along with the download history of all products that appeared in this category from January 2005 to August 2007. In total, these data cover 394 products over 32 months. The attributes include the price of the non-trial version of a product, backward compatibility with preceding operating systems (e.g., Windows 98 and Windows Server 2003), expert ratings provided by the site, and consumer ratings of the product. Trial versions typically come with a 30-day license to use the product for free, after which consumers are expected to pay for its use. Expert ratings at the site are collected from several industrial experts of these products. The consumer rating is based on the average feedback score about the product from consumers. Table 2 give summary statistics for all products as well as active bidders' products. Based on the compatibility information, we sum each product's operating system compatibility dummies and define this summation as a measure for that product's compatibility with older operating systems. This variable is later used in our estimation.

Overall, active bidders' products have higher prices, better ratings, and more frequent updates.

⁷Since some products were launched after January 2005, they were not observed in all periods.

Table 2: Product Attributes and Downloads

	Mean	Std. Dev.	Minimum	Maximum
All Products				
Non-trial Version Price \$	16.65	20.43	0	150
Expert Rating (if rated)	3.87	0.81	2	5
Average Consumer Rating (if rated)	3.89	1.31	1	5
Months Lapse Since Last Update	15.31	9.88	1	31
Compatibility Index	3.29	1.47	0	5
Number of Downloads/(Product×Month)	1367.29	9257.16	0	184442
Bidders' Products				
Non-trial Version Price \$	21.97	15.87	0	39.95
Expert Rating (if rated)	4	0.50	3	5
Average Consumer Rating (if rated)	4.06	0.91	2.5	5
Months Lapse Since Last Update	2.38	0.66	1	3
Compatibility Index	3.51	1.51	0	5
Number of Downloads/(Product×Month)	1992.12	6557.43	0	103454

3.1.3 Consumer File

The consumer file contains the log files of consumers from May 2007 to August 2007. This file contains each consumer's browsing log when they visit the search engine both within the search site and across Internet properties owned by the search site. The consumer file also has the registration information for those that register.

The browsing log of a consumer indicates whether the consumer made downloads and, if yes, which products she downloaded. Upon a user viewing the search results of software products, the search engine allowed the consumer to sort the results based on some attributes such as the ratings; consumers can also filter products based on some criteria such as whether a product's non-trial version is free. The browsing log records the sorting and filtering actions of each consumer. Prior to sorting and filtering, the top five search results are allocated to sponsored search slots and the remaining slots are ordered by how recently the software has been updated. There is a small, discrete label indicating whether a search result is sponsored, and sorting and filtering will often remove these links from the top five premium slots.

As the demographic information upon the registration is only optional, the dataset provides little if any reliable demographics of consumers. Hence we focus instead upon whether a consumer

is a registered user of the search engine and on their past search behavior at the other website properties, in particular whether they visited any music related site (which should control for the consumers' interests in music).

3.2 The Dynamics of Advertiser Bidding

As the search engine considers advertisers' past downloads when assigning current placements, there exists the potential for dynamic bidding behavior on the part of advertisers. Advertisers can bid lower amounts for the same placement with a large number of preceding period downloads.


To further illustrate dynamic bidding behavior in our data, we consider two non-parametric spline regressions. One regresses advertiser bids on past downloads and the update recency of the product (because the site returns a higher organic rank to more recently updated products). Another considers advertiser bids on past downloads and total past competing products' downloads. Figure 1 plots the results. For all levels of update recency and lagged competing downloads, there is a strong inverse relationship between bid levels and past downloads, suggesting that advertisers do account for past downloads when making bidding decisions.⁸ The second regression affords additional evidence of dynamic bidding; when competitors have a large number of lagged downloads, the advertiser bids more aggressively to offset its competitors' bidding advantage.⁹

By itself, the negative autocorrelation between downloads and bids does not necessarily imply advertisers are strategic; rather, advertisers may simply be myopic, reacting to their downloads in the preceding period. Accordingly, when we develop our model in the next section, we shall consider the possibility that advertisers are not forward-looking (Section 6.2.1). Results from that analysis are also consistent with dynamic bidding behavior.

⁸In the regression of bids on past downloads and update recency, this effect of past downloads is moderated slightly by update recency and is generally lowest for recently updated products. The moderating effect of update recency may be a consequence of advertisers having less of an incentive to promote a recently updated product, in light of its advantaged position in the organic search section.

⁹The results demonstrated in Figure 1 could be an artifact of pooling bidders' observations together. Thus, we consider analogous nonparametric analyses for three frequent advertisers, and find the results to be similar.

Figure 1: The Relationship between Bids, Past Downloads, Update and Competing Products



4 Model

The model incorporates behaviors of the agents interacting on the search engine platform: i) advertisers who bid to maximize their respective profits and ii) utility maximizing consumers who decide whether to click on the advertiser’s link. For any given policy applied by the search engine, this integrated model enables us to predict equilibrium revenues for the search engine (the consumer-advertiser interactions are analogous to a sub-game contingent on search engine behavior). The behavior of the bidder (advertiser) is dependent on the behavior of the consumer as consumer behavior affects advertiser expectations for downloads and, hence, their bids. The behavior of the consumer is dependent upon the advertiser because the rank of the advertisement affects the behavior of the consumer. Hence, the behaviors are interdependent. We first expost the consumer model and then solve the bidder problem conditioned on the consumers’ behaviors.

4.1 Consumer Model

Advertiser profit (and therefore bidding strategy) is incumbent upon their forecast of consumer downloads for their products $d_j^t(k, X_j^t; \Omega_c)$, where k denotes the position of the advertisement on the search engine results page, X_j^t indicate the vector of attributes of advertiser j ’s product at period t , and Ω_c are parameters to be estimated.¹⁰ Thus, we seek to develop a forecast for $d_j^t(k, X_j^t; \Omega_c)$


¹⁰In our application, we treat the periodicity of t as monthly because that is consistent with the bidding process. To explore the robustness of our findings to this treatment, we re-estimate the consumer model at a bi-weekly level and

and the attendant consequences for bidding. To be consistent with the advertisers information set, we base these forecasts of consumer behavior solely on statistics observed by the advertiser: the aggregate download data and the distribution of consumers characteristics. Later, in the policy section of the paper, we assess what happens to bidding behavior and platform revenues when disaggregate information is revealed to advertisers by the platform. We begin by describing the consumer’s download decision process and how it affects the overall number of downloads.

4.1.1 The Consumer Decision Process

Figure 2 overviews the decisions made by consumers. In any given period t , the consumer’s problem is whether and which software to select in order to maximize their utility. The resolution of this problem is addressed by a series of conditional decisions.

Figure 2: Consumer Decisions



First, the consumer decides whether she should search on the category considered in this analysis (C1). We presume that the consumer will search on the site if it maximizes her expected utility.¹¹

Conditioned upon engaging a search, the consumer next decides whether to sort and/or filter the results (C2). The two search options lead to the following 4 options for viewing the results: $\kappa = \{0 \equiv \text{neither}, 1 \equiv \text{sorting but not filtering}, 2 \equiv \text{not sorting but filtering}, 3 \equiv \text{sorting and filtering}\}$.¹² We assume that the consumer will choose the option that maximizes her expected utility. We assume that the consumer will choose the option that maximizes her expected utility. We assume that the consumer will choose the option that maximizes her expected utility.

¹¹Though we do not explicitly model the consumer’s decision to search across different terms, product categories or competitors, our model incorporates an "outside option" that can be interpreted as a composite of these alternative behaviors.

filtering}.¹² For each option, the set of products returned by the search engine differs in terms of the number and the order of products. Consumers choose the sorting/filtering option that maximizes their expected utility.

Finally, the consumer chooses which, if any products to download (C3). We presume that consumers choose to download software if it maximizes their expected utility. We discuss the modeling details for this process in a backward induction manner (C3–C1).

Download We assume that consumers exhibit heterogeneous preferences for products and download those alternatives that maximize their expected utility. We specify consumer i of preference segment g to have underlying latent utility $\tilde{u}_{ijt}^{g\kappa} = \tilde{w}_{ijt}^g - c_{ijt}^{g\kappa}$ for downloading software j in period t . In particular, \tilde{w}_{ijt}^g represents the expected benefit from the usage of the downloaded alternative j whereas $c_{ijt}^{g\kappa}$ can be interpreted as the opportunity cost (disutility) of time spent on locating the product. Letting a index product attributes, we have:

$$\tilde{w}_{ijt}^g = \tilde{\alpha}_j^g + \sum_a x_{jat} \tilde{\beta}_a^g + \tilde{e}_{ijt} \quad (1)$$

where

- $\tilde{\alpha}_j^g$ is the segment specific intercept for product j ;
- x_{jat} is the level of observed attribute a of product j ;
- $\tilde{\beta}_a^g$ is consumer i 's “taste” regarding product attribute a , which is segment specific;
- \tilde{e}_{ijt} is individual idiosyncratic preference shock, realized after the sorting/filtering decision.

The shocks are independently distributed over individuals, products and periods as zero mean normal random variables.

We assume that the search cost of locating a product $c_{ijt}^{g\kappa}$, is a function of its slot position, k_{jt}^κ , because consumers tend to view a webpage from the top down and may spend more time to locate

¹²We categorize sorting/filtering based on the most prevalent behaviors observed in the data. Sorting by ratings and/or filtering by price (free or not) account for 83% observations using sorting/filtering options. We also experiment a specification with all sorting/filtering options included but the model AIC deteriorates from -12491.2 to -12525.6 and our key insights are unaffected. As a result, we present the more parsimonious specification.

a product if the product is placed at the bottom of the page (Ansari and Mela, 2003).¹³ Specifically,

$$-\tilde{c}_{ijt}^{g\kappa} = \tilde{\theta}^g k_{jt}^\kappa + \tilde{e}_{ijt}^c \quad (2)$$

where $\tilde{\theta}^g$ is segment specific cost parameters on slot ranking; \tilde{e}_{ijt}^c is individual cost shock that is independently distributed across people, products and periods as a mean zero normal random variable.

Hence the net utility of product j becomes

$$\begin{aligned} \tilde{u}_{ijt}^{g\kappa} &= \tilde{w}_{ijt}^g - \tilde{c}_{ijt}^{g\kappa} \\ &= \tilde{\alpha}_j^g + \sum_a x_{jat} \tilde{\beta}_a^g + \tilde{\theta}_0^g k_{jt}^\kappa + \tilde{\varepsilon}_{ijt}^{g\kappa} \end{aligned} \quad (3)$$

where $\tilde{\varepsilon}_{ijt}^{g\kappa} = \tilde{e}_{ijt} + \tilde{e}_{ijt}^c$.¹⁴

To allow the variances of download errors ($\tilde{\varepsilon}_{ijt}^{g\kappa}$) and sorting/filtering errors ($\xi_{it}^{g\kappa}$, which will be detailed below) to differ, both must be properly scaled (cf., Train, 2003, Chapter 2). Hence we have the following assumption:

Assumption 1: $\tilde{\varepsilon}_{ijt}$'s are independently and identically distributed normal random variables with mean 0 and variance normalized to $(\delta^g)^2$. $\xi_{it}^{g\kappa}$'s are independently and identically distributed Type I extreme value random variables.

Under assumption 1, we may re-define the utility in Equation 3 as

$$u_{ijt}^{g\kappa} = \delta^g (\bar{u}_{ijt}^{g\kappa} + \varepsilon_{ijt}) \quad (4)$$

$$\bar{u}_{ijt}^{g\kappa} = \alpha_j^g + \sum_a x_{jat} \beta_a^g + \theta_0^g k_{jt}^\kappa \quad (5)$$

¹³With an additional dummy variable of “left vs. right” which interacts with k_{jt}^κ , this specification can be easily extended to accommodate search results that are sorted both from left to right and from top to bottom such as those at Google.

¹⁴We also consider a specification wherein we include a dummy variable for sponsored links to ascertain whether there is a signaling value of sponsorship over and above link order. Inconsistent with this conjecture, model fit decreases from -12491 to -12513 and the estimate is insignificant.

where $\{\alpha_j^g, \beta_a^g, \theta^g, \varepsilon_{ijt}\} = \{\tilde{\alpha}_j^g, \tilde{\beta}_a^g, \tilde{\theta}^g, \tilde{\varepsilon}_{ijt}\} / \delta^g$; $\bar{u}_{ijt}^{g\kappa}$ is the scaled “mean” net utility and $\varepsilon_{ijt} \sim N(0, 1)$. The resulting choice process is a multivariate probit choice model.¹⁵ Letting $d_{ijt} = 1$ indicate download (and $d_{ijt} = 0$ no download), we have

$$d_{ijt} = \begin{cases} 1 & \text{if } u_{ijt}^{g\kappa} \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

and the probability of downloading conditional on parameters $\{\alpha_j^g, \beta_a^g, \theta^g\}$ is

$$\begin{aligned} \Pr(d_{ijt} = 1) &= \Pr(u_{ijt}^{g\kappa} \geq 0) \\ &= \Pr(\delta^g(\bar{u}_{ijt}^{g\kappa} + \varepsilon_{ijt}) \geq 0) \\ &= \Pr(-\varepsilon_{ijt} \leq \bar{u}_{ijt}^{g\kappa}) \\ &= \Phi(\bar{u}_{ijt}^{g\kappa}) \end{aligned} \quad (7)$$

where $\Phi(\cdot)$ is the standard normal distribution CDF.

Although consumers know the distribution of the product utility error terms ($\tilde{\varepsilon}_{ijt}^{g\kappa}$), these error terms do not realize before the sorting/filtering (C2) and search (C1) decisions (cf. Hong and Shum, 2006; Hortacsu and Syverson, 2004; Kim et al., 2009).¹⁶ Hence, consumers can only form an expectation about the total utilities of all products under a given sorting/filtering option κ prior to choosing that option. Viewed in this light, the choice of a sorting and filtering strategy is informative about consumer preferences and provides an additional source of information to identify their preferences.

¹⁵We consider an alternative specification that allows the utilities across products to be correlated. Using a compound symmetric covariance structure for the product errors, we find decreased model fit (AIC: -12491 vs. -12503). It can be shown that, under the weak assumptions that (1) the consumer allocates her time between searching/browsing and the outside options (such as leisure time), and (2) it is not optimal to allocate all time to searching/browsing (i.e., there is no corner solution), the consumer download problem reduces to a multivariate independent choice probit model. The discussion as an Appendix can be requested from the authors.

¹⁶In an alternative model, we relax the assumption that consumers know the attributes and replace it with the less restrictive assumption that consumers only know the empirical *distribution* of the attribute levels. Hence, these consumers need to integrate over this uncertainty in their sort and filter decisions. The model fit deteriorates mainly because of the simulation errors (AIC: -12491.2 vs. -12550.2), but there is little impact on the models’ parameter estimates.

Sorting and Filtering Prior to making a download decision, consumers face several sorting and filtering decisions which are indexed as $\kappa = 0, 1, 2, 3$ – corresponding to no sorting or filtering, no sorting but filtering, sorting but no filtering and both sorting and filtering, respectively. We expect consumers to choose the option that maximizes their expected download utility.

Let $U_{it}^{g\kappa}$ denote the total expected utility from products under option κ , which can be calculated based on Equation 3:

$$U_{it}^{g\kappa} = \sum_j E_\varepsilon(u_{ijt}^{g\kappa} | u_{ijt}^{g\kappa} \geq 0) \Pr(u_{ijt}^{g\kappa} \geq 0). \quad (8)$$

This definition reflects that a product's utility is realized only when it is downloaded. Hence, the expected utility $E_\varepsilon(u_{ijt}^{g\kappa} | u_{ijt}^{g\kappa} \geq 0)$ is weighted by the download likelihood, $\Pr(u_{ijt}^{g\kappa} \geq 0)$. The expectation, $E_\varepsilon(\cdot)$, is taken over the random preference shocks $\varepsilon_{ijt}^{g\kappa}$.

In addition to the expected download utility, $U_{it}^{g\kappa}$, individuals may accrue additional benefits or costs for using sorting/filtering option κ that are known to the individuals but not observed by researchers. These benefits and costs might accrue through unobserved browsing experience or time constraints. Denote such unobserved benefits or costs of the sort/filter decision by $\eta^{g\kappa} + \xi_{it}^{g\kappa}$'s, where $\eta^{g\kappa}$ is an intercept term and $\xi_{it}^{g\kappa}$ is a random error term. The total utility of search option κ is thus given by

$$z_{it}^{g\kappa} = \eta^{g\kappa} + U_{it}^{g\kappa} + \xi_{it}^{g\kappa}. \quad (9)$$

Consumers choose the option of sorting/filtering that leads to the highest total utility $z_{it}^{g\kappa}$.

With $\xi_{it}^{g\kappa}$ following a Type I extreme value distribution (Assumption 1), the choice of sorting/filtering becomes a logit model such that

$$\Pr(\kappa)_{it}^g = \frac{\exp(\eta^{g\kappa} + U_{it}^{g\kappa})}{\sum_{\kappa'=0}^3 \exp(\eta^{g\kappa'} + U_{it}^{g\kappa'})} \quad (10)$$

To better appreciate the properties of this model, note that $U_{it}^{g\kappa}$ in Equation 8 can be written in

a closed form:¹⁷

$$\begin{aligned}
U_{it}^{g\kappa} &= \sum_j E_\varepsilon(u_{ijt}^{g\kappa} | u_{ijt}^{g\kappa} \geq 0) \cdot \Pr(u_{ijt}^{g\kappa} \geq 0) \\
&= \delta^g \sum_j \left(\bar{u}_{ijt}^{g\kappa} + \frac{\phi(\bar{u}_{ijt}^{g\kappa})}{\Phi(\bar{u}_{ijt}^{g\kappa})} \right) \cdot \Phi(\bar{u}_{ijt}^{g\kappa}).
\end{aligned} \tag{11}$$

With such a formulation, the factors driving the person's choice of filtering or sorting become more apparent:

- Filtering eliminates options with negative utility, such as highly priced products (because consumer price sensitivity is negative). As a result, the summation in Equation 11 for the filter option will increase as the negative $\bar{u}_{ijt}^{g\kappa}$ are removed. This raises the value of the filter option suggesting that price sensitive people are more likely to filter on price.
- Sorting re-orders products by their attribute levels. Products that appear low on a page will typically have lower utility regardless of their product content (because consumer slot rank sensitivity is negative). For example, suppose a consumer relies more on product ratings. By moving more desirable items that have high ratings up the list, sorting can increase the $\bar{u}_{ijt}^{g\kappa}$ for these items, thereby increasing the resulting summation in Equation 11 and the value of this sorting option.

¹⁷For a normal random variable x with mean μ , standard deviation σ and left truncated at a (Greene, 2003), $E(x|x \geq a) = \mu + \sigma \lambda(\frac{a-\mu}{\sigma})$, where $\lambda(\frac{a-\mu}{\sigma})$ is the hazard function such that $\lambda(\frac{a-\mu}{\sigma}) = \frac{\phi(\frac{a-\mu}{\sigma})}{1-\Phi(\frac{a-\mu}{\sigma})}$.

Hence with $u_{ijt}^{g\kappa} \sim N(\delta^g \bar{u}_{ijt}^{g\kappa}, (\delta^g)^2)$, we have

$$\begin{aligned}
&E(u_{ijt}^{g\kappa} | u_{ijt}^{g\kappa} \geq 0) \\
&= \left(\delta^g \cdot \bar{u}_{ijt}^{g\kappa} + \delta^g \cdot \frac{\phi(-\frac{\delta^g \cdot \bar{u}_{ijt}^{g\kappa}}{\delta^g})}{1 - \Phi(-\frac{\delta^g \cdot \bar{u}_{ijt}^{g\kappa}}{\delta^g})} \right) \\
&= \delta^g \left(\bar{u}_{ij}^{g\kappa} + \frac{\phi(\bar{u}_{ij}^{g\kappa})}{\Phi(\bar{u}_{ij}^{g\kappa})} \right)
\end{aligned}$$

Keyword Search The conditional probability of keyword search takes the form

$$\Pr(search_i^g) = \frac{\exp(\lambda_0^g + \lambda_1^g IV_{it}^g)}{1 + \exp(\lambda_0^g + \lambda_1^g IV_{it}^g)} \quad (12)$$

where IV_{it}^g is the inclusive value for searching conditional on the segment membership. IV_{it}^g is defined as

$$IV_{it}^g = \log\left[\sum_{\kappa} \exp(z_{it}^{g\kappa})\right]. \quad (13)$$

This specification can be interpreted as the consumer making a decision to use a keyword search based on the rational behavior of utility maximization (McFadden, 1977; Ben-Akiva and Lerman, 1985). A search term is more likely to be invoked if it yields higher expected utility.

Segment Membership Recognizing that consumers are heterogeneous in behaviors described above, we apply a latent class model in the spirit of Kamakura and Russell (1989) to capture heterogeneity in consumer preferences. Heterogeneity in preference can arise, for example, when some consumers prefer some features more than others. We assume G exogenously determined segments. Note that our specification implies a dependency across decisions that is not captured via the stage-specific decision errors, and therefore captures the effect of unobserved individual specific differences in search behavior.

The prior probability for user i being a member of segment g is defined as

$$pg_{it}^g = \exp(\gamma_0^g + Demo_{it}'\gamma^g) / \sum_{g'=1}^G \exp\left(\gamma_0^{g'} + Demo_{it}'\gamma^{g'}\right) \quad (14)$$

where $Demo_{it}'$ is a vector of attributes of user i such as demographics and past browsing history; vector $\{\gamma_0^g, (\gamma^g)'\}_{\forall g}$ contains parameters to be estimated. For the purpose of identification, one segment's parameters are normalized to zero.

4.1.2 Consumer Downloads

The search, sort/filter, and download models can be integrated over consumer preferences to obtain an expectation of the number of downloads that an advertiser receives for a given position of its

keyword advertisement. Advertisers must form this expectation predicated on observed *aggregate* download totals, d_j^t (in contrast to the search engine who observes y_{ijt}, κ_{it} and $Demo_{it}$).

To develop this aggregate download expectation, we begin by noting that the download utility $u_{ijt}^{g\kappa}$ is a function of consumer specific characteristics and decisions $\zeta_{ijt} = [\varepsilon_{ijt}^{g\kappa}, \xi_{it}^{g\kappa}, search_i^g, \text{segment } g \text{ membership}, Demo_{it}']$ and that an advertiser needs to develop an expectation of downloads over the distribution of these unobserved (to the advertiser) individual characteristics. Define

$$A_{ijt} = \{\zeta_{ijt} : u_{ijt}^{g\kappa} \geq 0\},$$

i.e., A_{ijt} is the set of values of ζ_{ijt} which will lead to the download of product j in period t .

Let $\mathcal{D}(\zeta_{ijt})$ denote the distribution of ζ_{ijt} . The likelihood of downloading product j in period t can be expressed as

$$P_j^t = \int_{\zeta_{ijt} \in A_{ijt}} \mathcal{D}(\zeta_{ijt}) \quad (15)$$

$$= \int_{Demo_{it}} \sum_g \sum_{\kappa} \left[\Phi(\bar{u}_{ijt}^{g\kappa}) \frac{\exp(U_{it}^{g\kappa})}{\sum_{\kappa'=0}^3 \exp(U_{it}^{g\kappa'})} \right] \Pr(search_{it}^g) pg_{it}^g d\mathcal{D}(Demo_{it}) \quad (16)$$

where the first term in the brackets captures the download likelihood, the second term captures the search strategy likelihood, and the first term outside the brackets captures the likelihood of search. pg_{it}^g is the probability of segment g membership and $\mathcal{D}(Demo_{it})$ is the distribution of demographics.

Correspondingly, the advertiser with attributes X_j^t has an expected number of downloads for appearing in slot k , $d_j^t(k, X_j^t; \Omega_c)$, which can be computed as follows

$$d_j^t(k, X_j^t; \Omega_c) = M_t P_j^t \quad (17)$$

where Ω_c is the set of consumer preference parameters; M_t is the market size in period t .

Product attributes are posted on the search engine and are therefore common knowledge to all

advertisers and consumers. We assume these X_j^t (including prices) are exogenous within the scope of our sponsored search analysis for several reasons. First, advertisers distribute and promote their products through multiple channels and they do so over longer periods of time than considered herein. Hence, product attributes and prices are more likely to be determined via broader strategic considerations than the particular auction game and time frame we consider. Second, the attribute and price levels for each product are stable over the duration of our data and analysis. We would expect more variation in attribute and price levels if they were endogenous to the particular advertiser and search engine decisions we consider. Third, keeping product attributes and prices stable may actually be strategic decisions of advertisers. However, because there is little or no variation in the data over time, it is not feasible to estimate endogenous attribute/price decision making with our data.

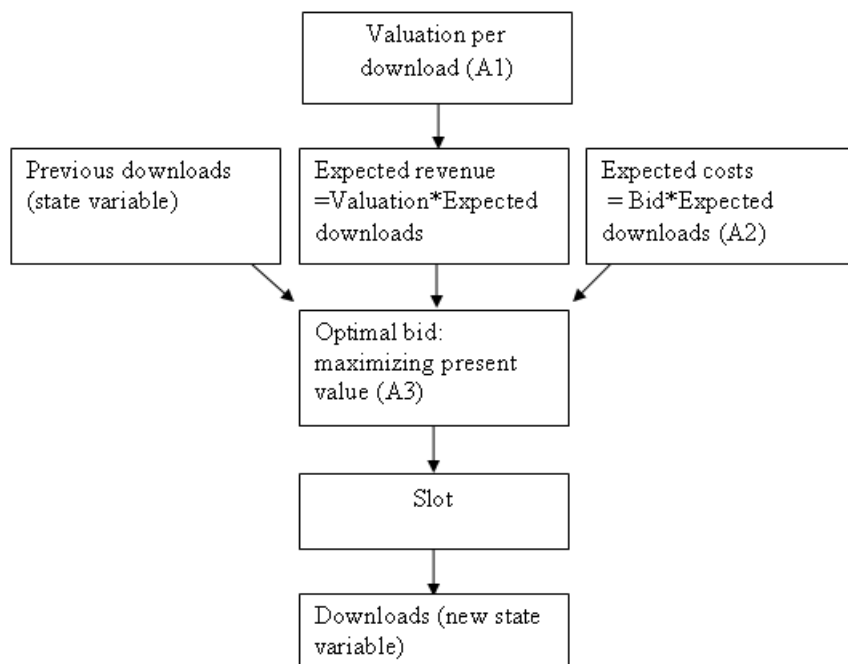
4.2 Advertiser Model

Figure 3 overviews the dynamic game played by the advertiser. Advertiser j 's problem is to decide the optimal bid amount b_j^t with the objective of maximizing discounted present value of payoffs.¹⁸ Higher bids lead to greater revenues because they yield more favorable positions on the search engine, thereby yielding more click-throughs for the advertiser. However, higher bids also increase costs (payments) leading to a trade-off between costs and revenues. The optimal decision of whether and how much to bid is incumbent upon the bidding mechanism, the characteristics of the advertiser, the information available at the time of bidding (including the state variables), and the nature of competitive interactions.

An advertiser's period profit for a download is the value it receives from the download less the costs (payments) of the download. Though we do not observe the value of a download, we infer this value by noting the observed bid can be rationalized only for a particular value accrued by the advertiser. We presume this value is drawn from a distribution known to all firms. The total period revenue for the advertiser is then the value per download times the expected number

¹⁸Because the search engine used in our application has the dominant market share in the considered category, we do not address advertiser bidding on other sites. Also, it would be difficult to obtain download data from these more minor competitors. We note this is an important issue and call for future research.

Figure 3: Advertiser Decisions



of downloads.¹⁹ The total period payment upon winning is the number of downloads times the advertiser’s bid. Hence, the total expected period profit is the number of downloads times the profit per download (i.e., the value per downloads less the payment per download).

Of course, the bid levels and expected download rates are affected by rules of the auction. Though we elaborate in further details on the specific rules of bidding below, at this point we simply note that the rules of the auction favor advertisers whose products were downloaded more frequently in the past since such products are more likely to lead to higher revenues for the platform.²⁰ Current period downloads are, in turn, affected by the position of the advertisement on the search engine. Because past downloads affect current placement, and thus current downloads, the advertiser’s problem is inherently dynamic; and past downloads are treated as a state variable.

Finally, given the rules of the auction, we note that all advertisers move simultaneously. While we presume a firm knows its own value, we assume competing firms know only the distribution of

¹⁹The expected number of downloads is inferred from the consumer model and we have derived this expression in section 4.1.2.

²⁰This is because the payment made to the search engine by an advertiser is the advertiser’s bid times its total downloads.

this value.

The process is depicted in Figure 3. We describe the process with more details as follows: Section 4.2.1 details the rules of the auction that affect the seller costs (A2), section 4.2.2 details the advertisers' value distribution (A1), and section 4.2.3 indicates how period values and costs translate to discounted profits and the resulting optimal bidding strategy (A3).

4.2.1 Seller Costs and the Bidding Mechanism

We begin by discussing how slot positions are allocated with respect to bids and the effect of these slot positions on consumer downloads (and thus advertiser revenue).

Upon a consumer completing a query, the search engine returns $k = 1, 2, \dots, K, \dots, \bar{N}$ slots covering the products of all firms. Only the top $K = 5$ slots are considered as premium slots. Auctions for these K premium slots are held every period ($t = 1, 2, \dots$). An advertiser seeks to appear in a more prominent slot because this may increase demand for the advertiser's product. Slots $K + 1$ to \bar{N} are non-premium slots which compose a section called organic search section.

There are N advertisers who are interested in the premium slots ($N \leq \bar{N}$). In order to procure a more favorable placement, advertiser j submits bid b_j^t in period t . These bids, submitted simultaneously, are summarized by the vector $\mathbf{b}^t = \{b_1^t, b_2^t, \dots, b_N^t\}$.²¹ Should an advertiser win slot k , the realized number of downloads d_j^t is a random draw from the distribution with the expectation $d_j^t(k, X_j^t; \Omega_c)$. The placement of advertisers into the K premium slots is determined by the ranking of their $\{b_j^t d_j^{t-1}\}_{\forall j}$, i.e., the product of current bid and last period realized downloads; the topmost bidder gets the best premium slot; the second bidder gets the second best premium slot; and so on. A winner of one premium slot pays its own bid b_j^t for each download in the current period. Hence, the total payment for winning the auction is $b_j^t d_j^t$.

Given that the winners are determined in part by the previous period's downloads, the auction game is inherently dynamic. Before submitting a bid, the commonly observed endogenous state

²¹For the purpose of a clear exposition, we sometimes use boldface notations or pairs of braces to indicate row vectors whose elements are variables across all bidders. For example, $\mathbf{d}^t = \{d_j^t\}_{\forall j}$ is a vector whose elements are $d_j^t, \forall j$.

variables at time t are the realized past downloads of all bidders from period $t - 1$,

$$\mathbf{s}^t = \mathbf{d}^{t-1} = \{d_1^{t-1}, d_2^{t-1}, \dots, d_N^{t-1}\}. \quad (18)$$

If an advertiser is not placed at one of the K premium slots, it will appear in the organic section; advertisers placed in the organic section *do not* pay for downloads from consumers. The ranking in the organic search section is determined by the product update recency at period t , which is a component of the attribute of products, \mathbf{X}^t . Other attributes include price, consumer ratings, and so on. In contrast to \mathbf{s}^t , \mathbf{X}^t can be considered as *exogenous* state variables, evolving according to some exogenously determined distribution. The endogenous state variables, in contrast, are affected by bidders' actions.²² All state variables \mathbf{s}^t and \mathbf{X}^t are commonly observed by all bidders before bidding.

4.2.2 Seller Value

The advertiser's bid determines the cost of advertising and must be weighed against the potential return when deciding how much to bid. We denote advertiser j 's valuation regarding one download of its product in period t as v_j^t . We assume that this valuation is private information but drawn from a normal distribution that is commonly known to all advertisers. Specifically,

$$\begin{aligned} v_j^t &= v(X_j^t; \theta) + f_j + r_j^t \\ &= X_j^t \theta + f_j + r_j^t \end{aligned} \quad (19)$$

where θ are parameters to be estimated and reflect the effect of product attributes on valuation. The f_j are firm-specific fixed effect terms assumed to be identically and independently distributed *across advertisers*. This fixed effect term captures heterogeneity in valuations that may arise from omitted firm-specific effects such as more efficient operations.²³ The $r_j^t \sim N(0, \psi^2)$ are private

²²Throughout the paper "state variables" is sometimes used implicitly to refer to the endogenous state variable, past downloads.

²³To capture unobserved heterogeneity of advertisers' valuations and the corresponding bidding strategies, we also consider a latent class advertiser model with segment specific θ , f and bidding policies (Arcidiacono and Miller, 2009;

shocks to an advertiser's valuation in period t , assumed to be identically and independently distributed *across advertisers and periods*. The sources of this private shock may include: (1) temporary increases in the advertiser's valuation due to some events such as a promotion campaign; (2) unexpected shocks to the advertiser's budget for financing the payments of the auction; (3) temporary production capacity constraint for delivering the product to users; and so on. The random shock r_j^t is realized at the beginning of period t . Although r_j^t is private knowledge, we assume the distribution of $r_j^t \sim N(0, \psi^2)$ is common knowledge among bidders. We further assume the fixed effect f_j of bidder j is known to all bidders but not to researchers. Given bidders may observe opponents' actions for many periods, the fixed effect can be inferred among bidders (Greene, 2003).

4.2.3 Seller Profits: A Markov Perfect Equilibrium (MPE)

Given v_j^t and state variable s^t , predicted downloads and search engine's auction rules, bidder j decides the optimal bid amount b_j^t with the objective of maximizing discounted present value of payoffs. In light of this, every advertiser has an expected period payoff, which is a function of s^t , \mathbf{X}^t , r_j^t and all advertisers' bids \mathbf{b}^t

$$\begin{aligned}
& \mathbf{E}\pi_j(\mathbf{b}^t, \mathbf{s}^t, \mathbf{X}^t, r_j^t; \theta, f_j) & (20) \\
= & \mathbf{E} \sum_{k=1}^K \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot (v_j^t - b_j^t) \cdot d_j^t(k, X_j^t; \Omega_c) \\
& + \mathbf{E} \sum_{k=K+1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot v_j^t \cdot d_j^t(k, X_j^t; \Omega_c) \\
= & \mathbf{E} \sum_{k=1}^K \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot (X_j^t \theta + f_j + r_j^t - b_j^t) \cdot d_j^t(k, X_j^t; \Omega_c) \\
& + \mathbf{E} \sum_{k=K+1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot (X_j^t \theta + f_j + r_j^t) \cdot d_j^t(k, X_j^t; \Omega_c)
\end{aligned}$$

where the expectation for profits is taken over other advertisers' bids \mathbf{b}_{-j}^t . $\Pr(k|\cdot)$ is the conditional probability of advertiser j getting slot k , $k = 1, 2, \dots, \bar{N}$. $\Pr(k|\cdot)$ depends not only on bids, but also on states \mathbf{s}^t (the previous period's downloads) and product attributes \mathbf{X}^t . This is because: i)

Chung et al., 2009). The first-step model fit for the bidding policies decreases (AIC changes from 2076 to 2110). The insights stay the same pertaining to the valuations of advertisers from the second-step estimation.

the premium slot allocation is determined by the ranking of $\{b_j^t d_j^{t-1}\}_{\forall j}$, where \mathbf{d}^{t-1} are the state variables and ii) the organic slot allocation is determined by product update recency, an element of \mathbf{X}^t .

In addition to the current period profit, an advertiser also takes its expected future payoffs over an infinite horizon into account when making decisions. In period t , given the state variables, advertiser j 's discounted expected future payoffs evaluated prior to the realization of the private shock r_j^t is given by

$$\mathbf{E} \left[\sum_{\tau=t}^{\infty} \rho^{\tau-t} \pi_j (\mathbf{b}^{\tau}, \mathbf{s}^{\tau}, \mathbf{X}^{\tau}, r_j^{\tau}; \Omega_a) \right] \quad (21)$$

where $\Omega_a = \{\theta, \psi, f'_{\forall j}\}$, with a denoting advertiser behavior (in contrast to the parameters Ω_c in the consumer model). The parameter ρ is a common discount factor. The expectation is taken over the random term r_j^t , bids in period t as well as all future realization of \mathbf{s} , \mathbf{X} , shocks, and bids. The endogenous state variables \mathbf{s}^{t+1} in period $t + 1$ is drawn from a probability distribution $P(\mathbf{s}^{t+1} | \mathbf{b}^t, \mathbf{s}^t, \mathbf{X}^t)$.

We use the concept of a pure strategy Markov perfect equilibrium (MPE) to model the bidder's problem of whether and how much to bid in order to maximize the discounted expected future profits (Bajari et al., 2007; Dubé et al., 2008 and others). The MPE implies that each bidder's bidding strategy only depends on the then-current profit-related information, including state, \mathbf{X}^t and its private shock r_j^t . Hence, we can describe the equilibrium bidding strategy of bidder j as a function $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t) = b_j^t$.²⁴ Given a state vector \mathbf{s} , product attributes \mathbf{X} and prior to the realization of current r_j (with the time index t suppressed), bidder j 's expected payoff under the equilibrium strategy profile $\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_N\}$ can be expressed recursively as:

$$V_j(\mathbf{s}, \mathbf{X}; \sigma) = \mathbf{E} \left[\pi_j(\sigma, \mathbf{s}, \mathbf{X}, r_j; \Omega_a) + \rho \int_{\mathbf{s}'} V_j(\mathbf{s}', \mathbf{X}'; \sigma) dP(\mathbf{s}' | \mathbf{b}, \mathbf{s}, \mathbf{X}) | \mathbf{s} \right] \quad (22)$$

where the expectation is taken over current and future realizations of random terms \mathbf{r} and \mathbf{X} . To

²⁴The bidding strategies are individual specific due to the fixed effect f_j (hence the subscript j). For the purpose of clear exposition, we use $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$ instead of $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t; f_j)$ throughout the paper. Multiple observations for each advertiser allows the identification of $\sigma_j, j = 1, 2, \dots, N$.

test the alternative theory that advertiser's may be myopic in their bidding, we will also solve the advertiser problem under the assumption that period profits are maximized independently over time.

The advertiser model can then be used in conjunction with the consumer model to forecast advertiser behavior as we shall discuss in the policy simulation section. In a nutshell, we presume advertisers will choose bids to maximize their expected profits. A change in information states, bidding mechanisms, or webpage design will lead to an attendant change in bids conditioned on the advertisers value function, which we estimate as described next.

5 Estimation

5.1 An Overview

Though it is standard to estimate dynamic MPE models via a dynamic programming approach such as a nested fixed point estimator (Rust, 1994), this requires one to repetitively evaluate the value function (Equation 22) through dynamic programming for each instance in which the parameters of the value function are updated. Even when feasible, it is computationally demanding to implement this approach. Instead, we consider the class of two-step estimators. Specifically, in this application we implement the two-step estimator proposed by Bajari et al. (2007) (BBL henceforth). In a technical appendix available from the authors, we also derive a Bayesian likelihood based estimator for the two-step model. This approach has the advantage that it does not rely on asymptotics for inference. The estimates are essentially identical though the posterior predictive 95% intervals for the Bayesian model parameters are slightly more narrow than the BBL confidence intervals, and their distribution is slightly skewed.

As can be seen in equation 22, the value function is parametrized by the primitives of the value distribution Ω_a . Under the assumption that advertisers are behaving rationally, these advertiser private values for clicks should be consistent with observed bidding strategies. Therefore, in the second step estimation, values of Ω_a are chosen so as to make the observed bidding strategies congruent with rational behavior. We detail this step in Section 5.3 below.

However, as can be observed in equations 22 and 20, computation of the value function is also incumbent upon i) the bidding policy function that maps bids to downloads, product attributes, and private shocks $\sigma_j(s^t, \mathbf{X}^t, r_j^t) = b_j^t$; ii) the expected downloads $d_j^t(k, X_j^t; \Omega_c)$; and iii) a function that maps the likelihood of future states as a function of current states and actions $P(s^{t+1} | \mathbf{b}^t, s^t, \mathbf{X}^t)$. These are estimated in the first step as detailed in Section 5.2 below and then substituted into the value function used in the second step estimation.

The identification of the consumer model follows the identification strategies of classical discrete choice models. The advertiser model’s identification follows BBL. We provide a more detailed discussion of its identification in Appendix A.3.

5.2 First Step Estimation

In the first step of the estimation we seek to obtain:

1. A “partial” policy function $\tilde{\sigma}_j(s, \mathbf{X})$ describing the equilibrium bidding strategies as a function of the observed state variables. We estimate the policy function by noting that players adopt equilibrium strategies (or decision rules) and that behaviors generated from these decision rules lead to correlations between i) the observed states and ii) advertiser decisions (i.e., bids). The partial policy function captures this correlation. In our case, we use a Tobit model with a flexible polynomial specification in state variables to link bids to downloads and product characteristics. Details are described in Section A.1.1 of the Appendix.²⁵ Subsequently, the full policy function $\sigma_j(s, \mathbf{X}, r_j^t)$ can be inferred based on $\tilde{\sigma}_j(s, \mathbf{X})$ by integrating out the private random shocks r_j^t . Hence the partial policy function can be thought of as the marginal distribution of the full policy function.
2. The expected downloads for a given firm at a given slot, $d_j^t(k, X_j; \Omega_c)$. The $d_j^t(k, X_j; \Omega_c)$ follows directly from the consumer model. Hence, the first step estimation involves i) estimating the parameters of the consumer model and then ii) using these estimates to compute

²⁵As a robustness check, we also consider a thin-plate spline function for the policy function. We obtain essentially the same the second step estimates under both specifications. We report the results of the polynomial specification since the identification of BBL with continuous control under nonparametric policy function is still not established (BBL, p.1346). We discuss the robustness check in the Appendix.

the expected number of downloads. The expected total number of downloads as a function of slot position and product attributes is obtained by using the results of the consumer model to forecast the likelihood of each person downloading the software and then integrating these probabilities across persons.²⁶ We discuss our approach for determining the expected downloads in Section A.1.2 of the Appendix.

3. The state transition probability $P(s'|\mathbf{b}, \mathbf{s}, \mathbf{X})$ which describes the distribution of future states (current period downloads) given observations of past downloads, product attributes and actions (current period bids). These state transitions can be derived by i) using the policy function to predict bids as a function of past downloads and product attributes, ii) determining the slot ranking as a function of these bids, past downloads and product attributes, and then iii) using the consumer model to predict the number of current downloads. Details regarding our approach to determining the state transition probabilities is outlined in Section A.1.3 of the Appendix.

With the first step estimates of $\sigma_j(\mathbf{s}, \mathbf{X}, r_j^t)$, $d_j^t(k, X_j; \Omega_c)$, and $P(s'|\mathbf{b}, \mathbf{s}, \mathbf{X})$, we can compute the value function in Equation 22 as a function with only Ω_a unknown. In the second step, we estimate these parameters.

5.3 Second Step Estimation

The goal of the second step estimation is to recover the primitives of the bidder value function, Ω_a . The intuition behind how the second-stage estimation works is that true parameters should rationalize the observed data. For bidders' data to be generated by rational plays, we need

$$V_j(\mathbf{s}, \mathbf{X}; \sigma_j, \sigma_{-j}; \Omega_a) \geq V_j(\mathbf{s}, \mathbf{X}; \sigma'_j, \sigma_{-j}; \Omega_a), \forall \sigma'_j \neq \sigma_j \quad (23)$$

²⁶As an aside, we note that advertisers have limited information from which to form expectations about total downloads because they observe the aggregate information of downloads but not the individual specific download decisions. Hence, advertisers must infer the distribution of consumer preferences from these aggregate statistics. In a subsequent policy simulation we allow the search engine to provide individual level information to advertisers in order to assess how it affects advertiser behavior and, therefore, search engine revenues.

where σ_j is the observed equilibrium policy function and σ'_j is some deviation from σ_j . This equation means that any deviations from the observed equilibrium bidding strategy will not result in more profits. Hence, we first simulate the value functions under the equilibrium policy σ_j and the deviated policy σ'_j (i.e., the left hand side and the right hand side of equation 23). Then we obtain Ω_a using a minimum distance GMM estimator as described in BBL. We describe the details of this second step estimation in Appendix A.2.

6 Results

6.1 First Step Estimation Results

Recall, the goal of the first step estimation is to determine the policy function, $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$, the expected downloads $d_j^t(k, X_j^t; \Omega_c)$, and the state transition probabilities $P(\mathbf{s}^{t+1} | \mathbf{b}^t, \mathbf{s}^t, \mathbf{X}^t)$. To determine $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$, we first estimate the partial policy function $\tilde{\sigma}_j(\mathbf{s}^t, \mathbf{X}^t)$ and then compute the full policy function. To determine $d_j^t(k, X_j^t; \Omega_c)$, we first estimate the consumer model and then compute the expected downloads. Last $P(\mathbf{s}^{t+1} | \mathbf{b}^t, \mathbf{s}^t, \mathbf{X}^t)$ is derived from the consumer model and the partial policy function. Thus, in the first stage we need only to estimate the partial policy function and the consumer model. With these estimates in hand, we compute $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$, $d_j^t(k, X_j^t; \Omega_c)$, and $P(\mathbf{s}^{t+1} | \mathbf{b}^t, \mathbf{s}^t, \mathbf{X}^t)$ for use in the second step. Thus, below, we report the estimates for the partial policy function and the consumer model on which these functions are all based.

6.1.1 Partial Policy Function $\tilde{\sigma}_j(\mathbf{s}, \mathbf{X})$

The vector of independent variables (\mathbf{s}, \mathbf{X}) for the partial policy function (i.e. the flexible polynomial function and the alternative thin plate spline function as outlined in Appendix A.1.1) contains the following variables:

- Product j 's state variable, last period download d_j^{t-1} and the square of this term. We reason that high past downloads increase the likelihood of a favorable placement and, therefore, affect bids. We introduce $(d_j^{t-1})^2$ to accommodate potential nonlinearity in the effect of past downloads on bids.

- Two market level variables (and their respective squares): the sum of last period downloads from all bidders and the number of bidders in last period. Since we only have 322 observations of bids, it is infeasible to estimate a parameter to reflect the effect of each opponent’s state (i.e., competition) on the optimal bid. Moreover, it is unlikely a bidder can monitor every opponent’s state in each period before bidding because such a strategy carries high cognitive and time costs. Hence, summary measures provide a reasonable approximation of competing states in a limited information context. Others in the literature who have invoked a similar approach include Jofre-Bonet and Pesendorfer (2003) and Ryan (2009). Like them, we find this provides a fair model fit. Another measure of competitive intensity is the number of opponents. Given that bidders cannot directly observe the number of competitors in the current period, we used a lagged measure of the number of bidders.
- The interaction term between past download d_j^{t-1} and update recency. This term is introduced to capture the interaction between the two variables observed in Section 3.2.
- Product j ’s attributes in period t (X_j^t), including its non-trial version price, expert rating, consumer rating, update recency, and compatibility with an older operating system. We expect that a higher quality product will yield greater downloads thereby affecting the bidding strategy.
- An advertiser specific constant term to capture the impact of the fixed effect f_j on bidding strategy.²⁷
- To control the possible effect of the growth of ownership of MP3 players, we also collect the average lagged price of all new MP3 players in the market from a major online retailing

²⁷An alternative, and more flexible approach to capture heterogeneity used by Misra and Nair (2009) estimates the two-step model agent by agent; this approach is feasible in contexts with large amounts of data for each agent, a moderate state and actions space, and a modicum of agent interactions. Given this is not the case in our context, we instead employ a fixed effect specification in both the valuation function and the bidding policy and assume that the fixed effects in the valuation function do not moderate the bidding policy function. Recently, Arcidiacono and Miller (2009) and Chung et al. (2009) have proposed a latent class approach to accommodate heterogeneity that is feasible to estimate in our context. As noted in section 4.2.2, our findings are robust to this approach. Accordingly, we believe the fixed effect assumption is of limited consequence in our context.

platform (www.pricegrabber.com).

Table 3 reports the estimation results.²⁸ As a measure of fit of the model, we simulated 10,000 bids from the estimated distribution. The probability of observing a positive simulated bid is 41.0%; the probability of observing a positive bid in the real data is 41.6%. Conditional on observing a positive simulated bid, these bids have a mean of \$0.19 with a standard deviation of \$0.09. In the data, the mean of observed positive bids is \$0.20 and the standard deviation is \$0.08. At the individual bids level, the within-sample bidding choice hit rate is 0.98. Conditional on observing a positive bid, the mean absolute percentage error (MAPE) is 0.05. To assess the out-of-sample fit, we also estimate the same model only using 70% (227/322) of the observations and use the remaining 30% as a holdout sample. The change in estimates is negligible. We then use the holdout to simulate 10,000 bids. The probability of observing a positive bid is 41.1%, while there are 42.4% positive bids in the holdout sample. Among the positive simulated bids, the mean is \$0.23 and the standard deviation is \$0.08. The corresponding statistics in the holdout are \$0.21 and \$0.07. The hit rate and MAPE for the holdout are 0.94 and 0.08, respectively. Overall, the fit is good.

Table 3: Bidding Function Estimates

	Parameters	Std. Err.
φ		
Lagged Downloads $_{jt}/10^3$	-0.32**	0.06
(Lagged Downloads $_{jt}/10^3$) ²	-0.09	0.07
Total Lagged Downloads $_t/10^3$	0.08**	0.04
(Total Lagged Downloads $_t/10^3$) ²	0.02**	0.01
Lagged Downloads $_{jt}/10^3 \times$ Lapse Since Last Update $_{jt}$	0.06**	0.03
Lagged Number of Bidders $_t$	0.02**	0.01
Lapse Since Last Update $_{jt}$	-0.55*	0.30
Non-trial Version Price $_{jt}$	0.40**	0.21
Expert Ratings $_{jt}$	0.46	0.56
Consumer Ratings $_{jt}$	0.82**	0.38
Compatibility Index $_{jt}$	-0.19**	0.03
Lagged MP3 Player Price $_t$	0.09**	0.03
τ	7.17**	1.06
Log Likelihood	-1002.9	
Note: ** p<0.05; * p<0.10		

²⁸To conserve space, we do not report the estimates of fixed effects.

The estimates yield several insights into the observed bidding strategy. First, the bidder’s state variable (d_j^{t-1}) is negatively correlated with its bid amount b_j^t because the ranking of the auction is determined by the product of b_j^t and d_j^{t-1} . All else being equal, a higher number of lagged downloads means a bidder can bid less to obtain the same slot. Second, the total number of lagged downloads in the previous period ($\sum_{j'} d_{j'}^{t-1}$) and the lagged number of bidders both have a positive impact on a bidder’s bid. We take this to mean increased competition leads to higher bids. Third, bids are increasing in the product price. One possible explanation is that a high priced product yields more value to the firm for each download, and hence the firm competes more aggressively for a top slot. Similarly and fourth, a high price for MP3 players reflects greater value for the downloads also leading to a positive effect on bids. Fifth, “Lapse Since Last Update” has a negative effect on bids. Older products are more likely obsolete, thereby generating lower value for consumers. If this is the case, firms can reasonably expect fewer final purchases after downloads and, therefore, bid less for these products. Likewise and sixth, higher compatibility with prior software versions reflects product age leading to a negative estimate for this variable. Seventh, though the effect is quite small, the interaction between update recency and lagged download is significant. This result may stem from older products appearing lower in the organic search results, thereby enhancing the incremental effect of securing a sponsored slot near the top, thus increasing the advertiser incentive to bid. Finally, ratings from consumers and experts (albeit not significant for experts) have a positive correlation with bid amounts – these again imply greater consumer value for the goods, making it more profitable to advertise them.

6.1.2 Consumer Model

The consumer model is estimated using MLE approach based on the likelihood function described in Appendix A.1.2. We consider the download decisions for each of the 21 products who entered auctions, plus the top 3 products who did not. Together these firms constitute over 80% of all downloads. The remaining number of downloads are scattered across 370 other firms, each of

whom has a negligible share. Hence, we exclude them from our analysis.²⁹

Table 4: Alternative Numbers of Latent Segments

	AIC
1 Segment	-12159.2
2 Segments*	-12491.2
3 Segments	-12571.1
4 Segments	-12551.4

Note: * indicates the model with the best fit.

We estimate an increasing number of latent segments until there is no improvement in model fit as defined by the AIC. Table 4 reports the AIC values for up to four segments. The two segment model with yields the best result, with an in-sample MAPE of 0.07 and a 10% of the sample holdout MAPE of 0.11. The overall fit is good.

Table 5 presents the estimates of the model with two segments. Conditional on the estimated segment parameters and demographic distribution, we calculate the segment sizes as 88% and 12%, respectively. Based on the parameter estimates in Table 5, Segment 1 is less likely to initiate a search (low λ_0^g and low download utility function intercept). The primary basis of segmentation is whether a customer has visited a music website at other properties owned by the download website; these customers are far more likely to be in the frequent download segment. Moreover, upon engaging a search, segment 1 appears to be less sensitive to slot ranking but more sensitive to consumer and expert ratings than segment 2. Segment 2, composed of those who search more frequently, relies more heavily on the slot order when downloading. Overall, we speculate that segment 1 are the occasional downloaders who base their download decisions on others' ratings and tend not to exclude goods of high price. In contrast, segment 2 contains the "experts" or frequent downloaders who tend to rely on their own assessments when downloading. Of interest is the finding that those in segment 2 rely more on advertising slot rank. This is consistent with a perspective that frequent downloaders might be more strategic; knowing that higher quality firms

²⁹As noted by Zanutto and Bradlow (2006), excluding products from the analysis might induce sample bias. As a robustness check, we re-estimate the model with a random sample of 5 additional products that were originally omitted. There is little change to the estimates but the model fit deteriorates (AIC: -12491 vs. -12517). Hence we retain the current specification.

Table 5: Consumer Model Estimates

	Segment 1 (88%) (Infrequent searcher)	Segment 2 (12%) (Frequent searcher and slot sensitive)
	Estimate (S.E.)	Estimate (S.E.)
β^g (utility parameters)		
Constant	-0.05 (0.03)	0.25 (0.10)
Slot Rank	-0.13 (0.05)	-0.65 (0.05)
Non-trial Version Price	0.01 (0.01)	-0.08 (0.03)
Expert Ratings	0.10 (0.05)	0.08 (0.03)
Consumer Ratings	0.11 (0.03)	0.05 (0.02)
Compatibility Index	-0.06 (0.03)	-0.15 (0.20)
Total Download Percentage	0.03 (0.02)	0.12 (0.04)
δ^g (sorting/filtering scaling)	1.44 (0.40)	1.50 (0.53)
$\eta^{g\kappa}$ (fixed effect, sorting/filtering)		
Sorting only	0.13 (0.06)	0.04 (0.15)
Filtering only	0.06 (0.11)	0.27 (1.21)
Sorting and filtering	-0.03 (0.18)	-0.21 (1.11)
λ^g (search probability)		
λ_0^g (base)	-8.13 (2.22)	-0.66 (0.30)
λ (1-correlation)	0.79 (0.31)	0.83 (0.12)
γ^g (segment parameters)		
Constant	—	-3.25 (1.20)
Music Site Visited	—	6.50 (2.11)
Registration Status	—	-0.15 (0.25)
Product Downloaded in Last Month	—	-0.35 (0.15)

tend to bid more and obtain higher ranks, those who download often place greater emphasis on this characteristic (Chen and He, 2006; Athey and Ellison, 2008). It could also reflect the greater opportunity cost of time for frequent searchers. Because these consumers conduct more searches, they search less “deeply” conditioned on a search. Otherwise, the total number of searches (i.e., the number of searches times the number of alternative considered per search), and hence the total cost of search, would be extremely large.

More insights on this difference in download behavior across segments can be gleaned by determining the predicted probabilities of searching and sorting/filtering by computing $\Pr(search_i^g) = \frac{\exp(\lambda_0^g + \lambda_1^g IV_{it}^g)}{1 + \exp(\lambda_0^g + \lambda_1^g IV_{it}^g)}$ and $\Pr(\kappa)_{it}^g = \frac{\exp(\eta^{g\kappa} + U_{it}^{g\kappa})}{\sum_{\kappa'=0}^3 \exp(\eta^{g\kappa'} + U_{it}^{g\kappa'})}$ in Equations 12 and 10, respectively. Table 6 reports these probabilities for both segments.

Table 6 confirms the tendency of those in segment 2 to be more likely to initiate a search in the

Table 6: Searching Behavior of Consumers

	Segment 1	Segment 2
Searching	0.09%	62.5%
No sorting or filtering	74.8%	85.5%
Sorting but no filtering	25.1%	7.9%
No sorting but filtering	→ 0	6.1%
Sorting and filtering	→ 0	0.5%

focal category. Though comprising only 12% of all consumers, they represent 90% of all searches. The increased searching frequency suggests that members of segment 2 are ideal customers to target because more searches lead to more downloads.

Moreover, segment 2 (heavy downloaders) is more likely to be influenced by sponsored advertising. To see this, note that segment 1 consumers put more weight on the ratings of products (e.g., expert and consumer ratings) than do segment 2 consumers. As a consequence segment 1 consumers engage in far more sorting. Sorting eliminates the advantage conferred by sponsored advertising because winners of the sponsored search auction may be sorted out of desirable slots on the page.

Table 6 also indicates consumers in segment 1 (occasional downloaders) seldom filter. Filtering occurs when consumers seek to exclude negative utility options from the choice set (e.g., omitting a product not compatible with a certain operating system). Given the high sensitivity to rank order, segment 2 consumers are more prone to eliminate advertised options by filtering. We suspect this segment, by virtue of being a frequent visitor, searches for very specific products that conform to a particular need. Overall, however, segment 1 is more likely to sort and/or filter than segment 2 (25.1% vs. 14.5%) suggesting that segment 2 is more valuable to advertisers. We will explore this conjecture in more detail in our policy analysis.

6.2 Second Step Estimation Results

6.2.1 Alternative Models

In addition to our proposed dynamic bidding model, we consider two alternative models of advertiser behavior: i) myopic bidding and ii) heterogeneous advertiser valuations across consumer

segments.³⁰ Table 7 reports the fit of each model. In the first alternative model, advertisers maximize period profits independently as opposed to solving the dynamic bidding problem given in Equation 22. This model yields a considerably poorer fit, with the average objective under the dynamic model of 1.1, compared to 3.2 under the myopic setting.³¹ Hence, we conclude that the data are consistent with a specification where advertisers are bidding strategically.³² This strategic behavior might result from dynamics in the bidding process coupled with nonlinearity in advertising response. Similar dynamic behavior has been evidenced in the face of non-linear advertising demand systems with dynamics in advertising carryover (Bronnenberg, 1998).

The second alternative model considers the case wherein advertiser valuations for clicks differ across segments. In this model, we augment Equation 19 by allowing these valuations to vary by segment and then integrate this heterogeneity into the seller profit function given by Equation 20. This model leads to only a negligible increase in fit. Closer inspection of the results indicates little difference in valuations across segments, implying advertisers perceive that the conversion rates of each segment are essentially the same. Hence, we adopt the more parsimonious single valuation model. It is further worth noting that all of our subsequent results and policy simulations evidence essentially no change across these two models.

Table 7: Alternative Models

Model	Average GMM Objective Functions
Base Model	1.11
Base Model Without Advertiser Dynamics	3.15
Base Model With Heterogeneous Customer Valuations	1.09

³⁰We do not estimate the discount factor ρ . As shown in Rust (1994), the discount factor is usually unidentified. We fix $\rho = 0.99$ for our estimation. We also consider $\rho = 0.90$ and $\rho = 0.95$ and observe minimal differences in the results.

³¹Specifically, we re-estimate the second step of BBL approach by bootstrapping across the empirical distribution of first stage estimates and computing the average of the GMM objective functions under the assumption of forward-looking. We then set the discount factor to zero and re-estimate the second step using bootstrapping and take the average of the GMM objective functions under the assumption of myopic bidding.

³²Owing to the inability of model fit alone to substantiate forward-looking behavior, techniques to disentangle myopic from dynamic behavior using field data have become an ongoing research problem of interest in marketing (Misra and Nair, 2009; Dubé et al., 2010).


6.2.2 Valuation Model Results

Table 8 shows the results of second step estimation for the favored model.³³ With respect to the advertiser value function, we find that newer, more expensive and better rated products yield greater values to the advertiser. This is consistent with our conjecture in Section 6.1.1 that firms bid more aggressively when having higher values for downloads. We find that, after controlling for observed product characteristics, 95% of the variation in valuations across firms is on the order of \$0.02. We attribute this variation in part due to differences in the operating efficiency of the firms.

Table 8: Value per Click Parameter Estimates

	Estimate	Std. Err.
θ		
Lapse Since Last Update $_{jt}$	-0.96*	0.27
Non-trial Version Price $_{jt}$	0.21*	0.10
Expert Ratings $_{jt}$	0.55*	0.08
Consumer Ratings $_{jt}$	0.88*	0.11
Compatibility Index $_{jt}$	-0.31*	0.03
Lagged MP3 Player Price $_t$	0.02*	0.01
ψ , Random shock std. dev.	1.44*	0.40

Figure 4: Distribution of Values per Download



Given the second step results, we can further estimate the value of a download to a firm in each period. In Figure 4, the kernel density estimator distribution of these estimates across time and

³³Advertiser specific constant terms f_j are not reported to conserve space.

advertisers is depicted. As indicated in the Figure there is substantial variation in the valuation of downloads. Table 8 explains some of this variation as a function of the characteristics of the software and firm specific effects. Results indicate that higher prices and quality correlate with higher valuations, presumably because these factors are associated with increased advertiser revenue and sales conversion rates. Overall, the mean value of a download to these advertisers is \$0.26. This compares to an average bid of \$0.20 as indicated in Table 1, suggesting that advertisers obtain a small surplus of about \$0.05. This surplus could arise from either i) the advertisers bidding less than their valuation due to the use of a generalized first price auction or ii) the benefit accruing from a high level of preceding downloads which would enable an advertiser to shade their bids further below their respective valuations. In our policy simulations, we will further explore the role of the auction mechanism on bids and whether it is possible to induce truth telling.

To our knowledge, this is the first paper to impute the advertiser’s return from a click in a keyword search context. One way to interpret these results is to consider the firm’s expected sales per download to rationalize the bid. The firm’s profit per click is roughly $CR_j^t \cdot P_j^t - b_j^t$, where CR_j^t indicates the download-sale conversion rate (or sales per download) and P_j^t is the non-trial version price. Ignoring dynamic effects and setting this profit per click equal to $v_j^t - b_j^t$ yields a rough approximation of the conversion rate as $CR_j^t = v_j^t / P_j^t$. Viewed in this light, the effect of higher quality software, which raises v_j^t , leads to a higher implied conversion rate. Noting that the average price of the software is \$22, this average per-click valuation implies that 1.2% of all clicks lead to a purchase (that is, the conversion rate is $0.26/22 = 1.2\%$). This estimate lies within the industry average conversion rate of 1 – 2% reported by Gamedaily.com, suggesting our findings have high face validity.³⁴

7 Policy Simulations

Given the behavior of consumers and advertisers, we can predict how changes in search engine policy affect overall bidding, downloads, consumer welfare, and revenues. The advertiser-consumer

³⁴“Casual Free to Pay Conversion Rate Too Low.” Gamedaily.com (<http://www.gamedaily.com/articles/features/magid-casual-free-to-pay-conversion-rate-too-low/70943/?biz=1>).

behaviors are analogous to a subgame conditioned on search engine policy. To assess the effect of changes in policy, we recompute the equilibrium behavior of consumers and advertisers conditioned on the new policy.³⁵ One might ask whether these deviations in policy are valid as the initial strategies might reflect optimal behavior on the part of the search engine. However, extensive interactions with the search site makes it clear that they have neither considered using these alternative policies nor have they tried them in the past in order to obtain a sense of the strategies' impacts. Hence, we do not conjecture that they are behaving strategically and thus we think these are reasonable policy simulations to consider. Alternatively, estimating a model incorporating the engine's behavior invokes rather strong assumptions of rationality due to the complexity and novelty of the problem. Also, we observe no variation in the considered behaviors of the search engine, meaning there is no means to identify the primitives driving such behaviors.

We describe three policy simulations: i) the effect of alternative webpage designs on search engine revenues, ii) the value of targeting (i.e., allowing advertisers to bid on keywords by segment), and iii) the effect of alternative pricing mechanisms on search engine revenue. As we can no longer assume the optimal advertiser policy function estimated in stage one of our two-step estimator remains invariant in the face of a change in search engine policy, the following policy simulations involve explicitly solving the infinite-horizon dynamic programming problem to recompute an updated (1) advertiser bidding function, (2) consumer download probability, and (3) set of state transitions. Owing to the complexities of solving this game, we develop an approximate dynamic programming approach to solve it.³⁶ More details regarding the implementation of the policy simulations are presented in Technical Appendix B.³⁷ Hence, we limit our discussion to the

³⁵The policy simulations assume that the parameters from the consumer utility function and the advertiser valuation for consumers' clicks are invariant to a change in website design or search engine's auction mechanisms.

³⁶Parallel to our research, a recent study by Farias et al. (2010) demonstrates the validity of the approximate DP algorithm predicated upon a non-parametric policy function. However, our application uses a parametric policy function; to the extent the parametric function is not sufficiently flexible to capture agent behavior, our results will be biased. It is worth noting we considered an array of different polynomial parametric models and our results were invariant to these alternative specifications. It is further worth noting that large action and state space coupled with complex interactions among bidders can complicate the implementation of a non-parametric approach.

³⁷The search for a revised parametric policy function (Appendix B.1) in the neighborhood of the original policy observed in the data, coupled with the assumption of the advertiser symmetry, mitigates the potential for a multiplicity of equilibrium (Bresnahan and Reiss, 1991; Dubé et al., 2005). Moreover, Jofre-Bonet and Pendorfer (2003) show the existence of pure strategy equilibrium in a dynamic procurement auction. If there exist multiple equilibria, the

objectives and insights from these simulations.

7.1 Policy Simulation I: Alternative Webpage Design

The goal of the search engine's sorting/filtering options is to provide consumers with easier access to price and rating information across different products. As shown in section 4.1 and evidenced by our results, sorting and filtering play a crucial role in consumer decision process. In light of this outcome, it is possible to consider an alternative webpage design of the search engine – eliminating the option of sorting and filtering for consumers – and assessing the resulting impact on consumer search, advertiser bidding, and the search engine's revenues. Because this change can have contrasting effects on consumer behavior (consumer should be less likely to search on the site because of the decrease in utility arising from fewer search options) and advertiser behavior (advertisers should bid more because of the decreased likelihood that their advertisements will be sorted or filtered out of the search results), the overall effect is unclear. Using our model, it can be tested which effect dominates. We do this by setting the probability of consumer choosing no sorting/filtering option in equation 10 to one. This manipulation mimics the scenario in which the sorting/filtering option is disabled. Under this new policy, we find that the search engine's revenue decreases by 2.9%, suggesting the consumer effect is larger.³⁸

Next, to more precisely measure these contrasting effects, we apportion the revenue change across consumers and advertisers. Let D_{j0}^t (D_{j1}^t) denote the number of downloads for product j in period t before (after) the change of the webpage. Let B_{j0}^t (B_{j1}^t) denote the bid from advertiser j in period t before (after) the new policy. Accordingly we can calculate (i) the revenue effect arising solely from changes in consumer behavior by holding advertiser behavior fixed, $(\sum_{j,t} B_{j1}^t D_{j1}^t - \sum_{j,t} B_{j1}^t D_{j0}^t)$ and (ii) the effect arising from changing advertiser behavior by holding consumer behavior fixed, $(\sum_{j,t} B_{j1}^t D_{j0}^t - \sum_{j,t} B_{j0}^t D_{j0}^t)$. Using this decomposition, we find the effect arising from consumers $(\sum_{j,t} B_{j1}^t D_{j1}^t - \sum_{j,t} B_{j1}^t D_{j0}^t) / \sum_{j,t} B_{j0}^t D_{j0}^t$ is -5.1% while the effect from advertisers $(\sum_{j,t} B_{j1}^t D_{j0}^t - \sum_{j,t} B_{j0}^t D_{j0}^t) / \sum_{j,t} B_{j0}^t D_{j0}^t$ is 2.2% . Consistent with

new functions can be interpreted as the policies that are the closest to the observed policy.

³⁸The bootstrapped 95% confidence interval for the revenue change of the search engine is $(-3.9\%, -1.0\%)$.

this result, consumer welfare as measured by their overall utility, *declines* 3.8% when the search tools are removed while advertiser profits *increase* 2.1%.³⁹ Thus, for the search engine, the disadvantage of this new policy to consumers outweighs the advantages resulting from more aggressive advertiser bidding.

7.2 Policy Simulation II: Segmentation and Targeting

Advertisers might realize notable dividends if they can capitalize upon the search engine's market intelligence about consumer preferences (Pancras and Sudhir (2007)). By sharing information on its consumers, the search engine can allow an advertiser to vary its bids across market segments. For example, consider two segments, *A* and *B*, wherein segment *B* is more sensitive to product price and segment *A* is more sensitive to product quality. Consider further, two firms, *X* and *Y*, where firm *X* purveys a lower price, but lower quality, product. Intuitively, firm *X* should bid more aggressively for segment *B* because quality sensitive segment *A* will not likely buy the low quality good *X*. This should lead to higher revenues for the search engine. On the other hand, there is less bidding competition for firm *X* within segment *B* because *Y* finds this segment unattractive – this dearth of competition can drive the bid of *X* down for segment *B*. This would place a downward pressure on search engine profits. Hence, the optimal revenue outcome for the search engine is likely to be incumbent upon the distribution of consumer preferences and the characteristics of the goods being advertised. Our approach can assess these effects of segmentation and targeting strategy on the search engine's revenue.

To implement this policy simulation, we enable the search engine to serve a different advertisement to each market segment and allow advertisers to bid differentially each period for these keyword slots across the two consumer segments (see Appendix B.2 for details). We find the search engine's resulting revenue increases by 1%. Using a similar decomposition mentioned in section 7.1, we find the revenue effect arising from the consumer side of the market is 1.4%. We attribute this effect mainly to the enhanced efficiency of advertisements under targeting. In other words,

³⁹The 95% confidence intervals for the welfare changes of consumers and advertisers are (-5.1%, -1.9%) and (0.6%, 4.0%), respectively.

targeting leads to more desirable advertisements for consumers thereby yielding increased downloads. In contrast, the effect arising from advertisers is -0.4% as a result of diminished competitive intensity. Overall, the consumer effect of targeting is dominant, and a net gain in profitability is indicated.⁴⁰

This policy also benefits advertisers in two ways: by increasing the efficiency of their advertising and reducing the competitive intensity of bidding within their respective segments. Overall, we project an 5.8% increase in advertiser revenue under the targeting policy. Consistent with this view of consumer gains, consumer welfare increases by 1.6% . In sum, every agent finds this new policy to be an improvement.

7.3 Policy Simulation III: Alternative Auction Mechanisms

Auction mechanism design has been an active domain of research since the seminal work of Vickrey (1961). Optimal mechanism design involves several aspects including the rules of the auction, the efficiency of the auction in terms of allocation surplus across players, new design to eliminate the dynamic bidding behavior, and so forth. We focus on the payment rules in this investigation. In particular, while the focal search engine currently charges winning advertisers their own bids, many major search engines such as `Google.com` and `Yahoo.com` are applying a “generalized second-price auction” (Edelman et al., 2007). Under the generalized second-price auction rules, winners are still determined by the ranking of $\{b_{j'}^t, d_{j'}^{t-1}\}_{\forall j'}$. However, instead of paying its own bid amount, the winner of a slot pays the highest losing bidder’s bid adjusted by their last period downloads. For example, suppose bidder j wins a slot with the bid of b_j^t and last period download d_j^{t-1} , its payment for each download will be $b_{j'}^t d_{j'}^{t-1} / d_j^{t-1}$, where j' is the highest losing bidders for the slot bidder j wins.

Though “generalized second-price auction” is widely adopted by major search engines, the optimality of such a mechanism has not been substantiated (Iyengar and Kumar, 2006; Katona and Sarvary, 2008). Further, whether the truth-telling equilibrium strategy still holds under a dynamic

⁴⁰The 95% confidence intervals for the revenue/welfare changes of the search engine, advertisers, and consumers are (0.2%, 1.5%), (4.8%, 6.4%), and (0.8%, 2.6%), respectively.

setting is unknown. By implementing a policy simulation that contrasts the search engine and advertiser revenues under the two different mechanisms, we find little difference in revenues for the advertiser or search engine (for example, search engine revenues increase 0.02%). However, the advertisers' bids for clicks approach their values for clicks. Under second price auction, the median ratio of bid/value is 0.98 compared to 0.77 under first price auction. This is consistent with the theory that in equilibrium bidders bid their true values under "generalized second-price auction" (Edelman et al., 2007). This offers empirical support for the contention that generalized second price auctions yield truth telling – though we find little practical consequence in terms of auction house revenue.

8 Conclusion

Given the \$9B firms annually spend on keyword advertising and its rapid growth, we contend that the topic is of central concern to advertisers and platforms that host advertising alike. In light of this growth, it is surprising that there is little extant empirical research pertaining to modeling the demand and pricing for keyword advertising in an integrated fashion across advertisers, searchers, and search engines. As a result, we develop a dynamic structural model of advertiser bidding behavior coupled with an attendant model of search behavior. Because we need to infer advertiser and consumer valuations and use these estimates to infer the effects of a change in search engine strategy, we develop a structural model of keyword search as a two-sided network. In particular, we consider i) how the platform or search engine should price its advertising via alternative auction mechanisms, ii) whether the platform should accommodate targeted bidding wherein advertisers bid not only on keywords, but also behavioral segments (e.g., those that purchase more often), and iii) how an alternative webpage design of the search engine with less product information would affect bidding behavior and the engine's revenues.

Our *model of advertiser bidding* behavior is predicated on the advertiser choosing its bids to maximize the net present value of its discounted profits. Specifically, we estimate advertiser valuations for clicks by choosing them such that, for an observed set of bids, the valuations rationalize

the bidding strategy. That is, their bids make advertisers' profits as high as possible. In this sense, our structural model "backs out" the advertiser's expectation for the profit per click. Given an estimate of these valuations, it becomes possible to ascertain how advertiser profits are affected by a change in the rules of the auction, a change in the webpage design, or a change in the information state of the advertiser.

We find that the estimated valuations for downloads/clicks are consistent with a download to sales ratio of 1.2%, well within industry estimates of 1% to 2%.

As noted above, a central component to the calculation of advertiser profits is the expectation of the number of clicks on its advertisement received from consumers. This expectation of clicks is imputed from our *consumer search and clicking model*. This model, which involves three steps (the choice of whether to search, whether to use search tools, and whether to download), follows from the standard random utility theory (McFadden, 1977).

Using the consumer and advertiser model, we conduct *policy simulations pertaining to search engine policy*. Relating to the consumer side, we explore the effect of changing the search engine's website design in order to reduce usability but increase advertising exposures. We manipulate usability by removing the sorting and filtering feature on the search engine site and find an overall reduction of 2.9% in search engine revenue, suggesting it would not be prudent to change the site. Second, we consider the possibility of allowing advertisers to bid by segment and allowing advertising slot rankings to differ by segment. Though this reduces competition within segments, targeting also enhances the expected number of downloads by increasing the relevance of the advertisements (suggesting larger search engine profits). Overall, the latter effect dominates, leading to an increase in search engine revenues of 1%. Third, we explore alternative auction designs. We find that a generalized second price auction leads to truth telling in advertiser bids and revenue equivalence for the search engine. This extends the work on generalized second price auction mechanisms to dynamic settings.

Several extensions are possible. First, we use a two-step estimator to model the dynamic bidding behavior of advertisers without explicitly solving for the equilibrium bidding strategy. Solving

explicitly for this strategy could provide more insights into bidder behavior in this new marketing phenomenon. For example, following the extant literature we assume that a bidder's return from advertising only comes from consumers' clicks. It is possible that advertisers also accrue some values from the exposures at the premium slots. Second, our analysis focuses upon a single category. The existence of multiple keywords auctions may present opportunities for collusion among bidders. By doing so, they can find a more profitable trade-off between payments to the search engine and clicks across keywords. One managerial implication is how to detect and discourage collusion and reduce its negative impact on search engine revenues. Third, competition between search engines over advertisers is not modeled. Though our data provider has a dominant role in this specific category, inter-engine competition is unattended in the literature. Fourth, the counterfactual policy functions reflect local equilibriums that are the closest to the observed policy (Doraszelski and Satterthwaite, 2009; Doraszelski and Escobar, 2009). Although this lessens the concern of multiplicity, we suggest that more rigid proof of the existence and uniqueness in the keyword auction context as a future research direction. Finally, our analysis is predicated on a relatively short duration of bidding behavior. Over the longer-term, there may be additional dynamics in bidding and download behavior that might arise from consumer learning or the penetration of search marketing into the market place, the so called "durable goods problem," (Horsky and Simon, 1983). Overall, we hope this study will inspire further work to enrich our knowledge of this new marketplace.

References

- Ansari, Asim, Carl F Mela. 2003. E-customization. *Journal of Marketing Research* **40**(2) 131–145.
- Arcidiacono, Peter, Robert Miller. 2009. CCP estimation of dynamic discrete choice models with unobserved heterogeneity. *Working Paper* .
- Athey, Susan, Glenn Ellison. 2008. Position auctions with consumer search. *Working Paper* .
- Bajari, Patrick, C. Lanier Benkard, Jonathan Levin. 2007. Estimating dynamic models of imperfect competition. *Econometrica* **75**(5) 1331–1370.
- Bajari, Patrick, Victor Chernozhukov, Han Hong, Denis Nekipelov. 2008. Nonparametric and semiparametric analysis of a dynamic game model. *Working Paper* .
- Ben-Akiva, Moshe, Steven Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press.
- Berry, Steven, James Levinsohn, Ariel Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* **63**(4) 841–890.
- Bradlow, Eric T., David C. Schmittlein. 2000. The little engines that could: Modeling the performance of world wide web search engines. *Marketing Science* **19**(1) 43–62.
- Bresnahan, Timothy F., Peter C. Reiss. 1991. Entry and competition in concentrated markets. *Journal of Political Economy* **99**(5) 977–1009.
- Bronnenberg, Bart J. 1998. Advertising frequency decisions in a discrete markov process under a budget constraint. *Journal of Marketing Research* **35**(3) 399–406.
- Bronnenberg, Bart J., Jean-Pierre Dubé, Carl F. Mela. 2009. Do DVRs moderate advertising effects? *Journal of Marketing Research* forthcoming.
- Chen, Yongmin, Chuan He. 2006. Paid placement: Advertising and search on the internet. *Working Paper* .
- Chung, Doug, Thomas Steenburgh, K. Sudhir. 2009. Do bonuses enhance sales productivity? a dynamic structural analysis of bonus-based compensation plans. *Working Paper* .
- Doraszelski, Ulrich, Juan Escobar. 2009. A theory of regular markov perfect equilibria in dynamic stochastic games: Genericity, stability, and purification. *Working Paper* .
- Doraszelski, Ulrich, Mark Satterthwaite. 2009. Computable markov-perfect industry dynamics. *Rand Journal of Economics* forthcoming.
- Dubé, Jean-Pierre, Günter J. Hitsch, Pradeep Chintagunta. 2008. Tipping and concentration in markets with indirect network effects. *Working Paper* .
- Dubé, Jean-Pierre, Günter J. Hitsch, Pranav Jindal. 2010. Estimating durable goods adoption decisions from stated preference data. *Working Paper* .

- Dubé, Jean-Pierre, K. Sudhir, Andrew Ching, Gregory Crawford, Michaela Draganska, Jeremy Fox, Wesley Hartmann, Günter Hitsch, V. Viard, Miguel Villas-Boas, Naufel Vilcassim. 2005. Recent advances in structural econometric modeling: Dynamics, product positioning and entry. *Marketing Letters* **16**(3) 209–224.
- Edelman, Benjamin, Michael Ostrovsky, Michael Schwarz. 2007. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *The American Economic Review* **97** 242–259.
- Farias, Vivek, Denis Saure, Gabriel Y. Weintraub. 2010. An approximate dynamic programming approach to solving dynamic oligopoly models. *Working Paper* .
- Feng, Juan. 2008. Optimal mechanism for selling a set of commonly-ranked objects. *Marketing Science* **27**(3) 501–512.
- Ghose, Anindya, Sha Yang. 2009. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science* **55**(10) 1605–1622.
- Goldfarb, Avi, Catherine Tucker. 2008. Search engine advertising: Pricing ads to context. *Working Paper* .
- Greene, William H. 2003. *Econometric analysis*. Prentice Hall.
- Hong, Han, Matthew Shum. 2006. Using price distributions to estimate search costs. *Rand Journal of Economics* **37**(2) 257–276.
- Horsky, Dan, Leonard S. Simon. 1983. Advertising and the diffusion of new products. *Marketing Science* **2**(1) 1–17.
- Hortacsu, Ali, Chad Syverson. 2004. Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. *Quarterly Journal of Economics* **119**(2) 403–456.
- Hotz, V. Joseph, Robert A. Miller. 1993. Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies* **60**(3) 497–529.
- Iyengar, Garud, Anuj Kumar. 2006. Characterizing optimal adword auctions. *Working Paper* .
- Jofre-Bonet, Mireia, Martin Pesendorfer. 2003. Estimation of a dynamic auction game. *Econometrica* **71**(5) 1443–1489.
- Judd, Kenneth L. 1998. *Numerical methods in economics*. MIT Press.
- Kamakura, Wagner A., Gary J. Russell. 1989. A probabilistic choice model for market segmentation and elasticity structure. *Journal of Marketing Research* **26**(4) 379–390.
- Katona, Zsolt, Miklos Sarvary. 2008. The race for sponsored links: A model of competition for search advertising. *Working Paper* .

- Kempe, David, Kenneth C. Wilbur. 2009. What can television networks learn from search engines? how to select, order, and price advertisements to maximize advertiser welfare. *Working Paper* .
- Kim, Jun, Paulo Albuquerque, Bart J. Bronnenberg. 2009. Online demand under limited consumer search. *Working Paper* .
- McFadden, Daniel L. 1977. Modeling the choice of residential location. *Cowles Foundation Discussion Paper No. 477* .
- Milgrom, Paul R., Robert J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* **50**(5) 1089–1122.
- Misra, Sanjog, Harikesh Nair. 2009. A structural model of sales-force compensation dynamics: Estimation and field implementation. *Working Paper* .
- Pakes, Ariel, Paul McGuire. 1994. Computing markov-perfect nash equilibria: Numerical implications of a dynamic differentiated product model. *The Rand Journal of Economics* **25**(4) 555–589.
- Pancras, Joseph, K. Sudhir. 2007. Optimal marketing strategies for a customer data intermediary. *Journal of Marketing Research* **44**(4) 560–578.
- Pesendorfer, Martin, Philipp Schmidt-Dengler. 2008. Asymptotic least squares estimators for dynamic games. *Review of Economic Studies* **75** 901–928.
- Rochet, Jean-Charles, Jean Tirole. 2006. Two-sided markets: A progress report. *Rand Journal of Economics* **37**(3) 645–667.
- Rust, John. 1994. Structural estimation of markov decision processes. Robert F. Engle, Daniel L. McFadden, eds., *Handbook of Econometrics*, vol. IV. Amsterdam: Elsevier Science.
- Rutz, Oliver J., Randolph E. Bucklin. 2007. A model of individual keyword performance in paid search advertising. *Working Paper* .
- Rutz, Oliver J., Randolph E. Bucklin. 2008. From generic to branded: A model of spillover dynamics in paid search advertising. *Working Paper* .
- Ryan, Stephen. 2009. The costs of environmental regulation in a concentrated industry. *Working Paper, Massachusetts Institute of Technology* .
- Train, Kenneth. 2003. *Discrete Choice Methods with Simulation*. Cambridge University Press.
- Varian, Hal R. 2007. Position auction. *International Journal of Industrial Organization* **25** 1163–1178.
- Vickrey, William. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* **16**(1) 8–37.
- Yao, Song, Carl F. Mela. 2008. Online auction demand. *Marketing Science* **27**(5) 861–885.
- Zanutto, Elaine, Eric Bradlow. 2006. The perils of data pruning in consumer choice model. *Quantitative Marketing and Economics* **4** 267 – 287.

Online Technical Appendix

A Two-step Estimator

A.1 First Step Estimation

A.1.1 Estimating the Advertiser’s Policy Function

The Partial Policy Function The partial policy function links states (\mathbf{s}) and characteristics (\mathbf{X}) to decisions (\mathbf{b}). Ideally this relation can be captured by a flexible parametric form and estimated via methods such as maximum likelihood or MCMC to obtain the partial policy function parameter estimates. The exact functional form is typically determined by model fit comparison among multiple specifications (e.g., Jofre-Bonet and Pesendorfer (2003)). We considered several different specifications for the distribution of bids and found the truncated normal distribution gives the best fit in terms of AIC.⁴¹ Specifically, we allow

$$b_j^t = \begin{cases} y_j^{t*} & \text{if } y_j^{t*} \geq \chi \\ 0 & \text{otherwise} \end{cases} \quad (\text{A1})$$
$$y_j^{t*} \sim N([\mathbf{s}^{t'}, \mathbf{X}_j^{t'}] \cdot \varphi + \varphi_j, \tau^2)$$

where $[\mathbf{s}^{t'}, \mathbf{X}_j^{t'}]$ is the vector of independent variables; τ is the standard deviation of y_j^{t*} ; φ_j is a bidder specific constant term due to the fixed effect f_j in valuations (Equation 19); χ is the truncation point, which is set at 15 to be consistent with the 15 cents minimum bid requirement of the search engine.⁴²

One possible concern when estimating the partial policy function $\tilde{\sigma}(\mathbf{s}, \mathbf{X})$ (and the full policy function $\sigma(\mathbf{s}, \mathbf{X}, r_j^t)$ next) is that there may be multiple equilibrium strategies; and the observed

⁴¹We experimented with alternative specifications including a Beta distribution and a Weibull distribution whose scale, shape, and location parameters are functions of (\mathbf{s}, \mathbf{X}) .

⁴²We also consider a semi-parametric specification by using thin-plate spline function. In particular, we model the bid level as non-parametric functions of lagged downloads, total past downloads, and update recency, but specify a linear function for the other covariates as the large dimension of the covariate space makes a fully non-parametric model infeasible. Results of this model are essentially identical to the polynomial tobit, though out of sample bid forecasts are slightly degraded (the MAPE increases from 0.05 to 0.09).

data are generated by multiple equilibriums. If this were the case, the policy function would not lead to a unique decision and would be of limited use in predicting advertiser behavior. It is therefore necessary to invoke the following assumption (BBL).

Assumption 2 (Equilibrium Selection): The data are generated by a single Markov perfect equilibrium profile σ .

Assumption 2 is relatively unrestrictive since our data is generated by auctions of *one* keyword and from *one* search engine. Given data are from a single market, the likelihood is diminished that different equilibriums from different markets are confounded. We note that this assumption is often employed in such contexts (e.g., Dubé et al. (2008)).

This partial policy function is then used to impute the full policy function $b_j = \sigma_j(\mathbf{s}, \mathbf{X}, r_j^t)$ as detailed below based on r_j^t 's distribution parameter ψ .

Full Policy Functions $\sigma_j^t(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$. To evaluate the value function of this dynamic game, we need to calculate bids as a function of not only $(\mathbf{s}^t, \mathbf{X}^t)$ but also the unobserved shocks r_j^t (see section 4.2.3). To infer this full policy function $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$ from the estimated partial policy function, $\tilde{\sigma}_j(\mathbf{s}^t, \mathbf{X}^t)$, we introduce one additional assumption.

Assumption 3 (Monotone Choice): For each bidder j , its equilibrium strategy $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$ is increasing in r_j^t (BBL).

Assumption 3 implies that bidders who draw higher private valuation shocks r_j^t will bid more aggressively.

To explore these two assumptions, note that the partial policy function $\tilde{\sigma}(\mathbf{s}^t, \mathbf{X}^t)$ presents distributions for bid b_j^t and the latent y_j^{t*} , whose CDF's we denote as $F_b(b_j^t | \mathbf{s}^t, \mathbf{X}^t)$ and $F(y_j^{t*} | \mathbf{s}^t, \mathbf{X}^t)$, respectively.⁴³ According to the model in equation A1, the population mean of y_j^{t*} across bidders and periods is $[\mathbf{s}^t, \mathbf{X}_j^t] \cdot \varphi + \varphi_j$. Around this mean, the variation across *bidders and periods* can be captured by the variance term τ^2 . With assumption 3, we can attribute τ^2 to the random shocks r_j^t .

Given the normal distribution assumption of the random shock $r_j^t \sim N(0, \psi^2)$, we may impute

⁴³To be more specific, we estimate a continuous distribution $F(y_j^{t*} | \mathbf{s}^t, \mathbf{X}^t)$ for y_j^{t*} from equation A1; then conditional on the truncation point χ , we can back out the (discontinuous) distribution $F_b(b_j^t | \mathbf{s}^t, \mathbf{X}^t)$ for b_j^t .

the y_j^{t*} (and hence b_j^t) for each combination of $(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$, i.e., the full policy function. To see this, note that since $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$ is increasing in r_j^t ,⁴⁴

$$F(y_j^{t*} | \mathbf{s}^t, \mathbf{X}^t) = \Pr(\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t) \leq y_j^{t*} | \mathbf{s}^t, \mathbf{X}^t) = \Phi(\sigma_j^{-1}(y_j^{t*}, \mathbf{s}^t, \mathbf{X}^t) / \psi)$$

where $\sigma_j^{-1}(y_j^{t*}, \mathbf{s}^t, \mathbf{X}^t)$ is the inverse function of $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$ with respect to r_j^t and $\Phi(\cdot)$ is the CDF of standard normal distribution. In equilibrium, we have $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t) = y_j^{t*}$. By substitution and rearrangement we get

$$\begin{aligned} y_j^{t*} &= \sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t) & (A2) \\ &= F^{-1}(\Phi(\sigma_j^{-1}(y_j^{t*}, \mathbf{s}^t, \mathbf{X}^t) / \psi) | \mathbf{s}^t, \mathbf{X}^t) \\ &= F^{-1}(\Phi(r_j^t / \psi) | \mathbf{s}^t, \mathbf{X}^t) \end{aligned}$$

where $\sigma_j^{-1}(y_j^{t*}, \mathbf{s}^t, \mathbf{X}^t) = r_j^t$; r_j^t / ψ has a standard normal distribution.

Therefore there is a unique mapping between the likelihood of observing a given valuation shock r_j^t and the y_j^{t*} . Each r_j^t drawn by a firm implies a corresponding quantile on the r_j^t 's distribution; this quantile in turn implies a y_j^{t*} from the distribution represented by that firm's partial bidding function $\tilde{\sigma}_j(\mathbf{s}^t, \mathbf{X}^t)$. However, because we do not know ψ and, thus, the distribution of r_j^t , we have to make draws from an alternative distribution r_j^t / ψ that has a one-one quantile mapping to r_j^t . To do this, we first draw a random shock r_j^t / ψ from $N(0, 1)$ for each advertiser i in period t . Next, we determine $F(y_j^{t*} | \mathbf{s}^t, \mathbf{X}^t)$ using results estimated in Equation A1 and looking at the distribution of its residuals to determine F . That is, for each value of y_j^{t*} , we should be able to compute its probability for a given \mathbf{s}^t and \mathbf{X}^t using F . Accordingly, F^{-1} links probabilities to y_j^{t*}

⁴⁴In this Appendix, we are abusing the notation of $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t)$. For the purpose of a clear exposition, we define $\sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t) = b_j^t$ in the paper. To match the bidding function estimated in equation A1, the more accurate definition should be

$$\begin{aligned} b_j^t &= \begin{cases} y_j^{t*} & \text{if } y_j^{t*} \geq \chi \\ 0 & \text{otherwise} \end{cases} \\ y_j^{t*} &= \sigma_j(\mathbf{s}^t, \mathbf{X}^t, r_j^t). \end{aligned}$$

(therefore b_j^t) for a given \mathbf{s}^t and \mathbf{X}^t . We then use F^{-1} to link the probability $\Phi(r_j^t/\psi)$ to b_j^t for a particular \mathbf{s}^t and \mathbf{X}^t . In this manner we ensure the bids and valuations in Equation A10 comport. In Appendix A.2.1, when evaluating the value function for a set of given parameter values of ψ in Equation A10 or evaluating base functions defined in Equation A11, we integrate out over the unobserved shocks r_j^t by drawing many r_j^t/ψ from $N(0, 1)$.

A.1.2 Consumer Model Estimation

We derive the consumer model conditioned on the information state of the advertiser as described in section 4.1. Given that advertisers do not observe what each person downloaded or the characteristics of these persons, they must infer consumer behavior from aggregate instead of individual level data.

Advertisers do observe the aggregate data in the form of download counts $d_j^t = \{d_1^t, d_2^t, \dots, d_N^t\}$ in period t . A single d_j^t follow a binomial distribution. Given the download probabilities P_j^t in Equation 15, a single d_j^t 's probability mass function is $\binom{M_t}{d_j^t} [P_j^t]^{d_j^t} [1 - P_j^t]^{M_t - d_j^t}$, where M_t is the consumer population size in period t . Hence the likelihood of observing d^t is

$$L(\mathbf{d}^t | \Omega_c) = \prod_j \binom{M_t}{d_j^t} [P_j^t]^{d_j^t} [1 - P_j^t]^{M_t - d_j^t}$$

where Ω_c are parameters to be estimated.

An advertiser's predicted downloads $d_j^t(k, X_j^t; \Omega_c)$ can readily be constructed using the parameter estimates as shown in equation 17

$$d_j^t(k, X_j^t; \hat{\Omega}_c) = M_t \hat{P}_j^t. \tag{A3}$$

This prediction is then used to forecast expectations of future downloads and slot positions in the firm's value function in the second step estimation.

A.1.3 State Transition Function $P(s'|b_j, \mathbf{b}_{-j}, \mathbf{s}, X)$

To compute the state transition, note that the marginal number of expected downloads is given by the expected downloads at a slot position multiplied by the probability of appearing in that slot position and then summed across all positions:

$$P(s'|b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}) = \sum_k d(k, X; \Omega_c) \Pr(k|b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}). \quad (\text{A4})$$

The expected downloads given a slot position in A4 is defined in 17. We can decompose the likelihood of appearing in slot k as follows

$$\begin{aligned} & \Pr(k|b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}) \\ &= \Pr_{\{k \leq K\}}(k|b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}) I\{k \leq K\} + \Pr_{\{k > K\}}(k|b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}) I\{k > K\} \end{aligned} \quad (\text{A5})$$

where $\Pr_{\{k \leq K\}}(k|b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X})$ is the probability of appearing in slot k of the sponsored search section (i.e., $k \leq K$), and $\Pr_{\{k > K\}}(k|b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X})$ is the likelihood of appearing in slot k of the organic search section (i.e., $k > K$). We discuss these two probabilities next.

Likelihood of Premium Slot $k \leq K$. Let us first consider the likelihood of winning one of the premium slots k ($k \leq K$), $\Pr_{\{k \leq K\}}(k|b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X})$ as an order statistic reflecting the relative quality of the advertiser's bid, which is defined as $b_j d_j^{(-1)}$. Higher quality bids are more likely to be assigned to better slots. Denote $\Psi_{bd}(b_{j'}, d_{j'}^{(-1)} | \mathbf{s}, \mathbf{X})$ as the distribution CDF of $b_{j'} d_{j'}^{(-1)}, \forall j'$, where $d_{j'}^{(-1)}$ is from the state vector and $b_{j'}$ has a distribution depending on the strategy profile $\sigma(\cdot)$.⁴⁵ For bidder j to win a premium slot k by bidding b_j , it implies that (1) among all of the other $N - 1$ competing bidders, there are $k - 1$ bidders who have a higher ranking than j in terms of $b_{j'} d_{j'}^{(-1)}$ and (2) the other ones have a lower ranking than j . The probability of having a higher ranking than j is $[1 - \Psi_{bd}(b_j d_j^{(-1)} | \mathbf{s}, X)]$. Thus the probability of bidder j winning slot k by bidding b_j is

⁴⁵It is difficult to write a closed form solution for Ψ_{bd} , but we may use the sample population distribution to approximate Ψ_{bd} .

simply an order statistics as shown below; note that the combination $\binom{N-1}{k-1}$ in the equation is because any $(k-1)$ out of the $(N-1)$ competing bidders can have a higher ranking than j .⁴⁶

$$\begin{aligned}
& \Pr_{\{k \leq K\}}(k | b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}) & (A6) \\
&= \binom{N-1}{k-1} [1 - \Psi_{bd}(b_j d_j^{(-1)} | \mathbf{s}, \mathbf{X})]^{k-1} [\Psi_{bd}(b_j d_j^{(-1)} | \mathbf{s}, \mathbf{X})]^{(N-1)-(k-1)} \\
&= \binom{N-1}{k-1} [1 - \Psi_{bd}(b_j d_j^{(-1)} | \mathbf{s}, \mathbf{X})]^{k-1} [\Psi_{bd}(b_j d_j^{(-1)} | \mathbf{s}, \mathbf{X})]^{N-k}
\end{aligned}$$

Likelihood of Organic Slot $k > K$. Next we consider what happens when an advertiser does not win this auction and is placed in the organic search section. In this case, by the rules of the auction, the bidder's slot is determined by its update recency compared to all products in the organic search section. For bidder j to be placed in organic slot $k > K$ it implies that (1) there are K bidders who have a higher ranking of $b_j d_j^{(-1)}$ than bidder j (i.e., j loses the auction) and (2) among the other $\bar{N} - K - 1$ products (i.e., all products at the search engine less those who win premium slots and j itself), there are $k - K - 1$ products that have a higher update recency than j and (3) the other ones have a lower ranking than j . Hence,

$$\begin{aligned}
& \Pr_{\{k > K\}}(k | b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}) \\
&= \Pr(k > K | b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}) \cdot \Pr(k | b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}, k > K) & (A7)
\end{aligned}$$

where the first term is the probability of losing the auction (condition 1) and the second term denotes the likelihood of appearing in position $k > K$ (condition 2 and 3). Note that the main reason for the difference between A6 and A7 is the change of ranking mechanisms. The ranking

⁴⁶An alternative interpretation of equation A6 is the probability mass function (PMF) of a binomial distribution. Among $N - 1$ competing bidders, there are $k - 1$ higher than bidder j and $(N - 1) - (k - 1)$ lower than j , and the probability of higher than j is $[1 - \Psi_{bd}(b_j d_j^{(-1)} | \mathbf{s}, \mathbf{X})]$. Hence, we may consider the expression in A6 as the PMF of a binomial distribution.

is based on $b_{j'}d_{j'}^{(-1)}$ for $k \leq K$ and update recency when $k > K$. The first term in A7 does not appear as an order statistics (as shown below) since when $k > K$ the order of $b_{j'}d_{j'}^{(-1)}$ becomes meaningless. Instead, the update recency is affecting the ranking. The two terms in A7 can be expressed as follows.

Losing the auction implies that among j 's $N - 1$ opponents, there are K bidders who have a higher ranking than j in terms of $b_{j'}d_{j'}^{(-1)}$. Hence,

$$\Pr(k > K | b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}) = \binom{N-1}{K} [1 - \Psi_{bd}(b_j d_j^{(-1)} | \mathbf{s}, \mathbf{X})]^K. \quad (\text{A8})$$

The conditional probability of being placed in an organic slot $k > K$ (condition 2 and 3) is, again, an order statistics.⁴⁷ This distribution is incumbent upon the update recency of all \bar{N} products exclusive of the K winners in the sponsored search section. Denoting the distribution of update recency of all products as Ψ_{up} , which can be approximated from the sample population distribution observed in the data, we obtain the following:

$$\begin{aligned} & \Pr(k | b_j, \mathbf{b}_{-j}, \mathbf{s}, \mathbf{X}, k > K) & (\text{A9}) \\ &= \binom{\bar{N} - K - 1}{k - K - 1} [1 - \Psi_{up}]^{k-K-1} [\Psi_{up}]^{(\bar{N}-K-1)-(k-K-1)} \\ &= \binom{\bar{N} - K - 1}{k - K - 1} [1 - \Psi_{up}]^{k-K-1} [\Psi_{up}]^{\bar{N}-k} \end{aligned}$$

Combining Equations A9 and A8 into A7, and then A7 and A6 into A5, yields the state transition equation.⁴⁸

Given that we have detailed the estimation of the first step functions ($\sigma_j(\mathbf{s}, \mathbf{X}, r_j^t)$, $d_j^t(k, X_j^t; \Omega_c)$, $P(\mathbf{s}' | \mathbf{b}, \mathbf{s}, \mathbf{X})$), we now turn to the second step estimator, which is incumbent upon these first step

⁴⁷This order statistics can again be interpreted as the PMF of a binomial distribution similar to A6.

⁴⁸Note that a change in the number of sponsored links has no appreciable effect on the computational burden implied by A5 and thus A4. Likewise, the number of sponsored links has no practical impact on the computation of equation 20. Hence our approach generalizes readily to a larger number of links.

functions.

A.2 Second Step Estimation of Bidder Model

In this Appendix we detail how to estimate the parameters in the value function. This is done in two phases: first, we simulate the value function conditioned on Ω_a , and second, we construct the likelihood using the simulated value function conditioned on Ω_a .

A.2.1 Phase 1: Simulation of Value Functions Given Ω_a

To construct the value function we first simplify its computation by linearization, and second using this simplification, we simulate the expected value function conditioned on Ω_a by integrating out over draws for \mathbf{s}^t , \mathbf{X}^t , and \tilde{r}_j^t .

Linearize the Value Function We simplify the estimation procedure by relying on the fact that Equation 20 is linear in the parameters Ω_a . We can rewrite Equation 20 by factoring out Ω_a .

$$\begin{aligned}
& \mathbf{E}\pi_j(\mathbf{b}^t, \mathbf{s}^t, \mathbf{X}^t, r_j^t; \Omega_a) && \text{(A10)} \\
&= \sum_{k=1}^K \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot (v(X_j^t; \theta) + f_j + r_j^t - b_j^t) \cdot d_j^t(k, X_j^t; \Omega_c) \\
&\quad + \sum_{k=K+1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot (v(X_j^t; \theta) + f_j + r_j^t) \cdot d_j^t(k, X_j^t; \Omega_c) \\
&= \left[\sum_{k=1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot d_j^t(k, X_j^t; \Omega_c) \cdot X_j^t \right] \cdot \theta \\
&\quad + \left[\sum_{k=1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot d_j^t(k, X_j^t; \Omega_c) \right] \cdot f_j \\
&\quad + \left[\sum_{k=1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot d_j^t(k, X_j^t; \Omega_c) \cdot \tilde{r}_j^t \right] \cdot \psi \\
&\quad - b_j^t \sum_{k=1}^K \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) d_j^t(k, X_j^t; \Omega_c) \\
&= Base_{j1}^t [\theta', f_j]' + Base_{j2}^t \psi - Base_{j3}^t
\end{aligned}$$

where

$$\begin{aligned}
Base_{j1}^t &\equiv \left[\frac{\sum_{k=1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot d_j^t(k, X_j^t; \Omega_c) \cdot X_j^t}{\sum_{k=1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot d_j^t(k, X_j^t; \Omega_c)} \right]' & (A11) \\
Base_{j2}^t &\equiv \left[\sum_{k=1}^{\bar{N}} \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) \cdot d_j^t(k, X_j^t; \Omega_c) \cdot \tilde{r}_j^t \right] \\
Base_{j3}^t &\equiv b_j^t \sum_{k=1}^K \Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t) d_j^t(k, X_j^t; \Omega_c) \\
\tilde{r}_j^t &= r_j^t / \psi \sim N(0, 1).
\end{aligned}$$

Note that the values of $\{Base_{j1}^t, Base_{j2}^t, Base_{j3}^t\}_{\forall t}$ are conditionally independent of θ, f_j and ψ . This enables us to first evaluate $\{Base_{j1}^t, Base_{j2}^t, Base_{j3}^t\}_{\forall t}$ and keep them constant when drawing θ, f_j and ψ from their posterior distributions. By doing so, we reduce the computational burden of estimation as described next.

Simulate the Value Functions Given Ω_a . After the linearization, given a set of advertiser parameters $\Omega_a = \{\theta, f_{j=1,2,\dots,N}, \psi\}$ and Equation A10, the value function depicted in Equation 22 can also be written as the following with period index t invoked:

$$\begin{aligned}
V_j(\mathbf{s}^0, \mathbf{X}^0; \sigma; \Omega_a) &= \mathbf{E}_{\mathbf{s}, \mathbf{X}, r} \left[\sum_{t=0}^{\infty} \rho^t \pi_j(\sigma, \mathbf{s}^t, \mathbf{X}^t, r_j^t; \Omega_a) \right] & (A12) \\
&= \mathbf{E} \left[\sum_{t=0}^{\infty} (\rho^t Base_{j1}^t \begin{bmatrix} \theta \\ f_j \end{bmatrix} + Base_{j2}^t \psi - Base_{j3}^t) \right] \\
&= \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j1}^t \right] \begin{bmatrix} \theta \\ f_j \end{bmatrix} + \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j2}^t \right] \psi - \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j3}^t \right]
\end{aligned}$$

where the expectation is taken over current and future private shocks, future states \mathbf{s}^t , future \mathbf{X}^t and R^t .

An estimated value function $\widehat{V}_j(\mathbf{s}^0, \mathbf{X}^0; \sigma; \Omega_a)$ can then be obtained by the following steps:

1. Draw private shocks \tilde{r}_j^t from $N(0, 1)$ for all bidders j in period 0; draw initial choice of \mathbf{s}^0 from the distribution of state variables derived from the observed data; draw \mathbf{X}^0 from the observed distribution of product attributes.
2. Starting with the initial state \mathbf{s}^0 , \mathbf{X}^0 and the \tilde{r}_j^t step 1, calculate \hat{b}_j^0 for all bidders using the inversion (equation A2) described in Appendix A.1.1.
3. Use \mathbf{s}^0 , \mathbf{X}^0 and $\hat{\mathbf{b}}^0$ to determine the slot ranking, whose distribution is $\Pr(k|b_j^t, \mathbf{b}_{-j}^t, \mathbf{s}^t, \mathbf{X}^t)$ in Equation A5 in Appendix A.1.3; using $d(k, X_j^0; \Omega_c)$ in Equation 17, obtain a new state vector \mathbf{s}^1 , whose distribution is $P(\mathbf{s}^1|\hat{\mathbf{b}}^0, \mathbf{s}^0, \mathbf{X}^0)$ in Equation A4 in Appendix A.1.3; draw \mathbf{X}^1 from the observed distribution of product attributes.
4. Repeat step 1-3 for T periods for all bidders to compute all \mathbf{s}^t , \mathbf{X}^t , \tilde{r}_j^t , and \mathbf{b}^t for all periods; T is large enough so that the discount factor ρ^T approaches 0.
5. Using \mathbf{s}^t , \mathbf{X}^t , \tilde{r}_j^t , $d_j^t(k, X_j^t; \Omega_c)$, and \mathbf{b}^t , evaluate $\{Base_{j1}^t, Base_{j2}^t, Base_{j3}^t\}_{t=0, \dots, T}$ and $\left\{ \left[\sum_{t=0}^T \rho^t Base_{j1}^t \right], \left[\sum_{t=0}^T \rho^t Base_{j2}^t \right], \left[\sum_{t=0}^T \rho^t Base_{j3}^t \right] \right\}$.
6. The resulting values of

$$\{Base_{j1}^t, Base_{j2}^t, Base_{j3}^t\}_{t=0, \dots, T}$$

and

$$\left\{ \left[\sum_{t=0}^T \rho^t Base_{j1}^t \right], \left[\sum_{t=0}^T \rho^t Base_{j2}^t \right], \left[\sum_{t=0}^T \rho^t Base_{j3}^t \right] \right\}$$

depend on the random draws of \mathbf{s}^t , \mathbf{X}^t , r^t . To compute

$$\left\{ \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j1}^t \right], \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j2}^t \right], \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j3}^t \right] \right\},$$

repeat step 1-6 for NR times so as to integrate out over the draws. Note that when T is large enough, $\left[E \sum_{t=0}^T \rho^t Base_{j.}^t \right]$ is a good approximation of $\left[E \sum_{t=0}^{\infty} \rho^t Base_{j.}^t \right]$ since ρ^T approaches 0.

7. Conditional on a set of parameters Ω_a and

$$\left\{ \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j1}^t \right], \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j2}^t \right], \left[\mathbf{E} \sum_{t=0}^{\infty} \rho^t Base_{j3}^t \right] \right\},$$

we may evaluate $\widehat{V}_j(\mathbf{s}^0, \mathbf{X}^0; \sigma; \Omega_a)$ from Equation A12.

An estimated deviation value function $\widehat{V}_j(\mathbf{s}^0, \mathbf{X}^0; \sigma'_j, \sigma_{-j}; \Omega_a)$ with an alternative strategy σ'_j other than σ_j can be constructed by following the same procedure. We draw a deviated strategy σ'_j by adding disturbance to the estimated policy function from Step 1. In particular, we add a normally distributed random variable ($mean = 0$; $s.d. = 0.3$) to each parameter.

We implement this process by first drawing initial states for each bidder and $\{X^t\}_{t=0,1,\dots,T}$ of all $T = 100$ periods. Then for each combination of bidder and initial state, we use this process to compute the base value functions and $ND = 100$ perturbed base functions. In Step 6, we use $NR = 100$. The discount factor ρ is fixed as 0.99.

The computational burden is reduced tremendously since we have linearized the value functions and factored out the parameters Ω_a . We do not need to re-evaluate the value functions for each set of parameters Ω_a . Instead, we only evaluate the base functions in Equation A11 once using step 1-6 and keep them fixed. Then for each draw of Ω_a from the posterior distribution we may evaluate the value functions (step 7) so as to recover Ω_a as described below.

A.2.2 Phase 2: Recover Ω_a

Recall our goal is to recover the Ω_a that satisfies equation 23. Expressing equation 23 in its simulated analog, we obtain

$$\left\{ \widehat{V}_j(\mathbf{s}^0, \mathbf{X}^0; \sigma_j, \sigma_{-j}; \Omega_a) \geq \widehat{V}_j(\mathbf{s}^0, \mathbf{X}^0; \sigma'_j, \sigma_{-j}; \Omega_a) \right\}, \quad (\text{A13})$$

This condition means that the estimated value function for any given initial state s^0 , with observed strategy σ_j , is greater than the estimated value function with any deviation σ'_j from that observed

strategy. Define

$$\hat{g}_1(\mathbf{s}^0, \mathbf{X}^0; \sigma, \sigma'_j; \Omega_a) \quad (\text{A14})$$

$$= \frac{1}{ND} \sum_{nd=1}^{ND} \min \left\{ 0, \widehat{V}_j(\mathbf{s}^0, \mathbf{X}^0; \sigma_j, \sigma_{-j}; \Omega_a) - \widehat{V}_j(\mathbf{s}^0, \mathbf{X}^0; \sigma'_j, \sigma_{-j}; \Omega_a)_{(nd)} \right\}$$

$$\hat{g}_2(\mathbf{s}^0, \mathbf{X}^0; \sigma, \sigma'_j; \Omega_a) \quad (\text{A15})$$

$$= \min \left\{ 0, \widehat{V}_j(\mathbf{s}^0, \mathbf{X}^0; \sigma_j, \sigma_{-j}; \Omega_a) \right\}$$

$$\hat{g}(\mathbf{s}^0, \mathbf{X}^0; \sigma, \sigma'_j; \Omega_a) = \begin{bmatrix} \hat{g}_1 \\ \hat{g}_2 \end{bmatrix}. \quad (\text{A16})$$

where (nd) denotes the nd -th simulated deviated value function under the deviated policy function.

g_1 is a negative number if the deviation leads to a greater value function than the observed strategy

and is 0 otherwise. g_2 is the Individual Rationality (IR) constraint, i.e., the value function should

be greater than zero. g_2 is a negative number if the IR constraint is violated and is 0 otherwise.

Since there are 21 advertisers, the dimension of \hat{g} is 42 by 1.

The minimum distance estimator is set up such that

$$\hat{\Omega}_a = \arg \min_{\Omega_a} \left\{ [\hat{g}(\mathbf{s}^0, \mathbf{X}^0; \sigma, \sigma'_j; \Omega_a)]' W \hat{g}(\mathbf{s}^0, \mathbf{X}^0; \sigma, \sigma'_j; \Omega_a) \right\} \quad (\text{A17})$$

where W is a 42 by 42 weighting matrix. As \hat{g} is not necessarily differentiable, the computation

of the covariance matrix as well as the optimal weighting matrix W^* is infeasible. Hence we use

an identity matrix for W . Another consideration pertains to the use of the first step estimates in

the second step estimation. Accordingly, the standard errors of these second step estimates need

to account for the empirical distribution of the first step estimates. Hence we use a bootstrapping

method in which we draw the first step parameters from their empirical distributions, repeat the

simulation of value functions, and re-estimate equation A17 for 100 times to obtain the standard

errors for these second step estimates.

A.3 Identification

We overview our identification strategy in this appendix, beginning with the advertiser model and concluding with the consumer model.

A.3.1 Advertiser Model

To achieve the point identification in BBL, we need to invoke two assumptions pertaining to the GMM estimator in Equation A17:

- The set of parameters Ω_a is compact and the true parameters minimize the objective function of the estimator in Equation A17.
- The objective function of the estimator in Equation A17 is twice differentiable in the parameters Ω_a .

Next we focus the discussion on the identification of heterogeneity across advertisers in bidding policy and in click valuations.

The policy function estimated in the first step is specified such that the action (bid levels) is a function of the state variables (past downloads and product attributes) and individual-specific fixed effects. We observe bid levels, product attributes and past downloads across advertisers and time. The correlation of bid levels and state variables across time and advertisers allows the identification of the effects of state variables on bids. The remaining variation in bids are caused by either the fixed effects or the error terms. The identification of the fixed effects is achieved through the variations of bid levels across advertisers. The variations of bids across both advertisers and time identify the variance of the errors.

In the second step estimation, several factors help to identify advertiser valuations. First, there is a one to one mapping between bid levels and the advertisers' click values (Milgrom and Weber, 1982). Accordingly, bid levels monotonically increase with click valuations. Hence, the variation in bids across advertisers and time is informative about advertisers' valuations. Observing multiple bids per advertisers, and multiple attribute levels across bids provides information about

both advertiser specific valuations and the attribute specific valuations. The variation of valuations across both advertisers and time identify the variance of random shocks. Second, the valuations should satisfy the equilibrium condition specified in equation 23 and the IR constraint such that $V_j(\mathbf{s}, \mathbf{X}; \sigma_j, \sigma_{-j}; \Omega_a) \geq 0$. These latter two constraints further bound the imputed valuations. Together, these three constraints yield the conditional distributions of click valuations across advertisers and time.

A.3.2 Consumer Model

The consumer log files include (1) product attributes across products and time, (2) slot rankings across products and time, (3) product downloads across products and time, and (4) consumer browsing information (clicks and search strategy) across time.

First, conditioned on the latent segments, the identification of the three decision models follows the argument of classical discrete choice models. In the latent class model, all parameters are segment-specific and fixed across time. As a result, the variations of product attributes, slot rankings and the corresponding product downloads across products and time enable us to identify these respective parameters. Further, as discussed in the consumer model, several normalizations facilitate identification: (1) the utility of outside good (not download) is normalized to zero as implied by the multivariate probit specification in the download model, (2) the variances of the search model and sorting/filtering model are fixed under the nested logit and logit specifications, and (3) the fixed effect of one of the sorting/filtering is normalized to zero (no sorting or filtering).

The segment specific parameter distribution is primarily identified from the mixture model as in the mixed logit with aggregate level data (e.g., Berry et al., 1995, Train, 2003). In particular, variation in attributes and slot rankings over time and products induce changes in browsing behavior (sorting/filtering) over time and demand over products and time. This variation identifies heterogeneity in preference across segments because individual level preference i.i.d. shocks have been integrated out across the population and time.

B Policy Simulations

It is reasonable to expect that advertisers will change their bidding strategy in response to changes in a search engine’s new policy. Thus, the advertiser bidding rules estimated in the first stage of our analysis are not likely to reflect advertiser behavior under the new policy. Hence, we need to solve the new optimal bidding strategy for advertisers conditional on the primitives estimated off the data. This requires explicitly solving the dynamic programming problem (DP) for advertisers. Because of the dimension of the state space and the interaction across advertisers, solving the infinite-horizon DP imposes a tremendous computational burden. In this appendix, we first outline our general approach to solving this dynamic advertiser bidding problem and then detail the manipulations underpinning the specific policy simulation of segmentation and targeting.

B.1 Computational Considerations

To solve the advertiser bidding problem we rely on maximizing the value function conditioned on the states and the expected response of competitors. More specifically, our solution to the game applies a modified version of the new advance in the literature of approximate dynamic programming, using iterated best response approach to solve oligopolistic dynamic games (Farias et al., 2010, Pakes and McGuire, 1994, Judd, 1998).

Denote the primitives of the consumers to be Ω_c . These primitives include all of the parameters in the consumer model. Further, denote the primitives of the advertiser model to be $\Omega_a = \{\theta, \psi, f'_{\forall j}\}$. These primitives include the advertiser valuations obtained from the consumers who click on their links. It is these primitives, Ω_c and Ω_a , that we presume to be invariant in the policy simulations. Next, denote the advertisers’ bidding policy parameters as $\Omega_Z = \{\varphi', \{\varphi_j\}'_{\forall j}, \tau\}$. These bidding function parameters can presumably change in response to changes in search engine policy and it is our objective below to find the set of parameters Ω_Z that maximize advertisers profits in response to a change in policy by the search engine. Also denote $\Omega_z = \{\varphi', \varphi_z, \tau\}$ and $\Omega_{-z} = \{\varphi_j\}'_{\forall j \neq z}$ whereas z is a given focal bidder. Specifically, the algorithm proceeds as follows.

To initiate the process, we first choose a given bidder as the focal bidder, indexed as z (we

will elaborate the choice criterion shortly). Second, we randomly draw a vector of state variables (past downloads and product attributes) from their empirical distributions. These state variables are used as the initial state $\mathbf{s}^0, \mathbf{X}^0$. Third, we specify a parametric form for policy as a function of state variables and the random shock, $b_z = \sigma(\mathbf{s}, \mathbf{X}, r_z; \Omega_z)$. Though the parameters of this decision rule can differ, we use the same functional form as the one estimated in the advertiser model (Equation A1). Thus $\Omega_z = \{\varphi', \varphi_z, \tau\}$ includes both the parameters common across bidders (φ', τ) and bidder-specific fixed effect (φ_z). Fourth, for all bidders j and periods t we draw a random shock r_j^t from the distribution $N(0, \psi)$. Define r^t as a vector of random shocks whose elements are $r_j^t, j = 1, \dots, z, \dots, N$ and $r = [r^0, r^1, \dots, r^T]$, $T = 100$. We repeat this step $NR = 50$ times yielding 50 draws for the sequence of r . Fifth, we find the parameters of this policy function that maximize firms' profits in an oligopolistic game where all firms move simultaneously, forming expectations about the bids of others. In this fifth step parameters are chosen for the bidders. The parameters $\Omega_{\mathbf{Z}} = \{\varphi', \{\varphi_j\}'_{\forall j}, \tau\}$ are calculated as follows:

(0) Initialize $\Omega_{\mathbf{Z}(old)} = \{\varphi'_{(old)}, \{\varphi_{j(old)}\}'_{\forall j}, \tau_{(old)}\}$ using the estimates of the advertiser model in Section 6.1.1.

(1a) For a given set of parameters $\Omega_z = \{\varphi', \varphi_z, \tau\}$, we may calculate the value of the following for bidder z :⁴⁹

$$\begin{aligned} & \mathcal{V}_z(\mathbf{s}^0, \mathbf{X}^0; \sigma_z, \sigma_{-z}, r; \Omega_z, \Omega_{-z(old)}) \\ = & \pi_z(\sigma_z, \sigma_{-z}, \mathbf{s}^0, \mathbf{X}^0, r^0; \Omega_z, \Omega_{-z(old)}) \\ & + \sum_{t=1}^T \rho^t \pi_z(\sigma_z, \sigma_{-z}, \mathbf{s}^t, \mathbf{X}^t, r^t; \Omega_z, \Omega_{-z(old)}) P(\mathbf{s}^t | \mathbf{b}^{t-1}, \mathbf{s}^{t-1}, \mathbf{X}^{t-1}) \end{aligned}$$

With the symmetry assumption, parameters in the bidding policy function $\{\varphi', \tau\}$ are common across all competing bidders (though competing bids can differ because their i) fixed effects in the decision rule, ii) errors (r_j^t) and iii) state variables (\mathbf{s}, \mathbf{X}) can differ). This symmetry assumption

⁴⁹The ensuing derivations are conditioned upon the invariant primitives Ω_c and Ω_a . However, to facilitate exposition, we omit Ω_c and Ω_a from the following equations.

enables us to compute expectations for bids of the competing bidders. Likewise, we can compute the transition probability $P(s'|b, s, \mathbf{X})$.

(1b) We repeat step (1a) for each draw r and average the \mathcal{V}_z . This yields an approximation of the value function for bidder z .

$$\hat{V}_z(s^0, X^0; \sigma_z, \sigma_{-z}; \Omega_z, \Omega_{-z(old)}) = \frac{1}{NR} \sum_{nr=1}^{NR} \mathcal{V}_z(s^0, X^0; \sigma_z, \sigma_{-z}, r^{(nr)}; \Omega_z, \Omega_{-z(old)})$$

(2) We then search for the $\Omega_{z(new)} = \{\varphi'_{(new)}, \varphi_{z(new)}, \tau_{(new)}\}$ that maximizes the $\hat{V}_z(\cdot; \Omega_z, \Omega_{-z(old)})$

$$\Omega_{z(new)} = \arg \max_{\Omega_z} \hat{V}_z(s^0, X^0; \sigma_z, \sigma_{-z}; \Omega_z, \Omega_{-z(old)})$$

thus yielding new $\Omega_{z(new)}$ conditioned on the imputed V .

(3) We then iterate through all competing bidders one by one using steps similar to (1) and (2) to search for updated fixed effects φ_{-z} , conditioned on the updated common parameters $\{\varphi_{(new)}, \tau_{(new)}\}$ and the fixed effects for other bidders.

(4) Conditional on new parameters $\{\varphi_{(new)}, \tau_{(new)}\}$ and new fixed effects, impute the updated value functions for all bidders:

$$\hat{V}_j(s^0, X^0; \sigma_j, \sigma_{-j}; \Omega_{j(new)}, \Omega_{-j(new)}) = \frac{1}{NR} \sum_{nr=1}^{NR} \mathcal{V}_j(s^0, X^0; \sigma_j, \sigma_{-j}, r^{(nr)}; \Omega_{j(new)}, \Omega_{-j(new)}), \forall j$$

If

$$\max_j \left| \hat{V}_j(s^0, X^0; \sigma_j, \sigma_{-j}; \Omega_{j(new)}, \Omega_{-j(new)}) - \hat{V}_j(s^0, X^0; \sigma_j, \sigma_{-j}; \Omega_{j(old)}, \Omega_{-j(old)}) \right| \leq 0.01,$$

stop the iteration. Otherwise, repeat from step (1a).

The updated policy rule (for each policy simulation) should not be sensitive to the choice of z , i.e., the choice of which bidder's value function to start the iteration. Nonetheless, as a robustness check, we experiment with several different bidders' value functions and find the results to all

be similar. In particular, we repeat the calculation for the most frequent bidder, the least frequent bidder, three other randomly chosen bidders. The reported counterfactuals are based on a randomly chosen bidder.

B.2 Policy Simulation: Segmentation and Targeting

Neither the search engine nor advertisers actually observes the segment memberships of consumers to help with targeting. However, it is possible for the advertiser to infer the posterior probability of consumer i 's segment membership conditional on its choices. These estimates can then be used to improve the accuracy and effectiveness of targeting.

More specifically, suppose the search engine observes consumer i in several periods. Let us consider consumer i 's binary choices over downloading, sorting/filtering, and searching in those periods. Denote these observations as $H_i(\{y_{ijt}\}_{j,t}, \{\kappa_{it}\}_t, \{search_{it}\}_t)$. The likelihood of observing $H_i(\{y_{ijt}\}_{j,t}, \{\kappa_{it}\}_t, \{search_{it}\}_t)$ is

$$\begin{aligned} & L(H_i(\{y_{ijt}\}_{j,t}, \{\kappa_{it}\}_t, \{search_{it}\}_t)) \\ &= \sum_g \prod_t L(H_i(\{y_{ijt}\}_j, \{\kappa_{it}\}_t, \{search_{it}\}_t) | g_{it}) \cdot pg_{it}^g \end{aligned} \quad (\text{A18})$$

where

$$\begin{aligned} & L(H_i(\{y_{ijt}\}_{j,t}, \{\kappa_{it}\}_t, \{search_{it}\}_t) | g_{it}) \\ &= \prod_j \int_{u_{ijt}^{g\kappa}} \int_{\mathbf{z}_{it}^g} \pi(y_{ijt} | u_{ijt}^{g\kappa}, \kappa_{it}^g, g_{it}) \pi(\kappa_{it}^g | \mathbf{z}_{it}^g, g_{it}) du_{ijt}^{g\kappa} d\mathbf{z}_{it}^g \Pr(search_{it}^g) \end{aligned} \quad (\text{A19})$$

Hence, the posterior probability of segment membership for consumer i can be updated in a Bayesian fashion,

$$\begin{aligned} \Pr(i \in g | H_i(\{y_{ijt}\}_{j,t}, \{\kappa_{it}\}_t, \{search_{it}\}_t)) \\ &= \frac{\prod_t L(H_i(\{y_{ijt}\}_j, \{\kappa_{it}\}_t, \{search_{it}\}_t) | g_{it}) \cdot pg_{it}^g}{\sum_{g'} \prod_t L(H_i(\{y_{ijt}\}_j, \{\kappa_{it}\}_t, \{search_{it}\}_t) | g'_{it}) \cdot pg_{it}^{g'}} \end{aligned} \quad (\text{A20})$$

As a consequence, the engine will have a more accurate evaluation about the segment membership of that consumer.

On the other hand, suppose some consumers only visit the engine once. Before they make the product choices, the search engine cannot obtain a posterior distribution outlined in Equation A20 since their choices of products are still unavailable. Still, it is possible to establish a more informative prediction about their memberships based on their κ_{it} 's before their product choices. Similar to Equation A20, the posterior in this case is

$$\Pr(i \in g | H_i(\kappa_{it})) = \frac{L(H_i(\kappa_{it}) | g_{it}) \cdot pg_{it}^g}{\sum_{g'} L(H_i(\kappa_{it}) | g'_{it}) \cdot pg_{it}^{g'}} \quad (\text{A21})$$

where

$$L(H_i(\kappa_{it}) | g_{it}) = \int_{z_{it}^{g_{it}}} \pi(\kappa_{it}^g | z_{it}^g, g_{it}) dz_{it}^{g_{it}}$$

We can construct an analysis to consider the benefits of targeting as follows. First, we compute the return to advertisers when advertisers can only bid on keywords for all segments. Second, we compute the return accruing to advertisers when they can bid for keywords at the segment level using the approach detailed in section B.1. The difference between the two returns can be considered as a measure for the benefits of targeting. At the same time, we may calculate the advertisers' returns under the two scenarios and the difference may be a measure for the value of market intelligence.