

Strategic Bidder Behavior in Sponsored Search Auctions

Benjamin Edelman
Harvard University
bedelman@fas.harvard.edu

Michael Ostrovsky*
Stanford University
ostrovsky@gsb.stanford.edu

Abstract

We examine sponsored search auctions run by Overture (now part of Yahoo!) and Google and present evidence of strategic bidder behavior in these auctions. Between June 15, 2002, and June 14, 2003, we estimate that Overture’s revenue from sponsored search might have been higher if it had been able to prevent this strategic behavior. We present a specific alternative mechanism that could reduce the amount of strategizing by bidders, raise search engines’ revenue, and also increase the overall efficiency of the market. We conclude by showing that advertisers’ strategic behavior has not disappeared over time; rather, such behavior remains present on both search engines.

1 Introduction

Search engines have considerable flexibility in designing systems to allocate their sponsored links. Modern information systems make it easy to receive information from advertisers about their desired ad placements and other goals, as well as to provide feedback about outcomes. But how exactly should search engines structure these sales?

Auctions are a natural choice. They relieve sellers of explicitly assigning a valuation to each item to be sold; instead, the market structure naturally assigns a valuation. This benefit is surely substantial in the context of thousands of search keywords, each independently valued by different would-be advertisers. Furthermore, auctions generally assure that each available item is sold—also an important benefit, since search advertising inventory is perishable, disappearing instantly if no advertisement is available to fill a placement spot.

Hence, it is not surprising that most search engines choose to auction off their sponsored link advertisements. But even after choosing an auction over alternative sales channels, important decisions remain to be made. How much information should advertisers receive about other advertisers’ bids? In what order should bids be shown in search results? What relationship should there be between an advertiser’s bid and the amount the advertiser is actually charged? More broadly, how should these auctions be designed?

*We are grateful to David Pennock and Yahoo! for data and advice. We also thank the editor, two anonymous referees, Drew Fudenberg, Louis Kaplow, Ariel Pakes, Al Roth, Michael Schwarz, and the participants of the First Workshop on Sponsored Search Auctions for helpful comments and suggestions.

While it is clear that differently designed auctions will have different properties, it is not immediately obvious that these differences are important enough for market designers to worry about. After all, if there are many bidders who compete for the same keywords, prices might be expected to be close to efficient, and revenue might be close to the highest possible level. In this paper, we show that this is not necessarily the case, and that auction design is a first-order question in these markets. More specifically, we estimate that during the period for which we have bidding data, an alternative design (based on the Vickrey auction) would have raised revenues for popular keywords by almost 10%. Limitations in our methodology make our estimates lower bounds of the true magnitude of potential improvements.

We emphasize search engine revenues conditional on users staying at search engines, and search engines remaining otherwise the same. We bracket the important issues raised in [2] and [13] as to a tradeoff between advertisement inclusion and a notion of quality (i.e., user satisfaction). We also put aside the relationship between relevance and advertiser valuations, as raised in [7]. Instead, we take advertiser valuations as given, conditional on other search engine characteristics (i.e., overall quality, user demographics, etc.), which we take to be unchanged across our counterfactuals.

The remainder of this article proceeds as follows. In Section 2, we describe the “cycling” patterns that we observe in the bidding data from Overture auctions. In Section 3, we present a brief theoretical example of losses from cycling, identify a specific improvement (the use of Vickrey auctions to induce truthful bidding and discourage strategic behavior), and compute counterfactual revenues. In Section 4, we present an analysis of current advertisement data, suggesting that strategic bidding behavior has not disappeared. In Section 5, we conclude by discussing practical considerations in changing a search engine’s auction mechanism.

2 Historical Evidence and Motivating Examples

In this section, we present evidence of strategic behavior of bidders using Overture’s paid placement auctions between June 2002 and June 2003.¹ During this period, Overture generally operated a first-price auction. That is, when a user clicked on an advertiser’s link, the advertiser paid Overture the amount of his bid.

This first-price auction structure is naturally unstable, in the sense that if bids can be adjusted frequently, bidders will not state their true valuations, and will keep changing their bids in response to other bidders’ behavior. For example, in a keyword market with two advertisers, suppose a click is worth \$0.60 to the first advertiser and \$0.80 to the second. If the first advertiser bids \$0.60, the second will bid \$0.61, thereby claiming the first spot.² Then, the first bidder is likely to lower his bid to the minimum bid (say, \$0.05), reducing his costs while still preserving his number two position, apparently the best position he can obtain given the other advertiser’s high bid. But then the second advertiser will lower his bid, e.g., to \$0.06, and the advertisers will raise each other in pennies until the second advertiser again outbids the first advertiser’s valuation (e.g., with a bid of

¹Overture was acquired by Yahoo! in the second half of 2003, after the end of our sample.

²Assuming, of course, that the top spot is more desirable than the second spot.

\$0.61), and the first advertiser again drops back to the minimum bid. Under these assumptions, their “cycling” will continue indefinitely.³

Moreover, some Overture advertisers apparently used an “autobid” system, whereby an Overture software system automatically adjusted an advertiser’s bid to achieve desired placement and to avoid overbidding. Advertisers began by telling the autobidder their maximum willingness to pay for a click, for a given keyword. The autobidder then automatically raised the advertiser’s bid in \$0.01 increments to obtain the highest possible position given the advertiser’s willingness to pay. The autobidder recognized opportunities to lower the advertiser’s bid without losses in rankings. For example, if a bidder was paying \$0.80 per click for the first place, while the next highest bidder was paying \$0.60, the autobidder would recognize that the first bidder could lower his bid to \$0.61 while remaining in the first place.

Overture’s (and, possibly, outside software developers’) autobidders exacerbated the instability of first-price auctions. When two or more advertisers activated autobidders, their bids tended to form a distinctive “sawtooth” pattern of gradual rises in price followed by sudden drops. Figure 1(a) shows an example of such behavior. The figure presents the top bids, in dollars, for a specific keyword (*phrase_id* = 24 in our sample), every 15 minutes from 12:15 AM to 2:15 PM on July 18, 2002. As indicated at point A, bidding starts with both advertisers below their maximum bids. The first advertiser’s autobidder recognizes an opportunity to obtain the first listing position by raising the second bidder’s bid by \$0.01, and the first advertiser’s autobidder does so. But then the second advertiser’s autobidder sees that it has been outbid, and raises its bid in turn. This process continues until the bids reach one advertiser’s maximum bid (say, the first advertiser’s maximum), as shown at B. The first advertiser’s autobidder can no longer increase its bid to obtain the first place, so it instead looks to avoid overspending, which it does by lowering its bid to \$0.01 more than the third-place bidder, yielding the bid at C. But then the second advertiser sees that it can still obtain first place by bidding \$0.01 more than the first advertiser’s newly-lowered bid. Bidding therefore begins to increase, yielding more and more “teeth” in the sawtooth pattern. Figure 1(b) shows the continuation of this pattern, for one week. Of course, there are occasional deviations from the simple pattern, as well as time differences between some subsequent peaks, but the general pattern is clearly present. In fact, this pattern is ubiquitous in the 2002-2003 Overture data, and it is present for most keywords and most time periods.

Clearly, this sawtooth pattern reduces market efficiency: the bidder who values the first spot the most spends only half the time at the top, and even less if there are more than two bidders competing for the top spot. Moreover, our analysis also indicates that this bidding pattern could have substantially reduced Overture’s revenue. In the next section, we estimate a lower bound on the size of this reduction, relative to outcomes under the Vickrey-Clarke-Groves mechanism, in which it is a dominant strategy for advertisers to bid truthfully.

³Of course, other types of strategic behavior are possible, but the cycling behavior is particularly natural in continuously repeated first-price auctions. As described below, we actually observe it in this context. The same type of behavior in sponsored search auctions is also reported in [1] and [8]. Cycling has also been reported in other contexts, e.g., retail gasoline markets [10].

3 Revenue Losses

3.1 Theoretical Example

To illustrate the potential loss in revenues due to the cycling behavior in the repeated first-price auction, consider the following simple example. There are two advertisers, with valuations V_a and V_b for a click ($V_a > V_b$) and one advertising slot. The minimum payment is ϵ ; assume $\epsilon < V_b$. Under Overture’s first-price bidding structure, prices increase in stair-steps from ϵ to V_b . Assuming that prices, on average, spend equal time at each step, the expected revenue conditional on a click is equal to $\frac{V_b + \epsilon}{2}$.

Now consider a second-price auction. If each advertiser bids its valuation, the bids will be V_a and V_b . Advertiser a will receive the slot. His payment per click will be V_b (the second advertiser’s bid). The expected search engine revenue conditional on a click will also be V_b , which is greater than $\frac{V_b + \epsilon}{2}$.

Of course, there is usually more than one advertising slot auctioned off by the search engines, there are more than two advertisers, their behavior in the first- and second-price auction may not correspond to the assumptions above, etc. It is therefore an empirical question whether (and by how much) first-price auctions underperform alternative mechanisms, and we now turn to answering this question.

3.2 Empirical Analysis: Methods and Implementation

To evaluate revenue losses, we need a counterfactual—what revenues would have been in a market without cycling, which requires devising an auction mechanism to stop cycling. A simple second-price auction, like the one used by Google and by more recent Yahoo! auctions, is not necessarily incentive compatible: a bidder may prefer to submit a lower valuation if he can get a reasonably good spot at a very low price, rather than submit his true valuation to get the top spot, but in the process pay a high price [6, 11].

Fortunately, the standard Vickrey-Clarke-Groves mechanism [12, 5, 9] works in this setting, at least in theory. However, it is more subtle than in the standard single-object case (where it is equivalent to the second-price auction), and it should be implemented carefully. The basic idea of the VCG mechanism is to maximize social welfare (given reported valuations), and to charge each player the negative externality he imposes on others. In the sponsored search setting, assuming that advertisers’ valuations of clicks do not vary by position (which is a close approximation, at least on average [4]), a search engine should ask for each bidder’s valuation of a click. It should then, based on each advertiser’s historical click-through rates in each position, compute the allocation of ads to maximize expected welfare; then, for each bidder, compute the maximum welfare without that bidder, and charge the bidder the difference between the two numbers, divided by the expected number of clicks (see [6] for the explicit derivation of VCG payments in this environment). By the standard VCG argument, this auction is incentive compatible, and it is also efficient. Note that it is different from both Yahoo!’s mechanism, which does not even take CTRs into account, and

Google’s mechanism, which does take CTRs into account, but charges fees that are not equal to VCG payments.

We received from Yahoo! Overture a dataset reporting times and amounts of all bids in 1000 top Overture markets, along with unique identifiers for bidders and keywords, between June 15, 2002, and June 14, 2003.⁴ Using this dataset, we compare Overture’s actual revenue from its first-price auctions during this period with the counterfactual revenue from auctions under a VCG auction regime.

In order to form counterfactuals of VCG revenue, we need a way to deduce advertisers’ valuations from their observed bids in Overture’s first-price auction. We take as our starting point the assumption that advertisers never bid above their valuation. Therefore, when we observe an advertiser’s maximum first-price bid, that is a lower bound for its valuation and for its bids under a second-price regime.

In backing out advertiser valuations from observed bids, we face a tradeoff. On one hand, we want to take valuations over a long period in order to make sure we catch the “peak” of a bidder’s sawtooth cycle. On the other hand, advertisers’ valuations may change over time, and so we hesitate to take maxima over unduly long periods. As a compromise, we take maxima over one-week periods. In particular, we take each bidder’s valuation for a given week to be its maximum bid during that week, in the actual first-price auctions we observe.

We assume that advertisers will truthfully bid their valuations in a VCG auction.⁵ Since we do not have data on individual clicks, we also assume that click-through rates are the same for all bidders on a given keyword, and that they decline, from the top spot to the bottom of a page, at the same rate for all keywords. We get data on the rate of decline, by position, from [3]. An ad ranked second is 77.7% as likely to be clicked as an ad ranked first for the same keyword. An ad ranked tenth is 7.8% as likely to be clicked. We assume that ads ranked below tenth do not get clicked.

Once we make these assumptions, for each keyword, we estimate the ratio of the revenue to Overture under the counterfactual VCG auction⁶ to the revenue under the actual mechanism. Note that while the revenues themselves are not identified (since we do not know the numbers of impressions and the actual click-through rates), the ratios of counterfactual revenues to actual revenues *are* identified, because the numbers of impressions and CTRs cancel out.

⁴All unique identifiers are arbitrary, giving us no individually-identifiable information about specific market participants and keywords.

⁵In a single-shot game, this is a unique dominant-strategy Nash equilibrium. Of course, in the actual sponsored search auctions, the game is, essentially, infinitely repeated, and thus there may be many other dynamic equilibria. However, truthful bidding is a particularly straightforward equilibrium, and thus we believe it provides the most accurate prediction of what would happen if search engines switched to the VCG mechanism.

⁶When computing Overture revenue, we bear in mind the canonical advertisement pricing policy: An advertiser is charged \$0.01 more than the bid of the next-highest advertiser, or \$0.05 if the advertiser is the lowest bidder for a given keyword.

3.3 Empirical Analysis: Results

We use the methods in the preceding section to back out bidder valuations, yielding VCG bids and VCG revenues for each keyword in our dataset.

Table 1 presents the estimated ratios of counterfactual VCG revenues to the actual revenues, using Overture’s 2002-2003 bid data and valuation estimates as described above. For most keywords, VCG yields lower revenue than Overture’s actual 2002-2003 auction rules, with mean and median revenue ratios (VCG revenue compared to actual revenue) of 0.66 and 0.68, respectively. However, restricting analysis to “popular” keywords (where the average number of bidders was greater than or equal to 10), the average and median revenue ratios were 1.07 and 1.06. We think these popular keywords are the most relevant comparison group, because the number of bidders on high-value keywords increases over time. By now many of these keywords have more bidders than even the densest markets in the 2002-2003 data.

To aggregate the ratios for individual keyword markets up to a total change in revenue, we need to know the contribution of each keyword market to Overture’s total revenues. Since we do not have this data, we use two different assumptions: all keyword markets contribute the same revenue to the search engine (“average ratio by keyword”), and all keywords have the same numbers of search requests and the same average click-through rates (“average ratio by click”). These numbers are relatively close (1.07 and 1.09, among popular keywords). We therefore conclude that if Overture had switched from its first-price auction regime to the VCG mechanism, during the period covered in our data, its revenues for popular keywords would have increased by more than 7%, assuming truthful bidding.

Finally, note that all the estimates in this subsection are lower bounds on the true values, since our estimates of bidder valuations were lower bounds.

4 Strategic Behavior: Current Evidence from Ranking Data

In the previous section, we presented evidence of strategic bidder behavior and corresponding revenue losses in the data from one search engine, covering the period from June 2002 to June 2003. That was a relatively early period for sponsored search auctions, and the search engines have since moved to mechanisms that are based on the second-price auction (which, despite not being incentive-compatible, are perhaps less susceptible to gaming than mechanisms based on the first-price auction). So one might wonder whether our results are peculiar to that early period—whether markets have subsequently stabilized and strategic behavior has disappeared. In this section, we show that this is only partly true: on both large search engines, Google and Yahoo!, strategic behavior still seems to be widespread.

Since we do not have current bidding data for either of the two largest search engines, we rely on publicly available information—the outcomes of the auctions, i.e., the rankings of advertisers displayed by the engines for various keywords. We have collected these rankings for ten keywords and phrases (“air travel”, “antivirus”, “arthritis”, “business card”, “car rent”, “contact lenses”,

“credit”, “mesothelioma”, “singles”, “travel agent”), once approximately every 7 minutes, from March 23, 2005, until April 4, 2005, from the main search pages at both Google and Yahoo!.

To assess the stability of each keyword market, we look at how often bidders’ rankings change. For each keyword, each search engine, and each search engine, we identified the highest-ranked website. We then looked at “streaks”—periods in the data when the highest-ranked website remained unchanged.⁷ If advertisers’ bids equal their true valuations of clicks, then rankings should change rarely, i.e., only when economic fundamentals change, and so streaks should be relatively long. However, if bidders track each other’s behavior and adjust their own bids accordingly, then we should see frequent changes in rankings, corresponding to frequent changes in bids; streaks would therefore tend to be short.

Table 2 presents the medians of streak lengths, in minutes, for the ten keywords and two search engines that we looked at. Several characteristics should be noted. First, over all keywords, the median length of a streak is 14 minutes at each of the search engines. This is a very short time interval, and it is hard to believe that economic fundamentals change so often, even if they depend on the time of the day and on the day of the week. Moreover, since we capture rankings only at discrete time intervals, we miss any changes that occurred between our requests, and so the true streak lengths are likely to be even shorter.

Second, streak lengths vary dramatically by keyword and search engine, with no clear pattern. Many keywords on both search engines (but especially on Google) have very short median streak lengths: 14 and 21 minutes. On the other hand, some keywords have very long streaks, and one had the same top bidder on Yahoo! throughout our sample,⁸ even though the highest-ranked bidder for the same keyword on Google changed very often. For “contact lenses,” Google and Yahoo! traded roles: streaks on Yahoo! were very short, while on Google they were comparatively very long—the median streak was longer than one day.

It is also instructive to look at particular examples of keywords with low streak lengths. Tables 3 and 4 present typical examples of advertisers competing for the top spot and constantly replacing each other there. Table 3 shows a sequence of top ads for keyword “mesothelioma” on Yahoo! and Table 4 presents such sequence for keyword “arthritis” on Google. Each line in the tables gives the URL associated with the top ad, the time stamp of the beginning of the streak with this ad at the top, and the length of the streak.

Unfortunately, we do not currently know the bids corresponding to these ads, and so we cannot be certain that these observations correspond to the bidding cycles that we described in Section 2 (rather than resulting from some other, technical issues). However, our observations suggest that it is likely that bidders are engaged in strategic behavior. Note that in Table 3, both bidders respond quickly—here too, probably even more quickly than our discrete data indicates. In contrast, in Table 4, only one bidder responds very quickly (perhaps almost instantly), while the other bidder

⁷Of course, since our requests had 7-minute gaps, it is possible that there were changes between the requests that we do not capture. However, this would imply that our estimates are biased upward, and would therefore only strengthen our conclusions.

⁸The keyword was “car rent” and the bidder was Travelocity.

takes a long time to respond. It is likely that the first bidder uses a fast robot to update his bids, while the second bidder either uses a slow robot or updates the bids manually. In both cases, it is hard to justify the observations by the constantly changing economic fundamentals, while the hypothesis of strategic behavior similar to that observed in Section 2 seems highly plausible.

5 Conclusions

Our theoretical and empirical analyses suggest that strategic behavior is widespread and costly, and that a switch to a VCG-based mechanism might stabilize auction outcomes with neutral or even positive effects on revenues, at least relative to the old Overture mechanism that was based on the first-price auction. It is important to note that search engines are not the only parties who would benefit from improvements in mechanism design. For one, under unstable mechanisms, the bidders who value top spots the most spend only a fraction of time there. Second, behaving strategically requires substantial investments: time, effort, bidding tools and software, consultant fees, etc. Finally, if an advertiser's willingness to pay for a click is generally positively correlated with the relevance of the advertiser's website, then users also suffer from suboptimal content.

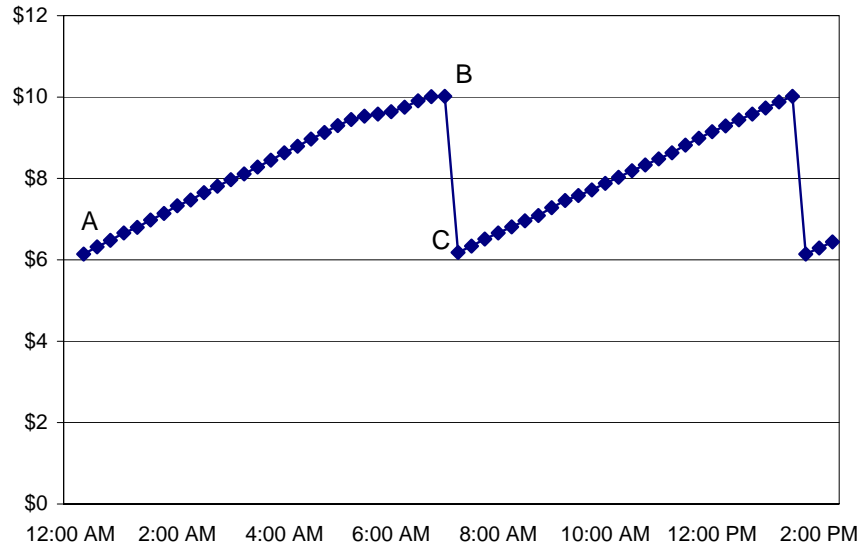
Beyond VCG, there may be alternative ways to reduce the problem of strategic behavior. For example, search engines might allow advertisers to update their bids only a few times a day. But such an approach is unlikely to eliminate strategic behavior completely. It might also have negative side effects—preventing advertisers from quickly responding to changes in economic fundamentals, or even facilitating collusion.

Of course, changing overnight from the current mechanisms to the efficient one would be too risky for the search engines. After all, advertisers have developed extensive institutional knowledge about the current design. However, given the magnitudes of revenue losses, as well as the costs to bidders and efficiency losses to search engine users, it would be worthwhile to experiment with alternative designs in small segments of sponsored search markets. Perhaps experimentation could begin with a number of keywords related to one specific geographical area, or a field with particularly savvy bidders. Useful insights can also be gained by analyzing richer data sets than what we have considered here, including position-specific click-through rates, conversion ratios, and changing bidding patterns over time.


Progress will most likely be achieved through a combination of auction-theoretical analysis, empirical work, and experimentation, which, in addition to being important, makes sponsored search an exciting field for future research.

References

- [1] K. Asdemir. Bidding patterns in search engine auctions. University of Alberta School of Business Working Paper, 2005.
- [2] H. K. Bhargava and J. Feng. Paid placement strategies for internet search engines. In *WWW '02: Proceedings of the eleventh international conference on World Wide Web*, pages 117–123, New York, NY, USA, 2002. ACM Press.
- [3] N. Brooks. The Atlas rank report: How search engine rank impacts traffic. Technical report, Atlas Institute, July 2004.
- [4] N. Brooks. The Atlas rank report, part 2: How search engine rank impacts conversions. Technical report, Atlas Institute, October 2004.
- [5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- [6] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. NBER Working Paper 11765, 2005.
- [7] J. Feng, H. K. Bhargava, and D. Pennock. Comparison of allocation rules for paid placement advertising in search engines. In *ICEC '03: Proceedings of the 5th international conference on Electronic commerce*, pages 294–299, New York, NY, USA, 2003. ACM Press.
- [8] J. Feng and X. Zhang. Price cycles in online advertising auctions. MIT Working Paper, 2006.
- [9] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
- [10] M. Noel. Edgeworth price cycles: Evidence from the Toronto retail gasoline market. UC San Diego Working Paper, January 2004.
- [11] H. R. Varian. Position auctions. UC Berkeley Working Paper, 2006.
- [12] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, March 1961.
- [13] T. A. Weber and Z. Zheng. A model of search intermediaries and paid referrals. Stanford University Working Paper, 2006.



(a) 14 hours



(b) 1 week

Figure 1: “Sawtooth” bidding pattern

Table 1: Distribution of VCG Counterfactual-to-Actual Revenue Ratios by Keyword

Statistic	Value (all keywords)	Value (“popular” keywords)
10th percentile	0.36	0.95
25th percentile	0.52	1.02
Median	0.68	1.06
75th percentile	0.81	1.12
90th percentile	0.92	1.13
Average ratio (by keyword)	0.66	1.07
Average ratio (by click)	0.76	1.09

Table 2: Median Streak Lengths (minutes)

Keyword	Google	Yahoo!
“air travel”	14	999
“antivirus”	14	14
“arthritis”	14	186
“business card”	21	5200
“car rent”	14	17049
“contact lenses”	1979	21
“credit”	14	91
“mesothelioma”	14	14
“singles”	98	314
“travel agent”	1286	364
All Keywords	14	14

Table 3: Keyword “Mesothelioma” on Yahoo!

Top Ad	Beginning	Length
www.mesotheliomaoptions.com	3/25 15:33	7
www.mesotheliomaweb.org	3/25 15:40	14
www.mesotheliomaoptions.com	3/25 15:54	21
www.mesotheliomaweb.org	3/25 16:15	21
www.mesotheliomaoptions.com	3/25 16:36	7
www.mesotheliomaweb.org	3/25 16:43	7
www.mesotheliomaoptions.com	3/25 16:50	7
www.mesotheliomaweb.org	3/25 16:57	7
www.mesotheliomaoptions.com	3/25 17:04	21
www.mesotheliomaweb.org	3/25 17:25	7
www.mesotheliomaoptions.com	3/25 17:32	21
www.mesotheliomaweb.org	3/25 17:53	7
www.mesotheliomaoptions.com	3/25 18:00	7
www.mesotheliomaweb.org	3/25 18:07	14
www.mesotheliomaoptions.com	3/25 18:21	14

Table 4: Keyword “Arthritis” on Google

Top Ad	Beginning	Length
www.StudyForKneePain.com	4/2 9:15	845
www.kaiserpermanente.org	4/2 23:20	7
www.StudyForKneePain.com	4/2 23:27	322
www.kaiserpermanente.org	4/3 4:49	7
www.StudyForKneePain.com	4/3 4:56	238
www.kaiserpermanente.org	4/3 8:54	7
www.StudyForKneePain.com	4/3 9:01	42
www.kaiserpermanente.org	4/3 9:43	7
www.StudyForKneePain.com	4/3 9:50	238
www.kaiserpermanente.org	4/3 13:48	7
www.StudyForKneePain.com	4/3 13:55	14
www.kaiserpermanente.org	4/3 14:09	7
www.StudyForKneePain.com	4/3 14:16	964
www.kaiserpermanente.org	4/4 6:20	7
www.StudyForKneePain.com	4/4 6:27	28