Empirical Price Modeling for Sponsored Search

Kuzman Ganchev
University of Pennsylvania
Philadelphia PA, USA
kuzman@seas.upenn.edu

Alex Kulesza
University of Pennsylvania
Philadelphia PA, USA
kulesza@seas.upenn.edu

Jinsong Tan
University of Pennsylvania
Philadelphia PA, USA
jinsong@seas.upenn.edu

Ryan Gabbard
University of Pennsylvania
Philadelphia PA, USA
gabbard@seas.upenn.edu

Qian Liu
University of Pennsylvania
Philadelphia PA, USA
qianliu@seas.upenn.edu

Michael Kearns
University of Pennsylvania
Philadelphia PA, USA
mkearns@cis.upenn.edu

ABSTRACT

We present a characterization of empirical price data from sponsored search auctions. We show that simple models drawing bid values independently from a fixed distribution can be tuned to match empirical data on average, but still fail to account for deviations observed in individual auctions. Hypothesizing that these deviations are due to strategic bidding, we define measures of "jamming" behavior and show that actual auctions exhibit significantly more jamming than predicted by such models. Correspondingly, removing the jamming bids from observed auction data yields a much closer fit. We demonstrate that this characterization is a revealing tool for analysis, using model parameter values and measures of jamming to summarize the effects of query modifers on a set of keyword auctions.

1. INTRODUCTION

Much of the academic literature on sponsored search to date has been theoretical in nature [7, 4, 5, 6], characterizing behavior or payoffs under strong assumptions that may fail in practice—especially when the markets in question are young, bidders are inexperienced, and relevant pieces of information (such as the exact rules of the game) are frequently kept secret. To our knowledge there has been no exploratory study of actual bidding data on a large scale to determine how real-world auctions can best be analyzed and understood. This paper provides a simple but needed first look at such questions. We utilize sponsored search data drawn from a wide array of Overture/Yahoo! auctions and examine how bids are distributed, what kinds of models of advertiser value can reasonably be proposed, and the evidence for strategic behavior.

Our analysis serves two immediate purposes. First, a better understanding of empricial bidding behavior improves the quality of data that can be synthetically generated for further study. We show that simple models used in practice fail to account for significant strategic effects, and suggest improvements that meaningfully enhance the "realism" of such models. Second, our characterization of sponsored search auctions includes measurable quantities and model parameters that can be used to summarize important features of an auction for further analysis. To demonstrate the insight provided by such summaries, we show how groups of query modifiers can influence bidding on a wide array of keyword auctions. We find, for example, that adding modifiers like "cheap" or "deal" to automobile brand names tends to increase the amount of bid jamming at the first slot of the corresponing sponsored search auctions.

The paper is structured as follows. Section 2 describes the methodology used to collect actual bids from the Overture bidview tool, and Section 3 presents a high-level, aggregate analysis using a subset of the resulting data. Section 4 demonstrates that the aggregate analysis fails to completely describe the character of individual auctions, and Section 5 proposes a measure of strategic behavior to explain the observed deviations. Section 6 examines an application of our analysis.

2. METHODOLOGY

Our data was obtained from the Overture bidview tool from approximately November 28 to December 2, 2006[1].

We collected bid data for two sets of queries. The first, smaller set includes the keywords used by Rusmevichientong et al [6] and comprises 859 queries related to travel. It is used for aggregate analysis in Section 3. The second set, used in the latter sections of the paper, comprises a wide array of 36,900 queries. For the purposes of further analysis (such as that carried out in Section 6), the second set is structured as a cross product of 450 base keywords—e.g., "lawyer"—intended to reflect basic searches that would generate advertiser interest, and 81 modifiers—e.g., "Philadelphia"—intended to capture the ways in which users might further specify searches. The base keywords and modifiers are further structured by placement in groups; there are nine groups of base keywords and six groups of modifiers. Table 1 and Table 2 give some summaries and examples for the groups of keywords in this data set. A complete query in the second set pairs one base keyword with zero or one modifiers (e.g., "lawyer" and "Philadelphia lawyer").

Although the complete cross-product of base keywords and modifiers results in a huge number of keyword phrases, a considerable portion of the phrases are not meaningful and thus have empty auctions or auctions with very few bids. We primarily focus on the auctions of significant size in this work.

[1] `http://www.overture.com/`; the bidview tool was discontinued shortly after we collected our data.

WWW2007, May 8–12, 2007, Banff, Canada.

group	#	examples
cars	41	BMW, Toyota
drugs	62	Xenical, Prozac
electronics	36	laptop, cell phone, camcorder
local-service	55	carpet cleaning, hair dresser
medical	50	anxiety, plastic surgery
non-local-service	27	car insurance, mortgage
software	67	Microsoft Windows, MySQL
subscription	91	cable, magazine
travel	21	cruise, hotel, vacation
total	450	

Table 1: Summary of the base keyword groups

group	#	examples
action	6	buy, purchase
info	11	information, review
location	40	New York, Ohio, Philadelphia
post	6	support, parts, repair
price	10	cheap, expensive, free, discount
quality	8	best, luxury, new, used
total	81	

Table 2: Summary of the modifier keyword groups

Due to resource constraints, we ran each query only once; our data provides no information on dynamic bidding behavior. The data returned by the bidview tool include up to 40 bids, ranked in order from highest to lowest. Advertiser names and ad text are provided, but not used for our analysis. Furthermore, we throw out the first bid in every auction. This is due to the method by which prices are determined: an advertiser pays a price equal to the bid of the next advertiser in bid order, so that the first bid is not relevant to money changing hands except insofar as it is higher than all other bids. For the remainder of the paper we use the term "price" assuming the convention that the k^{th} price is equal to the $(k+1)^{th}$ bid.

3. INITIAL LOOK AT THE DATA

Visualizing the bid books, it is apparent that the data are generally quite noisy. Consequently, we begin by examining the data in aggregate. A potential problem with this approach is the risk of observing artifacts of the mixing process; for example, if we have one market where only a small number of advertisers compete selling a very lucrative product, we might conclude that small auctions are likely to have a high prices. By contrast, we observe the opposite trend when we look at just one market (Figure 1). To limit such artifacts, we use only data from our smaller, travel domain query set in this section.

Previous research has looked at individual auctions over time [4] and at median prices [6], but never with the aim of modeling all the prices for a set of auctions. Because of this, we start by visualizing the data.

Figure 1 shows how the price paid by the top bidder is correlated with the total number of bidders in the auction. As can be seen in the figure, there is an almost linear relationship between the number of bidders and the mean price of the first position. An increasing trend is not surprising,

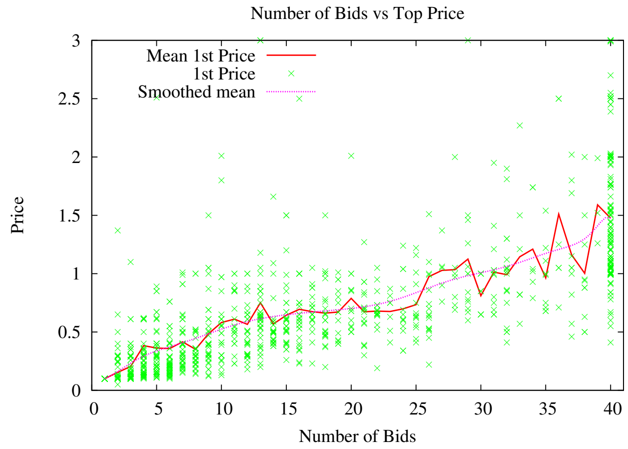

Figure 1: The correlation between first prices and the number of bids. The smoothed bids are a Bezier curve generated from the means.

but perhaps the (noisy) linearity in the mean is somewhat unexpected.

Figure 2 shows the correspondence between bid position and mean prices. Because prices increase as the number of bidders increases, all the auctions used in computing Figure 2 have at least 19 bids (i.e., 18 prices). We normalize the prices so that the first price is 1.0, and we have an equal contribution from all auctions. We find that an exponential decay fits the means surprisingly well (pink curve), and as a result an exponential decay model seems a natural choice for fitting individual auctions. We will see, however, that individual auctions show significantly different behavior than the aggregate.

Figure 3 shows the correspondence between bid position and median prices. Interestingly, the median prices are not as well explained by an exponential decay, especially in the top few slots (which are arguably most interesting).

Figure 4 shows the distribution of the first few prices for auctions where there are at least 19 bidders. As we can see, in most cases the first few prices are relatively low, with a few auctions where the first price is high. The histogram looks approximately like a binomial distribution.

Figure 5 shows the differences between the first and second price, the second and third price and the third and fourth price for auctions with at least 19 bids. We will see in Section 5 that the peak in the small price difference ranges may be due to "jamming" strategies used by bidding agents.

4. PARAMETRIC MODELS OF A SINGLE AUCTION

Section 3 motivates a simple approach to modeling individual auctions under the assumption of independence—i.e., that each bidder bids independently of every other bidder. If bid averages follow an exponential decay, independent bidders must be drawing from the distribution that yields this curve. In particular, let continuous function $p(\frac{k-1}{N-1})$ be the expected price of a bidder who is at position k out of N. (For example, $p(\cdot)$ might be an exponential decay.) Then we can sample prices $p(u)$ where u is drawn from the uniform dis-

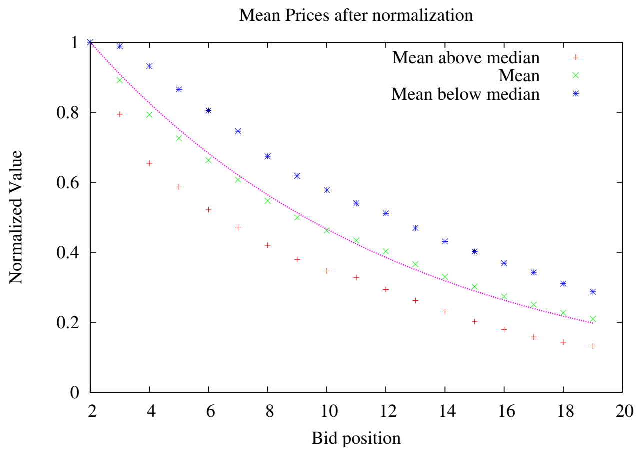

Figure 2: The correlation between prices and bid position, showing the mean price, the mean above the median, and the mean below the median. The curve is an exponential fit to the means.

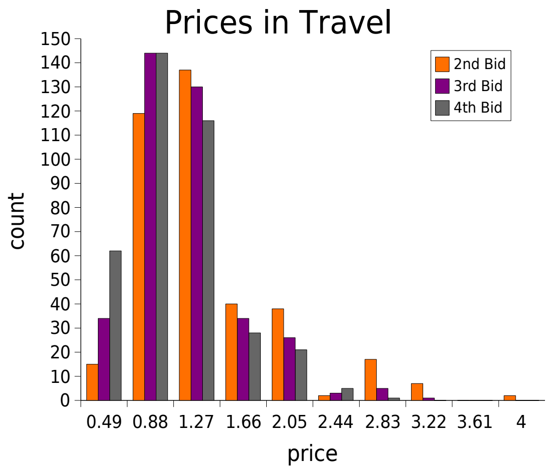

Figure 4: A profile of the first three prices. The horizontal axis is the value in dollars, the vertical is a count of the number of bidders with a price in the range up to that value.

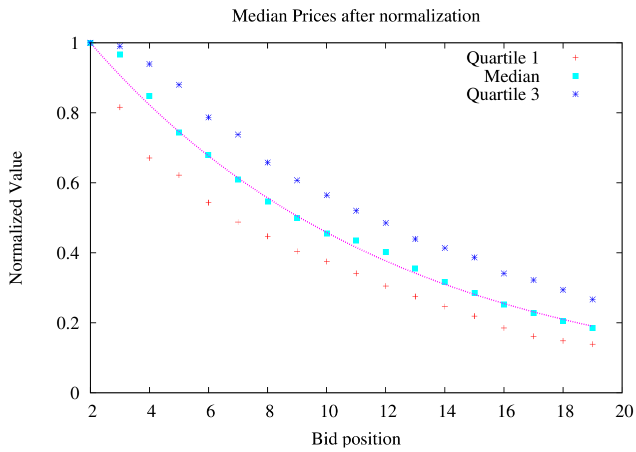

Figure 3: The correlation between prices and bid position, showing the median prices as well as the top and bottom quartiles. The curve is an exponential fit to the medians.

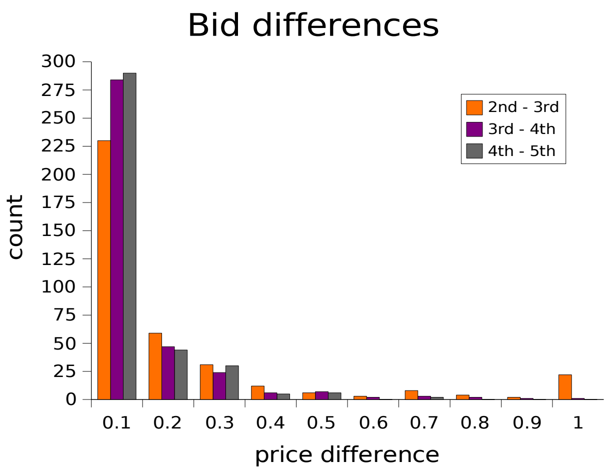

Figure 5: Bid differences corresponding to differences in price. The distribution looks smooth with the exception of the peak at small differences. See Section 5 for more on this peak.

tribution on $[0, 1]$. If the aggregate analysis is representative of individual auctions, data sampled in this manner should appear "similar" to real auction data. In fact, we show that this is not the case, motivating a more careful characterization of individual auctions in Section 5.

In order to measure the "similarity" of synthetic prices to real prices, we compare the abilities of various parametric models to fit the two types of data; i.e., we estimate $p(\cdot)$ from single auctions using both real data and synthetic data. If the synthetic generation procedure is accurate, then the data it produces should be fit equally as well as the empirical data. If there is some consistent difference in the quality of fits, then we can conclude that the synthetic data is somehow different.

We fit several different kinds of parametric models to the auctions. In all cases, the fits are performed to minimize the mean squared error (MSE) of the fit to the observed prices. MSE is computed as follows: for each price point, calculate the squared difference between the fit curve and the observed price. For each auction compute the mean of this value. Now compute the mean over all auctions of these per-auction averages. In this way each auction is given equal weight regardless of the number of bids it attracts.

To evaluate the quality of the fits, we examine not only MSE but also two other measures. The first, refered to as normalized MSE, involves normalizing the auctions so that the highest price point is always 1.0 and then computing MSE on the normalized prices. Due to the squaring operation, MSE of any kind can be difficult to interpret, so our third measure is mean absolute error (MAE). MAE is computed in the same manner as MSE but taking absolute values rather than sqaring the differences. The mean absolute error can be interpreted as the expected value of the difference (in cents) between the fit and the observed price if we choose an auction uniformly at random.

4.1 Exponential Models

Motivated by the exponential decay seen in Figure 3, our first model is an exponential curve. This model has two parameters a and b and is given by

$$price[i] = ab^i, \tag{1}$$

where i is an index of the price positions. An important possibility not considered by this model is that prices may converge to something other than 0. For example, we know that there is a reserve price in the auctions and that bids below the reserve price are not possible. This suggests a three parameter model where prices are given by

$$price[i] = ab^i + c. \tag{2}$$

Note that the second exponential model is more powerful than the first and always has a (weakly) better fit.

4.2 Linear Models

For the sake of a more robust similarity metric, we include a second model family that may capture different characteristics of the data: piecewise linear models. An n piece piecewise linear model has the form

$$price[i] = \begin{cases} a_1 i + b_1 & \text{if } i < l_1 \\ a_2 i + b_2 & \text{if } l_1 \le i < l_2 \\ \vdots & \\ a_n i + b_n & \text{if } l_{n-1} \le i < l_n \end{cases} . \tag{3}$$

Model	Real	Gaussian	Uniform	Exponential
ab^i	3.8	1.9	3.7	1.7
$ab^i + c$	2.9	1.2	1.3	1.0
linear	14.2	3.2	1.5	7.8
2-linear	1.2	0.6	0.5	0.6

Table 3: The MSE values of different model fits on real and synthetic data. All error values have been multiplied by 100 for ease of reading.

Model	Real	Gaussian	Uniform	Exponential
ab^i	9.5	5.7	8.0	7.4
$ab^i + c$	5.9	4.3	3.2	3.0
linear	36.8	15.2	4.2	34.1
2-linear	2.1	1.8	0.9	1.6

Table 4: The normalized MSE values of different model fits on real and synthetic data. All error values have been multiplied by 1000 for ease of reading.

Fitting such a model requires figuring out where to place the breaks l_i in 3 as well as fitting the line to the prices between breaks. This can be done using dynamic programming in a straightforward manner. We used up to 5 pieces for all auctions of size between 10 and 40. However, MSE consistently decreases by a factor of 10 when moving from 1-piece to 2-piece model, and again by a factor of 3 when moving to a 3-piece model. Since the MSE virtually vanishes at this point, we only give results for 1 and 2-piece models.

4.3 Results

For each auction in our second query set we generated parallel synthetic price data using the sampling technique described above, where $p(\cdot)$ was chosen to be the best exponential fit for the auction in question (Eq. (2)). Note that the synthetic data is *sampled* from the induced distribution and not generated directly from $p(\cdot)$, thus it is in no way guaranteed a better fit than the real auction data. We also generated two other types of synthetic prices corresponding to methods sometimes used in practice. Gaussian prices were generated by sampling from a Gaussian using the mean and variance of the empirical prices, and uniform prices were generated by sampling from a uniform distribution over the range of empirical prices.

Each of the four data sets (real, exponential, Gaussian, and uniform) was fit by each of the four models (exponential, exponential with offset, linear, and two-piece linear). The results are presented in Table 3, Table 4, and Table 5.

It is apparent that none of the synthetic methods for

Model	Real	Gaussian	Uniform	Exponential
ab^i	0.10	0.08	0.10	0.07
$ab^i + c$	0.06	0.06	0.06	0.05
linear	0.18	0.09	0.07	0.14
2-linear	0.05	0.04	0.04	0.04

Table 5: The MAE values of different model fits on real and synthetic data.

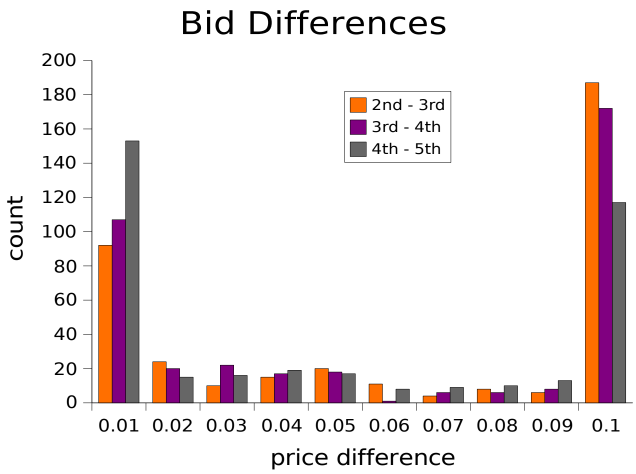

Figure 6: Bid differences in \$0.01 increments. We see that the peak in Figure 5 is due mostly to differences of \$0.01 or less. The bars marked "0.1" include all bid differences of \$0.10 or greater.

generating data displays fit error rates similar to those of real auction data. We conclude that these models are not fully capturing the relevant qualities of empirical auction data. This point is addressed further in the following section. However, it seems clear that the fits for exponential synthetic data vary across models in the way most similar to those of true auction data, though in magnitude the error measures are much smaller. In contrast, the linear fits for uniform and Gaussian data are much too accurate, and the exponential with offset fit for Gaussian data appears to give a reduced improvement over the plain exponential as compared with empirical data. We conclude, therefore, that the exponential model is the most accurate (this is also supported by the aggregate analysis), but that simple, independent-bidder models miss some key characteristic of sponsored search auctions.

5. JAMMING

Since our exponential model utilizes the only distribution that, sampled independently, gives rise to the exponential inverse cumulative density curve seen in our aggregate analysis (in particular, adding independent noise to each sample does not maintain the curve), there must be significant dependencies between bids in sponsored search auctions. While there are many possible reasons for such dependencies, we show here that one cause may be strategic behvaior on the part of advertisers.

Figure 6 is a zoom of the zero to ten cents region of Figure 5. It is clear that the unexpectedly large number of bid differences below ten cents is due to a peak at differences of one cent or less. We propose that the sharp peak around one cent in the price differences may be due, in part, to the use of a bidding tactic known as "jamming."

Jamming is defined as increasing one's bid, without increasing the price paid, in order to increase the price of the next highest bidder. This may be an attempt to deplete a competitor's budget or simply to increase their advertising costs in general. For example, if an auction contains three bidders A, B, and C, and they initially bid 1.00, .75, and .60, respectively, then B may jam A by changing its bid to 0.99. B will still pay .61, and so incurs no cost, but A's price will jump from .76 to 1.00. This tactic appears to be widely used and is indeed done automatically by many bidding packages [3, 2].

Of course, jamming is not the only possible explanation for the unusual number of one cent bid gaps. It is possible that bidders choose to play a strategy in which bids are set to one cent *above* the next lowest bidder in an effort to avoid being jammed, or to be "kind" to the the bidder above. It is also possible that other bidding strategies create the observed effects indirectly, or that collusion in somehow encouraging clumped bidding. Going forward we will continue to use the term "jamming," but we will define it as a purely statistical measure of bid closeness, with the understanding that further experiments, especially those using time series data for a dynamic analysis, will be necessary to draw conclusions about the true causes.

This inherent uncertainty also discourages us from attempting to directly evaluate generative bidding models in this section. Instead, we turn from modeling the bids directly to proposing some measures of the amount of jamming in an auction. We will then compare these measures on the real data as compared to uniformly, normally, and exponentially distributed synthetic price data.

First, we put forth a formal definition of jamming. There are a variety of options, but Figure 6 suggests one which is the simplest and most restrictive available: we will call a bid a jamming bid if it is one cent or less below the next highest bid. We will call a series of consecutive bids each of which (except the first) jams the previous one a jamming region. To reduce the influence of compression effects due to the reserve price, we do not consider bids within two cents of the reserve.

From this definition we propose various measures of jamming in an auction. In all of these, we seek to tease apart *true* jamming (jamming that results from dependencies between bidders) and apparent jamming due merely to the chance clustering of bids, which may be significant in a market with dozens of bids all less than a dollar.

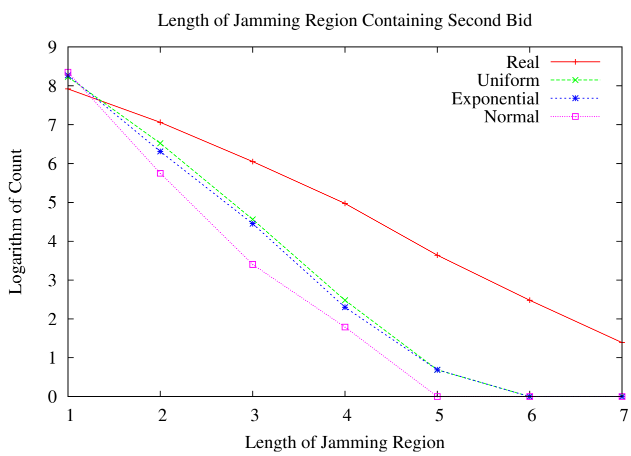

Figure 7: Counts of the length of the jamming region containing the second bid, on a logarithmic scale.

The first measure we propose is the size of the jamming region starting at the second bid. Intuitively, this corresponds to he number of people jamming the first price point. The results for this measure are shown in Figure 7. Note that since the counts fall off very sharply, the graph has been logarithmically scaled. Here we see that for the real data, the length of the first jamming region is very close to being exponentially distributed and that it falls off far more slowly than for the simulated data.[2] This slower drop–off is not unexpected: there is considerable motivation for multiple bidders to attempt to compete with each other for the top spot, whereas there is of course no such extra pressure for randomly distributed data.

The next measure is how frequently each price position is jammed, shown in Figure 8. There is a general upward trend for all distributions, probably due to the increasing compression of the range of possible bid values at lower positions creating more chance clusters. However, across a range of price points that there is a strong and statistically significant jamming effect in the real data over and above that seen in the independent-bidder models.

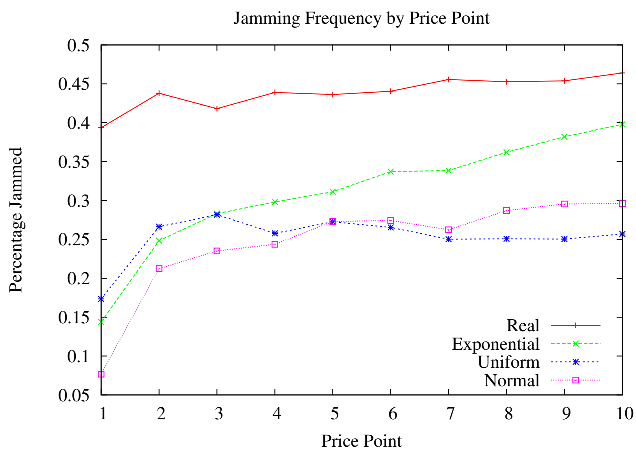

Figure 8: Frequency of jamming by price position.

The last measure we define is the jamming compression measure. This is defined to be the ratio of the number of jamming bids to the total number of bids and represents how much an auction would shrink if all the jamming bids were removed. In Figure 9 this measure is plotted against auction size. We can see that, as expected, it increases gradually for all distributions (though not much for the uniform), probably due again to more chance clusters occurring as more bids are crammed into a small space of possible bids. Although this data is noisy due to the relatively small number of auctions for some auction sizes, the jamming ratio is significantly higher over all auction sizes less than 27 except size 20.

Finally, Figure 10 shows how the jamming compression measure changes as we vary the jamming threshold (the maximum bid difference considered to be an instance of jamming). Of course, there is a general upward trend across all the distributions as the threshold is increased. Note that there is a sharp jump in jamming ratio for the real data when the threshold goes from zero cents to one cent, while the other curves are smooth. This is in line with our expectations since there is little motive to jam a competitor at a lower price than necessary since this just decreases the price they pay. Also interesting is that as the threshold increases the curve is actually slightly flatter for the real data compared to the reference distributions.

[2]All differences between the real and simulated data here are significant at 95% confidence.

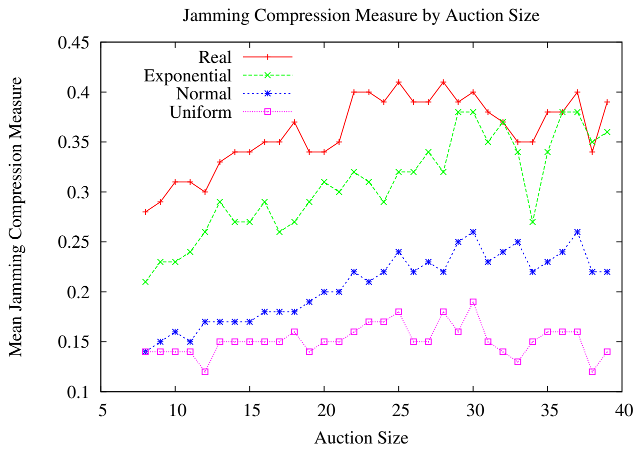

Figure 9: Jamming ratios by auction length.

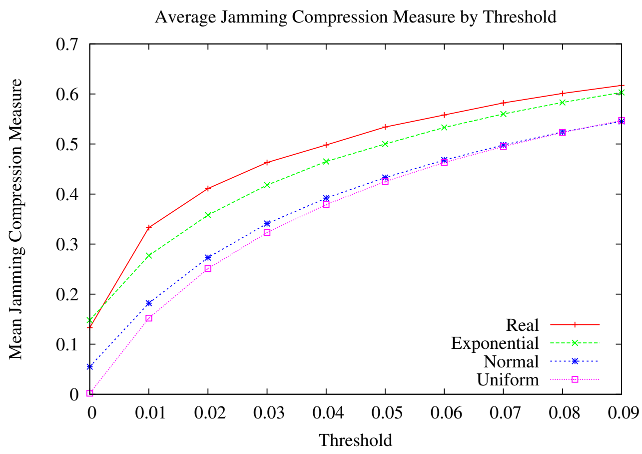

Figure 10: Jamming ratio averages by threshold.

5.1 Improvement of Fits by Unjamming

Having identified a characteristic of empirical sponsored search auctions that distinguishes them from simple, independent-bidder synthetic models, we recompute parametric model fits after removing jamming from the real auction data. To "unjam" an auction, we identify all jamming bids and remove them, as though the bidders never existed. The MSE results are presented in Table 6.

Though the gap between synthetic and real data still exists, it has been closed considerably. We hypothesize that further improvements might be obtained by considering other common strategies employed by advertisers. Using similar techniques, it should be possible to formalize and measure the prevalence of such strategies in real-world auction data.

Model	Real	Unjam	Exp
ab^i	3.8	2.8	1.7
$ab^i + c$	2.9	1.9	1.0
linear	14.2	13.1	7.8
2-linear	1.2	0.8	0.6

Table 6: The MSE values of different model fits on real, unjammed, and synthetic data. All error values have been multiplied by 100 for ease of reading.

6. MEASURING THE EFFECTS OF MODIFIERS

Since the queries in our second data set are structured as cross products of keyword/modifier groups, we can visualize some interesting trends using the characterization we have developed. In this section we examine the effects of adding modifiers to keyword auctions on the parameters of our model fits and our measures of jamming.

We first pre-process the data by removing auctions with fewer than ten price points. Including auctions with a small number of bidders greatly increases the variance of parameter values because it mixes the effects of the modifiers on the underlying valuations with effects that can be attributed to auction size. In general, as we increase the minimum number of price points we get more significant results and more dramatic differences.

We present a separate table for each quantity of interest. Each row in a table corresponds to a base group of keywords, each column corresponds to a modifier group or an empty ("null") modifier group, and each cell shows the mean of the parameter under the cross product of the corresponding groups. For example, the cell corresponding to "location" and "local services" contains the mean over queries consisting of keywords referring to local services combined with modifiers describing locations. If a cell is shaded, the number is statistically different from the "null" entry in the same row; red shading indicates an increase, and green shading indicates a decrease.[3] Significance levels were computed using a t-test[4] with confidence level 0.05. Abbreviations are defined in Table 7.

Figure 11 shows the effect of keyword modifiers on the first price parameter for the exponential model (the a in $price_i = ab^i + c$).

As we can see, the modifier group "location" significantly increases the highest price paid for the four base groups of "local service", "medical", "subscription," and "cars." This seems very natural, as all of these groups involve products that frequently depend on local providers, for example, "cosmetic surgery" and "hair implants" in "medical", and "gym membership" and "cable" in "subscription." Conversely, "non-local service" keywords have their prices reduced by the addition of a location, as expected. Perhaps surprisingly, the group of travel keywords did not see an increase in price (on average) with the addition of a location. This may be because the major travel bidders provide services to many destinations and are thus indifferent to whether a user is already interested in a specific one. Finally, we note that

[3]the red squares also have a border
[4]http://www.nmr.mgh.harvard.edu/Neural_Systems_Group/gary/python.html

Abbreviation	long name
act	action modifiers
info	info modifiers
loc	location modifiers
post	post purchase modifiers
price	price modifiers
quality	quality modifiers
cars	cars keywords
drugs	drugs keywords
electr	electronics keywords
local	local keywords
med	medical keywords
n-loc	non-local keywords
soft	software keywords
subscr	subscription keywords
travel	travel keywords

Table 7: Abbreviated row and column headings.

	null	act	info	loc	post	price	quality
cars	0.97	1.09	0.88	1.19	1.55	0.9	0.88
drugs	1.31	1.22	1.11	-	-	0.78	-
electr	1.04	0.85	0.39	1.9	0.7	0.78	0.8
local	1.45	0.94	0.99	3.05	1.36	1.18	1.32
med	1.93	-	2.03	3.66	2.19	0.65	-
n-loc	4.76	2.47	1.98	3.58	1.48	2.67	2.65
soft	0.86	0.48	0.42	-	-	-	0.48
subscr	0.87	0.97	0.59	1.47	1.05	0.9	0.83
travel	0.92	0.54	0.53	0.92	-	0.95	1.3

Figure 11: The means of the first parameter of the exponential model if we group keywords and modifiers. The figure shows auctions with at least 10 price points.

	null	act	info	loc	post	price	quality
cars	0.93	0.9	0.87	0.89	0.83	0.88	0.9
drugs	0.86	0.84	0.76	-	-	0.8	-
electr	0.91	0.87	0.88	0.8	0.86	0.88	0.86
local	0.9	0.83	0.82	0.8	0.83	0.86	0.82
med	0.84	-	0.75	0.73	0.81	0.84	-
n-loc	0.9	0.83	0.83	0.89	0.77	0.84	0.86
soft	0.87	0.84	0.88	-	-	-	0.87
subscr	0.88	0.85	0.85	0.8	0.81	0.86	0.86
travel	0.92	0.88	0.9	0.88	-	0.89	0.86

Figure 12: The means of the decay rate parameter of the exponential model if we group keywords and modifiers. The figure shows auctions with at least 10 price points.

	null	act	info	loc	post	price	quality
cars	32.4	21.1	17.3	21.1	21.1	18.8	21.2
drugs	17.8	14.5	14.0	-	-	12.3	-
electr	34.3	22.3	16.1	12.2	16.2	22.7	20.7
local	30.2	19.6	17.1	19.5	19.4	17.8	16.8
med	28.7	-	14.7	15.2	15.3	13.4	-
n-loc	35.5	21.3	22.9	32.1	17.7	25.3	26.5
soft	22.0	12.8	14.8	-	-	-	12.7
subscr	26.1	19.6	14.3	14.8	13.5	21.3	17.6
travel	34.3	15.1	22.0	24.5	-	28.0	23.8

Figure 13: The means of the number of bidders grouped by keyword and modifier groups. The figure shows auctions with at least 10 price points.

the conjunction of "car" and "post purchase modifiers" results in higher starting bids than car keywords alone. This is not surprising, since cars frequently require expensive post-purchase care (repairs, parts, etc.).

Figure 12 shows the decay rates for the exponential model (the b in $price_i = ab^i + c$). We see that the decay rates are always lowered (falloff speed increases) with the addition of modifiers. This is not surprising, given that the mean number of bidders decreases with the addition of modifiers as shown in Figure 13.

We can also visualize some other interesting trends. For example, we see in Figure 14 that while adding modifiers in general reduces jamming, adding a location or an action word like "buy" or "lease" actually increases the amount of jamming significantly for bids of at least \$0.40. Figure 15 shows that adding modifiers such as "cheap" to the name of an automobile manufacturer increases the expected length of the jamming region at bid 2.

7. CONCLUSIONS

We collected a large set of empirical sponsored search data and performed an exploratory analysis, attempting to characterize and understand real-world search auction data. We found an aggregate exponential decay of prices across many auctions, but showed that this model does not fully describe bidding behavior on a per-auction basis. We showed that jamming is more prevalent in real data than would be predicted by a model of independent bidders, and that removing jamming from empirical data (or, convesely, adding jamming to synthetic data) improves the similarity significantly. Future work will include studying effects other than jamming that contribute to this disparity. Finally, we demonstrated that our model parameters and measures of jamming provide useful summaries of important auction features, revealing trends in the ways modifiers influence the bids for search keywords.

	null	act	info	loc	post	price	quality
cars	0.19	0.25	0.2	0.26	0.08	0.23	0.2
drugs	0.16	0.16	0.09	-	-	0.08	-
electr	0.11	0.06	0.01	0.09	0.04	0.06	0.05
local	0.15	0.05	0.06	0.11	0.06	0.09	0.07
med	0.09	-	0.07	0.08	0.08	0.07	-
n-loc	0.27	0.12	0.14	0.23	0.04	0.17	0.19
soft	0.03	0.05	0.02	-	-	-	0.0
subscr	0.08	0.08	0.05	0.12	0.03	0.08	0.09
travel	0.11	0.02	0.05	0.12	-	0.11	0.13

Figure 14: Fraction of jammed bids over 40 cents grouped by keyword and modifier groups. The figure shows auctions with at least 10 price points.

	null	act	info	loc	post	price	quality
cars	1.98	2.49	2.34	2.47	1.85	2.65	2.11
drugs	2.23	2.1	1.94	-	-	2.04	-
electr	1.47	1.22	1.56	1.78	1.69	1.41	1.38
local	1.69	1.37	1.37	1.6	1.5	1.83	1.53
med	1.29	-	1.26	1.34	1.83	1.2	-
n-loc	1.52	1.39	1.57	1.61	1.0	1.35	1.43
soft	1.33	2.0	1.4	-	-	-	1.0
subscr	1.65	1.3	1.79	1.68	1.17	1.63	1.64
travel	1.68	1.8	2.06	1.57	-	1.5	1.47

Figure 15: The expected length of the jamming region starting t the first price grouped by keyword and modifier groups. The figure shows auctions with at least 10 price points.

8. REFERENCES

[1] A. Animesh, V. Ramachandran, and S. Viswanathan. Quality Uncertainty and the Performance of Online Sponsored Search Markets: An Empirical Investigation. *SSRN eLibrary*, 2006.

[2] ApexPacific. Overture bid management software & overture bid tool. http://www.apexpacific.com/ bidmaximizer/ overturebidding.html.

[3] Atlas. Rules-based bidding for pay-per-click management. http://www.atlasonepoint.com/ products/ bidmanager/ rulesbased.

[4] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Working Paper, http://rwj.berkeley.edu/schwarz/, 2005.

[5] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, New York, NY, USA, 2006. ACM Press.

[6] P. Rusmevichientong and D. P. Williamson. An adaptive algorithm for selecting profitable keywords for search-based advertising services. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 260–269, New York, NY, USA, 2006. ACM Press.

[7] H. R. Varian. Position auctions. *International Journal of Industrial Organization (to appear)*, 2006.