

Reserve Prices in Internet Advertising Auctions: A Field Experiment

Michael Ostrovsky
Stanford University
ostrovsky@gsb.stanford.edu

Michael Schwarz*
Yahoo! Labs
mschwarz@yahoo-inc.com

December 24, 2009

Abstract

We present the results of a large field experiment on setting reserve prices in auctions for online advertisements, guided by the theory of optimal auction design suitably adapted to the sponsored search setting. Consistent with the theory, following the introduction of new reserve prices revenues in these auctions have increased substantially.

*We are especially indebted to the engineers at Yahoo!, too numerous to list here, without whom this project could never have been implemented. We are also grateful to David Reiley and Preston McAfee for many helpful discussions and to Ken Hendricks, Jon Levin, and seminar participants at Chicago, Stanford, Yale, and the SIEPR-Microsoft Conference on Internet Economics for helpful comments and suggestions.

1 Introduction

Auctions are used to sell a wide variety of objects, ranging from flowers, paintings, and used cars to electromagnetic spectrum and Internet advertisements. One of the most natural questions about the design of an auction is revenue maximization: How should an auction be designed to generate the highest expected payoff to the seller? This question was answered by Myerson (1981) and Riley and Samuelson (1981) for the setting with one object for sale and independently distributed private bidder values. For the case with symmetric bidders, the answer is particularly elegant: the optimal mechanism can be implemented by a second-price auction with an appropriately chosen reserve price.

This theoretical work has been extended in many directions: e.g., Cremer and McLean (1988) and McAfee, McMillan, and Reny (1989) construct optimal auctions in settings with correlated and common bidder values; Maskin and Riley (1984) derive optimal mechanisms in settings with risk-averse bidders; and Maskin and Riley (1989), Armstrong (2000), and Avery and Hendershott (2000) study optimal design in settings with multiple objects.

Economists have also obtained empirical estimates of (and bounds on) optimal reserve prices for a variety of auctions (McAfee and Vincent, 1992; Paarsch, 1997; McAfee, Quan, and Vincent, 2002; Athey, Cramton, and Ingraham, 2002; Haile and Tamer, 2003; Tang, 2009). Notably, taken together, the results of these papers present a puzzle: all of them find that reserve prices actually observed in real-world auctions are substantially lower than the theoretically optimal ones. This raises the possibility that reserve prices are not a particularly important part of auction design, and sellers cannot use them to substantially raise revenues. Moreover, if participation in the auction is costly for bidders, increasing the reserve price may make the auction less attractive to them and fewer will bid, leading to lower revenues. Indeed, Bulow and Klemperer (1996) find that in symmetric single-object auctions, adding just one more bidder (and setting a zero reserve price) is always preferable to setting the optimal reserve price. So perhaps reserve prices are not important in practice?

In this paper, we address this question directly, by presenting the results of a large-scale field experiment on reserve prices in a particular setting: “sponsored search” auctions conducted by Yahoo! to sell advertisements. Reserve prices in the randomly selected “treatment” group were set based on the guidance provided by the theory of optimal auctions, while in the “control” group they were left at the old level of 10 cents per click. The revenues in the treatment group have increased substantially relative to the control group, showing that reserve prices in auctions can in fact play an important role and that theory provides a useful guide for setting them.¹ This increase is especially pronounced for keywords with relatively high search volumes, for keywords in which the theoretically optimal reserve price is relatively high, and for keywords with a relatively small number of bidders.

¹The reserve price of 10 cents per click was itself a result of experimentation by Yahoo! Search Marketing, which has over time raised the price from 1 cent in 1998 to 5 cents in 2001 and then to 10 cents in 2003. This experimentation, however, was not guided by theory.

Two prior studies have analyzed the results of field experiments on setting reserve prices in auctions. Reiley (2006) reports the results of a field experiment on reserve prices in a first-price online auction for trading cards for a popular collectible card game. His findings confirm several predictions of auction theory, such as the reduction in the probability of a sale when reserve prices are present and, more subtly, the increase in bids when they are present (which is a consequence of equilibrium behavior in first-price auctions). Brown and Morgan (2009) report the results of field experiments on auctions for collectible coins conducted on Yahoo! and eBay. The primary focus of the study is the competition between platforms and market tipping, but the authors also consider the effects of reserve prices. They find that positive reserve prices, set at the level of 70% of the purchase price of the coins from the dealer, lead to significantly higher revenues and lower numbers of bidders, relative to zero reserve prices.

Our paper makes several contributions relative to these studies. First, it analyzes a much larger and economically important setting, with thousands of keywords and millions (and potentially, given the size of the online advertising industry, billions) of dollars at stake. Consequently, many of the bidders in this setting spend considerable time and resources on optimizing their advertising campaigns. Second, the reserve prices in the experiment are guided by theory, based on the estimated distributions of bidder values. To the best of our knowledge, there are no other papers describing direct practical applications of the seminal results of Myerson (1981) and Riley and Samuelson (1981). Third, unlike the previous studies, the benchmark in our analysis is not a zero reserve price, but the existing reserve price set by the company after a long period of experimentation.

The paper is organized as follows. Section 2 provides an overview of the sponsored search setting. Section 3 extends theoretical results on optimal auction design (Myerson, 1981; Riley and Samuelson, 1981) to the current setting and discusses simulations of revenue impact of reserve prices. Section 4 describes the design of the experiment. Section 5 presents the experimental results. Section 6 concludes.

2 Sponsored Search Auctions

We start with a brief description of the sponsored search setting; for a detailed description, see Edelman, Ostrovsky, and Schwarz (2007, subsequently EOS). When an Internet user enters a search term (“query”) into a search engine, he gets back a page with results, containing both the links most relevant to the query and the sponsored links, i.e., paid advertisements. The ads are clearly distinguishable from the actual search results, and different searches yield different sponsored links: advertisers target their ads based on search keywords. For instance, if a travel agent buys the word “Hawaii,” then each time a user performs a search on this word, a link to the travel agent will appear on the search results page. When a user clicks on the sponsored link, he is sent to the advertiser’s Web page. The advertiser then pays the search engine for sending the user to its Web page. Different positions of an ad on the search results page have different desirability for

advertisers: an ad shown at the top of a page is more likely to be clicked than an ad shown at the bottom.

To allocate positions to advertisers, most search engines use variations of the “generalized second-price” (GSP) auction. In the simplest GSP auction, for a specific keyword, advertisers submit bids stating their maximum willingness to pay for a click. An advertiser’s bid remains active until he changes or disables it. When a user enters a keyword, she receives search results along with sponsored links, the latter shown in decreasing order of bids. In particular, the ad with the highest bid is displayed at the top, the ad with the next highest bid is displayed in the second position, and so on. If a user subsequently clicks on an ad in position i , that advertiser is charged by the search engine the amount equal to the next highest bid, i.e., the bid of the advertiser in position $(i + 1)$.

If a search engine offered only one ad position per result page, this mechanism would be equivalent to the standard second-price auction, coinciding with the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973). With multiple ad positions, GSP generalizes the second-price auction (hence the name). Here, each advertiser pays the next highest advertiser’s bid. Aggarwal, Goel, and Motvani (2006), EOS, and Varian (2007) show that with multiple positions, the GSP auction is no longer equivalent to the VCG auction. In particular, unlike the VCG mechanism, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP. Nevertheless, GSP has a natural equilibrium, with advertisers in general bidding less than their true values, in which the payoffs of advertisers and the search engine are the same as under VCG for every realization of bidder values.

3 Theory

The ideas of Myerson (1981) can be combined with the analysis of EOS and Varian (2007) to derive the optimal mechanism for the sponsored search setting and to show how it can be implemented with minimal changes to the existing GSP auction. Below is a sketch of the derivation, using the notation of EOS.² Suppose in an auction for a particular keyword there are K bidders and N positions on the screen. The (expected) number of clicks per period received by the advertiser whose ad was placed in position i is α_i . The value per click to advertiser k is s_k . These values are private information of the advertisers, drawn from distribution $F_k(\cdot)$ on $[0, \bar{s}_k]$. Values are independently distributed, so the distribution over vectors of values s is $F(s) = F_1(s_1) \times \cdots \times F_K(s_K)$ over $S = [0, \bar{s}_1] \times \cdots \times [0, \bar{s}_K]$. Advertisers are risk-neutral, and advertiser k ’s payoff from being in position i is equal to $\alpha_i s_k$ minus his payments to the search engine. Without loss of generality, positions are labeled in descending order ($\alpha_1 > \alpha_2 > \dots$).

Now consider an incentive-compatible direct revelation mechanism. Let $t_k(s_k)$ be the expected payment of bidder k with value s_k , let $x_k(s_k)$ be the expected number of clicks received by bidder k with expected value s_k , and, slightly abusing notation, let $x_k(s)$ be the expected number of clicks

²Similar derivations are contained in several recent papers, including Iyengar and Kumar (2006), Roughgarden and Sundararajan (2007), and Edelman and Schwarz (2010).

received by bidder k when the vector of bidder values is s . Then, using the same arguments as in the case of single-object optimal auctions (see, e.g., Krishna, 2002, for an exposition), except that the probability of receiving the object in the single-object case is replaced by the expected number of clicks in our case, we have the following equality for the expected payment of each bidder:

$$t_k(s_k) = t_k(0) + x_k(s_k)s_k - \int_0^{s_k} x_k(u_k)du_k.$$

This, following the standard argument, in turn implies that the expected payoff of the search engine is equal to

$$\sum_{1 \leq k \leq K} t_k(0) + \int_S \left(\sum_{1 \leq k \leq K} \psi_k(s_k)x_k(s) \right) f(s)ds,$$

where $\psi_k(s_k) = s_k - \frac{1-F_k(s_k)}{f_k(s_k)}$ is the virtual valuation of advertiser k with value s_k .

We now make two additional assumptions. First, assume that the virtual valuation is an increasing function. Second, assume that bidders are symmetric, i.e., have identical distributions of values. Then the revenues of the search engine are maximized when $t_k(0) = 0$ for any k and when $\left(\sum_{1 \leq k \leq K} \psi_k(s_k)x_k(s) \right)$ is maximized pointwise, for every s , which happens when (i) only bidders with positive virtual valuations are allocated clicks and (b) among them, bidders with higher virtual valuations (and thus, by assumption, with higher actual valuations) are allocated as many clicks as possible. Since each advertiser can only have one position on the screen, this simply means that the bidder with the highest value receives the top position, the bidder with the second highest value receives the second position, and so on.

Now consider an indirect mechanism: the generalized second-price auction with reserve price r^* such that $\psi(r^*) = 0$. By the argument analogous to that in EOS, in the bidder-optimal envy-free equilibrium of this auction (or, equivalently, in the unique equilibrium of the corresponding generalized English auction with reserve price r^*), bidders with values less than r^* (i.e., bidders with negative virtual valuations) will receive no clicks; among the bidders with values greater than r^* (i.e., with positive virtual valuations), the ones with higher values will receive higher positions; and finally bidders with value zero receive (and make) payments of zero. Hence, the allocations and expected payoffs in this mechanism are the same as those in the optimal direct mechanism, and thus GSP with reserve price r^* is a revenue-maximizing mechanism.

3.1 The Impact of Reserve Prices on Revenues

Remarkably, the optimal reserve price depends neither on the number of bidders nor on the number of available positions. The impact of reserve prices on revenues, however, depends critically on these parameters. In fact, in single-object auctions, reserve prices only play an important role if the number of bidders is small. To give a simple example (Table 1), suppose bidder values are distributed uniformly on $[0, 1]$, with the corresponding optimal reserve price $r^* = 0.5$. Then with just two bidders, the effect of setting the optimal reserve price rather than no reserve price is

substantial: it raises the expected revenues by 25%, from 0.33 to 0.42.³ With six bidders, however, the effect is trivial: moving from no reserve price to the optimal reserve price changes the expected revenues by less than one third of one percent. The intuition for this decline is straightforward: reserve price r only has a positive impact when one bidder’s realized value is above r and the other bidders’ values are all below r , and the probability of this event becomes small as the number of bidders increases.

Of course, the same effect holds for multi-unit auctions, like the sponsored search ones, if the number of slots is fixed but the number of bidders increases. However, reserve prices retain their power for much higher numbers of bidders, and are in general much more important. To see this, consider a generalized second-price auction with the decline factor of 0.7 (i.e., the top position expects to receive one click, the second position expects to receive 0.7 clicks, the third position expect to receive 0.49 clicks, and so on). With two bidders and no reserve price, the expected revenues of the auctioneer are only 0.1: in essence, both bidders get 0.7 clicks “for free,” and only compete for the remaining 0.3 clicks, thus generating the revenue of $0.3 \times 0.33 = 0.1$. Note, however, that it would be feasible for the search engine to “shut down” all positions below the top one, not allocating them to anyone, and so the revenues in the optimal auction have to be at least as high as in the optimal single-object one, i.e., 0.42. As we know from the theoretical analysis, the optimal auction does not in fact involve “shutting down” any positions: the auctioneer simply sets the reserve price equal to 0.5. The resulting expected revenue turns out to be 0.475, i.e., an improvement of 375% relative to the case of no reserve price.

Even with six bidders, reserve prices remain very important: the optimal reserve price improves the revenues by 45%. To see why the difference relative to the single-object case is so dramatic, consider what would have happened if the decline factor in the sponsored search auction was equal to 1 rather than 0.7 (i.e., all positions received the same number of clicks) and there were as many available positions as bidders. Without a reserve price, there would be no competition for positions and the auctioneer’s revenue would be equal to zero. With the optimal reserve price $r^* = 0.5$, revenue would be equal to the number of bidders times 0.25. Of course, with the discount factor of 0.7, the positions are no longer perfect substitutes, and the importance of reserve prices is not as dramatic, but the intuition is essentially the same. Note that in many other settings, substitutable objects are also auctioned off, either simultaneously or sequentially, and if bidders in these auctions have limited demands or each is restricted to one or only a small number of objects, then for the same reason, the analysis based on an individual single-object auction may severely understate the importance of reserve prices.

In order to estimate optimal reserve prices for the experiment, it was assumed that bidders’ values are drawn from lognormal distributions. Table 2 shows the impact of various levels of reserve prices on revenues in GSP under this assumption, with the parameters of the distribution chosen in such a way that its mean is equal to 0.5 and its standard deviation is also equal to

³The expected revenue without a reserve price is equal to $E[\min\{s_1, s_2\}] = 0.33$. The expected revenue with $r^* = 0.5$ is equal to $0.25 \times E[\min\{s_1, s_2\} | s_1 > 0.5, s_2 > 0.5] + 0.5 \times 0.5 = 0.42$.

0.5. The corresponding optimal reserve price is equal to .37. These parameters were chosen to give an illustration of a representative keyword; for instance, as we describe below, the optimal reserve price of 37 cents corresponds to the 75th percentile of estimated optimal reserve prices for the analyzed sample. The table presents the expected revenues for four levels of reserve prices: 0, 0.10 (corresponding to the old reserve price at Yahoo!, 10 cents), 0.235 (corresponding to the midpoint between the old reserve price and the theoretically optimal reserve price), and 0.37 (the theoretically optimal reserve price). Similar to the example with the uniform distribution of values, the impact of optimal reserve prices on revenues in the GSP auction is substantial: with six bidders, setting the reserve price at zero instead of the optimal level results in the loss of 25% of revenues; and even with ten bidders, the loss is noticeable: 9%.

While the reserve price of 0.10 produces much higher revenues than the reserve price of 0, it still falls far short of the optimal revenue. In contrast, moving to the midpoint between 0.10 and 0.37 allows the search engine to capture almost all of the upside from optimal reserve prices, likely because in that region, the derivative of expected revenues with respect to reserve prices is small (of course, at the optimal price level itself, that derivative equals zero).⁴ As discussed below, this observation plays an important role in the implementation of reserve price levels in the field experiment (and may in fact be the key to the reserve price puzzle).

4 Experiment

In practice, sponsored search auctions have a number of complicating features that make the model of Section 3 only a stylized representation of reality. Nevertheless, the model was viewed as a useful approximation and was used as the basis for the experiment. In this section, we outline the implementation of the experiment and discuss several of the complicating features. Broadly speaking, the implementation of the experiment involved two steps: estimating the distributions of bidder values and setting reserve prices.

4.1 Estimating the Distributions of Bidder Values

The company picked a set of criteria for choosing keywords suitable for the experiment (for instance, keywords that had very few bidders or very few searches were not included in the sample, because the distributions of bidder values for such keywords could not be reliably estimated). The resulting sample consisted of 461,648 keywords. It also picked a time interval of several weeks during which the data for estimation were collected.

For each keyword in the sample, they then computed the average number of advertisers bidding on this keyword, the average bid, and the average standard deviation of the bids, where the average was taken over all searches (i.e., every time the keyword was searched, the three statistics were

⁴Note that in this setting, one should be careful about various convexity and concavity statements. In fact, with two or more bidders, the function mapping reserve prices to revenues is neither convex nor concave: its derivative equals zero both at the optimal reserve price and at the reserve price of zero.

computed, and then the average over all searches for the given keyword was taken). The bid of the highest bidder in every auction was excluded from the statistics, because the theory does not allow us to pin it down (just like in a single-object second-price auction, under GSP, every bid of the highest bidder above a certain value results in the same vector of payoffs).

Next, it was assumed that bidders' values were drawn from a lognormal distribution with a mean and a standard deviation to be estimated. The number of potential bidders needed to be estimated as well: during the period when data were collected, Yahoo!'s sponsored search auctions had a uniform reserve price of 10 cents on all keywords, and so bidders with per-click values of less than 10 cents were excluded from the auctions and the data.

The next step was to simulate the three moments (observed number of bidders, average bid in positions 2 and below, and the standard deviation of the bids in positions 2 and below) for various true values of the number of potential bidders and the mean and the standard deviation of the lognormal distribution of values. To do that, for each combination of true values of the variables of interest, several hundred draws of the vectors of bidder values were drawn. For each draw, equilibrium bids were computed, taking into account the 10 cents reserve price and assuming that the bidders were playing the unique perfect Bayesian equilibrium of the Generalized English Auction in EOS. The moments of interest were then computed, and averaged over all draws of vectors of bidder values.

For each keyword, the number of bidders and the parameters of the distribution of bidder values were then estimated by matching the observed moments to the simulated ones. Note that the number of bidders is irrelevant for setting the optimal reserve price, but it needs to be estimated in order to get an accurate estimate of the mean and the standard deviation of the distribution of values. Finally, for each keyword, the theoretically optimal reserve price was computed using the formula in Section 3.

Table 3 presents the distribution of estimated optimal reserve prices for the sample. The median optimal reserve price is equal to 20 cents, the 10th percentile is 9 cents, and the 90th percentile is 72 cents. The 37 cents reserve price used in the example in Section 3.1 corresponds to the 75th percentile. Note that just like in the previous empirical studies of reserve prices in auctions, we find that for most of the sample (almost 90%), the estimated optimal reserve price exceeds the actual reserve price used in the auction (10 cents), and for much of the sample, the difference is substantial. Unlike the previous studies, however, in the current paper we can directly measure the importance of this difference, by conducting a controlled experiment.

Several details of the estimation procedure deserve additional attention.

First, as a simplification, in the simulation procedure it was assumed that each ad's probability of being clicked (conditional on where it is shown) is the same and the ads are ranked solely based on bids. This is an approximation: in practice, ads' "clickabilities" may differ, and auctions rank ads based not only on the bids, but also on their quality scores.

Second, in the simulations, the same "CTR curve" was used for all keywords. The click-through rate (CTR) curve is a function that estimates the ratios of the numbers of clicks the same ad would

receive in different positions on the screen. The CTR curve used in simulations was calibrated to the average estimated CTR curve for a number of auctions. Note also that this assumption implicitly rules out the possibility that the number of clicks that an ad receives, conditional on its position, is influenced by what other ads are shown on the screen (Jeziorski and Segal, 2009).

Third, it was assumed that the values that advertisers assigned to clicks did not depend on where on the screen the ads were shown or on which or how many other ads appeared on the screen. Athey and Ellison (2008) present an alternative model of sponsored search auctions that allows for this possibility by endogenizing advertiser values and discuss how the derivation of optimal reserve prices in that setting differs from the current one.

Fourth, in sponsored search auctions run by Yahoo!, advertisers can allow the company to “advanced match” their ads, by showing them not only for the keyword on which the advertiser bid, but also on other keywords that the company’s algorithms deem sufficiently closely related. For the purposes of the experiment, this possibility was ignored, which means that (1) it was assumed that the bids of the advertisers on a keyword reflect only their value of a click from this keyword and the values of other advertisers bidding for that same keyword and (2) the ads that were “advanced matched” into the keywords in the sample were not used for estimation.

Next, the “theoretical optimality” of the computed reserve prices ignores the dynamic aspects of the real-world sponsored search environment: if bidders know that their bids will be used to set reserve prices in the future, they will change their bids. This problem can in principle be circumvented by setting each advertiser’s reserve price based only on the bids of *other* advertisers. However, the company’s view was that all advertisers for a given keyword should face the same quality-score-adjusted reserve price (more on that below). In addition, with sufficiently many bidders, this dynamic effect becomes small. Hence, it was ignored.

Finally, note that while the estimation procedure is based “in spirit” on the method of moments, we cannot make any claims about its consistency, because that would require a large number of independent observations for each keyword. Nevertheless, for practical purposes, this procedure was viewed as sufficiently robust to various perturbations, such as, e.g., the choice of the lognormal distribution vs. other possible distributions of values, the calibration of the CTR curve, and so on.

4.2 Setting Reserve Prices

The theoretically optimal reserve prices were computed under the assumption that all bidders have the same quality scores and are ranked solely on the basis of their bids, which is a simplification. In practice, the ads on Yahoo! are ranked based largely on the product of their quality scores and bids, and the amount each advertiser pays is lower when his ad’s quality is higher. Thus, in order to keep the implementation of reserve prices consistent with the company’s ranking and pricing philosophy, the theoretical reserve prices were converted into advertiser-specific reserve prices that reflected the quality scores of the ads: ads with higher quality scores faced lower per-click reserve prices, and vice versa. Note that this is a deviation from the theoretically optimal auction design with asymmetric bidders.

In addition, the company wanted to be conservative when implementing the first large-scale application of a previously untested theory and also wanted to experiment with various levels of reserve prices to see how much they impact revenues in practice. For that reason, each keyword was randomly assigned an “adjustment factor,” with values .4, .5, or .6 for most keywords, and the final reserve price was set equal to (optimal reserve price)*(adjustment factor) + (10 cents)*(1 - adjustment factor); a number between the old reserve price of 10 cents and the theoretically optimal reserve price. Simulations like the one presented in Table 2 suggest that this “conservatism” need not be very costly, since most of the upside from the reserve prices is already obtained once the price is set at the midpoint between 10 cents and the optimal reserve price. Moreover, “overshooting” by the same amount may be considerably more costly, and therefore the seller facing uncertainty about the optimal reserve price will prefer to be conservative. This may, in fact, be a part of the explanation of the reserve price puzzle. The flatness of the revenue function around the optimal reserve price level also makes distinguishing between different “adjustment factors” statistically hard, and partly for that reason, the company allocated 95% of the keywords to the “treatment” group (with subgroups receiving different adjustment factors) and only 5% to the “control” group.^{5,6}

5 Experimental Results

The main results of the experiment are described in Table 4. The sample contains 438,632 keywords in the treatment group and 23,016 keywords in the control group. To measure the effect of new reserve prices on various magnitudes of interest, we consider differences-in-differences estimates: we compare the pre-intervention to post-intervention change in the average quantity of interest in the treatment group to that change in the control group.⁷ The introduction of new reserve prices has a strong negative effect on the number of advertisements shown on the page: on average, new reserve prices reduce this number by 0.91, i.e., almost one fewer ad per page is shown as a result of these reserve prices. This number is highly statistically significant. This effect should not be surprising: as Table 3 shows, most reserve prices were raised substantially, thus pricing out many advertisers.

Next, we consider the effect on average revenues per keyword.⁸ The new reserve prices raise revenues by almost 13%. However, this estimate is only marginally statistically significant (the p-value is 9.15%) and turns out not to be robust: for instance, by excluding just one keyword from the control sample, this number can be reduced to around 8%. The reason why this estimate is not

⁵This concern did turn out to be valid: the experiment did not find any systematic differences between different adjustment factors. Hence, when discussing the results, we put all of the treatment keywords in the same group.

⁶Another reason for the small size of the control group was that both theoretical considerations and a smaller pilot experiment were strongly suggestive that new reserve prices would substantially increase revenues, and so allocating more keywords to the control group would be costly for the company.

⁷More specifically, the pre-intervention data come from a 30-day period before the introduction of reserve prices, in May and June of 2008; then several weeks of data are skipped, because new reserve prices were phased in gradually and advertisers had grace periods before they became binding; and then the post-intervention data come from a 30-day period after all reserve prices were phased in and all grace periods ended, in August of 2008.

⁸We first compute the average effect on revenues in dollars, and then divide it by the average pre-intervention per-keyword revenues to get the percentage numbers. Unfortunately, we cannot disclose the absolute dollar amounts.

robust is that average revenues per keyword are affected not only by the bids of the advertisers, but also by the number of searches per keyword, and this number turns out to be highly volatile. For example, the numbers of searches for keywords related to Mother’s and Father’s Days and the NBA and NHL playoffs and finals have dropped dramatically from the pre-intervention period to the post-intervention one, while the opposite has happened to keywords related to the Beijing Olympic games, the NFL season, and the beginning of a new academic year. To control for this volatility and isolate the effect of reserve prices on revenues, instead of computing the difference between pre- and post-intervention revenues for each keyword, we compute the difference between pre- and post-intervention revenues per search for each keyword, and then multiply this difference by the number of searches for this keyword before the intervention. In other words, for each keyword, we compute the difference between its pre-intervention revenues and its hypothetical post-intervention revenues if the number of searches for that keyword remained unchanged. Conceptually, this is the same computation as the Laspeyres index for the rate of inflation, which averages the rates of inflation for individual goods with the weights proportional to the consumption of those goods in the initial period. Once we control for the changes in search volumes in this manner, the estimate of the impact of new reserve prices on revenues drops to 2.7%, but the precision of the estimator increases, and this number is highly statistically significant. Given the size of the sponsored search advertising market, this translates into potential improvements to search engine profits and revenues on the order of hundreds of millions of dollars per year. Moreover, by identifying the segments of keywords where new reserve prices perform relatively poorly and modifying them accordingly, the impact can be further improved. In the next three sections, we study the effect of reserve prices on the subsamples of rarely- and frequently-searched keywords, on the subsamples with high and low theoretically optimal reserve prices, and on the subsamples with many and few advertisers.⁹

5.1 Experimental Results by Search Volume

In the full sample, all keywords were searched by users at least once a day during the pre-intervention period. The distribution of search frequency, however, is highly skewed, with top 5% of keywords receiving 63% of searches and top 25% of keywords receiving 83% (in the pre-intervention period). In this section, we study how the search volume interacts with the impact of reserve prices. Specifically, we consider the impact of new reserve prices on revenues in two subsamples: the subsample of keywords that receive fewer than ten searches per day and the subsample of keywords that receive at least ten searches per day. The latter sample contains only 12.7% of keywords, but receives 74% of searches and generates 81% of revenues.

Tables 5 and 6 present the results for these subsamples. In the former, the impact is negative, with the new reserve prices reducing the revenues by 2.2%, while in the latter it is positive, increasing the revenues by 3.3%. Both numbers are statistically significant at the 5% level. As we mentioned earlier, despite being much smaller, the second subsample is responsible for most of the

⁹We present the results on only the most basic cuts of the data. In more granular cuts, samples become smaller and the results become noisier and less robust.

searches and revenues, which is why the net effect of +2.7% is much closer to the second number than to the first one. We can speculate that the effect of reserve prices on revenues for rarely searched keywords is negative because advertisers simply do not spend much time optimizing bids on these keywords. They may set bids less carefully or update them less frequently, thus making the theory less applicable. Another related possibility is that low search volumes result in less accurate estimates of bidder values, which in turn leads our methodology to set less accurate reserve prices for these keywords.

5.2 Experimental Results by Reserve Price Level

Another dimension along which keywords differ substantially is the estimated theoretically optimal reserve prices, varying from 9 cents for the 10th percentile of keywords to 72 cents for the 90th. Since the original reserve price for all keywords was equal to 10 cents, the shift from the old reserve prices to the new ones (set midway between the old and the theoretically optimal ones) is relatively much more important for the keywords with optimal reserve prices much larger than 10 cents than the keywords with optimal reserve prices close to 10 cents, and thus we expect the impact on revenues in these two groups to be different as well.

Tables 7 and 8 present the results for two subsamples of keywords: those with the optimal reserve price lower than 20 cents and those with the optimal reserve price greater than or equal to 20 cents. These two subsamples are of similar sizes, because 20 cents is approximately equal to the median optimal reserve price in the full sample. The first subsample, however, is responsible for only 7% of revenues, because the keywords in it have, on average, much lower revenues per search (this subsample receives 36% of searches).

For keywords with high theoretically optimal reserve prices, the intervention is very successful: the impact of new reserve prices on revenues is equal to 3.8% and is highly statistically significant. For the other group, however, the intervention reduced revenues by a large amount: 9.2%. This suggests that while the theory used to set prices in this experiment provided useful guidance when the required change was large, it was less successful at “fine-tuning” reserve prices, because of the various simplifications. For example, under the old regime, all bidders faced reserve prices of 10 cents. Under the new regime, consistent with the rest of the company’s pricing practices, bidders are rewarded for having ads that attract users and receive many clicks. Thus, even if the keyword-level reserve price is raised from 10 cents to, say, 12 cents, individual bidders may face reserve prices lower than 10 cents, thus lowering revenues, especially if bidders with highly clickable ads are also more likely to have higher willingness to pay for the clicks.

5.3 Experimental Results by the Number of Advertisers

The final dimension along which we split the data is the number of bidders placing ads on the keyword, also known as the keyword’s “depth.” The median keyword has the average (over time) depth of approximately 5.5 advertisers. The subsample of keywords with the average depth of less than 5.5 advertisers contains 228,292 keywords; each of the remaining 233,356 keywords has the

average depth of at least 5.5 advertisers. The former subsample receives 40% of searches and is responsible for 14% of revenues.

Tables 9 and 10 present the results for these two subsamples. The impact of reserve prices is positive and highly statistically significant in both subsamples. However, consistent with the theory, the revenue impact is much higher for low-depth keywords than for high-depth ones: While for the high-depth subsample the estimated increase in revenues is around 2.5%, for the low-depth subsample the impact is almost four times as large—approximately 10%.

6 Conclusion

The results of the experiment described in this paper show that setting appropriate reserve prices can lead to substantial increases in auction revenues. These results also show (to the best of our knowledge, for the first time) that the theory of optimal auction design is directly applicable in practice. We conclude with a quote from a Yahoo! executive, describing the overall impact of improved reserve prices on company revenues.

On the [revenue per search] front I mentioned we grew 11% year-over-year in the quarter [...], so thats north of a 20% gap search growth rate in the US and that is a factor of, attributed to rolling out a number of the product upgrades we've been doing. [Market Reserve Pricing] was probably the most significant in terms of its impact in the quarter. We had a full quarter impact of that in Q3, but we still have the benefit of rolling that around the world.

Sue Decker, President, Yahoo! Inc. Q3 2008 Earnings Call.^{10,11}

¹⁰<http://seekingalpha.com/article/101002-yahoo-inc-q3-2008-earnings-call-transcript?page=-1>.

¹¹The experiment described in this paper was one of several experiments comprising the Market Reserve Pricing project.

References

- [1] Aggarwal, Gagan, Ashish Goel, and Rajeev Motvani (2006), “Truthful Auctions for Pricing Search Keywords,” In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC06)*, pp. 1–7.
- [2] Armstrong, Mark (2000), “Optimal Multi-Object Auctions,” *Review of Economic Studies*, 67(3), pp. 455–481.
- [3] Athey, Susan, Peter Cramton, and Allan Ingraham (2002), “Setting the Upset Price in British Columbia Timber Auctions,” working paper.
- [4] Athey, Susan, and Glenn Ellison (2008), “Position Auctions with Consumer Search,” working paper.
- [5] Avery, Christopher, and Terrence Hendershott (2000), “Bundling and Optimal Auctions of Multiple Products,” *Review of Economic Studies*, 67(3), pp. 483–497.
- [6] Brown, Jennifer, and John Morgan (2009), “How Much Is a Dollar Worth? Tipping Versus Equilibrium Coexistence on Competing Online Auction Sites,” *Journal of Political Economy*, 117(4): 668–700.
- [7] Clarke, Edward H. (1971), “Multipart Pricing of Public Goods,” *Public Choice*, 11(1), pp. 17–33.
- [8] Cremer, Jacques, and Richard P. McLean (1988), “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56(6), pp. 1247–1257.
- [9] Edelman, Benjamin and Michael Ostrovsky (2007), “Strategic Bidder Behavior in Sponsored Search Auctions,” *Decision Support Systems*, 43(1), pp. 192–198.
- [10] Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz (2007), “Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review*, 97(1), pp. 242–259.
- [11] Edelman, Benjamin and Michael Schwarz (2010), “Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions,” *American Economic Review Papers and Proceedings*, forthcoming.
- [12] Groves, Theodore (1973), “Incentives in Teams,” *Econometrica*, 41(4), pp. 617–31.
- [13] Haile, Philip A., and Elie Tamer (2003), “Inference with an Incomplete Model of English Auctions,” *Journal of Political Economy*, 111(1), pp. 1–51.
- [14] Iyengar, Garud, and Anuj Kumar (2006), “Characterizing Optimal Keyword Auctions,” *Second Workshop on Sponsored Search Auctions, ACM EC’06*.
- [15] Jeziorskiy, Przemyslaw, and Ilya Segal (2009), “What Makes Them Click: Empirical Analysis of Consumer Demand for Search Advertising,” working paper.
- [16] Krishna, Vijay (2002), *Auction Theory*, 1st ed. (San Diego, CA: Academic Press).
- [17] Maskin, Eric, and John Riley (1984), “Optimal Auctions with Risk Averse Buyers,” *Econometrica*, 52(6), pp. 1473–1518.

- [18] Maskin, Eric, and John Riley (1989), “Optimal Multi-Unit Auctions,” in *The Economics of Missing Markets*, ed. by Frank Hahn (Oxford, U.K.: Oxford University Press).
- [19] McAfee, R. Preston, John McMillan, and Philip J. Reny (1989), “Extracting the Surplus in a Common Value Auction,” *Econometrica*, 57(6), pp. 1451–1459.
- [20] McAfee, R. Preston, Daniel C. Quan, and Daniel R. Vincent (2002), “How to Set Minimum Acceptable Bids, with Application to Real Estate Auctions,” *Journal of Industrial Economics*, 50(4), pp. 391–416.
- [21] McAfee, R. Preston, and Daniel R. Vincent (1992), “Updating the Reserve Price in Common Value Auctions,” *American Economic Review Papers and Proceedings*, 82(2), pp. 512–518.
- [22] Myerson, Roger B. (1981), “Optimal Auction Design,” *Mathematics of Operations Research*, 6(1), pp. 58–73.
- [23] Paarsch, Harry J. (1997), “Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales,” *Journal of Econometrics*, 78(2), pp. 333–357.
- [24] Reiley, David (2006), “Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet.” *RAND Journal of Economics*, 37(1), pp. 195–211.
- [25] Riley, John G., and William F. Samuelson (1981), “Optimal Auctions,” *American Economic Review*, 71(3), 381–392.
- [26] Roughgarden, Tim, and Mukund Sundararajan (2007), “Is Efficiency Expensive?,” *Third Workshop on Sponsored Search Auctions, WWW2007*.
- [27] Tang, Xun (2009), “Bounds on Revenue Distributions in Counterfactual Auctions with Reserve Prices,” working paper.
- [28] Varian, Hal R. (2007), “Position Auctions,” *International Journal of Industrial Organization*, 25(6), pp. 1163–1178.
- [29] Vickrey, William (1961), “Counterspeculation, Auctions, and Competitive Sealed Tenders,” *Journal of Finance*, 16(1), pp. 8–37.

Table 1: The impact of optimal reserve prices on expected revenues

Bidders	$r = 0$	$r = 0.5$	Impact
<i>Single-object second-price auction</i>			
$n = 2$	0.3333	0.4166	25%
$n = 6$	0.7143	0.7165	0.31%
<i>GSP with a “decay factor” of 0.7</i>			
$n = 2$	0.1	0.475	375%
$n = 6$	0.8123	1.1764	45%

Table 2: The impact of optimal reserve prices, lognormal distribution

Bidders	$r = 0$	$r = 0.10$	$r = 0.235$	$r = 0.37$
$n = 2$	0.08 24%	0.22 63%	0.32 93%	0.34 100%
$n = 6$	0.68 75%	0.78 86%	0.87 96%	0.91 100%
$n = 10$	1.24 91%	1.28 94%	1.33 98%	1.36 100%

Table 3: The distribution of estimated optimal reserve prices

Percentile	10%	25%	Median	75%	90%
Estimated r^*	9¢	12¢	20¢	37¢	72¢

Table 4: Full sample

Variable	Value	t -statistic	p -value
Number of keywords (T – treatment group)	438,632		
Number of keywords (C – control group)	23,016		
(Mean change in depth in T) – (mean change in depth in C)	−0.9125	−80.53	< 0.0001
(Mean change in revenue in T) – (mean change in revenue in C)	12.85%	1.69	0.0915
Estimated impact of reserve prices on revenues	2.71%	4.73	< 0.0001

Table 5: Restricted sample (keywords with fewer than 10 searches per day)

Variable	Value	<i>t</i> -statistic	<i>p</i> -value
Number of keywords (T – treatment group)	382,860		
Number of keywords (C – control group)	20,133		
(Mean change in depth in T)–(mean change in depth in C)	–0.9039	–75.53	< 0.0001
(Mean change in revenue in T)–(mean change in revenue in C)	10.33%	1.19	0.2354
Estimated impact of reserve prices on revenues	–2.19%	–2.36	0.0183

Table 6: Restricted sample (keywords with at least 10 searches per day)

Variable	Value	<i>t</i> -statistic	<i>p</i> -value
Number of keywords (T – treatment group)	55,772		
Number of keywords (C – control group)	2,883		
(Mean change in depth in T)–(mean change in depth in C)	–0.971	–28.07	< 0.0001
(Mean change in revenue in T)–(mean change in revenue in C)	14.03%	1.51	0.1316
Estimated impact of reserve prices on revenues	3.30%	2.32	0.0201

Table 7: Restricted sample (optimal reserve price < 20¢)

Variable	Value	<i>t</i> -statistic	<i>p</i> -value
Number of keywords (T – treatment group)	222,249		
Number of keywords (C – control group)	11,615		
(Mean change in depth in T)–(mean change in depth in C)	–0.8612	–60.29	< 0.0001
(Mean change in revenue in T)–(mean change in revenue in C)	–11.88%	–2.45	0.0144
Estimated impact of reserve prices on revenues	–9.19%	–11.1	< 0.0001

Table 8: Restricted sample (optimal reserve price ≥ 20¢)

Variable	Value	<i>t</i> -statistic	<i>p</i> -value
Number of keywords (T – treatment group)	216,383		
Number of keywords (C – control group)	11,401		
(Mean change in depth in T)–(mean change in depth in C)	–0.9664	–55.09	< 0.0001
(Mean change in revenue in T)–(mean change in revenue in C)	14.59%	1.79	0.0736
Estimated impact of reserve prices on revenues	3.80%	5.41	< 0.0001

Table 9: Restricted sample (average depth < 5.5)

Variable	Value	t -statistic	p -value
Number of keywords (T – treatment group)	217,087		
Number of keywords (C – control group)	11,205		
(Mean change in depth in T)–(mean change in depth in C)	–0.9332	–77.66	< 0.0001
(Mean change in revenue in T)–(mean change in revenue in C)	25.50%	1.58	0.1142
Estimated impact of reserve prices on revenues	10.06%	7.29	< 0.0001

Table 10: Restricted sample (average depth ≥ 5.5)

Variable	Value	t -statistic	p -value
Number of keywords (T – treatment group)	221,545		
Number of keywords (C – control group)	11,811		
(Mean change in depth in T)–(mean change in depth in C)	–0.9113	–52.93	< 0.0001
(Mean change in revenue in T)–(mean change in revenue in C)	10.44%	1.25	0.2102
Estimated impact of reserve prices on revenues	2.54%	3.59	< 0.0003