# Sponsored Search with Contexts

Eyal Even-Dar, Michael Kearns, and Jennifer Wortman
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104

## ABSTRACT

We examine a formal model of sponsored search in which advertisers can bid not only on search terms, but on search terms under specific *contexts*. A context is any auxiliary information that might accompany a search, and might include information that is factual, estimated or inferred. Natural examples of contexts include the zip code, gender, or abstract "intentions" (such as researching a vacation) of the searcher. After introducing a natural probabilistic model for context-based auctions, we prove several theoretical results, including the fact that under rather general circumstances, the overall social welfare of the advertisers and auctioneer together can only increase when moving from standard to context-based mechanisms. In contrast, we also provide and discuss specific examples in which only one party (advertisers or auctioneer) benefit at the expense of the other in moving to context-based search, and we give extensive simulations contrasting standard and context-based mechanisms in light of these observations.

## 1. INTRODUCTION

In the standard model of sponsored search, advertisers may place bids on individual terms or keywords. When a user searches for a given keyword, an auction is held to determine which advertisers' ads will appear on the search results page presented to the user and in what order they will appear. Generally the auction mechanism used is a generalized second-price (or *next-price*) mechanism [4] in which advertisers are ranked either by bid alone (the so-called "Rank By Bid" allocation, or RBB) or by the product of bid and a numerical measure of the quality of the ad ("Rank By Revenue," or RBR) [7]. An advertiser pays a fee to the auctioneer (in this case, the search engine) only when the search results in a click on the ad.

In this paper we investigate the extension of standard sponsored search auction models to incorporate what we shall call *contexts*. Broadly speaking, a context is any piece of auxiliary information that might modify the interpretation or expected value of a specific search query. Contexts may be "factual" information, or may be based on (possibly noisy) inferences. For example:

- It is often possible to infer a user's zip code from their IP address. Advertisers providing only local services (such as dentistry or child care) might value searches originating from certain zip codes much more highly than others. Certain zip codes might also correlate with various demographic factors such as income, and thus also be more valuable to advertisers.

- An online retailer specializing in maternity clothing might place a high value on clicks from women in their twenties and thirties, while a website selling dorm room supplies might prefer clicks from teenagers of either gender. Search engines are often able to collect such demographic information about searchers directly via site accounts and could potentially use this information to select more relevant ads.

- In some cases it may be possible to estimate a user's abstract "intention" from their recent search or web activity. For example, we might infer a user's interest in planning a vacation from a series of searches for travel web sites. Though their true goal might be to open a competing site, a hotel in Istanbul would probably be willing to pay more per click on a search for the keyword "turkey" if it were simply *more likely* that the searcher was planning a vacation than a home-cooked meal. (Currently, a search for "turkey" on Google yields sponsored ads for both tour groups in Istanbul and a brand of deli meat.)

Search engines already include limited abilities for advertisers to modify bids based on searcher context. Both Google's AdWords program[1] and Yahoo's Search Marketing program[2] allow advertisers to bid on searches limited to specified geographic areas. Microsoft's adCenter[3] allows bidders to target searchers by location, age, or gender by specifying an additional bid amount for targeted searches on top of a base keyword bid. Aside from these formal context-based mechanisms, advertisers may informally implement their own by bidding on more specific search terms (e.g. "Philadelphia dentist" rather than "dentist" alone). Indeed, there are now companies (see for instance `www.natpal.com`) offering proxy bidding to businesses that provide geographically localized services. Such proxy bidding optimizes the modifiers or contexts (e.g. "Philadelphia") that should be added to the basic search terms (e.g. "dentist") in order to maximize click-through rates and conversions.

Despite these various existing forms of context-based bidding, to our knowledge no formal study has been published


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

[1] `http://adwords.google.com`
[2] `http://searchmarketing.yahoo.com`
[3] `http://adcenter.microsoft.com`

showing that these additional bidding capabilities are beneficial either to the advertisers or to the search engines themselves. The closest existing work is the economic literature on *bundling*, the strategy of offering multiple distinct goods for sale together as a single item. (In our setting, these goods are clicks from searches on the same keyword in different contexts, and we are interested in what happens when they are *unbundled*.) Palfrey first analyzed the effects of bundling on bidder and auctioneer welfare in VCG auctions [8]. He showed that under certain statistical assumptions about bidders' values for each item, when the number of bidders is relatively small, bundling will be helpful to the auctioneer but will decrease the revenue of the bidders. On the other hand, under the same assumptions, when the number of bidders is large, auctioneers will be better off holding separate auctions for each item. In this case, only bidders with high values across multiple items would profit from bundling. Further analysis by Chakraborty [3] showed that there always exists a critical threshold on the number of bidders. When the number of bidders is lower than this threshold, it is in the seller's best interest to bundle; when the number of bidders is higher, it is in the seller's best interest to hold separate auctions. While none of these results carry over directly to the more complex multiple-slot sponsored search auction setting, the underlying intuition is similar.

The idea of allowing increased expressiveness in sponsored search auctions, including context-specific bidding, has also been suggested by Parkes and Sandholm in the context of efficient solutions for the winner determination problem [9]. However, their work does not address the effects of increased expressiveness on revenue. Here we examine sponsored search auctions in which advertisers may place explicit bids on pairs of keywords and contexts, and compare the welfare of both the advertisers and the search engines in this setting to the welfare when bids are restricted to words alone.

In Section 2, we introduce standard sponsored search terminology and definitions, along with our probabilistic model for incorporating contexts. In Section 3, we prove our main theoretical result, which states that the total profits of all bidders and the auctioneer taken together can only increase in moving to a context-based mechanism. This result, whose proof is straightforward, holds under fairly weak assumptions about the original underlying mechanism and the behavior of the advertisers. These assumptions are met, for example, when advertisers bid according to a symmetric Nash equilibrium [10] in a next-price auction, or under truthful bidding in a Vickrey-Clarke-Groves (VCG) auction. Section 4 is devoted to the discussion of simple examples of trade-offs in revenue between advertisers and auctioneer, including cases in which only one party benefits at the other's expense. In Section 5 we generalize the theory of Section 3 to the case in which contexts are estimated noisily, and in Section 6 we provide extensive simulation results that elucidate both the theory and our observations about trade-offs.

## 2. PRELIMINARIES

Without loss of generality we will limit our analysis to auctions on a single fixed keyword (or search term) $w$. We assume that there exists a fixed and known distribution $P$ over the set of user contexts $C$ for searches on $w$, and that for all $c \in C$, each advertiser $a \in \{1, \cdots, A\}$ has a known (expected) value $v_{a,c}$ for a click from a user with context $c$.

For each query, we assume there are $S$ advertiser slots available. We make the standard assumption [10, 1] that the click-through rate (CTR) of an ad shown in slot $s$ can be factored into two parts, a slot-specific base click-through rate $x_s$ that is monotonically decreasing in the $s$, and a quality effect $e_{a,c}$ that can depend on the advertiser $a$ and context $c$ in an arbitrary way. We can then write the click-through rate of advertiser $a$ in slot $s$ for context $c$ as the product $e_{a,c}x_s$, and the expected click-through rate over all contexts as $e_a x_s$ where $e_a = \sum_{c \in C} P(c) e_{a,c}$. For convenience, we define $x_s = 0$ for $s > S$.

In order to compute the expected value over all contexts of a click to an advertiser, we must take into account the advertiser's quality since the distribution of clicks that an advertiser receives will be affected by his quality scores. We can compute the expected value of a click to advertiser $a$ as $v_a = \sum_{c \in C} P(c) v_{a,c} e_{a,c} / e_a$. Note that if the quality of the advertiser is constant over all contexts then $v_a = \sum_{c \in C} P(c) v_{a,c}$ as expected.

We examine both the standard VCG mechanism (see, for example, Edelman et al. [4]) and the generalized second-price auction mechanism using a rank by revenue (RBR) allocation scheme [7]. Under RBR, advertisers are ranked by the product of their quality effect ($e_{a,c}$ or $e_a$) and their bid ($b_{a,c}$ or $b_a$) rather than bid alone. This approximately models the allocation methods currently used by both Google and Yahoo. In the RBR generalized second-price auction, the payment of advertiser $i$ for a click in slot $s$ is calculated as $b_j e_j / e_i$ (or $b_{j,c} e_{j,c} / e_{i,c}$ in a context-based model) where $j$ is the advertiser in slot $s + 1$. Notice that this payment is the minimum amount that advertiser $i$ must bid to remain in slot $i$, i.e. the minimum value of $b_i$ for which $b_i e_i \geq b_j e_j$.

We will analyze three quantities of interest: the advertiser profit, the combined social welfare, and the auctioneer revenue. We define the advertiser profit as the sum over all advertisers of the expected value received from clicks on a given user search minus the expected price paid. The combined social welfare is simply the advertiser profit plus the expected revenue of a user search to the auctioneer, i.e. the social welfare if we think of the search engine as a player. Intuitively, this can be thought of as a measure of the economic efficiency of the auction. Note that since the revenue of the auctioneer is by definition equal to the total amount paid by all bidders, the combined social welfare is equivalent to the sum over all advertisers of the expected value of a search.

We will examine sponsored search mechanisms under various equilibrium concepts. For next-price auctions, it is appropriate to consider the concept of symmetric Nash equilibrium introduced by Varian [10].[4] While Varian's equilibria concepts were originally defined in the RBB setting, they can naturally be extended to the RBR setting. Letting $v_s$ and $e_s$ denote the value and quality of the bidder in slot $s$, we can express the bids under high SNE as

$$b_s = \frac{1}{e_s x_{s-1}} \left( \sum_{t=s}^{S} (x_{t-1} - x_t) e_{t-1} v_{t-1} + x_S e_{S+1} v_{S+1} \right)$$

[4]Edelman et al. [4] independently developed the idea of symmetric Nash equilibria around the same time. They called these equilibria *locally envy-free* and showed that in next-price auctions there exists one particular locally envy-free equilibrium (which Varian calls the *low SNE*) for which revenue to both the auctioneer and the bidders is the same as under truthful bidding in a VCG auction.

and the bids under low SNE as

$$b_s = \frac{1}{e_s x_{s-1}} \sum_{t=s}^{S+1} (x_{t-1} - x_t) e_t v_t$$

where for both we have $b_{S+1} = v_{S+1}$. For VCG, it is appropriate to examine the dominant-strategy equilibrium in which all advertisers bid truthfully. Of course under truthful bidding we simply have $b_s = v_s$ for all $s$.

Given a sponsored search mechanism and an equilibrium concept (such as Nash, symmetric Nash, or dominant strategy equilibria), we will say that the mechanism is *efficient at equilibrium* under the specified equilibrium concept if the mechanism maximizes combined social welfare any time such an equilibrium is played. It is easy to see that an RBR generalized second-price auction is efficient under symmetric Nash equilibria. (The proof, which is similar to Varian's "monotone values" proof for SNE of next-price auctions without quality scores [10], relies on the fact that in sponsored search auctions, social welfare is always maximized when ads are ranked in decreasing order by the product of value and quality.) Additionally, VCG is efficient under the dominant-strategy equilibrium.

To start we will assume that the context of a search query is always known. This assumption can be relaxed as we show in Section 5.

# 3. COMBINED SOCIAL WELFARE

We begin by examining the shift in combined social welfare that occurs when we move from the standard keyword auction to a context-based auction. We will show that under a broad variety of conditions, the combined social welfare can only increase when context-based bidding is introduced. This result is not surprising given similar results from bundling theory (see, for example, Theorem 4 of Palfrey [8]), but is nice in its generality. In particular, our result generalizes Palfrey's to the multiple-slot auction setting. It holds for any mechanism that is efficient at equilibrium for an appropriate equilibrium concept.

Before stating the main theorem of this section, we give the following short and intuitive lemma. The proof is by straight-forward induction and can be found in standard textbooks [6].

Lemma 3.1 (Rearrangement Inequality). *Suppose we are given two sets of ordered values, $\{y_1, y_2, \cdots, y_k\}$ and $\{z_1, z_2, \cdots, z_k\}$ such that for all $\ell \in \{1, \cdots, k-1\}$, $y_\ell \geq y_{\ell+1}$ and $z_\ell \geq z_{\ell+1}$. The permutation $\pi$ on the numbers $\{1, \cdots, k\}$ maximizing the value $\sum_{\ell=1}^{k} y_{\pi(\ell)} z_\ell$ is the identity $\pi(\ell) = \ell$.*

We are now ready to prove the main theorem of this section. We state the conditions as generally as possible, and then immediately apply the result to some specific standard mechanisms.

Theorem 3.2. *Consider any sponsored search mechanism that is efficient at equilibrium for a given equilibrium concept. For this mechanism, the combined social welfare at equilibrium under context-based bidding is at least as high as the combined social welfare at equilibrium under standard keyword bidding.*

Proof. Let $\pi$ be a permutation capturing the slot ordering of advertisers at equilibrium under standard keyword bidding. In other words, if advertiser $a$ is in slot $s$ at equilibrium, then $\pi(s) = a$. Similarly, for all $c$, let $\pi_c$ be a permutation on advertisers capturing their slot ordering for keyword $c$ at equilibrium under context-based bidding.

At equilibrium, the expected total value to the advertisers of a given query for the word $w$ in the standard word-based auction is

$$\sum_{c \in C} P(c) \sum_{s=1}^{S} v_{\pi(s),c} e_{\pi(s),c} x_s \tag{1}$$

By the Rearrangement Inequality, it is clear that this sum will be maximized when $v_{\pi(s)} e_{\pi(s)} \geq v_{\pi(s+1)} e_{\pi(s+1)}$ for all $s$, i.e. when ads are sorted according to the product of expected value and quality.

Similarly, the expected total value to the advertisers of a given query for the word $w$ in the context-based auction can be calculated as

$$\sum_{c \in C} P(c) \sum_{s=1}^{S} v_{\pi_c(s),c} e_{\pi_c(s),c} x_s \tag{2}$$

This is maximized when $v_{\pi_c(s),c} e_{\pi_c(s),c} \geq v_{\pi_c(s+1),c} e_{\pi_c(s+1),c}$ for all $s$ and $c$, i.e. when ads are sorted according to the product of value and quality for the current context.

Consider the inner sums in Equations 1 and 2 above for a given context $c$. Since we have established that the permutation $\pi_c$ sorts advertisers according to the product of value and quality for context $c$, and since $x_s$ is decreasing in $s$, we can again apply Lemma 3.1 to see that

$$\sum_{s=1}^{S} v_{\pi_c(s),c} e_{\pi_c(s),c} x_s \geq \sum_{s=1}^{S} v_{\pi(s),c} e_{\pi(s),c} x_s$$

Thus the social welfare will always be at least as high when advertisers may bid on contexts than it would have been when advertisers must bid on words alone. □

The following corollaries illustrate some specific examples of sponsored search mechanisms and equilibrium concepts for which Theorem 3.2 applies. They follow immediately from the efficiency of the mechanisms under the specified equilibrium concepts.

Corollary 3.3. *In a RBR generalized second-price auction, the combined social welfare at a symmetric Nash equilibrium under context-based bidding is at least as high as the combined social welfare at a symmetric Nash equilibrium under standard keyword bidding.*

Corollary 3.4. *In a VCG auction, the combined social welfare at the dominant-strategy truthful equilibrium under context-based bidding is at least as high as the combined social welfare at the dominant-strategy truthful equilibrium under standard keyword bidding.*

# 4. TRADE-OFFS IN REVENUE

In the previous section, we saw that under a wide variety of auction mechanisms and bidding assumptions, it will always be more efficient in terms of combined social welfare to allow context-based bidding. However, this does not necessarily imply that context-based bidding always produces higher revenue for the auctioneer or that it always increases the total revenue of the advertisers. Indeed there are situations in which context-based auctions result in lower revenue for the auctioneer or for the bidders as a whole. In this section we examine scenarios in which decreases in revenue might occur and analyze why this is the case.

For simplicity, we consider next-price auctions over a single ad slot and assume that all advertisers bid truthfully.[5] Similar examples can be shown in multiple-slot models. In each example, the word $w$ will have two possible contexts, $c_1$ and $c_2$, with $P(c_1) = P(c_2) = 0.5$, and all advertisers will have a uniform quality effect of 1.

## 4.1 A Decrease in Auctioneer Revenue

It is possible to construct a simple example with only two bidders in which the revenue of the auctioneer will decrease if context-based bidding is allowed. The values of the advertisers for each context are given in the following table.

| $a$ | $v_{a,c_1}$ | $v_{a,c_2}$ | $v_a$ |
|---|---|---|---|
| 1 | 10 | 1 | 5.5 |
| 2 | 1 | 10 | 5.5 |

Under standard word-based bidding, the expected auctioneer revenue will be

$$x_1 e_{r(1)}(b_{r(2)}e_{r(2)}/e_{r(1)}) = x_1 v_{r(2)}e_{r(2)} = 5.5x_1$$

If context-based bidding is introduced, the expected revenue of the auctioneer will be

$$\begin{aligned} & \sum_{c\in C} P(c)x_1 e_{r_c(1),c}(b_{r_c(2),c}e_{r_c(2),c}/e_{r_c(1),c}) \\ = & \sum_{c\in C} P(c)x_1 v_{r_c(2),c}e_{r_c(2),c} \\ = & \; x_1(0.5*1 + 0.5*1) = x_1 \end{aligned}$$

The problem that arises in this simple example is the general problem of splitting the competition. When the advertisers are forced to bid on both contexts of $w$, they are placed in direct competition with each other. Because of the nature of second-price auctions, this competition is enough to drive up the price per click. However, when the advertisers are free to bid separately for each context, they are not in direct competition for the contexts they each prefer most and are thus able to pay less per click, reducing the revenue to the auctioneer.

As a hypothetical example of where this situation might be seen in reality, consider the effect of introducing zip code-based bidding on advertising by cable service providers. While all cable providers are likely to value clicks from sponsored ads appearing on general keywords like "cable" or "cable tv", there is typically only one cable provider available in a given zip code. It is likely that context-based bidding would thus reduce competition between cable providers, lowering prices per click which would in turn lower the overall revenue to the auctioneer. In this case, it would be in the search engine's best interest to stick with standard keyword bidding.

This point is further illustrated in Figure 1. As we will illustrate through simulations in Section 6, when the number of bidders is close to the number of contexts and many bidders exhibit "singleton" behavior, preferring one context above the others, it can be harmful for the auctioneer to allow context-based bidding. However, we will see that this is often not the case, and under other circumstances it can be quite beneficial.

[5]Note that when there is only a single slot, the next-price auction mechanism is equivalent to VCG. Truthful bidding is therefore a dominant-strategy equilibrium that maximizes combined social welfare.

## 4.2 A Decrease in Advertiser Profit

The introduction of context-based bidding can sometimes lead to a *concentration* of competition, yielding higher revenue for the auctioneer at the expense of the advertisers. Suppose we have three advertisers with values as given in the following table.

| $a$ | $v_{a,c_1}$ | $v_{a,c_2}$ | $v_a$ |
|---|---|---|---|
| 1 | 10 | 10 | 10 |
| 2 | 9 | 1 | 5 |
| 3 | 1 | 9 | 5 |

Clearly advertiser 1 will be the high bidder for both the word-based and the context-based auctions. Because advertisers 2 and 3 each value only one context highly, they will not cause a high price for advertiser 1 in the word-based setting. However, we will see that moving to a context-based auction will result in increased competition for each context, raising the price for advertiser 1.

In the word-based auction, the advertiser profit can be calculated as

$$\begin{aligned} & x_1 e_{r(1)}\left(v_{r(1)} - (b_{r(2)}e_{r(2)}/e_{r(1)})\right) \\ = & \; x_1\left(v_{r(1)}e_{r(1)} - v_{r(2)}e_{r(2)}\right) \\ = & \; x_1(10-5) = 5x_1 \end{aligned}$$

In the context-based auction, the advertiser profit will be

$$\begin{aligned} & \sum_{c\in C} P(c)x_1 e_{r_c(1),c}\left(v_{r_c(1),c} - (b_{r_c(2),c}e_{r_c(2),c}/e_{r_c(1),c})\right) \\ = & \sum_{c\in C} P(c)x_1(v_{r_c(1),c}e_{r_c(1),c} - v_{r_c(2),c}e_{r_c(2),c}) \\ = & \; x_1\left(0.5*(10-9) + 0.5*(10-9)\right) = x_1 \end{aligned}$$

This example is especially striking because the allocation of clicks is the same under both models; only the payment scheme has changed. The introduction of contexts has concentrated the competition forcing advertiser 1 to pay a higher price per click.

Such a scenario can occur when there are a mix of large corporations and smaller local services in competition for ads on a given keyword. Consider the market for ads on the term "pizza." Most local pizzerias would not bother to place ads on this term in a word-based setting as the majority of the clicks they would receive would have no value. This would allow larger nationwide pizza chains to purchase these ads at moderate prices. However, introducing zip code-based bidding could motivate many smaller pizza chains to begin placing ads to attract locals, driving up the prices of ads for the large chains and increasing the revenue to the auctioneer.

Indeed we will see in Section 6 that when a portion of bidders exhibit the singleton behavior yet enough bidders exist such that there is still competition for each context, prices will be driven up and the advertiser profit will decrease when context-based bidding is allowed.

## 4.3 Increased Revenue for Everyone

Finally, it is often the case that introducing context-based bidding will simultaneously allow advertisers to reach their target audience while still allowing enough competition for the auctioneer to benefit. Consider four advertisers with the values shown in the following table.

| $a$ | $v_{a,c_1}$ | $v_{a,c_2}$ | $v_a$ |
|---|---|---|---|
| 1 | 10 | 1 | 5.5 |
| 2 | 7 | 1 | 4 |
| 3 | 1 | 10 | 5.5 |
| 4 | 1 | 7 | 4 |

As in the first example, under standard word-based bidding, the expected auctioneer revenue is

$$x_1 e_{r(1)}(b_{r(2)}e_{r(2)}/e_{r(1)}) = x_1 v_{r(2)}e_{r(2)} = 5.5x_1$$

Now when context-based bidding is introduced, the expected revenue of the auctioneer increase to

$$\begin{aligned}&\sum_{c\in C} P(c)x_1 e_{r_c(1),c}(b_{r_c(2),c}e_{r_c(2),c}/e_{r_c(1),c})\\ &= \sum_{c\in C} P(c)x_1 v_{r_c(2),c}e_{r_c(2),c}\\ &= x_1(0.5*7+0.5*7) = 7x_1\end{aligned}$$

Meanwhile, due to tightly packed competition for the only slot, the advertiser profit under word-based bidding will be

$$x_1 e_{r(1)}v_{r(1)} - 5.5x_1 = 5.5x_1 - 5.5x_1 = 0$$

while in the context-based auction, the advertiser profit will be

$$\begin{aligned}&\sum_{c\in C} P(c)x_1 e_{r_c(1),c}v_{r_c(1),c} - 7x_1\\ &= x_1\,(0.5*10+0.5*10) - 7x_1 = 3x_1\end{aligned}$$

In this example, introducing context-based auctions allows the advertisers to focus ads on their target audiences, raising advertiser profit, but without completely splitting up the competition, enabling the auctioneer to profit as well. This scenario could arise when contexts are once again zip codes, and advertisers are local businesses in competition with each other. Consider a set of dentist offices bidding on ads for the keyword "dentist." Each dentist office would be interested only in local clicks and as such would be happy to have the option to bid by context. However, since there are often multiple dentists servicing patients in any given zip code, there would still be enough competition that the search engine would be better off as well. We suspect that this scenario is likely to fit the bidding patterns of advertisers on most common keywords.

## 5. THE NOISY CONTEXT SETTING

In this section we relax our assumptions and analyze the case in which the context reported by the auctioneer is not necessarily the true user context, but rather an estimate. Such a situation might arise when we treat context not as simple environmental facts about the user, but as something more difficult to infer such as the intent to buy a new car versus the intent only to browse.

In the previous sections we assumed that the advertisers and the auctioneer know the prior over the different contexts. Here we assume that they also know how to compute the posterior, i.e. the probability that the true context of a user is $c$ given that the auctioneer has predicted that the user's context is $p$. We denote this posterior as $P(c|p)$. As before, we use $P(c)$ to denote the prior probability that the *true* search is context $c$. We use $Q(p)$ to denote the prior probability that the *predicted* context is $p$.

For each advertiser $a$, we can compute the expected value of a click from a user with predicted context $p$:

$$\bar{v}_{a,p} = \frac{\sum_{c\in C} P(c|p)e_{a,c}v_{a,c}}{\sum_{c\in C} P(c|p)e_{a,c}}$$

Similarly, we can compute the expected quality of advertiser $a$ on predicted context $p$:

$$\bar{e}_{a,p} = \sum_{c\in C} P(c|p)e_{a,c}$$

Recall that in the deterministic setting, advertisers were ranked according to the value of $b_a e_a$ for word-based auctions, or $b_{a,c}e_{a,c}$ for context-based auctions. In the noisy context setting, they will still be ranked according to $b_a e_a$ for word-based auctions, but will now be ranked by $b_{a,c}\bar{e}_{a,c}$ for context-based auctions. Let $r$ be the ranking of ads for word-based auctions, and $r_p$ be the ranking for context-based auctions when the *predicted* context is $p$.

Like Theorem 3.2, the following result is stated quite generally. It will again hold when advertisers bid according to any symmetric Nash equilibrium in a RBR next-price auction, or when advertisers bid truthfully in a VCG auction.

THEOREM 5.1. *Consider any sponsored search mechanism that is efficient at equilibrium for a given equilibrium concept. For this mechanism, the combined social welfare at equilibrium under noisy context-based bidding is at least as high as the combined social welfare at equilibrium under standard keyword bidding.*

PROOF. Define the permutations $\pi$ and $\pi_c$ as in the proof of Theorem 3.2. As before, it can be shown using the Rearrangement Inequality that since the auction mechanism we are considering is efficient at equilibrium, it must be the case that $v_{\pi(s)}e_{\pi(s)} \geq v_{\pi(s+1)}e_{\pi(s+1)}$ for all $s$ and that $v_{\pi(s),c}\bar{e}_{\pi(s),c} \geq v_{\pi(s+1),c}\bar{e}_{\pi(s+1),c}$ for all $s$ and all contexts $c$. Thus we have

$$\begin{aligned}&\text{Word-Based Combined Social Welfare}\\ &= \sum_{c\in C} P(c)\sum_{s=1}^{S} v_{\pi(s),c}e_{\pi(s),c}x_s\\ &= \sum_{c\in C}\sum_{p\in C} Q(p)P(c|p)\sum_{s=1}^{S} v_{\pi(s),c}e_{\pi(s),c}x_s\\ &= \sum_{p\in C} Q(p)\sum_{s=1}^{S}\sum_{c\in C} P(c|p)v_{\pi(s),c}e_{\pi(s),c}x_s\\ &\leq \sum_{p\in C} Q(p)\sum_{s=1}^{S}\sum_{c\in C} P(c|p)v_{\pi_p(s),c}e_{\pi_p(s),c}x_s\\ &= \text{Noisy Context-Based Combined Social Welfare}\end{aligned}$$

The inequality is another application of Lemma 3.1. □

## 6. SIMULATIONS

This section contains results from a set of preliminary simulations investigating the circumstances under which context-based bidding is superior to standard keyword bidding and vice versa. These simulations study the effects of altering a wide range of parameters including the number of bidders, bidder types, and quality effects.

A bidder's values for each context are generated randomly, based on the type of the bidder:

- A *constant bidder* has the same value for clicks from all contexts. This value is drawn from a normal distribution centered around an underlying *base word value*.
- A *singleton bidder* has non-zero value only for one context. The value for this context is drawn from a normal distribution centered around an underlying *base context value*.
- A *random bidder* may have non-zero value for every context. Each value is again drawn from a normal distribution centered around the underlying base context value.

Except where otherwise specified, base context values are set to be equal to the base word value, making the implicit assumption that all contexts are inherently of equal value.

Once values have been set for each bidder for every context, quality effects are generated in one of three ways:

- *Uniform quality* implies that all advertisers have a quality score of 1 across all contexts.
- When quality is *correlated*, the advertiser with the $i$th highest value for a context is assigned a quality of $0.95^{i-1}$ for this context.
- Similarly, when quality is *anti-correlated*, the advertiser with the $i$th lowest value for a context has quality $0.95^{i-1}$ for the context.

All simulations assume a model with 8 ad slots with base click-through rates set according to the empirical numerical values reported by the Atlas Institute's rank report [2]. All results except for those in Section 6.4 have been averaged over 10,000 random trials.

## 6.1 Varying the Number of Bidders

We chose to start by examining the case in which most advertisers have a high value for only a single context, studying what happens to auctioneer revenue and advertiser profit as the number of bidders grows. As we saw in Section 4, when only a single advertiser is interested in each context, allowing context-based bidding can be harmful to the auctioneer, whereas when many advertisers have a singleton interest in each, auctioneer revenue can improve greatly. Figure 1 can be seen as a visualization of this shift, illustrating what happens between these extreme points.

In this set of experiments, the number of contexts is fixed at 20. 95% of advertisers are singleton bidders while the other 5% are constant bidders, and quality effects are assigned uniformly. The total number of bidders ranges from 1 to 500.

The upper left and upper right plots of Figure 1 show the revenue of the auctioneer as a function of the number of advertisers when advertisers bid according to Varian's low and high symmetric Nash equilibria. The bottom left plot shows the advertiser profit when advertisers bid according to the low SNE. (The plot for high SNE is similar.) Finally, the bottom right plot shows the combined social welfare when advertisers bid according to any symmetric Nash equilibrium (or truthfully, or more generally, whenever bidders bid in such a way that their ranking by the product of bid and quality is the same as their ranking by the product of value and quality).

It is clear from these plots that as the number of bidders grows, the revenue of the auctioneer in the context-based auction grows more and more superior to the revenue in the word-based auction. This happens very quickly; by the time there are twice as many bidders as contexts, the revenue of the auctioneer is already higher in the context-based setting. The advertiser profit also tends to be higher when context-based auctions are used up until the point in which there is too much competition between bidders for each context and the revenue for advertisers decreases. There is a clear sweet spot in between these two extremes in which both auctioneer revenue and advertiser profit grow. This can be seen as a generalization of the final example in Section 4 in which advertisers are able to focus ads on their target audiences, but the competition for ad slots is not completely split.

## 6.2 Varying the Bidders' Types

Next we chose to examine the effect of bidder types, varying the percentage of singleton bidders. In these experiments, the number of bidders is fixed at 100 and the number of contexts at 10. The percentage of advertisers who are singletons ranges from 0 to 100; the remaining advertisers are assumed to be constant.

Figure 2 shows the the auctioneer revenue under low and high SNE, advertiser profit under low SNE (again, high SNE is similar), and combined social welfare, all as a function of the percent of bidders with singleton values. We can see that auctioneer revenue improves when context-based bidding is introduced, except in the extreme case when nearly all of the bidders are singletons. When this is the case, context-based bidding results in too much splitting of the competition and a decrease in revenue occurs. On the other hand, advertiser profit is superior in the context-based model only when a sufficient number of advertisers are singletons. Again we can see a sweet spot at which both advertiser profit and auctioneer revenue increase. This happens when the percentage of bidders exhibiting singleton behavior is between 70 and 80.

## 6.3 Varying Other Model Parameters

We subsequently tried varying other parameters of the model in order to see how these might affect the gap in revenue between the context-based auction and the standard keyword auction. First, we tried setting the quality effects in various ways. Figure 3 shows the auctioneer revenue, advertiser profit, and combined social welfare as a function of number of advertisers for the three quality effect generation methods. As in Section 6.1, the number of contexts is fixed at 20, with 95% singleton bidders and 5% constant.

It is clear that under all three types of quality generation, the auctioneer revenue is always improved through the use of context-based bidding as long as the number of advertisers is not too small with respect to the number of contexts. The difference in revenue between context-based auctions and standard auctions is most dramatic when quality scores are positively correlated with value or are uniformly 1. (When quality scores are negatively correlated with value, the auctioneer revenue tends to be low either way.)

Finally, we tried to model the situation in which some contexts are inherently more valuable than others. For example, the context of “IP address from a well-off neighborhood” might be more valuable to retailers on the whole than the context of “IP address from an average neighborhood.” The number of bidders is fixed at 30, and the number of contexts at 10. Now 10%, 50%, or 90% of the bidders are singletons, and the rest are random, all with uniform quality.

Figure 4 shows the auctioneer revenue, advertiser profit, and combined social welfare as a function of the number of contexts with inherently low values. In all cases both the revenue to the auctioneer and the advertiser profit are higher under context-based bidding. The biggest increase in advertiser profit occurs when most bidders are singletons, while bigger increases in auctioneer revenue occur when more are random.

### 6.4 An Experiment on Real Data

In our final simulation, we generated parameter values based on data from Yahoo's Search Marketing program. This data was gathered as part of a separate project [5] in November 2006, and contains advertiser bids on the term “used car” in conjunction with the names of 18 states, e.g. bids on the search terms “used car Pennsylvania” and “used car New York.” Using this data, we estimated the true value of each bidder for each state context, setting the value to 0 if the advertiser chose not to bid on “used car” in conjunction with a particular state. We set the probability $P(c)$ of each context $c$ to be proportional to the product of the state's population and its median income. At the time when the data was recorded, Yahoo did not yet incorporate quality into their rating system, so we chose to use a constant quality effect for all advertisers.

While one should be careful not to read too much into a single simulation in which many parameter values must be set somewhat arbitrarily, the effects of allowing context-based bidding are striking and demonstrate the potential of context-based sponsored search. The revenue to the auctioneer increased by 79% under low SNE and by 80% under high SNE when context-based bidding was introduced. The advertiser profit increased by 22%. Overall, the combined social welfare increased by 99%.

## 7. ACKNOWLEDGMENTS

We would like to thank our anonymous reviewers for many helpful comments and in particular for pointing out the connection between our work and the existing economic literature on bundling.

## 8. REFERENCES

[1] C. Borgs, J. Chayes, O. Etesami, N. Immorlica, K. Jain, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th International World Wide Web Conference*, 2007.

[2] N. Brooks. The Atlas rank report: How search engine rank impacts traffic. Accessed at `http://atlassolutions.com/pdf/RankReport.pdf` in February 2007.

[3] I. Chakraborty. Bundling decisions for selling multiple objects. *Economic Theory*, 13:723–733, 1999.

[4] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 2006. Forthcoming.

[5] K. Ganchev, A. Kulesza, J. Tan, R. Gabbard, Q. Liu, and M. Kearns. Empirical price modeling for sponsored search. Paper to be presented at the Third Workshop on Sponsored Search Auctions, Banff, Canada, May, 2007.

[6] G. Hardy, J. E. Littlewood, and G. Polya. *Inequalities*. Cambridge University Press, 2nd edition, 1952.

[7] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, 2006.

[8] T. Palfrey. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51(2):463–483, 1983.

[9] D. Parkes and T. Sandholm. Optimize-and-dispatch architecture for expressive ad auctions. Paper presented at the First Workshop on Sponsored Search Auctions, Vancouver, Canada, June, 2005.

[10] H. Varian. Position auctions. *International Journal of Industrial Organization*, 2006. Forthcoming.

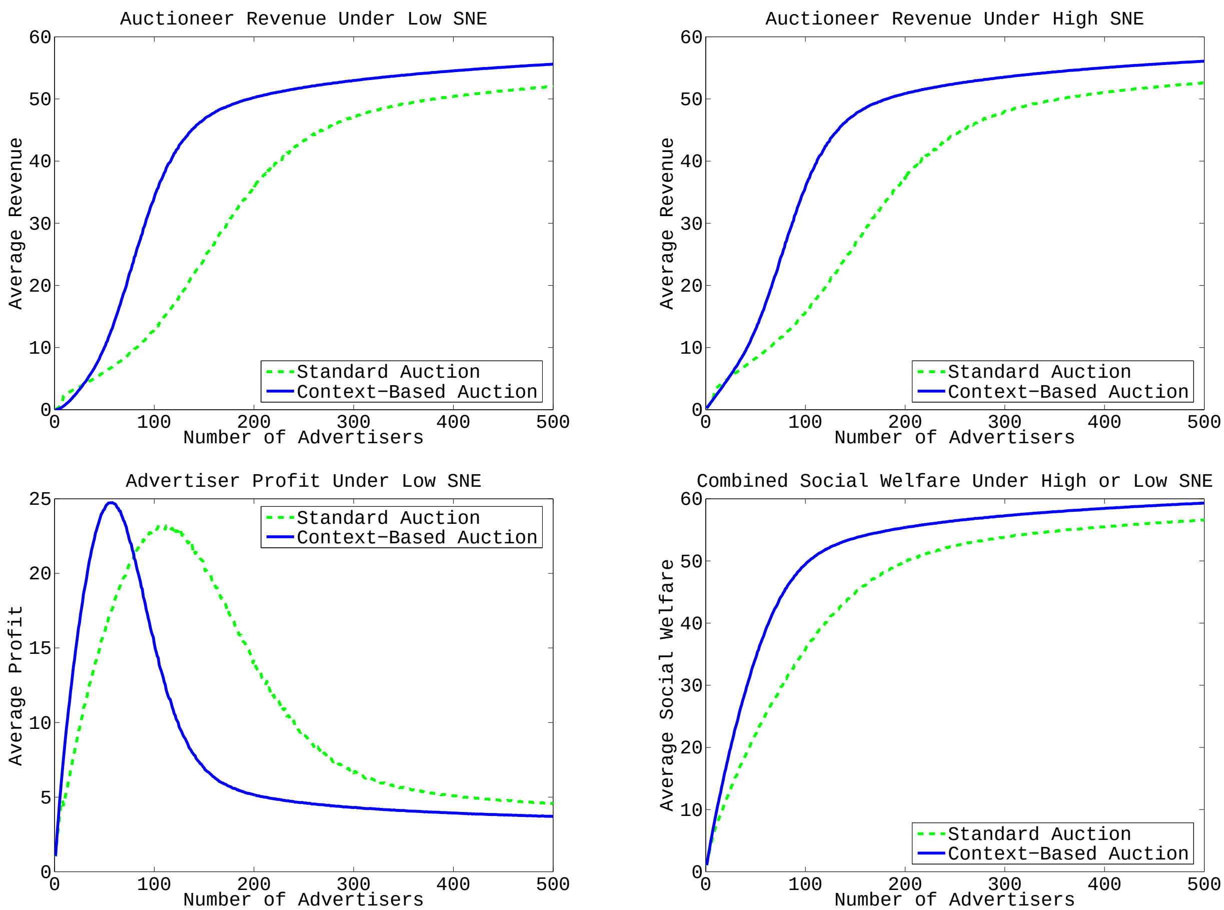


**Figure 1: Varying the number of bidders**

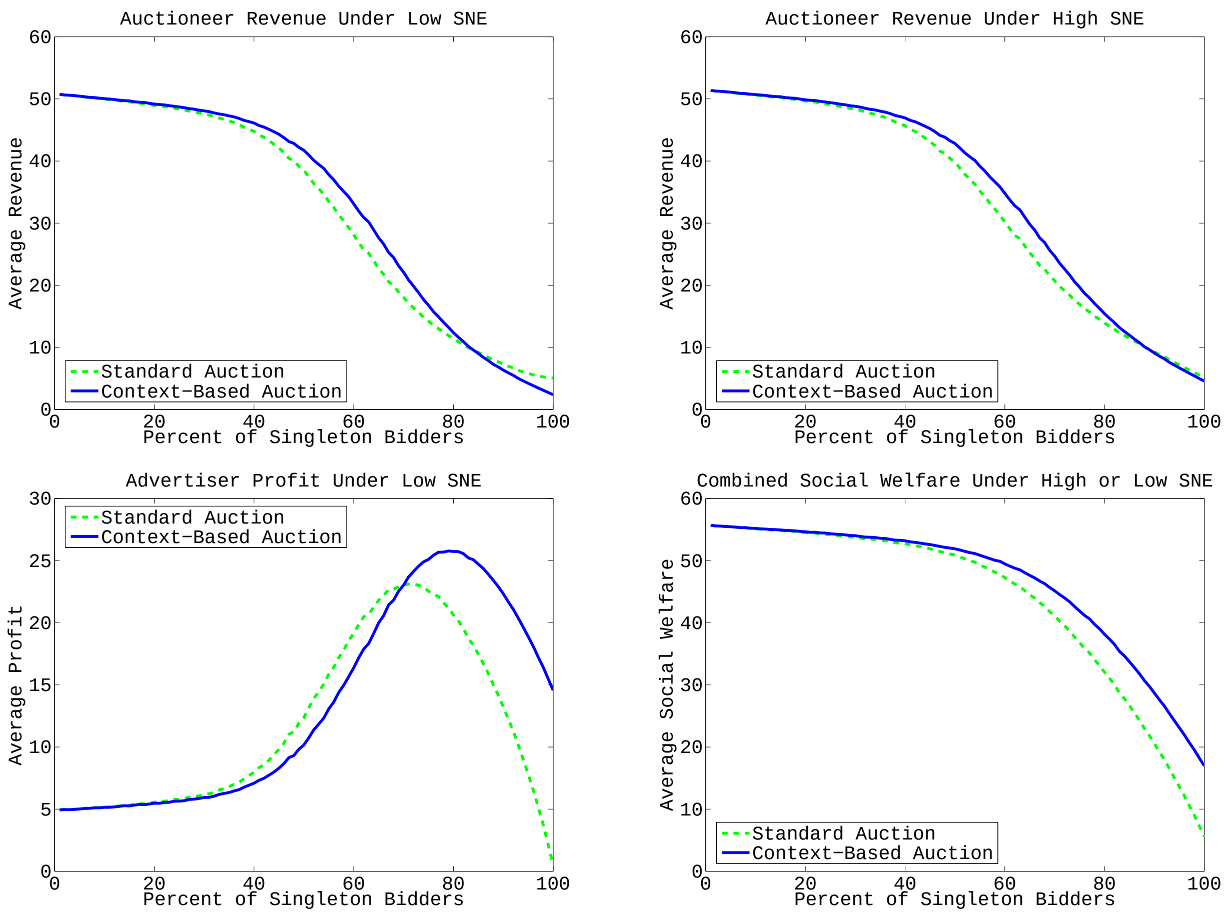


**Figure 2: Varying the fraction of singleton bidders**

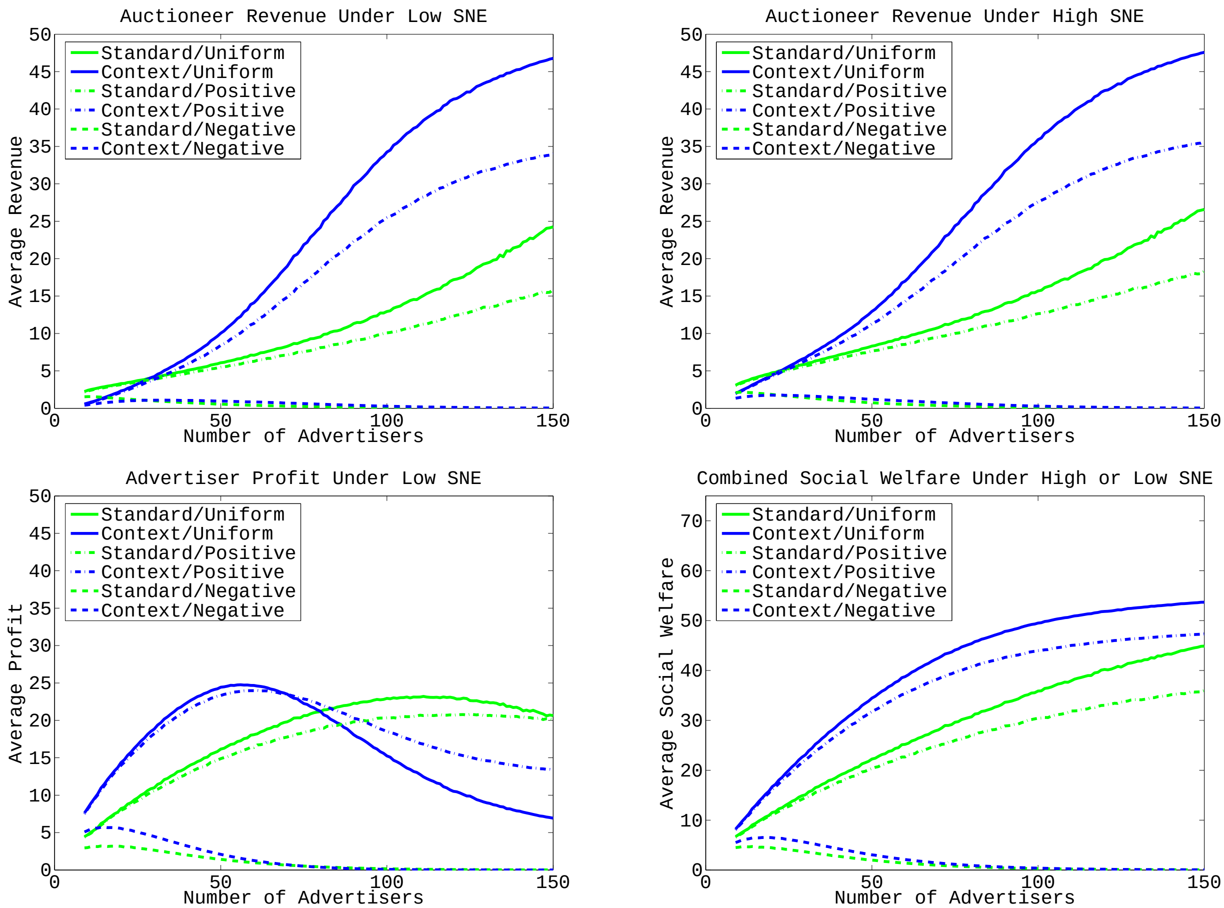


**Figure 3: Varying the way in which quality scores are generated**

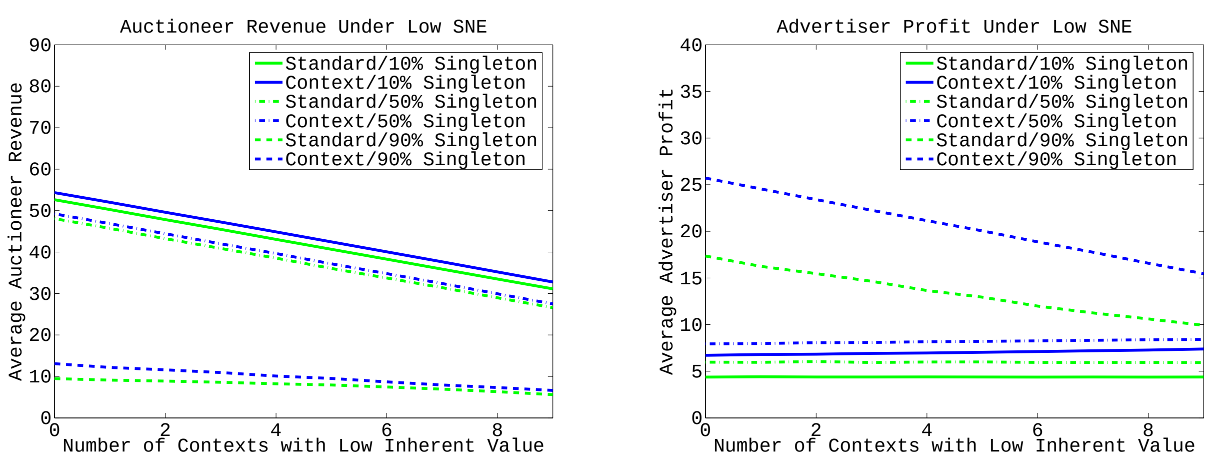


**Figure 4: Varying the inherent values of contexts**