

# Is Efficiency Expensive?

Tim Roughgarden \*  
Department of Computer Science, Stanford  
University, 462 Gates Building, 353 Serra Mall,  
Stanford, CA 94305.

tim@cs.stanford.edu

Mukund Sundararajan †  
Department of Computer Science, Stanford  
University, 470 Gates Building, 353 Serra Mall,  
Stanford, CA 94305.

mukunds@cs.stanford.edu

## ABSTRACT

We study the simultaneous optimization of efficiency and revenue in pay-per-click keyword auctions in a Bayesian setting. Our main result is that the efficient keyword auction yields near-optimal revenue even under modest competition. In the process, we build on classical results in auction theory to prove that increasing the number of bidders by the number of slots outweighs the benefit of employing the optimal reserve price.

## 1. INTRODUCTION

What objective should a keyword auction optimize? As most search engines are controlled by public companies, their primary responsibility is to maximize revenue and create value for their stockholders. On the other hand, from a social standpoint, we prefer an auction that optimizes social efficiency—that is, an auction that maximizes the total value to its participants. Can both objectives be optimized simultaneously?

Consider the following motivating example. Suppose Alice has an object, such as a cell phone, for which she has no value. There is one potential customer, Bob, with a non-negative value  $v_b$  for the cell phone. For Alice to optimize social efficiency, she must allocate the cellphone to Bob whenever Bob’s value is positive. If Alice does not know  $v_b$ , the only incentive-compatible efficient auction offers the cell phone to Bob for free. Of course, Alice makes no revenue from such an auction.

This example suggests that optimizing efficiency may lead to sacrificing revenue completely. Is the friction between these objectives typically so severe? In this paper we study the tension between revenue and efficiency in keyword auctions. Call the revenue extracted by the revenue-maximizing auction the *optimal revenue*. We ask:

What fraction of the optimal revenue does an efficient search auction extract?

\*Supported in part by ONR grant N00014-04-1-0725, an NSF CAREER Award, and an Alfred P. Sloan Fellowship.

†Supported in part by OSD/ONR CIP/SW URI “Software Quality and Infrastructure Protection for Diffuse Computing” through ONR Grant N00014-01-1-0795 and by OSD/ONR CIP/SW URI “Trustworthy Infrastructure, Mechanisms, and Experimentation for Diffuse Computing” through ONR Grant N00014-04-1-0725.

We adopt the efficient auction as our protagonist (over the optimal one) for two reasons. First, an optimal auction typically sets a reserve price based on information about the distribution of bidder valuations. Such information is not always available, and even when it is, collecting and processing it involves non-trivial effort. Efficient auctions require no such information and are simpler to run. Second, optimal auctions only make sense in monopoly settings. This assumption does not hold, for instance, in the search market. On the flip side, the obvious case *against* an efficient keyword auction is that it does not optimize revenue. This paper shows that, besides their other laudable properties, efficient keyword auctions often guarantee *near-optimal revenue*.

## 1.1 Results and Techniques

We assume that the bidders’ values are drawn i.i.d. from a known distribution  $D$ . The value represents the maximum amount that a bidder is willing to pay for a click on its advertisement. Suppose that the auction has  $n$  bidders and  $k$  slots. (The slots are not identical and are parameterized by *click-through-rates*; see Section 2.1.) Under fairly general assumptions on the distribution  $D$ , we give two guarantees on the expected revenue achieved by an efficient auction, relative to the revenue-maximizing (or *optimal*) auction. First, we show that the revenue earned by an efficient auction *with  $k$  additional bidders* exceeds the revenue from the optimal auction. This shows that modest extra competition is as valuable as precisely learning the distribution  $D$  and employing the optimal reserve price. We also give an explicit comparison between the revenue of the two auctions—we prove that the efficient auction gives a  $(1 - \frac{k}{n})$ -approximation of the optimal revenue.

The first result builds on techniques from Bulow and Klemperer [2], who studied single and multi-item auctions. Our second result hinges on quantitative bounds on the increase in the revenue of efficient keyword auctions as the number of bidders increases. This, in turn, leads us to study the behavior of order statistics of the distribution  $D$ . The challenging aspect is to prove bounds that are meaningful for a modest number of bidders and also do not make strong assumptions about the distribution  $D$ . Qualitatively, our analysis shows that among distributions that satisfy the well-known monotone hazard condition, exponential distributions exhibit worst-case behavior.

## 1.2 Related Work

There have been several previous theoretical analyses studying various aspects of keyword auctions. Much of this work

is surveyed by Lahaie et al. [8]. For example, work by Varian [15], Edelman et al. [3, 4] and Agarwal et al. [1] study bidding strategies and equilibria. Work by Mehta et al. [12] studies revenue maximization with budgets but without incentive constraints. Feng et al. [5] show how to calculate reserve prices that increase revenue and quantify this improvement when bidders’ valuations are drawn i.i.d. from a uniform distribution. Iyengar and Kumar [7] discuss the format of the optimal auction under various assumptions on the click through rates. Lambert and Shoham [9] show, in a repeated setting, that a certain class of auctions are asymptotically optimal. Liu, Chen, and Whinston [11] view keyword auctions as weighted unit price auctions, study bidding equilibria, prove revenue equivalence results and study the form of the efficient and the optimal auctions.

The following results are more closely related to ours. Likhodedov and Sandholm [10] design multi-item auctions that maximize efficiency subject to achieving specified revenue targets. The focus of [10] is on computing the optimal constrained auction, rather than on deriving analytical bounds on feasible efficiency-revenue trade-offs. The paper of Neeman [14] is philosophically closer to ours, and quantifies revenue loss in an efficient auction in a Bayesian setting. However, Neeman [14] studies only single-item auctions. Also, the distributional assumptions and trade-offs of [14] are incomparable to ours. While our results require a regularity or monotone hazard rate assumption, Neeman [14] assumes that valuations are drawn from a distribution with bounded support, and his results are parameterized by the ratio of the expected value to the maximum value (which is zero for most of the distributions we study).

Finally, we also build on techniques from two classic papers in auction theory—Myerson [13] and Bulow and Klemperer [2].

## 2. PRELIMINARIES

### 2.1 The Model

We examine revenue properties of efficient auctions in a classical Bayesian auction setting studied by Myerson [13]). We study standard pay-per-click keyword auctions (see for instance [1]). We use the terms *bidder* and *advertiser* interchangeably. Details of the model follow.

In the standard pay-per-click model, the  $n$  bidders are advertisers who compete to have their advertisement displayed in one of  $k$  slots. Such an auction is run by the search engine on the event of a search query. When the advertisement of a bidder  $i$  is placed in a slot  $j$ , the estimate of the probability of a future click on the advertisement is modeled by the *click-through rate*  $CTR_{i,j}$ . We assume that these estimates are accurate. A common assumption in both theory and practice is that the CTRs are *separable* [1]—that is, for all bidders  $i$  and slots  $j$ ,  $CTR_{i,j}$  is  $\mu_i \cdot \Theta_j$ . We also make the natural assumption that  $\Theta_j \geq \Theta_{j+1}$  for every slot  $j$ . In this paper, we primarily focus on the case where all ads are equally relevant, in the sense that all  $\mu_i$ ’s are equal. (The effect of relaxing this assumption is discussed in Section 4.2.)

We assume that  $n$  symmetric bidders with non-negative valuations  $v_i$  are drawn independently from a known distribution  $D$ . Here  $v_i$  represents the value that the advertiser has for a click on its advertisement—this value may represent the profit that the advertiser expects to make on a subsequent sale, with the probability of a sale appropriately

factored in. The value  $v_i$  is private to the bidder. The distribution  $D$  is described by a probability density function  $f$  and a cumulative distribution function  $F$ . We assume that bidders’ utility functions are quasilinear—if a bidder  $i$  is allocated  $x_i$  clicks, its utility is  $x_i \cdot v_i - p_i$ , where  $p_i$  is the total amount it pays. (The amount paid per click is  $p_i/x_i$ .) See Section 4.2 for further discussion of these assumptions.

As stated in the introduction, we look at two auction objectives. The first is efficiency, defined as the sum of the value served:  $\sum_i v_i \cdot x_i$ . The second objective is revenue, defined as the amount that the auctioneer receives:  $\sum_i p_i$ . As in Myerson [13], we are interested in the expected revenue, where the expectation is over the product distribution arising from the  $n$  valuations drawn i.i.d. from the distribution  $D$ . Also, we restrict our attention to auctions that are Bayesian incentive-compatible, where truth-telling forms a Bayes-Nash equilibrium.

We assume throughout that the distribution  $D$  satisfies the following *regularity condition*. Following Myerson [13], we define the virtual valuation  $\mathfrak{v}$  of the distribution  $D$  as:

$$\mathfrak{v}(v) = v - \frac{1 - F(v)}{f(v)}. \quad (1)$$

By definition, our regularity condition asserts that the virtual valuation is strictly increasing in  $v$ . Under this assumption, all of the auctions that we consider are strategyproof. We will sometimes strengthen this regularity assumption and require that  $D$  satisfies the *monotone hazard rate condition*. Define the hazard rate  $h$  as

$$h(v) = \frac{f(v)}{1 - F(v)}. \quad (2)$$

The monotone hazard rate condition states that the hazard rate is non-decreasing in  $v$ . Both assumptions are common in auction theory [2, 13]. For example, Gaussian, uniform, and exponential distributions all satisfy the monotone hazard rate condition.

### 2.2 Useful Results from Auction Theory

Myerson [13] discusses the form of the revenue-maximizing single-item auction. An intermediate result of [13] applies to general single-parameter settings including ours. An auction can be viewed as an allocation rule—which maps bids to allocations—and a payment rule—which maps bids to payments. We define a *monotone* allocation rule as one in which the expected number of clicks that a bidder receives is non-decreasing in its bid, where the expectation is over the values of the other bidders.

*Definition 1.* An allocation rule is *monotone* if for every bidder  $i$  and values  $v_i \geq v'_i$ ,

$$\int_{v_{-i}} x_i(v_i, v_{-i}) f_{-i}(v_{-i}) dv_{-i} \geq \int_{v_{-i}} x_i(v'_i, v_{-i}) f_{-i}(v_{-i}) dv_{-i} \quad (3)$$

Myerson [13, Lemma 3.1] shows that every Bayesian incentive-compatible auction must satisfy the following three properties.

- It has a monotone allocation rule.
- The allocation rule determines the payments (up to a pivot term), with player  $i$ ’s payment given by

$$p_i(v_i, v_{-i}) = x_i(v_i, v_{-i}) \cdot v_i - \int_0^{v_i} x_i(v, v_{-i}) dv. \quad (4)$$

- The expected revenue of the auctioneer can be expressed as a function of the allocation rule:

$$\int_{\mathcal{V}} \left( \sum_i \mathbf{v}(v_i) x_i(v_1, \dots, v_n) \right) f(v_1, \dots, v_n) d(v_1, \dots, v_n) \quad (5)$$

Rewriting the expression in (5) in terms of virtual valuations (1) shows that the expected revenue of an auction is the expected virtual valuation served. The optimal auction selects a monotone allocation rule that maximizes the expected virtual valuation.

### 3. REVENUE PROPERTIES OF EFFICIENT KEYWORD AUCTIONS

Before we compare the revenue of the optimal and efficient auctions, we briefly consider their allocation rules. The efficient auction is discussed in detail in [1], while the optimal auction is discussed in [7]. As the auctions that we discuss are Bayesian incentive-compatible, we use the terms *bids* and *valuations* interchangeably.

Sort the bidders in non-increasing order of the bids. By the regularity condition, sorting by virtual valuations gives the same result. Recall that for all slots  $j$ , the CTR  $\Theta_j$  is at least  $\Theta_{j+1}$ . As the efficient auction attempts to maximize the total value— $\sum_{1 \leq i \leq k} \Theta_i \cdot v_i$ —it assigns the  $i^{\text{th}}$  bidder to the  $i^{\text{th}}$  slot. The optimal auction, which maximizes the total virtual valuation, allocates in the same way *after* ignoring bids with negative virtual valuations. Ignoring bids with negative virtual valuations corresponds to ignoring bids less than a reserve price  $\mathbf{v}^{-1}(0)$ . (The inverse function  $\mathbf{v}^{-1}$  is well defined as  $\mathbf{v}$  is strictly increasing.) The next two facts follow from this discussion.

*Fact 1.* The expected revenue of the efficient auction is the weighted sum of the top  $k$  virtual valuations. It is the expected value of  $\sum_{1 \leq i \leq k} \mathbf{v}(v_{m_i}) \cdot \Theta_i$ , where  $m_i$  denotes the index of the bidder with the  $i^{\text{th}}$  highest (virtual) valuation.

*Fact 2.* The expected revenue of the optimal auction is the expected value of  $\sum_{1 \leq i \leq \min(k, l)} \mathbf{v}(v_{m_i}) \cdot \Theta_i$ , where  $m_i$  denotes the index of the bidder with the  $i^{\text{th}}$  highest (virtual) valuation, and  $l$  is the largest value such that  $\mathbf{v}(v_{m_l}) \geq 0$ .

#### 3.1 Increased Competition vs. an Optimal Reserve Price

We next use Facts 1 and 2 to prove the following theorem.

**THEOREM 1.** *The expected revenue of the efficient keyword auction with  $n + k$  bidders is at least the expected revenue of the optimal auction with  $n$  bidders.*

Before we prove the theorem, we discuss its implications. Theorem 1 compares the efficacy of two ways that a search engine can improve its revenue. The search engine can collect information to learn the distribution  $D$  and calculate the optimal reserve price  $\mathbf{v}^{-1}(0)$ . Alternatively, it can expand its market and run the efficient auction, which does not require knowledge of  $D$ . Theorem 1 implies that enlarging the market by the number of slots outweighs the benefit of running an optimal auction. Bulow and Klemperer [2] proved an analogous theorem for auctions with identical goods.

**PROOF.** First, we show that the expected virtual valuation of a bidder is 0. From (1), the expected virtual valuation of a bidder is

$$\int_0^\infty \mathbf{v}(v) \cdot f(v) \cdot dv = \int_0^\infty \left( v - \frac{1 - F(v)}{f(v)} \right) \cdot f(v) \cdot dv = 0,$$

where in the second equality we use the identity  $\int_0^\infty v \cdot f(v) \cdot dv = \int_0^\infty (1 - F(v)) \cdot dv$  for non-negative random variables.

Let  $W(a_1 \cdots a_{n+k})$  denote the weighted sum of the top  $k$  numbers among  $a_1 \cdots a_{n+k}$ , where the weight associated with the  $i^{\text{th}}$  highest number is  $\Theta_i$ . By Fact 1, the expected revenue of the efficient auction with  $n + k$  bidders is

$$E[W(\mathbf{v}(v_1) \cdots \mathbf{v}(v_{n+k}))] = E[E[W(\mathbf{v}(v_1) \cdots \mathbf{v}(v_{n+k})) | v_1 \cdots v_n]]. \quad (6)$$

Let  $l$  denote the random variable equal to the number of the first  $n$  bidders with non-negative virtual valuations. For  $i \in \{1, 2, \dots, n\}$ , let  $m_i \in \{1, 2, \dots, n\}$  denote the index of the bidder with the  $i^{\text{th}}$  highest (virtual) valuation (among the first  $n$  bidders). We can lower bound the right-hand side of (6) by assuming that the efficient auction allocates the first  $\min\{k, l\}$  slots to the  $\min\{k, l\}$  highest bidders among the first  $n$ , and the remaining slots to the bidders  $n + 1, n + 2, \dots, n + (k - \min\{k, l\})$ ; the efficient auction always chooses an allocation that is at least this good. Precisely, we have  $E[W(\mathbf{v}(v_1) \cdots \mathbf{v}(v_{n+k}))] \geq$

$$E \left[ E \left[ \left( \sum_{i=1}^{\min\{k, l\}} \Theta_i \cdot \mathbf{v}(v_{m_i}) + \sum_{i=1}^{k - \min\{k, l\}} \Theta_{l+i} \cdot \mathbf{v}(v_{n+i}) \right) \middle| v_1 \cdots v_n \right] \right]$$

By the mutual independence of the different valuations, the fact that the expected virtual valuation of a single bidder is 0, and linearity of expectations, the right-hand side of the above inequality is equal to

$$E \left[ E \left[ \sum_{i=1}^{\min\{k, l\}} \Theta_i \cdot \mathbf{v}(v_{m_i}) \middle| v_1 \cdots v_n \right] \right]$$

By Fact 2, this is the expected revenue of the optimal  $k$ -slot keyword auction with  $n$  bidders. This completes the proof.  $\square$

One drawback of Theorem 1 is that it does not directly compare the revenue obtained by the efficient and optimal auctions in the same environment. We address this issue in the next section.

#### 3.2 Approximation Bounds

How can we directly compare the revenue generated by the efficient and optimal auctions? Theorem 1, which shows that the revenue of the efficient auction with  $n + k$  players is at least the revenue of the optimal auction with  $n$  bidders, suggests a potential approach. If the efficient auction with  $n$  players collects at least a  $c$  fraction of the revenue of the efficient auction with  $n + k$  bidders, then the efficient auction also  $c$ -approximates the optimal revenue (with  $n$  bidders). We use this idea to prove the following revenue guarantees for the efficient keyword auction.

**THEOREM 2.** *Suppose the distribution  $D$  of valuations is regular. Then the expected revenue of the efficient keyword auction with  $k$  slots and  $n$  bidders is at least  $(1 - k \cdot (k + 1) / n)$  times that of the optimal auction.*

**THEOREM 3.** *Suppose the distribution  $D$  of valuations satisfies the monotone hazard rate. Then the expected revenue of the efficient keyword auction with  $k$  slots and  $n$  bidders is at least  $(1 - k/n)$  times that of the optimal auction.*

As expected, these theorems confirm the intuition that the revenue of the efficient auction approaches that of the optimal one as the number of bidders tends to infinity. But Theorems 2 and 3 show something much stronger: the efficient auction obtains near-optimal revenue even in the practically important case of a modest number of bidders—as long as the number of bidders is a small multiple of the number of slots, the revenue is close to optimal. Qualitatively, these theorems imply that distributional knowledge and reserve prices have a negligible effect on auction revenue when there is at least moderate competition (as for popular keywords such as “camera” and “laptop”).

We now provide proofs for Theorems 2 and 3. Thus far, we viewed the expected revenue of the efficient auction as the expected value of the weighted sum of the top  $k$  virtual valuations (Fact 1). In this section, we use equation (4) instead. Assume that the bidders are sorted in non-decreasing order of values. The allocation rule of the efficient auction and equation (4) imply that the payment of the  $i^{\text{th}}$  bidder is

$$p_i = \sum_{j=i}^K (\Theta_j - \Theta_{j+1}) \cdot v_{j+1}. \quad (7)$$

A similar expression was established in [1]. Equation (7) shows that for all  $i$  and  $j \geq i$ , bidder  $i$  pays the  $j + 1^{\text{th}}$  highest bid for the marginal clicks  $\Theta_j - \Theta_{j+1}$ . Using the above equation, the total revenue of the auctioneer is

$$\begin{aligned} \sum_{1 \leq i \leq k} p_i &= \sum_{1 \leq i \leq k} \sum_{j=i}^k (\Theta_j - \Theta_{j+1}) \cdot v_{j+1} \\ &= \sum_{1 \leq j \leq k} (\Theta_j - \Theta_{j+1}) \cdot j \cdot v_{j+1}. \end{aligned}$$

*Fact 3.* The expected revenue of the efficient auction is a weighted sum of the expected values of  $k$  order statistics of  $n$  i.i.d. samples from the distribution  $D$ . The  $k$  order statistics range from the  $2^{\text{nd}}$  highest to the  $k + 1^{\text{th}}$  highest number.

We aim to show that the efficient auction with  $n$  players collects at least a  $c$  fraction of the revenue of the efficient auction with  $n + k$  players. By Fact 3, we need only show that the expected value of the  $i^{\text{th}}$  highest number of  $n$  samples is at least  $c$  times that of the  $i^{\text{th}}$  highest number of  $n + k$  samples (where all samples are drawn i.i.d. from the distribution  $D$ ). We begin with the regular case (Theorem 2).

**LEMMA 1.** *Let  $D$  be a regular distribution and  $l \in \{2, \dots, k+1\}$ . The expected value of the  $l^{\text{th}}$ -highest number among  $n$  samples is at least  $(1 - \frac{k-l}{n})$  times that of the  $l^{\text{th}}$ -highest number among  $n + k$  samples, when both sets of samples are drawn i.i.d. from distribution  $D$ .*

**PROOF.** We view the two sample sets in the following way. We first draw  $n + k$  i.i.d. samples from  $D$ . We then permute the indices of these samples randomly (since the samples are i.i.d., this does not affect the distribution). Define  $X$  to be the random variable equal to the  $l^{\text{th}}$  highest value among the

$n + k$  samples. Define  $Y$  to be zero whenever any of the  $l$  highest values occur among the final  $k$  samples (after the random permutation), and equal to  $X$  otherwise. Since  $Y$  is either the  $l^{\text{th}}$  highest value among the first  $n$  samples or zero, its expectation is a lower bound on the expectation of the  $l^{\text{th}}$  highest value among  $n$  i.i.d. samples.

Condition on the outcome of the first step of the random process, thereby fixing the value of  $X$ . The conditional expected value of  $Y$  is  $X \cdot \Pr[\mathcal{E}]$ , where  $\mathcal{E}$  is the event that none of the  $l$  highest samples are mapped to the last  $k$  indices. By the Union Bound, this occurs with probability at least  $1 - (kl/n)$ . Taking expectations gives  $E[Y] \geq (1 - kl/n) \cdot E[X]$ , which completes the proof.  $\square$

Theorem 1, Fact 3, and Lemma 1 now give Theorem 2.

We next discuss the proof of Theorem 3. We start by showing that among distributions that satisfy the monotone hazard condition, the increase in revenue from additional bidders is maximized by exponential distributions.

**LEMMA 2.** *Let  $D$  be a distribution that satisfies the monotone hazard condition. The ratio between the expected value of the  $l^{\text{th}}$ -largest of  $n$  samples and the expected value of the  $l^{\text{th}}$ -largest of  $n + k$  samples (all i.i.d. from  $D$ ) is minimized when  $D$  is an exponential distribution.*

The proof of this lemma is technical and can be found in the Appendix. The intuition behind the proof is that distributions with long tails minimize the ratio. We can write the distribution function  $F$  in terms of the hazard rate of the distribution, via  $F(v) = 1 - \exp\{-\int_0^v h(v)dv\}$ . Thus, increasing the hazard rate effectively reduces the mass in the tail. Exponential distributions, which have constant hazard rates, have the longest tails among distributions that satisfy the monotone hazard rate condition.

Lemma 2 justifies restricting attention to exponential distributions, for which there are closed form formulas for order statistics.

*FactEfi[16, 6].* The expected value of the  $k^{\text{th}}$ -largest value of  $n$  samples drawn i.i.d. from an exponential distribution with rate  $\lambda$  is  $(H_n - H_{k-1})/\lambda$ , where  $H_i = \sum_{j=1}^i 1/j$  denotes the  $i^{\text{th}}$  Harmonic number.

Lemma 2 and Fact 4 now imply the following.

**LEMMA 3.** *Let  $D$  be a distribution that satisfies the monotone hazard condition. The expected value of the  $l^{\text{th}}$ -largest value of  $n$  samples is at least  $(H_n - H_{l-1})/(H_{n+k} - H_{l-1})$  times that of the  $l^{\text{th}}$ -largest value of  $n + k$  samples (all i.i.d. draws from  $D$ ).*

Theorem 3 now follows easily from Theorem 1, Fact 3, and Lemma 3.

We finish this section by stating a stronger version of Theorem 3. The theorem uses the optimal efficiency as a benchmark, instead of the optimal revenue (which can only be less, assuming individual rationality).

**THEOREM 4.** *Suppose the distribution  $D$  of valuations satisfies the monotone hazard rate. Then the expected revenue of the efficient keyword auction with  $k$  slots and  $n$  bidders is at least  $(1 - k/n)$  times the optimal efficiency.*

By Fact 3, the expected value of the  $k + 1^{\text{th}}$  highest of  $n$  numbers times the total click-through-rate is a lower bound on the expected revenue. The proof of the above theorem is now an easy implication of Lemma 3 and the following fact.

*Fact 5.* The optimal efficiency is a weighted sum of the expected values of  $k$  order statistics of  $n$  samples taken i.i.d. from the distribution  $D$ . The  $k$  order statistics range from the  $1^{\text{st}}$  highest to the  $k^{\text{th}}$  highest number. The weight corresponding to the  $i^{\text{th}}$  highest number is the CTR  $\Theta_i$ .

## 4. DISCUSSION

### 4.1 Efficiency Properties of Optimal Keyword Auctions

We briefly consider efficiency properties of optimal keyword auctions. The following theorem asserts that under modest competition, revenue maximizing auctions yield near optimal efficiency. Taken together with Theorem 1, Theorem 2, and Theorem 3, it asserts that under modest competition, the objectives of efficiency and revenue are indeed well aligned. The theorem is an easy implication of results from the previous section.

*THEOREM 5.* Suppose the distribution  $D$  of valuations satisfies the monotone hazard rate. Then the expected efficiency of the optimal keyword auction with  $k$  slots and  $n$  bidders is at least  $(1 - k/n)$  times the optimal efficiency.

*PROOF.* The revenue of the efficient auction is at most the revenue of the optimal auction, which is the at most the efficiency of the optimal auction. Applying Theorem 4 completes the proof.  $\square$

### 4.2 Modeling Assumptions

In this section we discuss the sensitivity of our results to various modeling assumptions.

First, we restrict attention to Bayesian incentive-compatible auctions—we compare the revenue of the efficient auction to the optimal one in this class. Most real-world keyword auctions are not incentive compatible (see [1, 15, 4]) and generally can have multiple Nash equilibria. On the other hand, Aggarwal et al. [1] show that, in these non-truthful auctions, there is always an equilibrium that is revenue-equivalent to the truthful auction. While this only shows that the efficient auction is competitive with *some* equilibrium of these auctions, we expect that other equilibria will have similar revenue properties. Second, we have assumed that the CTR of a slot is independent of the advertiser that is awarded it. We can extend our results to advertiser-dependent CTRs, although we lose a factor of the ratio  $\mu_{\min}/\mu_{\max}$  in our bounds. Third, our results depend on the number of bidders being larger than the number of slots. On the other hand, in practice there are many keywords with very few bidders (although it is not clear that these auctions account for a large fraction of a search engine’s revenue). Conceivably, reserve prices play a more significant role in auctions for these unpopular keywords. Fourth, we assume that the valuations are drawn identically from a distribution  $D$  (which may or may not be known)—for a fixed keyword, we think that this assumption is a reasonable approximation of reality. Finally, all of our results are for a single-shot setting—it seems challenging to derive similar results in a repeated setting.

## 4.3 Multi-item Auctions

The results of Section 3.1 imply approximation bounds for single and multi-item auctions as a special case.

*THEOREM 6.* If the valuations are picked i.i.d. from a distribution  $D$  that satisfies the monotone hazard rate condition, the expected revenue earned by an efficient  $k$ -item auction with  $n$  bidders is at least a  $(1 - k/n)$  fraction of the expected revenue generated by the optimal auction with  $n$  bidders. If we relax the condition to regularity, then the auction obtains a  $1 - (k \cdot (k + 1))/n$  fraction of the optimal revenue.

The above theorem can be viewed as a companion to a theorem by Bulow and Klemperer [2], which states that the revenue of the efficient  $k$ -item auction with  $n + k$  bidders is at least the revenue of the optimal  $k$ -item auction with  $n$  bidders. Theorem 6 clarifies the effect of the addition of  $k$  extra bidders—it shows that even without these additional bidders, if the number of bidders modestly exceeds the number of items, then the efficient auction approximates the optimal revenue.

We conclude with a correspondence between multi-item and keyword auctions. We show that the optimal keyword auction is essentially the superposition of multiple optimal  $k$ -item auctions, while the efficient keyword auction is the superposition of multiple efficient  $k$ -item auctions. We prove the lemma generally for any auction that sets a reserve price  $r$  and allocates the slots in order of non-increasing bid to the bidders that bid at least  $r$ . The efficient auction uses a reserve  $r = 0$ , while the optimal auction uses a reserve  $r = v^{-1}(0)$ .

*LEMMA 4.* The expected revenue from a keyword auction with reserve price  $r$ ,  $k$  slots, and  $n$  bidders with valuations drawn i.i.d. from  $D$ , is equal to the weighted sum of revenues from  $k$  multi-item auctions, each with reserve price  $r$  and  $n$  bidders with valuations drawn i.i.d. from  $D$ . The  $i^{\text{th}}$  multi-item auction sells  $i$  objects.

*PROOF.* Fix the valuations of the  $n$  bidders. Let  $l$  be the last bidder in this sequence with  $v_l \geq r$ ; if there is no such bidder,  $l = 0$ . Let  $k' = \min(k, l)$ . By (7), the total revenue of the auction is

$$\begin{aligned} \sum_{1 \leq i \leq k'} p_i &= \sum_{1 \leq i \leq k'} \sum_{j=i}^k (\Theta_j - \Theta_{j+1}) \max(r, v_{j+1}) \\ &= \sum_{1 \leq j \leq k} (\Theta_j - \Theta_{j+1}) \min(l, j) \max(r, v_{j+1}). \end{aligned}$$

Note that the  $j^{\text{th}}$  summand in the above expression is precisely  $(\Theta_j - \Theta_{j+1})$  times the revenue of a  $j$ -item auction with reserve price  $r$ .  $\square$

## 5. REFERENCES

- [1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7, New York, NY, USA, 2006. ACM Press.
- [2] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–94, March 1996. available at <http://ideas.repec.org/a/aea/aecrev/v86y1996i1p180-94.html>.

- [3] B. Edelman and M. Ostrovsky. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *Decision Support Systems, forthcoming*, 2006.
- [4] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *Stanford GSB Research Papers: Paper 1917*, 2005.
- [5] J. Feng, R. Zhan, and M. Shen. Ranked items auctions and online advertisement. *Special Issue about Auctions at POMS*, 2006.
- [6] J. S. Huang. Characterizations of the exponential distribution by order statistics. *Annals of the Institute of Statistical Mathematics*, 11(3):605–608, sep 1974.
- [7] G. Iyengar and A. Kumar. Characterizing optimal keyword auctions. *Second Workshop on Sponsored Search Auctions*, 2006.
- [8] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. Sponsored search auctions. In N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 28. Cambridge University Press, 2007. To appear.
- [9] N. Lambert and Y. Shoham. Asymptotically optimal repeated auctions for sponsored search. *submitted*, 2007.
- [10] A. Likhodedov and T. Sandholm. Auction mechanism for optimally trading off revenue and efficiency. In *ACM Conference on Electronic Commerce*, pages 212–213, 2003.
- [11] D. Liu, J. Chen, and A. B. Whinston. Weighted unit-price auctions. *submitted*, 2006.
- [12] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized on-line matching. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 264–273, Washington, DC, USA, 2005. IEEE Computer Society.
- [13] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [14] Z. Neeman. The effectiveness of english auctions. *Games and Economic Behavior*, 43(2):214–238, May 2003.
- [15] H. Varian. Position auctions. *To appear in International Journal of Industrial Organization*, 2006.
- [16] M. Xiangwei, G. Jian, and W. Xinzheng. Order statistics for negative exponential distribution and its applications. In *Proceedings of the 2003 International Conference on Neural Networks and Signal Processing, 2003.*, volume 1, pages 720–722, 2003.

## APPENDIX

### A. MISSING PROOFS

Fix a distribution  $D$  with a p.d.f.  $f$  and c.d.f.  $F$  and a hazard function  $h(t) = f(t)/1 - F(t)$ . We assume that  $F$  is strictly increasing and that  $D$  satisfies the monotone hazard condition. We write the c.d.f  $F$  in terms of the hazard function:

$$F(t) = 1 - e^{-\int_0^t h(x)dx} \quad (8)$$

Let the random variable  $D(n, k)$  denote the value of the  $k^{\text{th}}$  highest of  $n$  samples taken i.i.d. from distribution  $D$ . Let

$U$  denote the uniform distribution with support  $[0, 1]$ . We first prove a lemma that allows us to relate order statistics of  $U$  to the order statistics of  $D$ .

LEMMA 5.  $E[D(n, k)] = E[F^{-1}(U(n, k))]$ . The first expectation is over the joint distribution resulting from  $n$  samples drawn i.i.d. from distribution  $D$ , while the second is over the joint distribution resulting from  $n$  samples drawn i.i.d. from  $U$ .

PROOF. Let  $Y_1, \dots, Y_n$  be  $n$  random variables drawn i.i.d. from  $D$ . Set  $X_i = F(Y_i)$ . This defines the random variables  $X_1, \dots, X_n$ . We define the random variables  $Y_1^* \dots Y_n^*$  such that  $Y_k^*$  denotes the  $k^{\text{th}}$  smallest number among  $Y_1, \dots, Y_n$ .  $X_1^* \dots X_n^*$  are defined similarly for the  $X$ 's.

As  $F$  is strictly increasing, the relative order between the  $Y_i$ 's is preserved by the transformation. Also,  $F^{-1}$  is defined because  $F$  is strictly increasing. So  $E[Y_k^*] = E[F^{-1}(X_k^*)]$ , where the second expectation is taken over the induced distribution on the  $X$ 's. We now show that the distribution on the  $X_i$ 's is uniform in the interval  $[0, 1]$ . For any  $x \leq 1, x \geq 0$ ,  $Pr(X_i < x) = Pr(F(Y_i) < x) = Pr(Y_i < F^{-1}(x)) = F(F^{-1}(x)) = x$ . As  $F$  is a c.d.f, for  $x > 1$ ,  $Pr(X < x)$  is 1 and for any  $x < 0$ ,  $Pr(X < x)$  is 0. Noting that  $X_k^* = U(n, k)$  and  $Y_k^* = D(n, k)$  completes the proof.  $\square$

We are now ready to prove Lemma 2. Let  $P_\lambda$  denote an exponential distribution with rate parameter  $\lambda$  such that  $0 < \lambda < \infty$ , with the c.d.f is  $F_\lambda$ . Formally, we show that for all  $\lambda$  such that  $0 < \lambda < \infty$ ,

$$\frac{E[F_\lambda^{-1}(U(n, k))]}{E[F_\lambda^{-1}(U(n+l, k))]} \leq \frac{E[F^{-1}(U(n, k))]}{E[F^{-1}(U(n+l, k))]}$$

PROOF. Let  $X_1 \dots X_n$  denote  $n$  random variables drawn from  $U$ . We condition on values for these random variables. This fixes  $P_\lambda(n, k)$  and  $D(n, k)$ . Now add  $l > 0$  samples  $X_{n+1}, \dots, X_{n+l}$ . We now show that  $P_\lambda(n, k)/E[P_\lambda(n+l, k)] \leq D(n, k)/E[D(n+l, k)]$  for suitable  $\lambda$ . We then show that fixing  $\lambda$  is without loss of generality. We select  $\lambda$  such that  $P_\lambda(n, k) = D(n, k)$ . It suffices to show that  $E[P_\lambda(n+l, k)] \geq E[D(n+l, k)]$ .

We now condition on values of  $X_{n+1}, \dots, X_{n+l}$ . Any values below  $P_\lambda(n+1, k) = D(n+l, k)$ , do not change the value of the  $k^{\text{th}}$  highest number in either distribution. It suffices for us to show for all  $x > U(n, k)$ ,  $F^{-1}(x) \leq F_\lambda^{-1}(x)$ . Alternatively, we can show that for all  $Y \geq F^{-1}(U(n, k))$ ,  $F(Y) \geq F_\lambda(Y)$ . We now prove this claim.

First, by definition of the hazard rate, we can write  $F(t) = 1 - e^{-\int_0^t h(x)dx}$ . As  $F(D(n, k)) = F(P_\lambda(n, k)) = U(n, k)$ ,  $\int_0^{D(n, k)} h(x)dx = \lambda \cdot D(n, k)$ . Also as  $h$  is monotone increasing,  $h(D(n, k)) \geq \lambda$ . Further, for any  $Y \geq D(n, k)$ , as  $h$  is monotone increasing,  $\int_0^Y h(x)dx \geq \int_0^{D(n, k)} h(x)dx + \int_{D(n, k)}^Y h(x)dx \geq \lambda \cdot D(n, k) + \lambda(Y - D(n, k)) = \lambda Y$ . Applying Equation 8 proves the claim.

We now show that fixing  $\lambda$  is without loss of generality. For an exponential distribution with rate  $\lambda$ , it is easy to see that  $F^{-1}(X) = -\log(1 - X)/\lambda$ . We prove the claim pointwise—fix values for random variables  $X_1 \dots X_n$  and  $l$  additional variables  $X_{n+1}, \dots, X_{n+l}$ . This fixes values,  $P_{\lambda_1}(n, k) = F_{\lambda_1}^{-1}(U(n, k))$ ,  $P_{\lambda_2}(n, k) = F_{\lambda_2}^{-1}(U(n, k))$ ,  $P_{\lambda_1}(n+l, k) = F_{\lambda_1}^{-1}(U(n+l, k))$ ,  $P_{\lambda_2}(n+l, k) = F_{\lambda_2}^{-1}(U(n+l, k))$ . We can see that  $P_{\lambda_1}(n, k)/P_{\lambda_1}(n+l, k) = P_{\lambda_2}(n, k)/P_{\lambda_2}(n+l, k)$ . This completes the proof.  $\square$