

# Hybrid Keyword Search Auctions

Ashish Goel<sup>\*</sup>  
 Department of Management Sci. and Engg.  
 Stanford University  
 ashishg@stanford.edu

Kamesh Munagala<sup>†</sup>  
 Department of Computer Science  
 Duke University  
 kamesh@cs.duke.edu

## ABSTRACT

Search auctions have become a dominant source of revenue generation on the Internet. Such auctions have typically used per-click bidding and pricing. We propose the use of hybrid auctions where an advertiser can make a per-impression as well as a per-click bid, and the auctioneer then chooses one of the two as the pricing mechanism. We assume that the advertiser and the auctioneer both have separate beliefs (called priors) on the click-probability of an advertisement. We first prove that the hybrid auction is truthful, assuming that the advertisers are risk-neutral. We then show that this auction is superior to the existing per-click auction in multiple ways:

1. We show that risk-seeking advertisers will choose only a per-impression bid whereas risk-averse advertisers will choose only a per-click bid, and argue that both kind of advertisers arise naturally. Hence, the ability to bid in a hybrid fashion is important to account for the risk characteristics of the advertisers.
2. For obscure keywords, the auctioneer is unlikely to have a very sharp prior on the click-probabilities. In such situations, we show that having the extra information from the advertisers in the form of a per-impression bid can result in significantly higher revenue.
3. An advertiser who believes that its click-probability is much higher than the auctioneer's estimate can use per-impression bids to correct the auctioneer's prior *without incurring any extra cost*.
4. The hybrid auction can allow the advertiser and auctioneer to implement complex dynamic programming strategies to deal with the uncertainty in the click-probability using the same basic auction. The per-click and per-impression bidding schemes can only be used to implement two extreme cases of these strategies.

<sup>\*</sup>Research supported in part by NSF ITR grant 0428868, by gifts from Google, Microsoft, and Cisco, and by the Stanford-KAUST alliance.

<sup>†</sup>Research supported by NSF via a CAREER award and grant CNS-0540347.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2009, April 20–24, 2009, Madrid, Spain.  
 ACM 978-1-60558-487-4/09/04.

As Internet commerce matures, we need more sophisticated pricing models to exploit all the information held by each of the participants. We believe that hybrid auctions could be an important step in this direction. The hybrid auction easily extends to multiple slots, and is also applicable to scenarios where the hybrid bidding is per-impression and per-action (i.e. CPM and CPA), or per-click and per-action (i.e. CPC and CPA).

## Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems; J.4 [Social and Behavioral Sciences]: Economics

## General Terms

Algorithms, Economics, Theory

## Keywords

Keyword auctions, Mechanism Design, Internet

## 1. INTRODUCTION

While search engines had a transformational effect on Internet use and indeed, on human interaction, it was only with the advent of keyword auctions that these search engines became commercially viable. Most of the major search engines display advertisements along with search results; the revenue from these advertisements drives much of the innovation that occurs in search in particular, and Internet applications in general. Cost-per-click (CPC) auctions have evolved to be the dominant means by which such advertisements are sold [11]. An advertiser places a bid on a specific keyword or keyword group. The auctioneer (i.e. the search engine) maintains an estimate of the click-through probability (CTR) of each advertiser for each keyword. When a user searches for a keyword, the first advertising spot is sold to the advertiser which has the highest product of the bid and the CTR; in the event that this advertisement is clicked upon by the user, this advertiser is charged the minimum bid it would have to make to retain its position. The same process is repeated for the next slot, and so on. A full description of the tremendous amount of work related to keyword auctions is beyond the scope of this paper; the reader is referred to the excellent survey by Lahaie et al [6]. Two other alternatives to CPC auctions are widely used:

1. CPM, or Cost Per (thousand) Impressions: The publisher charges the advertiser for every instance of an

advertisement shown to a user, regardless of the click. This is widely used to sell banner advertisements.

2. CPA, or Cost Per Action (also known as Cost Per Acquisition): The publisher charges the advertiser when an actual sale happens. This is widely used in associate programs such as the ones run by Amazon, and by lead generation intermediaries.

The three models are equivalent when precise estimates of the click-through-probability and click-to-sale-conversion probability are known. In the absence of such estimates, CPC has emerged as a good way of informally dividing the risk between the auctioneer and the advertiser: the auctioneer is vulnerable only to uncertainty in its own estimates of CTR, whereas the advertiser is vulnerable only to uncertainty in its own estimates of the click-to-sale-conversion probability, assuming its advertisement gets displayed in a favorable spot. A great deal of effort has gone into obtaining good predictions of the CTR. The problem is made harder by the fact that many keywords are searched for only a few times, and typical CTRs are low. Advertisers often want to also bid by customer demographics, which further exacerbates the sparsity of the data. Hence, there has to be a great reliance on predictive models of user behavior and new ads (e.g. see [10]). Arguably, another approach would be to devise pricing models that explicitly take the uncertainty of the CTR estimates into account, and allow advertisers and auctioneers to jointly optimize over this uncertainty. In general, we believe that as Internet commerce matures, we need not just better estimation and learning algorithms but also more sophisticated pricing models to exploit all the information held by each of the participants.

In this paper, we propose the use of hybrid auctions for search keywords where an advertiser can make a per-impression as well as a per-click bid, and the auctioneer then chooses one of the two as the pricing mechanism. Informally, the per-impression bids can be thought of as an additional signal which indicates the advertiser's belief of the CTR. This signal may be quite valuable when the keyword is obscure, when the advertiser is aggregating data from multiple publishers or has a good predictive model based on domain knowledge, and when the advertiser is willing to pay a higher amount in order to perform internal experiments/keyword-selection. We assume that the advertiser and the auctioneer both have separate beliefs (called priors) on the click-probability of an advertisement.

## 1.1 Our results

We describe the hybrid auction in section 2, where we also outline the strategic model (that of discounted rewards) used by the auctioneer and the advertiser. We introduce the multi-armed bandit problem as it occurs naturally in this context. Our results, which we have already described at a high level in the abstract, are split into two parts.

**Myopic advertisers.** We first study (section 3) the case of myopic advertisers, which only optimize the expected profit at the current step. When these advertisers are risk-neutral, we show that truth-telling is a strongly dominant strategy: the advertiser bids the expected profit from an impression as a per-impression bid, and the value it expects from a click as the per-click bid, regardless of the auctioneer's prior or optimization strategy. Further, if the advertiser is certain about its CTR, and if this CTR is drawn from the auctioneer's

prior which follows the natural Beta distribution (defined later), then the worst case loss in revenue of the auctioneer over pure per-click bidding is at most  $1/e \approx 37\%$ . In contrast, the revenue-loss for the auctioneer when he uses the per-click scheme as opposed to the hybrid auction *can approach 100% for a fairly natural scenario, one that corresponds to obscure keywords*. We finally consider risk taking behavior of the advertisers when they are not certain about their CTR. We show that per-click bidding is dominant when the advertisers are risk-averse, but per-impression bidding is desirable when they are risk-seeking. Thus, the hybrid auction

1. Naturally extends the truthfulness of the single-slot per-click bidding auctions currently in use, for the case of myopic, risk-neutral advertisers (which is the situation under which the properties of the per-click auction are typically analyzed [11, 4, 1, 8]).
2. Results in bounded possible revenue loss but unbounded possible revenue gain for the auctioneer in the natural setting of risk-neutral, myopic advertisers, and where the auctioneer uses the Gittins index. The revenue gain occurs in the common setting of obscure keywords.
3. Naturally takes the risk posture of the advertiser into account, which neither per-click nor per-impression bidding could have done on its own (both risk averse and risk seeking advertisers occur naturally).

The result bounding the possible revenue loss of the auctioneer under the hybrid auction is for an arbitrary discount factor used by the auctioneer; the results about the possible revenue gain and the risk posture assume a myopic auctioneer. We believe these are the most appropriate assumptions, since we want to provide bounds on the revenue loss using hybrid auctions under the most general scenario, and want to illustrate the benefits of using the hybrid auction under natural, non-pathological scenarios.

**Semi-myopic advertisers.** In section 4, we remove the assumption that the advertisers are only optimizing some function of the profit at the current step. We generalize to the case where the advertisers optimize revenue over a time-horizon. We develop a tractable model for the advertisers, and show a simple dominant strategy for the advertisers, based on what we call the *bidding index*. Though this strategy does not have a closed form in general, we show that in many natural cases (detailed later), it reduces to a natural pure per-click or pure per-impression strategy that is socially optimal. Thus, our hybrid auctions are flexible enough to allow the auctioneer and the advertiser to implement complex dynamic programming strategies collaboratively, under a wide range of scenarios. Neither per-impression nor per-click bidding can exhaustively mimic the bidding index in these natural scenarios.

Finally, we show a simple bidding strategy for a certain (i.e. well-informed) advertiser to make the auctioneer's prior converge to the true CTR, while incurring no extra cost for the advertiser; per-click bidding would have resulted in the advertiser incurring a large cost. This is our final argument in support of hybrid auctions, and may be the most convincing from an advertiser's point of view.

We explain throughout the paper why the scenarios we consider are not arbitrarily chosen, but are quite natural

(indeed, we believe the most natural ones) to analyze. In the process, we obtain many interesting properties of the hybrid auction, which are described in the technical sections once we have the benefit of additional notation.

**Multiple Slots.** The main focus of the paper is analyzing the properties of the HYBRID scheme on a single ad slot. However, the auction itself can be easily generalized to multiple slots in two different ways; before describing these, we need to note that the hybrid auction assigns an “effective bid” to each advertiser based on the per-impression bid, the per-click bid, and the expected CTR or quality measure. The first generalization is akin to the widely used generalized second price auction [4, 11] (also referred to as the “next-price” auction [1]) for CPC-only bidding: the advertisers are ranked in decreasing order of effective bids, and the “effective charge” made to each advertiser is the effective bid of the next advertiser. We do not discuss this variant in the rest of this paper, since the computation methodology is no different from single slot auctions. Note that we can not expect this multi-slot generalization of the hybrid auction to be truthful because even the CPC-only next-price auction is not truthful [11, 4, 1]. However, given the immense popularity of the next-price auction, we believe that this generalization of the hybrid auction is the most likely to be used in real-life settings.

The second generalization mirrors VCG [4] (or equivalently, a laddered CPC auction [1]). This generalization assumes that the CTR is multiplicatively separable into a position dependent term and an advertiser dependent term, and under this assumption, guarantees truthfulness (on both the per-impression and per-click bids) for myopic, risk-neutral advertisers. Details of this auction are in section 6. The proof follows by extending the proof of theorem 3.1 exactly along the lines of [1] and is omitted.

The hybrid auction is also applicable to scenarios where the hybrid bidding is per-impression and per-action (i.e. CPM and CPA), or per-click and per-action (i.e. CPC and CPA).

## 2. THE HYBRID AUCTION SCHEME

As mentioned before, we will assume that there is a single slot that is being auctioned. There are  $n$  advertisers interested in a single keyword. When an advertiser  $j$  arrives at time  $t = 0$ , it submits a bid  $(m_{jt}, c_{jt})$  to the advertiser at time-slot  $t \geq 1$ . The interpretation of this bid is that the advertiser is willing to pay at most  $m_{jt}$  per impression or at most  $c_{jt}$  per click. These values are possibly conditioned on the outcomes at the previous time slots. The auctioneer chooses a publicly known value<sup>1</sup>  $q_{jt}$  which we term the *auctioneer index*, which can possibly depend on the outcomes for this advertiser at the previous time slots when its advertisement was shown, but is independent of all the bids.

The HYBRID auction scheme mimics VCG on the quantity  $R_{jt} = \max(m_{jt}, c_{jt}q_{jt})$ . We will call  $R_{jt}$  the effective bid of user  $j$  at time  $t$ . Let  $j^*$  denote the advertiser with highest  $R_j$  value, and  $R_{-j^*}$  denote the second highest  $R_j$ . There are two cases. First, suppose  $m_{j^*} > c_{j^*}q_{j^*}$ , then  $j^*$  gets the slot at per-impression price  $R_{-j^*}$ . In the other case,  $j^*$  gets the slot at per-click price  $\frac{R_{-j^*}}{q_{j^*}}$ . It is clear that the HYBRID

<sup>1</sup>It is conceivable that an auctioneer may strategically decide to not reveal its true prior; this would be an interesting direction to consider in future work.

scheme is feasible, since the per-impression price charged to  $j$  is at most  $m_j$ , and the per-click price is at most  $c_j$ . The auction generalizes in a natural way to multiple slots, but we will focus on the single slot case in this paper. If the auctioneer chooses  $q_{j,t}$  to be an estimate of the click-through-rate (CTR) and the advertiser submits only a per-click bid, then this reduces to the traditional next-price auction currently in use.

In order to analyze properties of the HYBRID auction, we need to make modeling assumptions about the advertiser, about the auctioneer index  $q$ , and about time durations.

### 2.1 Time Horizon and the Discount Factor

To model the time scale over which the auction is run, we assume there is a global discount factor  $\gamma$ . Informally, this corresponds to the present value of revenue/profit/cost that will be realized in the next step, and is an essential parameter in determining the correct tradeoff between maximizing present expected reward (exploitation) vs. obtaining more information with a view towards improving future rewards (exploration). The expected revenue at time step  $t$  gets multiplied by a factor of  $\gamma^t$ . Note that  $\gamma = 0$  corresponds to optimizing for the current step (the *myopic* case). In the discussion below, we assume the auctioneer and advertiser behave strategically in optimizing their own revenues, and can compute parameters and bids based on their own discount factors which could be different from the global discount factor  $\gamma$  used for discussing social optimality.

### 2.2 Auctioneer Model and the Gittins Index

For the purpose of designing an auctioneer index  $q_{jt}$ , we assume the auctioneer starts with a prior distribution  $\mathcal{Q}_j$  on the CTR of advertiser  $j$ . We assume further that he announces this publicly, so that the advertiser is aware of this distribution. Therefore, the a priori expected value of the CTR of advertiser  $i$  from the point of view of the auctioneer is  $\mathbf{E}[\mathcal{Q}_j]$ . Suppose at some time instant  $t$ ,  $T_{jt}$  impressions have been offered to advertiser  $j$ , and  $n_{jt}$  clicks have been observed. The natural posterior distribution  $\mathcal{Q}_{jt}$  for the advertiser is given by:

$$\Pr[\mathcal{Q}_{jt} = x] \propto x^{n_{jt}}(1-x)^{T_{jt}-n_{jt}} \cdot \Pr[\mathcal{Q}_j = x];$$

this corresponds to Bayesian updates, and when initialized with the uniform continuous prior, corresponds to the natural Beta distribution, defined later. The auctioneer chooses a function  $f(\mathcal{Q}_{jt})$  that maps a posterior distribution  $\mathcal{Q}_{jt}$  to a  $q$  value.

The function  $f$  is chosen based on the revenue guarantees the auctioneer desires. We will use the following idealized scenario to *illustrate* a concrete choice of  $f$ ; our results apply broadly and are not limited to the scenario we describe. Suppose the auctioneer wishes to optimize over a time horizon given by discount factor  $\gamma_a$ . Then, if the auctioneer were to ignore the per impression bids, and uses a first price auction on  $c_j q_{jt}$  assuming that  $c_j$  is the true per-click valuation, then his choice of  $q_{jt}$  should maximize his expected discounted revenue. Let  $v_j$  denote the true per-click valuation of the advertiser. By assumption,  $v_j = c_j$ . At time  $t$ , the auctioneer offers the slot to the bidder  $j^*$  with highest  $v_j q_{jt}$  at per-click price  $v_{j^*}$ , earning  $v_{j^*} \mathbf{E}[\mathcal{Q}_{jt}]$  in expectation. It is well-known that the discounted reward of this scheme is maximized when  $q_{jt}$  is set to the *Gittins index* (described next) of  $\mathcal{Q}_{jt}$  with discount factor  $\gamma_a$ , and hence

setting  $f(Q_{j,t}) =$  the Gittins index of  $Q_{j,t}$  would be a natural choice.

The *Gittins index* [5, 12] of a distribution  $Q$  for discount factor  $\gamma_a$  is defined as follows: Consider a coin with this prior distribution on probability of heads  $Q$ , and that yields reward 1 on heads. The Gittins index is  $(1-\gamma)M$ , where  $M$  is the largest number satisfying the following condition: Some optimal discounted reward tossing policy that is allowed to retire at any time point and collect a retirement reward of  $M$  will toss the coin at least once. It is well-known that the Gittins index is at least the mean  $\mathbf{E}[Q]$  of the prior, and for a given mean, the Gittins index increases with the variance of the prior, taking the lowest value equal to the mean only when the prior has zero variance. Further, the Gittins index also increases with the discount factor  $\gamma_a$ , being equal to the mean when  $\gamma_a = 0$ .

An equivalent definition will be useful: Consider a coin with the prior distribution on probability of heads  $Q_{j,t}$  that yields reward 1 on heads. Suppose the coin is charged  $G$  amount whenever it is tossed, but is allowed to retire any-time. The Gittins index is the largest  $G$  for which the expected discounted difference between the reward from tossing minus the amount charged in the optimal tossing policy is non-negative<sup>2</sup>.

Typically, the distribution  $Q$  is set to be the conjugate of the Bernoulli distribution, called the Beta distribution [5]. The distribution  $\text{Beta}(\alpha, \beta)$  corresponds to starting with a uniform distribution over the CTR and observing  $\alpha - 1$  clicks in  $\alpha + \beta - 2$  impressions. Therefore, if the initial prior is  $\text{Beta}(\alpha, \beta)$ , and  $n$  clicks are then observed in  $T$  impressions, the posterior distribution is  $\text{Beta}(\alpha+n, \beta+T-n)$ .  $\text{Beta}(1, 1)$  corresponds to the uniform distribution. Beta distributions are widely used mainly because they are easy to update. However, unless otherwise stated, our results will not depend on the distribution  $Q$  being a Beta distribution.

### 2.3 Advertiser Model

The value bid by the advertisers will depend on their optimization criteria. The true per-click value of advertiser  $j$  is  $v_j$ . The advertiser  $j$  maintains a time-indexed distribution  $\mathcal{P}_{j,t}$  over the possible values of the actual CTR  $p$  that is updated whenever he receives an impression. We assume advertiser  $j$ 's prior is updated based on the observed clicks in a fashion similar to the auctioneer's prior, but again, this is not essential to our results except where specifically mentioned.

The advertiser's bid will depend on its optimization criteria. In the next section, we consider the case where the advertisers only optimize their revenue at the current step, and could possibly take risk. In later sections, we consider the case where the advertisers attempt to optimize long-term revenue by bidding strategically over time.

In each case, the advertiser could be well-informed (or certain) about its CTR, so that  $\mathcal{P}_{j,t}$  is a point distribution, or uninformed about its CTR, so that it trusts the auctioneer's prior, *i.e.*,  $\mathcal{P}_{j,t} = Q_{j,t}$ , or somewhere in between. Depending on the optimization criterion of the advertiser, these cases lead to different revenue properties for the auctioneer and advertiser, and show the advantages of the HYBRID scheme over pure per-click bidding and pure per-impression bidding.

<sup>2</sup>The Gittins index is usually defined as  $M$  (*i.e.*,  $G/(1-\gamma)$ ) but the alternate definition  $(1-\gamma)M$  (*i.e.*,  $G$ ) is more convenient for this paper.

## 3. MYOPIC ADVERTISERS

In this section, we analyze single time-step properties of the auction. Specifically, we assume that the advertisers are myopic, meaning that they optimize some function of the revenue at the current time step.

Since we consider myopic properties, we drop the subscript  $t$  denoting the time step from this section. The auctioneer's prior is therefore  $Q_j$ , and the advertiser's prior is  $\mathcal{P}_j$ . Let  $p_j = \mathbf{E}[\mathcal{P}_j]$ .

We first show that when the advertisers are risk-neutral, then bidding  $(v_j p_j, v_j)$  is the dominant strategy, *independent of the auctioneer's prior or the choice of  $f$* . Further, if the advertiser is certain about its CTR, and if this CTR is drawn from the auctioneer's prior which follows a Beta distribution, then the worst case loss in revenue of the auctioneer over pure per-click bidding is at most  $1/e \approx 37\%$ . In contrast, the revenue-loss for the auctioneer when he uses the per-click scheme as opposed to the hybrid auction *can approach 100% for a fairly natural scenario, one that corresponds to obscure keywords*. We finally consider risk taking behavior of the advertisers when they are not certain about their CTR. We show that per-click bidding is dominant when the advertisers are risk-averse, but per-impression bidding is desirable when they are risk-seeking. Thus, the hybrid auction naturally extends the truthfulness of the single-slot per-click bidding auctions currently in use, results in bounded possible revenue loss but unbounded possible revenue gain, and naturally takes the risk posture of the advertiser into account; the precise qualitative conclusions are detailed in the introduction and the formal statements are proved below.

The result bounding the possible revenue loss of the auctioneer under the hybrid auction is for an arbitrary discount factor used by the auctioneer; the results about the possible revenue gain and the risk posture assume a myopic auctioneer. We believe these are the most appropriate assumptions, since we want to provide bounds on the revenue loss using hybrid auctions under the most general scenario, and want to illustrate the benefits of using the hybrid auction under natural, non-pathological scenarios.

### 3.1 Truthfulness

We first show that the dominant strategy involves truthfully revealing the expected CTR,  $p_j$ . Recall that the advertiser bids  $(m_j, c_j)$ . Further, the auctioneer computes an index  $q_j$  based on the distribution  $Q_j$ , and does VCG on the quantity  $R_j = \max(m_j, c_j q_j)$ .

**THEOREM 3.1.** *If  $p_j = \mathbf{E}[\mathcal{P}_j]$  and the advertiser is myopic and risk-neutral, then regardless of the choice of  $q_j$ , the (strongly) dominant strategy is to bid  $(v_j p_j, v_j)$ .*

**PROOF.** First, consider the case where  $q_j \leq p_j$ . Suppose the advertiser bids  $(m_j, c_j)$  and wins the auction. Then, the expected profit of this advertiser is at most  $p_j v_j - \min\{R_{-j^*}, R_{-j^*} \cdot (p_j/q_j)\}$  which is at most  $p_j v_j - R_{-j^*}$ . Thus, the maximum profit of the advertiser can be at most  $\max\{0, p_j v_j - R_{-j^*}\}$  which is obtained by bidding  $(p_j v_j, v_j)$ .

Next, consider the case where  $q_j > p_j$ . Suppose the advertiser bids  $(m_j, c_j)$  and wins the auction. Then, the expected profit of this advertiser is at most  $p_j v_j - \min\{R_{-j^*}, R_{-j^*} \cdot (p_j/q_j)\}$  which is at most  $p_j v_j - R_{-j^*} \cdot (p_j/q_j)$ . Thus, the maximum profit of the advertiser can be at most  $\max\{0, p_j v_j - R_{-j^*} \cdot (p_j/q_j)\}$  which is again obtained by bidding  $(p_j v_j, v_j)$ .

Thus, it is never suboptimal to bid truthfully. Let  $R^*$

denote the second highest value of  $\max(m_j, c_j q_j)$ . In order to show that  $(p_j v_j, v_j)$  is a (strongly) dominant strategy, we need to show that for any other bid-pair  $(m_j, c_j)$ , there exist values of  $q_j, R^*$  such that the profit obtained by bidding  $(m_j, c_j)$  is strictly less than that obtained by truthful bids. Suppose  $\epsilon$  is an arbitrary small but positive number. First consider the scenario where  $q_j = p_j$ , i.e., the auctioneer has a perfect prior. In this scenario, bidding  $(m_j, c_j)$  with either  $m_j > p_j v_j + \epsilon$  or  $c_j > v_j + \epsilon/p_j$  results in a negative profit when  $p_j v_j < R^* < p_j v_j + \epsilon$ ; truthful bidding would have resulted in zero profit. Further, if  $m_j < p_j v_j - \epsilon$  and  $c_j < v_j - \epsilon/p_j$ , then the advertiser obtains zero profit in the case where  $p_j v_j > R^* > p_j v_j - \epsilon$ ; truthful bidding would have obtained positive profit.

This leaves the cases where  $m_j = p_j v_j, c_j < v_j/(1 + \epsilon)$  or where  $m_j < p_j v_j/(1 + \epsilon), c_j = v_j$ . In the former case, the advertiser obtains zero profit in the situation where  $q_j > p_j > q_j/(1 + \epsilon)$  and  $R^* = p_j v_j$ ; truthful bidding would have obtained positive profit. In the latter case, the advertiser obtains zero profit when  $q_j < p_j/(1 + \epsilon)$  and  $R^* = p_j v_j/(1 + \epsilon)$ ; truthful bidding would have obtained positive profit.  $\square$

### 3.2 Well-Informed Advertisers: Loss in Auctioneer's Revenue

We now consider the case where the advertisers are certain about their CTR  $p_j$  and risk-neutral; by the results of the previous section, we will assume that they bid truthfully. More formally, we assume the prior  $\mathcal{P}_j$  is the point distribution at  $p_j$ . We suppose that the  $p_j$  are drawn from the auctioneer's prior that is of the form  $\mathcal{Q}_j = \text{Beta}(\alpha_j, \beta_j)$ . We now show that for  $q_j$  being the Gittins index of  $\text{Beta}(\alpha_j, \beta_j)$  for any discount factor  $\gamma_a$ , the expected revenue of the auctioneer at the current step is at least  $1 - 1/e$  times the revenue had he ignored the per-impression bids.

**THEOREM 3.2.** *In the above mentioned scenario, the expected revenue of the auctioneer at the current step is at least  $1 - 1/e \approx 63\%$  of the corresponding auction that ignores the per-impression bid.*

**PROOF.** Let  $q_j$  denote the Gittins index of  $\mathcal{Q}_j = \text{Beta}(\alpha_j, \beta_j)$ . Let advertiser 1 have the highest  $v_j q_j$ , and advertiser 2 the next highest. Let  $R^* = v_2 q_2$ . If the per-impression bids are ignored, advertiser 1 gets the impression at a per-click price of  $v_2 q_2/q_1$ , so that the expected revenue is  $R^* \frac{\mathbf{E}[\mathcal{Q}_1]}{q_1}$ .

In the HYBRID scheme,  $v_1 q_1$  and  $v_2 q_2$  are both at least as large as  $R^*$ . Hence, if the auctioneer makes a per-impression charge, then this charge must be at least  $R^*$  per impression. If the advertiser makes a per-click charge (which must be to advertiser 1), then the expected revenue is at least  $R^* \mathcal{Q}_1/q_1$ . Hence the expected revenue of the HYBRID scheme is at least  $R^* \mathbf{E} \left[ \min \left( 1, \frac{\mathcal{Q}_1}{q_1} \right) \right]$  and the ratio of the revenue of the HYBRID scheme to the per-click scheme is at least  $\frac{\mathbf{E}[\min(q_1, \mathcal{Q}_1)]}{\mathbf{E}[\mathcal{Q}_1]}$ .

For  $p$  drawn from the distribution  $\mathcal{Q}_1$ , we now need to show that  $\frac{\mathbf{E}[\min(q_1, \mathcal{Q}_1)]}{\mathbf{E}[\mathcal{Q}_1]} \geq 1 - 1/e$ . To show this, observe that for a fixed  $\mathcal{Q}_1$ , this ratio is smallest when  $q_1$  is as small as possible. This implies we should choose  $q_1 = \mathbf{E}[\mathcal{Q}_1] = \frac{\alpha}{\alpha + \beta}$ , which corresponds to a discount factor of 0. Denote  $\mu = \mathbf{E}[\mathcal{Q}_1]$ . Then, the goal is to minimize the ratio  $\frac{1}{\mu} \mathbf{E}[\min(\mu, \mathcal{Q}_1)]$  as a function of  $\alpha, \beta$ . Lemma 3.3 shows that this ratio is  $1 - 1/e$ , completing the proof.  $\square$

**LEMMA 3.3.** *If  $w$  is drawn from the Beta distribution with parameters  $\alpha, \beta \geq 1$ , and  $\mu = \alpha/(\alpha + \beta)$  is the mean of  $w$ , then  $\mathbf{E}[\min(\mu, w)] \geq \mu(1 - 1/e)$ .*

**PROOF.** We will allow  $\alpha, \beta$  to take on fractional values as long as they are both at least 1. Suppose  $\alpha, \beta$  are both strictly bigger than 1. Let  $z$  denote the random variable drawn from the Beta distribution with parameters  $\alpha' = \alpha - \mu\theta, \beta' = \beta - (1 - \mu)\theta$ , where  $\theta > 0$  is chosen such that  $\alpha', \beta' \geq 1$  and at least one of  $\alpha', \beta'$  is exactly 1. The mean of  $z$  is  $\frac{\alpha - \mu\theta}{\alpha + \beta - \theta} = \mu$ , which is the same as the mean of  $w$ .

Let  $f_w, f_z$  denote the probability density functions of  $w, z$  respectively, and let  $F_w(x)$  (resp.  $F_z(x)$ ) denote  $\Pr[w \geq x]$  (resp.  $\Pr[z \geq x]$ ). Consider the ratio  $r(x) = f_w(x)/f_z(x) = \phi x^{\mu\theta}(1-x)^{(1-\mu)\theta}$ , where  $\phi$  is the ratio of the corresponding normalizing terms and hence does not depend on  $x$ .

Observe that  $r(x)$  is uni-modal (since the derivative of  $r$  is 0 exactly once in the interval  $[0, 1]$ ); and that  $r(x) \rightarrow 0$  as  $x \rightarrow 0^+$  and as  $x \rightarrow 1^-$ . Since both  $F_w(x)$  and  $F_z(x)$  are monotonically decreasing curves connecting  $(0, 1)$  and  $(1, 0)$ , the above properties of  $r(x) = \frac{F'_w(x)}{F'_z(x)}$  easily imply that for some  $s \in (0, 1)$ , over the interval  $x \in [0, s]$ ,  $F_w(x) \geq F_z(x)$ , and over  $x \in [s, 1]$ ,  $F_w(x) \leq F_z(x)$ . This combined with the fact that  $\mathbf{E}[w] = \mathbf{E}[z] = \mu$  implies  $w$  Lorenz-dominates  $z$ , so that for all concave functions  $g$ , we have  $\mathbf{E}[g(w)] \geq \mathbf{E}[g(z)]$  [7].

Since  $g(w) = \min(w, y)$  is concave in  $w$  for fixed  $y$ , we have  $\mathbf{E}[\min(w, \mu)] \geq \mathbf{E}[\min(z, \mu)]$ . Therefore, it is sufficient to analyze  $\mathbf{E}[\min(\mu, z)]/\mu$ , i.e. the case where either  $\alpha$  or  $\beta$  is exactly 1, and the other is at least 1. It is easy to explicitly verify both these cases, and show that the worst case is when  $\alpha = 1$  and  $\beta \rightarrow \infty$  when  $\mathbf{E}[\min(\mu, z)] = (1 - 1/e)\mu$ .  $\square$

**A Typical Case.** Though the HYBRID scheme is not revenue dominant over the pure per-click scheme in pathological cases, the key advantage is in the following typical situation. There are  $n$  advertisers whose CTRs  $p_1 \geq p_2 \geq \dots \geq p_n$  are drawn from a common prior  $\mathcal{Q} = \text{Beta}(1, K)$ , whose mean is  $\mu = O(\frac{1}{K})$ . Assume further that  $n = 4^K$  or  $K = \frac{\log n}{2}$ . We have:

$$\begin{aligned} \Pr[p_2 \geq 1/2] &\geq \left( 1 - \left( 1 - \frac{2}{\log n} \left( \frac{1}{2} \right)^{\frac{\log n}{2}} \right)^{\frac{n}{2}} \right)^2 \\ &= \left( 1 - \left( 1 - \frac{2}{\sqrt{n} \log n} \right)^{\frac{n}{2}} \right)^2 \\ &= 1 - o(1) \end{aligned}$$

Recall that the advertisers are aware of their CTR, but the auctioneer is only aware of the prior. Suppose the per-click value for all the advertisers is  $v$ , and these are truthfully revealed. In the per-click scheme, the auctioneer sells the impression to an arbitrary advertiser at per-click price  $v$ , and in expectation earns  $\mu v$ . If the auctioneer is myopic ( $\gamma_a = 0$ ), then  $q = \mathbf{E}[\mathcal{Q}] < p_2$  w.h.p, and the HYBRID scheme charges per-impression. Here, the auctioneer sells to advertiser 1 at per-impression price  $vp_2$ . From the above,  $\mathbf{E}[p_2] = \Omega(1)$ , so that  $\mathbf{E}[p_2]/\mu = \Omega(\log n)$ . Therefore, for  $n$  advertisers with diffuse priors of the form  $\text{Beta}\left(1, \frac{1}{\log n}\right)$ , the auctioneer gains a factor  $\Omega(\log n)$  in revenue. This is particular relevant for obscure keywords, where the auctioneer will have very diffuse priors.

### 3.3 Uninformed Advertisers and Risk

So far, we have assumed that the advertiser is risk neutral and certain about the CTR, so that he is optimizing his expected profit. We now suppose that the advertiser is uncertain and trying to maximize a utility function  $U$  on his profit. The function  $U(x)$  is monotone with monotone derivative, and  $U(0) = 0$ . If  $U$  is convex, the advertiser is said to be risk-seeking, and if it is concave, the advertiser is said to be risk-averse. We show that for risk-averse advertisers, pure per-click bidding is dominant, whereas pure per-impression bidding is dominant when the advertisers are risk-seeking.

It is natural to assume that some advertisers may be either risk-averse or risk-seeking. Risk-aversion models advertisers with tight budget constraints. Risk-seeking advertisers also occur naturally in many settings; one example is when advertisers are conducting experiments to identify high performance advertising channels and keywords. Finding a high reward keyword may result in a higher budget allocated to this keyword and more revenue in the future, making the present utility function of winning this ad slot appear convex at the present time.

We assume the advertisers are uninformed, which is equivalent to assuming the advertiser and the auctioneer share a common prior, so that  $\mathcal{P}_j = \mathcal{Q}_j$ . Essentially, the advertiser has no information and simply trusts the auctioneer's prior<sup>3</sup>. In this section, we focus on a single advertiser, and drop the subscript corresponding to it. Let  $p = \mathbf{E}[\mathcal{P}] = \mathbf{E}[\mathcal{Q}]$ . As mentioned earlier, we assume that the auctioneer is myopic as well ( $\gamma_a = 0$ ), so that  $q = \mathbf{E}[\mathcal{Q}] = p$ .

Let  $(m, c)$  denote the advertiser's bid, and let  $v$  denote the true per-click valuation. Let  $I_R$  be the indicator corresponding to whether the bidder gets the impression if  $R_{-j} = R$ . The bidding strategy of the advertiser will attempt to maximize:

$$I_R \cdot \max \left( \mathbf{E}[U(v\mathcal{P} - R)], \mathbf{E} \left[ U \left( v\mathcal{P} - \frac{R\mathcal{P}}{p} \right) \right] \right)$$

In the above expression, the first term is the expected profit if the impression is obtained based on the per-impression bid; and the latter term is the expected profit if the impression is obtained based on the per-click bid. Our next lemma captures the structure of the dominant strategy.

**LEMMA 3.4.** *If  $U$  is concave, bidding  $(0, v)$  is a dominant strategy. If  $U$  is convex, the dominant strategy is of the form  $(m, 0)$  for a suitably chosen  $m$ .*

**PROOF.** First consider the case when  $vp < R$ . In this case, regardless of  $U$ ,  $\mathbf{E} \left[ U \left( v\mathcal{P} - \frac{R\mathcal{P}}{p} \right) \right] \leq 0$ . Therefore, to obtain positive profit, the bidder has to obtain the impression based on his per-impression bid. In this case, the expected profit is  $\mathbf{E}[U(v\mathcal{P} - R)]$ . Note that  $\mathbf{E}[v\mathcal{P} - R] < 0$ . Therefore if  $g$  is concave:

$$\mathbf{E}[U(v\mathcal{P} - R)] \leq U(\mathbf{E}[v\mathcal{P} - R]) \leq U(0) = 0$$

Therefore, if  $U$  is concave and  $vp < R$ , then bidding  $(0, v)$  is a dominant strategy. Since obtaining the impression based

<sup>3</sup>For the other extreme case of well-informed advertisers, there is no uncertainty, and hence the risk-averse and risk-seeking cases collapse to risk-neutral.

on the per-click bid does not yield positive profit, if  $U$  is convex, bidding  $(m, 0)$  with appropriately chosen  $m$  is a dominant strategy.

The next and most interesting case is when  $vp \geq R$ . Define random variable  $X = v\mathcal{P} - R$  and  $Y = v\mathcal{P} - \frac{R\mathcal{P}}{p}$ . First note that  $\mathbf{E}[X] = \mathbf{E}[Y] = vp - R \geq 0$ . Further, the cumulative distribution functions (CDFs) of  $X$  and  $Y$  cross exactly once, with the CDF of  $X$  being initially larger than the CDF of  $Y$ . This is sufficient to show that  $Y$  Lorenz-dominates  $X$ . This implies that for  $U$  being concave,  $\mathbf{E}[U(Y)] \geq \mathbf{E}[U(X)]$  [7], so that the advertiser only bids per click. Further, if  $U$  is convex,  $\mathbf{E}[U(X)] \geq \mathbf{E}[U(Y)]$ , so that the advertiser only bids per impression.  $\square$

Our main result in this sub-section is the following property which gives a single natural characterization of the optimum hybrid bid for both risk-averse and risk-seeking advertisers. We will then show that for risk-seeking advertisers ( $U$  is strictly increasing and convex), the expected myopic revenue of the auctioneer is larger in the HYBRID auction compared to the pure per-click auction, and for risk-averse advertisers, the HYBRID and per-click auctions coincide.

**THEOREM 3.5.** *Let  $m^* = \max\{y \mid \mathbf{E}[U(v\mathcal{P} - y)] \geq 0\}$ . Bidding  $(m^*, v)$  is a dominant strategy. Further, the auctioneer's revenue in the HYBRID scheme dominates the revenue in the pure per-click scheme.*

**PROOF.** First consider the case when  $U$  is concave. Then,  $\mathbf{E}[U(v\mathcal{P} - p)] \leq U(\mathbf{E}[v\mathcal{P} - p]) = 0$ . Therefore,  $m^* \leq vp$ , so that bidding  $(m^*, v)$  is equivalent to bidding  $(0, v)$ , which is a dominant strategy.

Next, when  $U$  is convex, we have  $m^* \geq vp$ , so that bidding  $(m^*, v)$  is equivalent to bidding  $(m^*, 0)$ . The previous lemma shows that the dominant strategy is of the form  $(m, 0)$ . Since  $m^*$  is the largest value of  $R_{-j}$  for which the advertiser makes a non-negative profit, bidding  $(m^*, 0)$  must be the dominant strategy.

The auctioneer's revenue in the HYBRID scheme is the second largest value of  $\max(m_j^*, v_j p_j)$  while that in the per click scheme is the second largest value of  $v_j p_j$ , which cannot be larger.  $\square$

## 4. SEMI-MYOPIC ADVERTISERS

In this section, we remove the assumption that the advertisers are optimizing some function of the profit at the current step. We now generalize to the case where the advertisers optimize revenue over a time-horizon. We develop a tractable model for the advertisers, and show a simple dominant strategy for the advertisers, based on what we call the *bidding index*. Though this strategy does not have a closed form in general, we show that in many natural cases (detailed later) cases, it reduces to a natural pure per-click or pure per-impression strategy that is socially optimal. Thus, our hybrid auctions are flexible enough to allow the auctioneer and the advertiser to implement complex dynamic programming strategies collaboratively, under a wide range of scenarios. Neither per-impression nor per-click bidding can exhaustively mimic the bidding index in these natural scenarios.

Recall that the true per-click value of advertiser  $j$  is  $v_j$ , and that the advertiser  $j$  maintains a time-indexed distribution  $\mathcal{P}_{jt}$  over the possible values of the actual CTR  $p$  that is updated whenever he receives an impression. We assume

advertiser  $j$ 's prior is updated based on the observed clicks in a fashion similar to the auctioneer's prior,  $Q_{jt}$ .

We assume the bidding strategy of the advertiser is **semi-myopic**, which we define as follows: The advertiser has a discount factor  $\gamma_b$ . The bid of an advertiser  $j$  depends on its current state  $\langle v_j, \mathcal{P}_{jt}, Q_{jt} \rangle$ , and on  $R_{-j}$  in a fashion described next. At every step, the value of  $R_{-j}$  is revealed. If the advertiser  $j$  got the impression the previous time step, the value of  $R_{-j}$  remains the same since the states of the other advertisers remains the same, else it changes adversarially. The optimization criterion of the advertiser is to maximize its discounted expected gain (using discount factor  $\gamma_b$ ) in the contiguous time that it receives impressions (so that the value of  $R_{-j}$  remains the same). We make the reasonable assumption that the advertiser cannot optimize for a horizon beyond that, since the value of  $R_{-j}$  changes in an unknown fashion. Finally note that a myopic advertiser is equivalent to assuming  $\gamma_b = 0$ .

**Discussion.** The semi-myopic model is closely related to the MDP approach of analyzing repeated auctions; see for instance [2, 3]. These works make the assumption that the priors  $\mathcal{P}_j$  of the advertisers are public knowledge. However, this leads to somewhat perverse incentives in which the optimal strategy for an advertiser could be to underbid at the current time step in the hope that the priors of the other advertisers resolve to low values, and he then wins the auction on the remaining time steps at a lower price. However, note that if there are sufficiently many bidders, this scenario is unlikely to happen, and the bidder will attempt to win the auction at the current time slot. We make this explicit by making the following assumptions:

1. The bidder  $j$  is aware of the revealed  $R_{-j}$  values of the other bidders, but may not be aware of their prior distributions, which are usually private information.
2. The bidder only optimizes over the contiguous time horizon in which he receives the impressions. In this horizon,  $R_{-j}$  is fixed, and further, this removes the perverse incentive to under-bid described above.

Note that in our model, the bidder is indeed *aware* of the current bids  $R_{-j}$  of the other bidders. However, unlike the model in [2, 3], the optimization time-horizon of the bidder leads to the existence of a nicely specified dominant strategy. We hope that our modeling, that is only slightly more restrictive than ones considered in literature, but which have nice analytic properties, will be of independent interest in this and other contexts.

#### 4.1 The Dominant Bidding Index Strategy

We first show a bidding strategy that we term the *bidding index strategy*, and show that it is weakly dominant in the class of *semi-myopic* strategies. The **bidding index**  $\mathcal{B}(v, \mathcal{P}, \mathcal{Q})$  is defined as follows: Suppose the advertiser's current prior is  $\mathcal{P}$  and the auctioneer's current prior is  $\mathcal{Q}$ . Denote the current time instant as  $t = 0$ . Since the advertiser computes this index, we assume the advertiser trusts his own prior but not the auctioneer's. For a parameter  $W$ , define the following game between the advertiser and the auctioneer with discount factor  $\gamma_b$ : At step  $t \geq 0$ , suppose the advertiser has prior  $\mathcal{P}_t$  (with mean  $\mathbf{E}[\mathcal{P}_t] = p_t$ ) and the auctioneer,  $Q_t$  (with  $q_t = f(Q_t)$  being the auctioneer's index), the advertiser can either stop the game, or continue. If

he continues, he gains  $vp_t$  in expectation and pays the auctioneer  $W \min\left(1, \frac{p_t}{q_t}\right)$ ; the difference is his gain. The advertiser's value for the game is the expected discounted (according to  $\gamma_b$ ) gain for the optimal strategy. Define  $\mathcal{W}(v, \mathcal{P}, \mathcal{Q})$  as the largest value of  $W$  for which the value of the game with initial priors  $\mathcal{P}$  and  $\mathcal{Q}$ , is positive. This value can easily be computed by dynamic programming, much like the Gittins index.

The **bidding index**  $\mathcal{B}(v, \mathcal{P}, \mathcal{Q})$  is defined as:

$$\mathcal{B}(v, \mathcal{P}, \mathcal{Q}) = \mathcal{W}(v, \mathcal{P}, \mathcal{Q}) \min\left(1, \frac{p_0}{q_0}\right)$$

This is the largest per impression price at time  $t = 0$  for which the value of the above game is positive.

**The Strategy:** At any time step, when the advertiser  $j$ 's prior is  $\mathcal{P}_{jt}$  with mean  $p_{jt}$ , and the auctioneer's prior is  $Q_{jt}$ , with  $q_{jt} = f(Q_{jt})$ , let  $W_{jt} = \mathcal{W}(v_j, \mathcal{P}_{jt}, Q_{jt})$  and  $B_{jt} = \mathcal{B}(v_j, \mathcal{P}_{jt}, Q_{jt})$ . The **bidding index strategy** involves bidding  $(B_{jt}, \frac{B_{jt}}{p_{jt}})$ .

It is clear that the bidding index strategy is well-defined for  $q_{jt}$  being an arbitrary function  $f(Q_{jt})$  chosen by the auctioneer, and not just for  $f$  being the Gittins index of  $Q_{jt}$  using discount factor  $\gamma_a$ .

**THEOREM 4.1.** *The bidding index strategy is (weakly) dominant in the class of semi-myopic strategies.*

**PROOF.** Consider a sequence of time steps when advertiser  $j$  gets the impression; call this a phase. During this time, the value  $R_{-j}$  used in the VCG scheme is fixed; denote this value  $R^*$ . Suppose at a certain time step, the mean of the advertiser's prior is  $p_{jt}$  and the auctioneer computes  $q_{jt}$ . If the advertiser gets the impression, the price he is charged in the VCG scheme is either  $R^*$  per impression or  $R^*/q_{jt}$  per click. The advertiser optimizes this by paying  $R^* \min(1, p_{jt}/q_{jt})$  in expectation per impression. The state evolution is only conditioned on getting the impression and not on the price paid for it.

Since the advertiser's strategy is semi-myopic, at any time step, the bid should fetch him a non-negative expected profit for the rest of the phase. This implies that  $R^* \leq W_{jt}$ . There are two cases.

First, if  $p_{jt} < q_{jt}$ , the advertiser essentially bids  $R_{jt} = B_{jt} \frac{q_{jt}}{p_{jt}} = W_{jt} \geq R^*$ , and receives the impression at a per-click price of  $\frac{R^*}{q_{jt}}$ . Therefore, the expected per impression price is  $R^* \frac{p_{jt}}{q_{jt}} = R^* \min\left(1, \frac{p_{jt}}{q_{jt}}\right)$ .

Next, if  $p_{jt} > q_{jt}$ , the advertiser essentially bids  $R_{jt} = B_{jt} = W_{jt} \geq R^*$ , and receives the impression at a per-impression price of  $R^* = R^* \min\left(1, \frac{p_{jt}}{q_{jt}}\right)$ .

Therefore, the bidding scheme ensures that the advertiser receives the impression and makes the most possible profit in the rest of the phase. Note finally that if  $W_{jt} < R^*$ , the maximum possible profit in the rest of the phase is negative, and the bidding scheme ensures the advertiser does not receive the impression.  $\square$

#### 4.2 Social Optimality of Bidding Index

Suppose the global discount factor is  $\gamma$ . We define the socially optimal strategy as follows: Suppose at time  $t$ , advertiser  $j$  with prior  $\mathcal{P}_{jt}$  receives the impression resulting in value  $v_j p_{jt}$  for him. The socially optimal solution maximizes

the infinite horizon expected discounted value with discount factor  $\gamma$ .

We show that the bidding index strategy implements the socially optimal solution in each of the following two situations: (1) The advertiser and the auctioneer share the same prior ( $\mathcal{P}_{jt} = \mathcal{Q}_{jt}$ ), and either (1a) only the advertisers are strategic ( $\gamma_a = 0$  and  $\gamma_b = \gamma$ ) or (1b) only the auctioneer is strategic ( $\gamma_a = \gamma$  and  $\gamma_b = 0$ ); and (2) The advertisers are certain about their CTR ( $\mathcal{P}_j = p_j$ ) and (2a) the auctioneer's index  $q_{jt}$  is always at most  $p_j$ . The bidding index also has a particularly simple form when the advertisers are certain, and (2b) the auctioneer's  $q_{j,t}$  is monotonically decreasing with  $t$  and always larger than  $p_{j,t}$ . In both (2a) and (2b), the bidding index strategy reduces to bidding  $(v_j p_j, v_j)$ .

These scenarios are not arbitrarily chosen, and are the most illustrative scenarios we could find. Scenario (1) corresponds to an advertiser and an auctioneer that trust each other and hence have a common prior; in case (1a), the auctioneer merely discloses its current estimate and trusts the advertisers to bid in an optimal fashion, whereas in (1b) the advertisers delegate the strategic decision making to the auctioneer. In scenario 2, the advertisers have a definitive model of the CTR in (2a), we model the case where the auctioneer starts with an underestimate of the click-through rate and hence the  $q_{j,t}$  are always smaller than  $p_j$  to which they will hopefully converge as this advertisement is shown more times and the auctioneer's prior gets sharpened, and in (2b) we model the mirror situation where the  $q_{j,t}$ 's are always an over-estimate. It will be interesting to find a general theorem about the bidding index that unifies all these diverse scenarios.

In each of these cases, the bidding strategy can be implemented using either per-impression or per-click bidding or both, but neither per-impression nor per-click bidding can exhaustively mimic the bidding index in all scenarios.

**Shared Priors.** When the advertisers are uncertain and simply share the auctioneer's prior, we have  $\mathcal{P}_{jt} = \mathcal{Q}_{jt}$ . Let  $G_{jt}$  denote the Gittins index of  $\mathcal{P}_{jt}$  with discount factor  $\gamma$ . The socially optimal solution always gives the impression to the advertiser with highest  $v_j G_{jt}$  at time  $t$ .

**THEOREM 4.2.** *For shared priors, the bidding index strategy implements the socially optimal solution in the following two cases:*

1. *The advertisers are strategic, i.e.,  $\gamma_b = \gamma$ , and the auctioneer is myopic, i.e.,  $\gamma_a = 0$ .*
2. *The advertisers are myopic, i.e.,  $\gamma_b = 0$ , and the auctioneer is strategic, i.e.,  $\gamma_a = \gamma$ .*

**PROOF.** For the first part, we have  $q_{jt} = p_{jt}$  since  $\gamma_a = 0$ . Therefore,  $\min\left(1, \frac{p_{jt}}{q_{jt}}\right) = 1$ , so that the value  $W_{jt}$  is the largest charge per impression so that the advertiser's discounted revenue is non-negative. This is precisely the definition of the Gittins index with discount factor  $\gamma_b = \gamma$ . Therefore, the bidding index strategy involves bidding  $(v_j G_{jt}, v_j \frac{G_{jt}}{p_{jt}})$ . This can easily be seen to be equivalent to bidding either  $(v_j G_{jt}, 0)$  or  $(0, v_j \frac{G_{jt}}{p_{jt}})$ , and hence can be mimicked with either pure-impression or pure-click bidding. Also, we have  $R_{jt} = v_j G_{jt}$ , so that the bidding index implements the socially optimal strategy.

For the second part, since the advertiser is myopic, the bidding index reduces to bidding  $(v_j p_{jt}, v_j)$ . Since  $\gamma_a = \gamma$ , we have  $q_{jt} = G_{jt} \geq p_{jt}$ . Therefore,  $R_{jt} = v_j G_{jt}$ , so that the bidding index implements the socially optimal solution; this can also be mimicked using per-click bidding but not per-impression bidding.  $\square$

**Well -Informed Advertisers.** We next consider the case where the advertisers are certain about their CTR  $p_j$ , so that  $\mathcal{P}_{jt} = p_j$ . The socially optimal solution always gives the impression to the advertiser with the largest  $v_j p_j$ . We show the following theorem:

**THEOREM 4.3.** *When  $\mathcal{P}_{jt} = p_j$ , then the bidding index strategy reduces to bidding  $(v_j p_j, v_j)$  in the following two scenarios:*

1. *The auctioneer's  $q_{jt}$  is always at most  $p_j$ . In this case, the strategy is equivalent to bidding  $(v_j p_j, 0)$  and is socially optimal.*
2. *The auctioneer's  $q_{jt}$  is at least  $p_j$ , and is monotonically decreasing with  $t$ .*

**PROOF.** When  $p_j > q_{jt}$  for all  $t$ , we have  $\min\left(1, \frac{p_j}{q_{jt}}\right) = 1$ , so that the value  $W_{jt}$  is the largest per-impression price for which the advertiser's discounted revenue is non-negative. This is precisely  $v_j p_j$ , so that the bidding index strategy reduces to bidding  $(v_j p_j, v_j)$ . This is clearly socially optimal. Since  $q_{jt} < p_j$ , this is equivalent to bidding  $(v_j p_j, 0)$ , but can not be simulated using per-click bids.

When  $p_j \leq q_{jt}$  for all  $t$  and  $q_{jt}$  is monotonically decreasing with  $t$ , the expected price  $W_{jt} \frac{p_j}{q_{jt}}$  charged to the advertiser increases with time. At any time  $t$ , the advertiser maximizes  $W_{jt}$  by setting it to  $v_j q_{jt}$  and stopping after the first step. Therefore,  $W_{jt} = v_j q_{jt}$ , and  $B_{jt} = v_j q_{jt} \frac{p_j}{q_{jt}} = v_j p_j$ . Therefore, the bidding index strategy reduces to bidding  $(v_j p_j, v_j)$ ; this is also equivalent to  $(0, v_j)$ , but can not be simulated using per-impression bids.  $\square$

## 5. EXPLORATION BY ADVERTISERS

We now show a simple bidding strategy for a certain (i.e. well-informed) advertiser to make the auctioneer's prior converge to the true CTR, while incurring no extra cost for the advertiser; per-click bidding would have resulted in the advertiser incurring a large cost. More concretely, this models the scenario where the advertiser has side information about the advertisement's CTR but the auctioneer does not have a good prior, for example, because the keyword may be obscure. The advertiser has incentive to help the auctioneer "learn" the true CTR because this improves the advertiser's chance of winning an ad slot in a pure per-click scheme.

To motivate why this is important, imagine a situation where the advertiser would not get the slot if the scheme were pure per-click, and he were to bid truthfully per-click, letting the auctioneer use his own estimate  $q_j$  of the CTR. Therefore, in the pure per-click scheme, the advertiser has to overbid on the per-click valuation to get the slot enough number of times to make the CTR used by the auctioneer converge to the true value; we show that this results in loss in revenue for the advertiser. However, allowing for per impression bids preserves truthfulness, and furthermore, helps



the auctioneer “learn” the true CTR, while incurring no revenue loss to the advertiser. This is our final argument in support of hybrid auctions, and may be the most convincing from an advertiser’s point of view.

Formally, we consider an advertiser that is certain about its CTR  $p_j$ , where  $v_j p_j > R_{-j}$  so that the advertiser can (and would like to) win the auction but where  $q_{jt} < p_j$ , and where the goal of the advertiser is to make the auctioneer’s prior converge to the true CTR. We show that the advertiser can achieve this goal without any loss in revenue, whereas achieving the same objective using per-click bidding would have resulted in a large revenue-loss. We assume the auctioneer’s prior is a Beta distribution.

We show a candidate strategy for an advertiser to make the Gittins index of the auctioneer’s distribution,  $\mathcal{Q}_{jt} = \text{Beta}(\alpha_{jt}, \beta_{jt})$  converge close to its true CTR  $p_j$  while incurring no loss in revenue. The loss is defined as the value earned from actual clicks minus the amount paid to the auctioneer.

We focus on a single advertiser and drop its subscript. For any  $\epsilon > 0$ , suppose the advertiser’s strategy is as follows: During an “explore” phase, submit a bid of  $(vp', v)$  where  $p' = p(1 - \epsilon)$ , and then switch to bidding  $(0, v)$ . During the explore phase, suppose the advertiser gets  $T$  impressions on a price per impression basis resulting in  $n$  clicks. Then the worst-case loss in revenue of the advertiser during the explore phase is  $v(Tp' - n)$ . The “explore” phase stops when the auctioneer’s posterior mean of the distribution  $\text{Beta}(\alpha + n, \beta + T - n)$  is at least  $p(1 - \epsilon)$ . Note that this also implies that the Gittins index for  $\text{Beta}(\alpha + n, \beta + T - n)$  is at least  $p(1 - \epsilon)$  irrespective of the discount factor  $\gamma$ ; this in turn implies that by switching to pure per-click bidding, the advertiser is assured that  $q \geq p(1 - \epsilon)$ , so that bidding  $(0, v)$  yields  $R_j \geq vp(1 - \epsilon)$ .

**CLAIM 5.1.** *Suppose the advertiser knows its true CTR is  $p$ , and the auctioneer’s initial prior is  $\text{Beta}(\alpha, \beta)$ . For any  $\epsilon > 0$ , the explore phase incurs no loss in revenue for the advertiser.*

**PROOF.** Let  $T$  denote the random stopping time of the explore phase and suppose it results in  $N$  clicks. First note that if  $T > 0$ , then  $p(1 - \epsilon) > \frac{\alpha}{\alpha + \beta}$ . The posterior mean on stopping is  $\frac{\alpha + N}{\alpha + \beta + T} \geq p(1 - \epsilon)$ , which implies  $N/T > p(1 - \epsilon)$ . Therefore,  $Tp(1 - \epsilon) - N < 0$ , which shows there is no loss in revenue (provided  $T$  is finite with probability 1, which follows from the law of large numbers in this case).  $\square$

Suppose  $R_{-j} = v_j p_j(1 - \epsilon)$ . In a pure per-click bidding scheme, the advertiser would have to bid at least  $v_j(1 - \epsilon)p_j/q_{jt}$  at time  $t < T$  with an expected loss (i.e. profit – cost) of  $p_j v_j((1 - \epsilon)p_j/q_{jt} - 1)$ . For a myopic auctioneer with initial prior  $(1, \beta)$ , the total loss of revenue for the advertiser till time  $T$  is  $\Omega(v_j p_j \beta)$  which can be arbitrarily large.

## 6. MULTI-SLOT AUCTION

In this section, we generalize the hybrid auction to multiple slots under the standard separable CTR assumption, such that the resulting generalization is truthful in a myopic setting analogous to Section 3. Assume there are  $K$  slots, where slot  $i$  is associated with a CTR multiplier  $\theta_i \in [0, 1]$ . Slot 1 is the topmost slot; since the CTRs decrease with slot number, we have  $1 = \theta_1 \geq \theta_2 \geq \dots \geq \theta_K \geq 0$ . We will also define  $\theta_{K+1} = 0$ .

As before, advertiser  $j$  and the auctioneer maintain priors on the CTR value for this advertiser in ad slot 1. As before, we denote these priors as  $\mathcal{P}_j$  and  $\mathcal{Q}_j$  respectively. Let  $p_j = \mathbf{E}[\mathcal{P}_j]$  be the expected CTR estimated by the advertiser, and let  $q_j = f(\mathcal{Q}_j)$  denote the Gittins index (or for that matter, any other function) of the auctioneer’s prior. Let  $v_j$  denote the true per-click valuation of advertiser  $j$ . Note that the priors  $\mathcal{P}_j$  and  $\mathcal{Q}_j$  correspond to the estimated CTR for advertiser  $j$  in the *first ad slot*, so that the expected CTR for the  $i^{\text{th}}$  slot based on the advertiser’s estimate is  $\theta_i p_j$ .

Advertiser  $j$  bids  $(m_j, c_j)$ , which is interpreted as the per-impression and per-click bids for obtaining the *first slot*. The auction is modeled after the laddered auction in [1], which is equivalent to VCG under the separability assumption [4]. First, compute the effective bid  $R_j = \max\{m_j, c_j q_j\}$  for every advertiser as described in section 2. Assume without loss of generality that there are  $K + 1$  advertisers, and that  $R_1 \geq R_2 \geq \dots \geq R_{K+1}$ . Then, the auction proceeds as follows:

1. Advertiser  $j$  is placed in slot  $j$ , for  $1 \leq j \leq K$ .
2. An “effective charge”,  $e_j$  is computed for advertiser  $j$  as  $e_j = \sum_{i=j}^K \left( \frac{\theta_i - \theta_{i+1}}{\theta_j} \right) R_{i+1}$ .
3. If  $m_j > c_j q_j$  then the advertiser is charged  $e_j$  per impression; else it is charged  $e_j/q_j$  per click.

It is easy to see that  $e_j \theta_j = R_j(\theta_j - \theta_{j+1}) + e_{j+1} \theta_{j+1}$ . Informally, advertiser  $j$ ’s effective charge is the same as the effective bid of the  $(j + 1)$ -th advertiser for the additional click-rate at the  $j$ -th position, and the same as the effective charge of the  $(j + 1)$ -th advertiser for the click-rate that would have already been realized at the  $(j + 1)$ -th position.

**THEOREM 6.1.** *If  $p_j = \mathbf{E}[\mathcal{P}_j]$  and the advertiser is myopic and risk-neutral, then regardless of the choice of  $q_j$ , the (strongly) dominant strategy is to bid  $(v_j p_j, v_j)$ .*

The proof of the above theorem is obtained by extending the proof of theorem 3.1 exactly along the line of the proof of truthfulness of the laddered auction in [1], and is omitted from this version. This proof can also be obtained using the analysis of VCG with probabilistic allocations, due to Myerson [9].

## 7. CONCLUSION

Advertising is a major source of revenue for search engines and other web-sites, and a major driver of innovation in web technology and services. Advertising spots are typically sold on the web using auctions, and these auctions have typically been either Cost-Per-Click (CPC), Cost-Per-Impression (CPM), or Cost-Per-Action (CPA). We defined a single-slot hybrid auction, which allows advertisers to enter per-impression as well as per-click bids. We showed that this auction is truthful for risk-neutral, myopic advertisers, the setting under which such auctions have typically been analyzed. When advertisers are risk-seeking, or non-myopic, or when the advertiser has much better information about the Click-Through-Rate (CTR) than the auctioneer, we show that the hybrid auction offers stronger revenue guarantees and advertiser flexibility than either pure CPC or CPM. The hybrid auction generalizes naturally to multi-slot scenarios

and is equally applicable to (CPM,CPA) or (CPC,CPA) bidding. Finally, the hybrid auction is fully backwards compatible with a CPC auction, in the sense that advertisers entering (optional) per-impression bids in addition to per-click bids can seamlessly co-exist with advertisers entering only per-click bids in the same auction.

## 8. REFERENCES

- [1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. *Proceedings of the seventh ACM conference on Electronic Commerce*, pages 1–7, June 2006.
- [2] S. Athey and I. Segal. An efficient dynamic mechanism. 2007. Available at <http://kuznets.fas.harvard.edu/~athey>.
- [3] A. Bapna and T. Weber. Efficient dynamic allocation with uncertain valuation. *Working Paper, Department of Management Science and Engineering, Stanford University*, 2005.
- [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*, 97(1):242–259, 2007.
- [5] J. C. Gittins. Bandit processes and dynamic allocation indices. *J Royal Statistical Societe Series B*, 14:148–167, 1979.
- [6] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. Sponsored search auctions. In *Algorithmic Game Theory*, edited by Nisan, Roughgarden, Tardos, and Vazirani, 2007.
- [7] A.W. Marshall and I. Olkin. *Inequalities: theory of majorization and its applications*. Academic Press (Volume 143 of Mathematics in Science and Engineering), 1979.
- [8] A. Mehta, A. Saberi, U. Vazirani, , and V. Vazirani. Adwords and generalized online matching. *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, 2005.
- [9] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [10] M. Richardson, E. Dominowska, and R. Ragno. Predicting clicks: estimating the click-through rate for new ads. *Proceedings of the 16th international conference on World Wide Web*, pages 521–530, 2007.
- [11] H. Varian. Position auctions. *International Journal of Industrial Organization*, October 2006.
- [12] R. Weber. On the gittins index for multiarmed bandits. *Annals of applied probability*, 2(4):1024–1033, 1992.