



Auctions for online ad space among advertisers sensitive to both views and clicks

Patrick Maillé, Bruno Tuffin

► To cite this version:

Patrick Maillé, Bruno Tuffin. Auctions for online ad space among advertisers sensitive to both views and clicks. *Electronic Commerce Research*, Springer Verlag, 2017, 10.1007/s10660-017-9267-6 . hal-01241308v2

HAL Id: hal-01241308

<https://hal.inria.fr/hal-01241308v2>

Submitted on 14 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Auctions for online ad space among advertisers sensitive to both views and clicks

Patrick Maillé

IMT Atlantique

2, rue de la Châtaigneraie, 35576 Cesson Sévigné Cedex, France,
`patrick.maille@imt.fr`

Bruno Tuffin

Inria Rennes Bretagne Atlantique

Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France
`bruno.tuffin@inria.fr`

July 19, 2017

Abstract

Advertisement in dedicated webpage spaces or in search engines sponsored slots is usually sold using auctions, with a payment rule that is either per impression or per click. But advertisers can be both sensitive to being viewed (brand awareness effect) and being clicked (conversion into sales).

In this paper, we generalize the auction mechanism by including both pricing components: the advertisers are charged when their ad is displayed, and pay an additional price if the ad is clicked. Applying the results for Vickrey-Clarke-Groves (VCG) auctions, we show how to compute payments to ensure incentive compatibility from advertisers as well as maximize the total value extracted from the advertisement slot(s). We provide tight upper bounds for the loss of efficiency due to applying only pay-per-click (or pay-per-view) pricing instead of our scheme. Those bounds depend on the joint distribution of advertisement visibility and population likelihood to click on ads, and can help identify situations where our mechanism yields significant improvements. We also describe how the commonly used generalized second price (GSP) auction can be extended to this context.

1 Introduction

An increasing number of online services (search engines, news, e-mail, video, maps, ...) are provided free of charge, with a business model based on advertising: upon accessing the service, the user is presented with advertisements,

usually next to (or before, for videos) the data of initial interest. The available advertisement space is limited (among others, to maintain the service attractiveness for users), and is generally allocated among advertisers using auction mechanisms: advertisers submit *bids*, in the form of prices, and advertisement slots are decided based on those bids and possibly other factors such as the advertisers' likelihood of being of interest (clicked) to users. The price paid by each advertiser may depend on all submitted bids, and would ideally satisfy a property called *incentive compatibility* or *truthfulness*, meaning that even strategic bidders cannot do better than truthfully declare the price they are willing to pay to get a slot.

In practice, two types of payment mechanisms are applied: either a price is paid each time the ad is displayed (*pay-per-impression* scheme), or each time the ad is clicked by a user (*pay-per-click* scheme). The former mechanism seems appropriate for advertisers aiming at building a "brand awareness effect" (getting the name of a brand or product known), while the latter rather corresponds to advertisers aiming at converting the ad into a more tangible outcome such as a product sale or a registration. Typically, Coca-Cola should be more interested in views, since clicks will hardly directly result in sales. Pay-per-click on the other hand more likely corresponds to online shops displaying their products. But many advertisers could be interested in both options.

For both of those payment types, appropriate pricing rules exist, and have been the subject of extensive research (see for example [11] and references therein). Note that some services use a fixed payment scheme (pay-per-click for most search engines like Google or Bing) while others (e.g., Facebook) leave the choice to the advertiser. However, to the best of our knowledge no platform proposes hybrid schemes, where a price would be paid per view and an additional charge would be applied if the ad is clicked. Actually the two models are mathematically equivalent when there are known values of the the click-through-rate (CTR) and click-to-sale-conversion probabilities because a view or a click can then be translated into a financial gain [7], even if it might actually be difficult to quantify the money gain associated to a view into in the case of brand awareness.

A hybrid mechanism similar to what we suggest has been defined and studied in [7, 8], under the name *impression plus click* pricing: the equivalence result is not satisfied there because the beliefs about the CTR are considered potentially different between the publisher and the advertiser, hence the need for a richer mechanism. An interesting feature of the proposed auction scheme is that despite the relaxation of this common knowledge assumption, the mechanism is still incentive-compatible. On the other hand, it requires strong assumptions on bounded valuations that we do not need here.

In this paper, we do not consider such discrepancies in CTR estimation—we assume it is the same for the advertiser and the publisher—but we rather go in a different direction in terms of assumption relaxation. We indeed investigate the situation where a common bid is used for a (possibly large) *set of webpages*, with different ad visibilities but also different likelihoods that an ad is clicked: typically, a page about fashion is more likely to bring clicks and then sales than

a page about news; similarly, for a given page the click likelihood can depend on the population viewing it, which can be determined using different means: knowing how the user landed on this page, having users connect using a social network account, etc. In this paper we will treat that case as distinct pages, without loss of generality.

Assuming such variability in ad visibility and click likelihood also implies that the equivalence result is not satisfied: we cannot simply and directly relate the valuation per impression and the valuation per click, since the click/impression ratio differs between pages. An advantage is that instead of requiring an estimation of the CTR for each page, the bids have to be provided once and for all independently of the web page (but under a separability assumption).

In such a context, it appears necessary to have advertisers declare *two* values, namely their willingness-to-pay for being clicked, and for being viewed. We show on simple examples that using one-dimensional bids leads to significant losses of efficiency (in the sense of the value obtained from the ad slots). On the other hand, with a Vickrey-Clarke-Groves (VCG)-based scheme applied for each webpage to 2-dimensional bids, we ensure truthful bidding (as a dominant bidding strategy) and maximized efficiency. We present bounds on the relative performance of one-dimensional bids with respect to our proposed scheme, and also explain how the commonly-used generalized second price (GSP) auctions can be extended to this context, with their pros and cons.

The remainder of this paper is organized as follows. Section 2 introduces our mathematical model, specifying our assumptions for ad visibility and click attractiveness. Section 3 then treats the case when each page only has one advertisement slot, describing the scheme and its social surplus performance bounds. Some of those results are generalized to the multiple-slot case in Section 4. Section 5 discusses our findings and proposes directions for future work.

2 Mathematical model

Consider advertisement slots to be sold on a web page or on a search engine. Those slots are sold through auctions; for more about such auctions, the reader is advised to look at [11]. In this paper, we consider that advertisers do not bid on each webpage separately: this is particularly true for ad banners in user-generated webpages such as forums; advertisers cannot treat separately each forum and estimate how much to bid to place an ad. Instead, we assume that advertisement slots are sold for *sets* of webpages (or combinations of keywords, in the search engine case). This paper proposes an auction scheme that can be applied on such a set: the characteristics of the set are revealed by the auctioneer, advertisers submit one bid for each set, and the auctioneer decides what ads to display and how much to charge for each visited page on the set. This section specifies our assumptions in terms of ad slot (random) visibility and user (random) click behavior, and advertiser preferences.

2.1 Advertisement visibility and click attractiveness

Inspired by the literature [6, 10] considering separable click-through rates, we describe here our model for the likelihood of an ad being *viewed* (depending on the page and the ad slot) and for the likelihood of that ad being *clicked* if viewed (depending on the page, and the ad itself). The mathematical decomposition of those effects is as follows:

- The **ad visibility** depends on the type of page visited, the size and location of the advertising slot, etc. It can be interpreted as the *probability* that the user actually *sees* the ad (although this may be difficult to measure), in particular we assume it is independent of the ads possibly displayed in the other ad slots. We denote by α_s the visibility of the ad displayed in the s^{th} ad slot of a given page: we model it as a random variable to take into account the variability in pages, sizes and locations. But for any given visited page, the vector $\vec{\alpha} = (\alpha_1, \dots, \alpha_m)$ is assumed known to the auctioneer, where m is the (max) number of slots available on each page. Also, without loss of generality we order ad slots on each page so that $\alpha_1 \geq \alpha_2 \geq \dots$, i.e., ad slot 1 being more visible than ad slot 2, etc. Note that the case when some pages have less than m ad slots is also covered: for such spaces one just has to set $\alpha_s = 0$ if there is no ad slot s .
- Once an ad is viewed, there is a probability that it is clicked by the user. We assume this probability is multiplicatively separable into
 - a **page effect**, which reflects the fact that the likelihood of a click depends on the type of page visited. For example, the likelihood of a click (to products related to the page) is larger for e-commerce webpages than for news webpages; or similarly, general knowledge about the audience of a webpage (age, gender, location) can give information about the likelihood to get a click. We represent that effect through another random variable β .
 - an **ad effect**, which represents the intrinsic attractiveness of the ad, and which we denote by CTR_j for advertiser j .

To summarize, in our model each webpage of the auctioned set is characterized by a pair $(\vec{\alpha}, \beta)$, with $\vec{\alpha}$ a vector giving the visibility value α_s for each slot s in the page. When displayed in slot s on such a page, the ad from an advertiser j is:

- viewed with probability α_s ;
- clicked with probability $\alpha_s \beta \text{CTR}_j$.

Given that advertisers do not submit a bid per page but only a global bid for the whole set of pages, they do not need to know the values of β and $\vec{\alpha}$ for each individual page: only the *joint distribution* of $(\vec{\alpha}, \beta)$ is needed by advertisers to decide how to bid; we assume that the auctioneer reveals this information to

advertisers before the auction starts, so that advertisers would not need to carry out statistical estimation campaigns on their own for the set of webpages. A natural incentive for the auctioneer is that without this entry cost, advertisers are also more likely to join the system.

Assuming β and $\vec{\alpha}$ to be known to the auctioneer raises some implementability questions:

- β characterizes the clicking behavior on a page, which we believe stems from the *population* visiting that page (which includes, e.g., the purpose of their visit such as e-shopping versus news catching). We think such aspects can be estimated by the auctioneer, correlating browsing and click behaviors;
- the estimation of $\vec{\alpha}$ may be hard in practice since visibility cannot be directly measured. However, we believe such an estimation is possible when β is known: by displaying the same ad on several pages (or slots) and comparing the numbers of clicks, one should be able to estimate the visibility of the particular page-slot pair.

Designing the most efficient estimation methods for β and $\vec{\alpha}$ based on observed browsing and clicking behavior is beyond the scope of this paper, but deserves some attention if our scheme is to be implemented in practice. Also, strategic auctioneer behaviors consisting in revealing false distributions are not considered in this paper; we think they may be detected through independent estimations so that auctioneers would not risk implementing them.

2.2 Advertiser preferences

Consider a set \mathcal{J} of advertisers. Each advertiser is potentially interested in both views and clicks. More exactly, Advertiser $j \in \mathcal{J}$ is characterized by a pair

$$(b_j^v, b_j^c)$$

where b_j^v is its valuation for a view (the brand awareness effect) and b_j^c its valuation for a click (related to potential sales). Due to the large number of page visits, it is reasonable to assume that advertisers are *risk-neutral*, i.e., they are interested in the value they can extract over the whole set of webpages and of visits.

Under the assumptions of the previous subsection, for an Advertiser j characterized by valuations (b_j^v, b_j^c) , the expected value of ad slot s at web page $(\vec{\alpha}, \beta)$ is

$$U_j(\vec{\alpha}, \beta) = \alpha_s (b_j^v + \beta \text{CTR}_j b_j^c). \quad (1)$$

That value quantifies how much the advertiser thinks its ad in slot s is worth, i.e., how much the advertiser would be willing to pay to obtain this slot (instead of not being displayed at all). Following the usage in game-theory literature, we will indifferently refer to that value as the advertiser *valuation* or *willingness-to-pay*. The corresponding *net utility* would then be the difference between the willingness-to-pay and the price actually paid.

Advertisers being sensitive to the *aggregated* effects of their ad behavior, each advertiser $j \in \mathcal{J}$ is assumed to behave (here, bid) in order to maximize its *expected net utility*, that is, the expectation

$$\mathbb{E}_{(\vec{\alpha}, \beta)}[U_j(\vec{\alpha}, \beta) - p_j(\vec{\alpha}, \beta)] = \mathbb{E}_{(\vec{\alpha}, \beta)}[\alpha_{s_j} (b_j^y + \beta \text{CTR}_j b_j^c) - p_j(\vec{\alpha}, \beta)], \quad (2)$$

where s_j is the slot that Advertiser j obtains when page $(\vec{\alpha}, \beta)$ is visited, and $p_j(\vec{\alpha}, \beta)$ is the (expected) amount that Advertiser j is charged because of that visit.

2.3 Auctioning for ad space over a set of webpages

In this paper, we suggest that publishers offer a simple interface to advertisers: advertisers are asked to declare how much they are willing to pay to be viewed, and—when displayed—what *additional* price they would be willing to pay to have their ad clicked. Those bids are then used for a possibly large set of webpages which has been agreed upon beforehand between the auctioneer and each particular advertiser. Note that the set of webpages can differ among advertisers: each advertiser j may define a set of webpages W_j it is interested in, although we do not develop that aspect in this paper because it is not its purpose to discuss the strategies on the slots to target. The difference for our analysis would be a decomposition of the set of ads into $2^{|\mathcal{J}|} - 1$ subsets, corresponding to all potential combinations of bidders, and the surpluses and expected revenues would just be a sum over all those subsets. We therefore simplify the notations by assuming without much loss of generality that the advertisers submit their bid on all slots.

The auctioneer allocates the ad slot(s) of a page in order to maximize the generated value, based on bids, i.e., on declared values $(\tilde{b}_j^y)_{j \in \mathcal{J}}$ and $(\tilde{b}_j^c)_{j \in \mathcal{J}}$. Focusing on a single slot, it should be given to the advertiser with the highest value for $\tilde{b}_j^y + \beta \text{CTR}_j \tilde{b}_j^c$. Intuitively, if a webpage is not likely to attract clicks (β small), its ad slot will rather be allocated to advertisers focusing on building a brand effect; on the contrary if β is large then the slot should be given to an advertiser with a high valuation for clicks.

2.4 Performance metrics

As is classical in mechanism design, a key property will be *dominant-strategy incentive compatibility* (or *incentive compatibility* in short), meaning that for any bid profile of one's competitors, one's best bidding strategy is to reveal one's willingness-to-pay for the ad slot resource. In this paper we will directly apply some well-known results on second-price auctions and their generalizations, which ensure incentive compatibility.

The other performance metric will be relative to the value that is extracted from the whole set of ad slots, and is generally called *social welfare*, to highlight the fact that we consider the society as a whole. Nevertheless, here we do not incorporate users in the metric, while they may also have gains or losses because

of those ad slots: modeling on the one hand the ad nuisance effect and on the other hand the potential gains from interesting ads is beyond the scope of this paper. As a result, what we will call social welfare—and aim at maximizing—is the sum of the net utilities for advertisers (given in (2)) and for the auctioneer (that is, the auctioneer’s revenue). Note that it is a standard definition when studying sponsored search auctions [10, 14, 15]. Since the payments from advertisers to the auctioneer cancel out, social welfare will come down to the sum of the expected values of advertisers’ valuations.

Mathematically, social welfare will therefore be

$$\text{SW} = \sum_{\text{Adv. } j} \mathbb{E}_{(\vec{\alpha}, \beta)} [\alpha_{s_j} (b_j^y + \beta \text{CTR}_j)],$$

with again s_j the slot allocated to Advertiser j on webpages characterized by $(\vec{\alpha}, \beta)$.

Social welfare will of course depend on the slot allocation scheme and the advertiser bidding behavior. An efficient auction scheme would incentivize advertisers to bid so that the resulting social welfare is (not far from) optimal.

3 The one-slot case

For simplicity of exposition, we first treat here the case when only one advertisement slot is available on each webpage. In that case, the vector $\vec{\alpha}$ is just a scalar $\alpha \geq 0$.

3.1 The *impression plus click* second-price auction scheme

We suggest that *on each webpage* w (with given values for (α, β)), the publisher performs a second-price auction [16] to allocate the advertising slot. The multiple-slot case will be discussed in the next section.

3.1.1 2D-auction rules

Definition 1 (2D-auction scheme). *Based on the submitted two-dimensional bids $(\tilde{b}_j^y, \tilde{b}_j^c)_{j \in \mathcal{J}}$, on a webpage with ad visibility α and click-likelihood β ,*

- *the slot is given to an advertiser j_1 with the highest declared expected value from being displayed*

$$j_1 \in \arg \max_{j \in \mathcal{J}} \underbrace{\tilde{b}_j^y + \beta \text{CTR}_j \tilde{b}_j^c}_{:=m_j(\beta)}$$

- *that advertiser is charged αp for each impression (or equivalently, p per expected view) where p is the second-highest declared valuation*

$$p := \max_{j \in \mathcal{J} \setminus j_1} \{\tilde{b}_j^y + \beta \text{CTR}_j \tilde{b}_j^c\}. \quad (3)$$

Note from the definition that the winner identity depends on β , and the price charged depends on β and α .

3.1.2 2D-auction incentive and surplus properties

As a second-price auction, the 2D-auction has good incentive properties, as stated below.

Proposition 1. *The 2D-auction scheme of Definition 1 satisfies the truthfulness property: for each Advertiser j , submitting its true pair of valuations $(\tilde{b}_j^v, \tilde{b}_j^c) = (b_j^v, b_j^c)$ is a dominant strategy.*

Moreover, the average social surplus per visited page is

$$S_{2D} = \mathbb{E} \left[\max_{j \in \mathcal{J}} \alpha (b_j^v + \beta CTR_j b_j^c) \right]. \quad (4)$$

Proof. On any page (α, β) , the ad slot is allocated to the advertiser who values it most in the sense of the declared valuation $\tilde{U}_j := \alpha (\tilde{b}_j^v + \beta CTR_j \tilde{b}_j^c)$, which corresponds to the real valuation when advertisers bid truthfully. The price paid then corresponds to the so-called *social opportunity cost* of the winner's presence, that is here simply the highest valuation among losing advertisers. We therefore exactly have a Vickrey auction using the declared values $(\tilde{U}_j)_{j \in \mathcal{J}}$, which is known to be incentive compatible [12]: for each webpage the advertisers cannot do better than being truthful about their valuation U_j . Since the value of β is unknown to advertisers, the only way to ensure that the value considered by the auctioneer is the true one for each page is for advertisers to truthfully declare their value-per-view and value-per-click. \square

Note that another implementation of the 2D pricing scheme can be in the form of impression+click pricing, as suggested in [8], instead of the per-impression price of (3) charged each time: Advertiser j_1 would be charged only $\alpha \tilde{b}_{j_2}^v$ when displayed, and an additional $\frac{CTR_{j_2}}{CTR_{j_1}} \tilde{b}_{j_2}^c$ if the ad is clicked. That implementation gives the same expected payment αp with p given in (3), and therefore will be considered equivalent in this paper. Nevertheless, it may be preferred by advertisers since being closer to the actual user behavior (the extra charge for a click being applied only when a click occurs instead of being averaged for each impression). But note on the other hand that it induces more variance in prices.

Given that the parameter β differs among pages (a difference with [7, 8]), there is no direct equivalence between a price per view and a price per click when submitting bids. Hence, as we illustrate later, using one-dimensional bids (interpreted as per impression, or per click) leads to an inevitable loss of efficiency: the ad slots cannot always be allocated to the advertiser valuing it most. At least a two-dimensional bid is needed, and we argue it is sufficient to ensure optimality.

Of course, separate auctions could be run independently on each webpage. The advantage of our method is that the bidding process is used once, for

a possibly very large set of webpages, and remains very simple (only a two-dimensional bid is sent by each advertiser at the beginning of its campaign) while still permitting advertisers to play their dominant-strategy bid on *each* page. Said differently, even if advertisers were allowed to submit a different bid on each page, they could not get a higher utility.

3.2 1D-auction schemes

For comparison purposes, we define the (classical) one-dimensional bidding auction: even if advertiser utilities are bi-dimensional, they are only allowed to submit a one-dimensional bid, representing either their valuation for being clicked or for being displayed.

3.2.1 1D-auction rules

Let us first define the allocation and pricing rules in a general manner.

Definition 2 (1D-auction scheme). *Based on values $(\tilde{v}_j)_{j \in \mathcal{J}}$ representing the average value of being displayed, and computed from one-dimensional bids:*

- the slot is given to an advertiser j with maximum value \tilde{v} :

$$j \in \arg \max_{i \in \mathcal{J}} \tilde{v}_i;$$

- the (average) price paid per impression is the second-highest value $p_j = \max_{i \in \mathcal{J} \setminus \{j\}} \tilde{v}_i$.

This definition is independent of the bid interpretation:

- if a bid represents the value-per-impression, then \tilde{v}_j is directly the submitted bid of Advertiser j ;
- if bids are interpreted as representing the value-per-view, then \tilde{v}_j can be obtained from the bid b_j as $\tilde{v}_j = \mathbb{E}[\alpha]b_j$;
- if bids are interpreted as a value-per-click, then \tilde{v}_j is computed from bid b_j as: $\tilde{v}_j = b_j \cdot \mathbb{E}[\alpha\beta]CTR_j$.

Also, both price implementations are covered:

- in a pay-per-impression scheme, the displayed Advertiser j is charged p_j at each impression;
- in a pay-per-click scheme, Advertiser j is charged $\frac{p_j}{\mathbb{E}[\alpha\beta]CTR_j}$ when clicked, hence on average p_j per impression.

Note that the definition above allows some implementations which are counter-intuitive or make little practical sense, such as, for example, bids being interpreted as value-per-impression but advertisers being charged only when clicks

occur. Actually we have 6 possible combinations for the (*bid interpretation, price implementation*) pair, which are all equivalent when the distribution of (α, β) is known to the auctioneer and to advertisers, as we illustrate now. We first show that all bid interpretations are equivalent, since in all cases Advertiser j can ensure that any chosen value v_j is considered by the auctioneer:

- with *value-per-impression* bids, Advertiser j just has to bid v_j ;
- with *value-per-view* bids, Advertiser j knows that to compute \tilde{v}_j , its bid will be multiplied by $\mathbb{E}[\alpha]$ by the system. It should then declare $\frac{v_j}{\mathbb{E}[\alpha]}$;
- with *value-per-click* bids, Advertiser j knows that \tilde{v}_j will be obtained from the bid through a multiplication by $CTR_j \mathbb{E}[\alpha\beta]$, i.e., the expected probability of being clicked when displayed. It just has to bid $\frac{v_j}{CTR_j \mathbb{E}[\alpha\beta]}$.

As a result, since we assumed the joint distribution of α and β to be common knowledge among all actors, advertisers can have the mechanism consider any value v_j they want, through very simple manipulations for any interpretation of the bids.

Similarly, as regards the price implementation, once the mechanism has computed the (average) price per display p_j charged to Advertiser j :

- under *pay-per-click* the mechanism will charge $\frac{p_j}{CTR_j \mathbb{E}[\alpha\beta]}$ for each click,
- under *pay-per-impression* the charge p_j will be applied at each impression.

Hence again, because of the distribution of (α, β) being known, both implementations are equivalent: Advertiser j will pay on average p_j for each impression. Therefore, all combinations are mathematically equivalent, even if those with non-matching bid interpretation and price implementation are not likely to be chosen by the auctioneer.

3.2.2 1D-auction properties

Let us formalize the equivalence between all the variants of the 1D-auction, as well as their incentive properties.

Lemma 1. *For the 1D-auction scheme of Definition 2, both pay-per-impression and pay-per-click implementations are equivalent in average, in the following sense:*

1. *for any interpretation of the bid (as the willingness-to-pay for an impression, for a view, or for a click), truthfully declaring one's (corresponding) expected value is the only dominant strategy;*
2. *the same advertiser is displayed in the ad slot of all webpages, and is an advertiser j_0 with maximum expected value for being displayed*

$$j_0 \in \arg \max_{j \in \mathcal{J}} \mathbb{E} [\alpha (b_j^v + \beta CTR_j b_j^c)]$$

3. the resulting average surplus (per visited webpage) is the maximum expected value

$$S_{1D} = \max_{j \in \mathcal{J}} \mathbb{E} [\alpha(b_j^y + \beta \text{CTR}_j b_j^c)] \quad (5)$$

4. the average revenue per visited webpage is the second-highest maximum expected value

$$R_{1D} = \max_{j \in \mathcal{J} \setminus j_0} \mathbb{E} [\alpha(b_j^y + \beta \text{CTR}_j b_j^c)]. \quad (6)$$

Proof. The 1D-auction actually allocates only one indivisible good (the right to be displayed in the ad slot of *all* webpages), and allocates it using a classical Vickrey auction: bidders are asked (directly or indirectly, depending on the bid interpretation) to declare how much they are willing to pay to obtain that good, which is then allocated to the highest bidder who is charged (on average) the second-highest declared value. It is well-known and easy to check that bidding truthfully is the only dominant strategy.

In our setting, this means that an advertiser j characterized by its valuations (b_j^c, b_j^y) should bid so that the value considered by the auction equals its truthful (expected) valuation for being displayed, that is $\mathbb{E} [\alpha(b_j^y + \beta \text{CTR}_j b_j^c)] = \mathbb{E}[\alpha]b_j^y + \text{CTR}_j \mathbb{E}[\alpha\beta]b_j^c$. If bids are interpreted as values per impression, that bidder should then directly declare that value, while it should divide it by the click probability $\text{CTR}_j \mathbb{E}[\alpha\beta]$ (resp., $\mathbb{E}[\alpha]$) if bids are interpreted as values per click (resp., per view).

The two other points of the lemma are then direct consequences of the truthfulness property: the ad slot goes to an advertiser with maximal expected valuation, and who is charged on average the second-highest expected valuation for each display. \square

The main consequence of Lemma 1 can be summarized as follows: for a given set of bids, the 1D-auction schemes defined above will select only one winning advertiser, which will be displayed on *all pages*. In our model where advertisers' valuations for views and clicks differ, and where clicking behavior varies among pages, this leaves some space for improvement, which our 2D-scheme accomplishes by displaying *on each page* the advertiser with the highest valuation.

3.3 Relative performance of 1D and 2D auctions

The following proposition compares the performance of one-dimensional bidding (pay-per-click or pay-per-impression) with our two-dimensional bidding scheme.

Proposition 2. *Recall that S_{1D} (resp. S_{2D}) are for the expected surplus over the set of webpages over which the auction is performed for the 1D (resp. 2D) auction. Then if advertisers are assumed to play their (unique) dominant strategy, the respective surpluses raised by both auction schemes verify:*

$$1 \leq \frac{S_{2D}}{S_{1D}} \leq \mathbb{E} \left[\alpha \max \left(\frac{1}{\bar{\alpha}}, \frac{\beta}{\alpha\bar{\beta}} \right) \right] \quad (7)$$

with $\bar{\alpha} := \mathbb{E}[\alpha]$ and $\overline{\alpha\beta} := \mathbb{E}[\alpha\beta]$.

Moreover, those bounds are tight, i.e., for any distribution of (α, β) there exists a valuation profile such that $S_{2D} = S_{1D}$ and a valuation profile such that $S_{2D} = S_{1D} \mathbb{E} \left[\alpha \max \left(\frac{1}{\bar{\alpha}}, \frac{\beta}{\overline{\alpha\beta}} \right) \right]$.

While the lower bound is expected since 2D auctions make use of more information, the upper bound characterizes the maximum gain in terms of α and β . In particular, if that bound is sufficiently close to 1 then the efficiency improvement of the 2D-scheme is not worth the extra complexity: it is sufficient to implement a 1D-scheme and display the same ad on all the webpages of the set.

Proof. To simplify notations, define for each $i \in \mathcal{J}$

$$a_i := \text{CTR}_i b_i^c.$$

The expressions for S_{1D} and S_{2D} are then:

$$\begin{aligned} S_{1D} &= \max_{i \in \mathcal{J}} \mathbb{E} [\alpha \{b_j^y + \beta a_j\}] = \max_{i \in \mathcal{J}} \{\bar{\alpha} b_j^y + \overline{\alpha\beta} a_j\}, \\ S_{2D} &= \mathbb{E} \left[\alpha \max_{i \in \mathcal{J}} \{b_j^y + \beta a_j\} \right]. \end{aligned}$$

The inequality $S_{2D} \geq S_{1D}$ is trivial, as well as the fact it is tight (just assume all bidders in \mathcal{J} have the same values for a_j and b_j^y).

Now, notice that for all $j \in \mathcal{J}$, $\bar{\alpha} b_j^y + \overline{\alpha\beta} a_j \leq S_{1D}$, hence $a_j \leq \frac{S_{1D} - \bar{\alpha} b_j^y}{\overline{\alpha\beta}}$, and thus for any page (α, β) we have

$$b_j^y + \beta a_j \leq \frac{\beta}{\overline{\alpha\beta}} S_{1D} + b_j^y \left(1 - \frac{\beta \bar{\alpha}}{\overline{\alpha\beta}} \right).$$

Therefore, since $b_j^y \geq 0$ we have when $\beta > \frac{\overline{\alpha\beta}}{\bar{\alpha}}$

$$b_j^y + \beta a_j \leq \frac{\beta}{\overline{\alpha\beta}} S_{1D} \quad \text{for all } j \in \mathcal{J}.$$

Consequently,

$$\begin{aligned} S_{2D} &= \mathbb{E} \left[\mathbb{1}_{\{\beta \leq \frac{\overline{\alpha\beta}}{\bar{\alpha}}\}} \left(\alpha \max_{j \in \mathcal{J}} \{b_j^y + \beta a_j\} \right) + \mathbb{1}_{\{\beta > \frac{\overline{\alpha\beta}}{\bar{\alpha}}\}} \left(\alpha \max_{j \in \mathcal{J}} \{b_j^y + \beta a_j\} \right) \right] \\ &\leq \mathbb{E} \left[\mathbb{1}_{\{\beta \leq \frac{\overline{\alpha\beta}}{\bar{\alpha}}\}} \underbrace{\left(\alpha \max_{j \in \mathcal{J}} \left\{ b_j^y + \frac{\overline{\alpha\beta}}{\bar{\alpha}} a_j \right\} \right)}_{= \frac{\alpha}{\bar{\alpha}} S_{1D}} + \mathbb{1}_{\{\beta > \frac{\overline{\alpha\beta}}{\bar{\alpha}}\}} \frac{\alpha \beta}{\overline{\alpha\beta}} S_{1D} \right] \\ &= S_{1D} \mathbb{E} \left[\alpha \max \left(\frac{1}{\bar{\alpha}}, \frac{\beta}{\overline{\alpha\beta}} \right) \right], \end{aligned}$$

which gives the right-hand side of (7).

There remains to show that this last inequality is tight: let us fix a joint distribution for (α, β) and consider a setting with two advertisers, both having the same expected value $\bar{\alpha}b_1^v + \bar{\alpha}\beta a_1 = \bar{\alpha}b_2^v + \bar{\alpha}\beta a_2$ per visit for the slot (hence that common value is S_{1D}), but where:

- Advertiser 1 is only interested in clicks, i.e., $b_1^v = 0$ and thus $a_1 = \frac{S_{1D}}{\alpha\beta}$;
- Advertiser 2 is only interested in views, i.e., $a_2 = 0$ and $b_2^v = \frac{S_{1D}}{\alpha}$.

Then with 2D-bidding, the slot goes to Advertiser 1 if and only if $\beta > \frac{\alpha\beta}{\alpha}$, and

$$\begin{aligned} S_{2D} &= \mathbb{E} \left[\mathbb{1}_{\{\beta \leq \frac{\alpha\beta}{\alpha}\}} \underbrace{\alpha \left(\max_{i \in \mathcal{J}} \{b_i^v + \beta a_i\} \right)}_{=b_2^v + \beta a_2 = S_{1D}/\alpha} + \mathbb{1}_{\{\beta > \frac{\alpha\beta}{\alpha}\}} \underbrace{\alpha \left(\max_{i \in \mathcal{J}} b_i^v + \beta a_i \right)}_{=b_1^v + \beta a_1 = \frac{\beta}{\alpha\beta} S_{1D}} \right] \\ &= S_{1D} \mathbb{E} \left[\mathbb{1}_{\{\beta \leq \frac{\alpha\beta}{\alpha}\}} \frac{\alpha}{\alpha} + \mathbb{1}_{\{\beta > \frac{\alpha\beta}{\alpha}\}} \frac{\alpha\beta}{\alpha\beta} \right] \\ &= S_{1D} \mathbb{E} \left[\alpha \max \left(\frac{1}{\alpha}, \frac{\beta}{\alpha\beta} \right) \right]. \end{aligned}$$

□

3.4 An illustrative example

A particular case is when α and β are independent. For example, if in that case β is uniformly distributed over an interval $[0, \beta_{\max}]$, then from Proposition 2 we have

$$1 \leq \frac{S_{2D}}{S_{1D}} \leq \frac{5}{4},$$

i.e., using 2D-bidding instead of 1D-bidding can yield a 25% surplus improvement.

Let us for example assume that ad slots on all webpages have the same visibility α , and take $\alpha = 1$ without loss of generality. Hence webpages differ only in the likelihood β of ads attracting clicks; in this example, we consider that β is uniformly distributed over the interval $[0, 1]$.

Consider two advertisers with the same CTR, with Advertiser 1 only interested in clicks, i.e., characterized by valuations $(0, b_1^c)$, and Advertiser 2 only interested in views, i.e., characterized by valuations $(b_2^v, 0)$.

Behavior of our two-dimensional auction

Under our impression+click (2D-auction) scheme, from Proposition 1, advertisers submit truthful bids. Hence,

- if $\beta > \frac{b_2^v}{b_1^c \text{CTR}}$ the slot goes to Advertiser 1, who is charged b_2^v per impression;

- otherwise the slot goes to Advertiser 2, who is charged $b_1^c \beta \text{CTR}$ on average per impression (either as a price per impression, or as a price b_1^c per click).

Then, assuming $b_2^y < b_1^c \text{CTR}$, the expected value (over the random β) from the slot for each impression is

$$\begin{aligned}
\mathbb{E}[\max(b_1^c \beta \text{CTR}, b_2^y)] &= b_2^y \frac{b_2^y}{b_1^c \text{CTR}} + \int_{x=\frac{b_2^y}{b_1^c \text{CTR}}}^1 b_1^c x \text{CTR} dx \\
&= \frac{(b_2^y)^2}{b_1^c \text{CTR}} + \frac{b_1^c \text{CTR}}{2} \left(1 - \left(\frac{b_2^y}{b_1^c \text{CTR}}\right)^2\right) \\
&= \frac{(b_2^y)^2}{2b_1^c \text{CTR}} + \frac{b_1^c \text{CTR}}{2}. \tag{8}
\end{aligned}$$

When $b_2^y \geq b_1^c \text{CTR}$ the slot always goes to Advertiser 2 so that the expected slot value is simply b_2^y .

Behavior of a one-dimensional auction

On the other hand, let us investigate what would happen under a 1D auction, where each advertiser could only submit a one-dimensional bid, for example indicating its willingness-to-pay for a click (from Lemma 1, a pay-per-impression scheme would be equivalent).

Denoting by b_i the (one-dimensional) bid of Advertiser i for $i = 1, 2$, the auctioneer would for each page:

- sort bids in terms of their product *click-through-rate* \times *bid*—here since we assumed both advertisers have the same CTR, the ranking is just based on bids—;
- give the slot to the highest bidder (in terms of this product);
- charge (per impression, on average) the second-highest price (still in terms of the *click-through-rate* \times *bid* product), hence for example $b_2 \frac{\text{CTR}_2}{\text{CTR}_1}$ per click if Advertiser 1 obtains the slot [11]. Here, with both advertisers having the same CTR, the price per click is just the second-highest bid.

The problem here is that the click likelihood is a random variable, so that as evoked before there is no predictable equivalence between the value for being displayed and the value for being clicked. While Advertiser 1 can declare its true valuation $b_1 = b_1^c$, bidding in such a scheme is more tricky for the view-sensitive Advertiser 2, since it cannot bid so that its truthful valuation be taken into account for each page. Since bidders are assumed risk-neutral, and as stated in Lemma 1, Advertiser 2 will declare its truthful *expected* valuation for the slot (in terms of value per click): since winning the auction will bring some (random) number N of views, and a corresponding expected number of clicks $N\bar{\beta}$, that bidder should declare it is willing to pay $b_2^y N$ for those views, and hence $b_2 = b_2^y / \bar{\beta}$ per click.

For our example, if $b_2 < b_1$ then the slot always goes to Advertiser 1, and hence has an expected value per display

$$\mathbb{E}[\beta \text{CTR} b_1^c] = \frac{\text{CTR} b_1^c}{2}.$$

On the other hand, if $b_1 < b_2$ the slot goes to Advertiser 2, for which it has an expected value per display b_2^y .

As a result, if advertisers bid truthfully, the expected value per impression is

$$\max\left(b_2^y, \frac{\text{CTR} b_1^c}{2}\right). \quad (9)$$

Behavior of a pay-per-impression auction

Similarly, a pay-per-impression auction would ask for a one-dimensional bid and always give the slot to the same bidder, yielding the same maximum expected value as the one in (9).

Figure 1 shows that expected value per ad impression for our 2D-bid scheme and for the two one-dimensional bidding schemes (click-based or impression-based) when $b_2^y = 0.01$, $\text{CTR} = 0.01$ and b_1^c varies. The second graph displays the ratio of allocated slot values. Our 2D-auction reaching maximum efficiency, a value 1 means an efficient outcome. A loss of efficiency can be observed for some per-click valuations b_1^c of Advertiser 1. Interestingly, this (relative) loss of efficiency is not monotonous as b_1^c increases. We also observe that the bound obtained in Proposition 2 is attained: when $b_1^c = \frac{b_2^y}{\beta \text{CTR}}$ we have $\frac{S_{2D}}{S_{1D}} = \frac{5}{4}$, or equivalently an 80% efficiency.

4 Extension to multi-slot pages

Our 2D-auction mechanism can quite easily be extended to the case where each page has a number m of ad slots. Let us index slots with the subscript s , $1 \leq s \leq m$. Recall that the slot visibilities are $(\alpha_s)_{1 \leq s \leq m}$ where $\alpha_1 \geq \dots \geq \alpha_m$, and the page effect β is the same for all slots of a given page. As before we also have the ad effect on clicks, that we still denote by CTR_j for advertiser j . Note that, since β is the same for all slots, the preference ranking of all advertisers is the same: they all prefer Slot 1 over Slot 2, etc.

As pointed out in Section 2, assuming m to be the same over all pages is without loss of generality.

Similar to what is suggested in [7], our auction scheme can be extended in two different ways.

- A first possibility is to apply the VCG principle to allocate and price ad slots on each visited page (see [16, 3, 9] or [12, 13] for a definition of VCG auctions). More precisely, for each page one would allocate slots to ads in order to maximize the “total declared value” of that allocation; and each

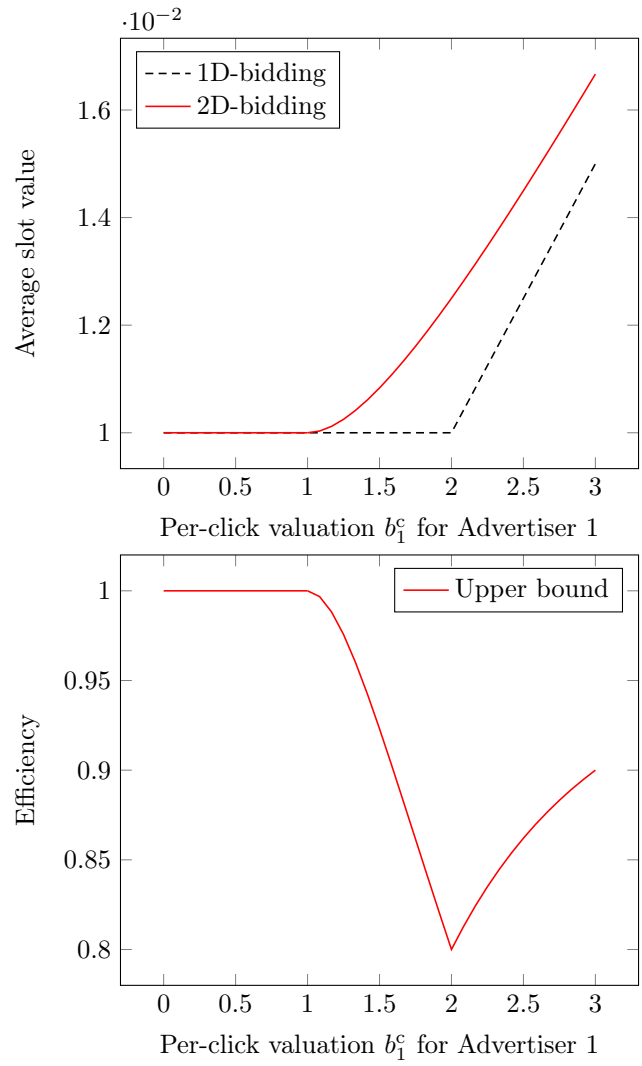


Figure 1: Efficiency of a 1D-bidding scheme

advertiser would be charged the cost its presence imposes on the others' total declared value. This approach is detailed in Subsection 4.1.

- The second approach is inspired by the Generalized Second Price mechanism, which is extensively applied in sponsored search auctions [11, 12]: as for VCG, the allocation is the most efficient one (in the sense of declared values); but the price each displayed advertiser pays is here the minimum one she could have declared and would have earned her the same slot. This approach is detailed in Subsection 4.2.

4.1 Applying VCG on each page

In VCG auctions, each bidder should declare how much she is willing to pay for all the possible outcomes. In our setting, each bidder is sensitive only to her allocated slot (if any) since our model assumes no interdependency among the view/click probabilities of jointly displayed ads.

Hence, the mechanism needs to know the willingness-to-pay of each bidder for all slots, which depend on the page characteristics $(\vec{\alpha}, \beta)$. As in the one-slot case, it is sufficient that each bidder j declares her per-view valuation b_j^v and per-click valuation b_j^c . Therefore the mechanism can simply ask each bidder j to declare a per-view bid \tilde{b}_j^v and a per-valuation bid \tilde{b}_j^c .

Vickrey-Clarke-Groves mechanisms are well-known for having three desirable properties:

- *Incentive Compatibility*, meaning that bidding truthfully is a dominant strategy (i.e., the best each bidder could do, whatever the others' actions);
- *Efficiency*, meaning that the equilibrium (in dominant strategies) of bidding strategies leads to a value-maximizing allocation;
- *Individual Rationality*, meaning that each bidder gets a non-negative utility (i.e., one never loses by entering the game).

In practice, once the bid values $(\tilde{b}_j^v, \tilde{b}_j^c)$ are collected from all bidders j , the mechanism follows the same method, independently, on each webpage, as described below.

1. The slots $1, \dots, m$ are allocated in order to maximize the declared valuation. More specifically, the advertisers $(\pi(s))_{1 \leq s \leq m}$, where $\pi(i)$ is the advertiser whose ad is displayed in slot i , are a solution of

$$\arg \max_{\pi(1), \dots, \pi(m)} \sum_{i=1}^m \alpha_i \left(\tilde{b}_{\pi(i)}^v + \beta \text{CTR}_{\pi(i)} \tilde{b}_{\pi(i)}^c \right) \quad (10)$$

s.t. $\pi(i) \neq \pi(i') \quad \forall i \neq i'$.

This is not computationally hard: since slots are such that $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m$, advertisers just need to be sorted by decreasing values of $\tilde{b}_j^v +$

$\beta \text{CTR}_j \tilde{b}_j^c$, the i^{th} advertiser of that ranking being then displayed in slot i for $i = 1, \dots, m$. Note that, as in the one-slot case, the ranking depends on the page only through the click-likelihood β .

2. VCG payments are so that each advertiser pays the loss of (declared) value her presence imposes on others: here only providers ranked lower than Advertiser j (and in the first $m + 1$ positions) are affected by Advertiser j 's presence. Had Advertiser j not been here, each of them would be one slot ahead, hence the price paid by Advertiser j is

$$p_j^{\text{VCG}} = \sum_{k=s(j)+1}^{m+1} (\alpha_k - \alpha_{k+1}) \left(\tilde{b}_{j(k)}^v + \beta \text{CTR}_{j(k)} \tilde{b}_{j(k)}^c \right),$$

where $s(j)$ is the slot allocated to Advertiser j on this page.

4.2 Using Generalized Second Price

In GSP, the allocation rule is the same as in VCG—hence, determined by (10)—, but the pricing rule is different, following a “next-price” principle [1]: Advertiser j pays the minimum bid (actually, the declared valuation) she could have bid while keeping the same ad slot, i.e., her price depends on the bid of the advertiser ranked just after her. Denoting by $s(j)$ the slot of Advertiser j and by $\pi(s(j)+1)$ the advertiser ranked just below her, we have an average price per display

$$p_j^{\text{GSP}} = \alpha_{s(j)} \left(b_{\pi(s(j)+1)}^v + \beta \text{CTR}_{\pi(s(j)+1)} b_{\pi(s(j)+1)}^c \right),$$

which can be decomposed in an *impression plus click* manner, with a price per impression

$$p_j^{\text{GSP imp}} = \alpha_{s(j)} b_{\pi(s(j)+1)}^v$$

and an additional price per click

$$p_j^{\text{GSP click}} = \frac{\text{CTR}_{\pi(s(j)+1)}}{\text{CTR}_j} b_{\pi(s(j)+1)}^c.$$

Indeed, on a visited webpage (α, β) such a pricing yields an average (per visit) payment due to clicks

$$\begin{aligned} \mathbb{E}[p_j^{\text{GSP click}} \mathbb{1}_{\{\text{ad clicked}\}}] &= \alpha_{s(j)} \beta \text{CTR}_j p_j^{\text{GSP click}} \\ &= \alpha_{s(j)} \beta \text{CTR}_{\pi(s(j)+1)} b_{\pi(s(j)+1)}^c. \end{aligned}$$

Despite not incentive compatible (see below for an illustration), this pricing scheme is more likely to be used than VCG for two reasons, as in traditional sponsored search auctions:

- it is claimed to be simpler to understand for advertisers (instead of depending on many bids, the price paid only depends on the next-high bid);
- it yields a larger revenue (still, with respect to VCG) to the slots owner [12].

4.3 A numerical example

The mechanisms being described, we now provide a comparison of their output, as well as a comparison with 1D schemes, by means of a numerical example.

Let us consider an example with m slots, for which $N^c + 1$ advertisers compete: we assume N^c advertisers are interested in clicks, while one advertiser (Advertiser 0) is interested in views. For simplicity, we assume that

- all click-interested advertisers have the same click-through-rate
- the willingness-to-pay values for a click are linearly distributed among click-interested advertisers, between b_{\min}^c and $b_{\max}^c > b_{\min}^c$: Advertiser k is willing to pay a price $b_{\min}^c + \frac{k-1}{N^c-1}(b_{\max}^c - b_{\min}^c)$
- the view probabilities are the same on all pages. More precisely, at each page the probability that an ad displayed in slot s is viewed is given in Table 1. Those numbers are taken from [4], and indicate the relative numbers of clicks on ranked results in a search-engine result page: even if the context is different here, this gives an idea of the influence of the position on the click likelihood.
- the page effect on the click likelihood is such that β is uniformly distributed on an interval of the form $[0, \beta_{\max}]$.

α_1	α_2	α_3	α_4	α_5	α_6	α_7	α_8	α_9	α_{10}
0.364	0.125	0.095	0.079	0.061	0.041	0.038	0.035	0.03	0.022

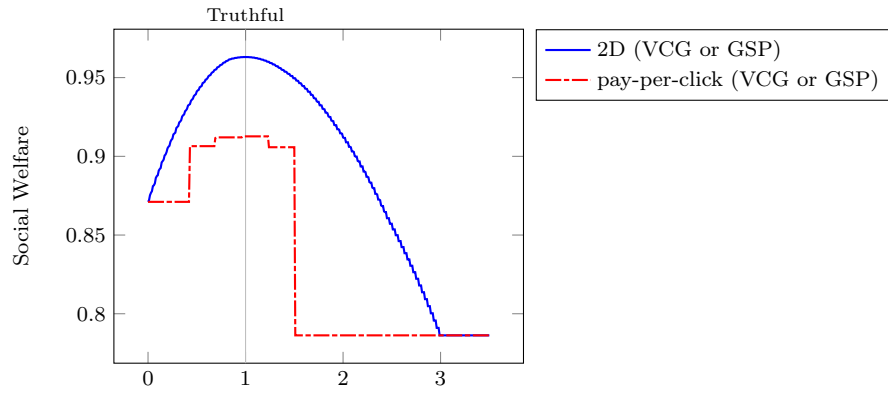
Table 1: Values for α_s used in the example.

Let us investigate a specific example, for which the parameter values are given in Table 2.

m	N^c	$\text{CTR}_k, k > 0$	CTR_0	b_{\min}^c	b_{\max}^c	b_0^v	β
5	6	0.1	0.05	10	100	1	$\mathcal{U}[0, 0.3]$

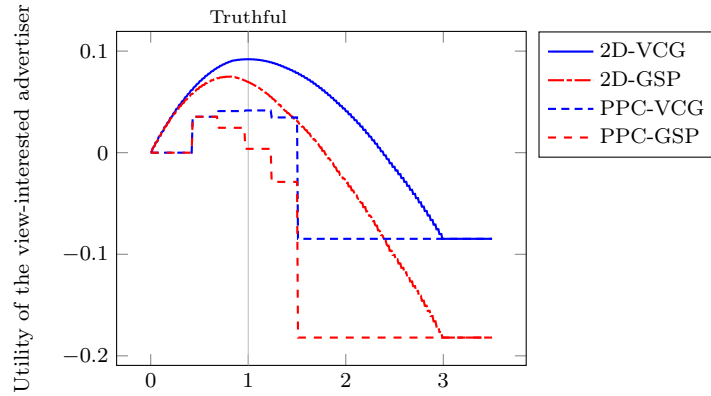
Table 2: Parameter values for the m -slot scenario considered.

We compare in Figures 2-4 our 2D-auction scheme—with VCG or GSP pricing—to a 1D-auction scheme, which can be seen as the currently used pay-per-click (PPC) scheme—implementing also the VCG and GSP pricing rules. We vary the bid (per click) of the view-interested advertiser, assuming the other advertisers (which are all click-interested) truthfully declare their willingness-to-pay for a click. Figures 2 and 3 respectively show the per-display expected surplus (that is, the total expected value that slots have for advertisers) and the view-interested valuation for the resulting allocation. The allocation rules being the same for VCG and GSP, the social welfare curves in Figure 2 coincide for both



Bid (per "equivalent view") of the view-interested advertiser

Figure 2: Average value extracted from the ad slots



Bid (per "equivalent view") of the view-interested advertiser

Figure 3: Average utility (value minus price) per visit for the view-interested advertiser

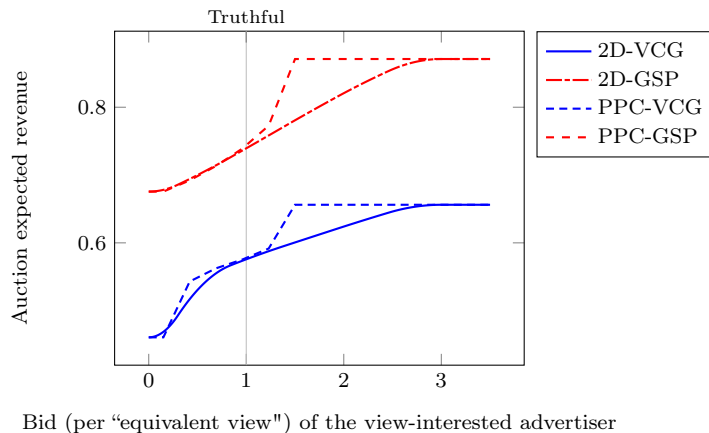


Figure 4: Auctioneer revenue for the different mechanisms

pricing rules. Figure 3 on the other hand distinguishes both schemes, and illustrates the incentive-compatibility property of the VCG pricing rule (both for 1D and 2D auctions): bidding truthfully maximizes the view-interested advertiser utility. Also, bidding truthfully ensures a non-negative utility (individual rationality property of the VCG scheme). Finally, both for 1D and 2D auction schemes, the reachable social welfare is maximized when bidders bid truthfully.

On the other hand, with GSP pricing truthfulness is not ensured: here Advertiser 0 (the view-interested advertiser) would be better off bidding below its actual (average) per-click valuation. However, determining the best bidding strategies with GSP pricing is a complex problem, that is beyond the scope of this paper: we rather refer to [6, 11, 15] for works related to this issue of optimal bidding when GSP is used. When such *bid shading* (bidding below truthful valuation) is implemented by advertisers, different ads may be displayed, which can result in a loss in the match between viewers and the displayed ads. Hence, as the authors of [2] recommend for the 1D-case, we think VCG should be preferred to GSP because of the VCG incentive properties.

Finally, Figure 4 shows the corresponding revenues. As stated before, if bidders bid truthfully, the GSP pricing rule yields a higher revenue than the VCG one. Somewhat counterintuitively, in this numerical experiment the 2D-auction revenues are lower than the 1D-auction revenues: even if more value (the social welfare) is extracted from the ad slots in the 2D-case, GSP pricing favors the auctioneer to an extent that exceeds this difference since the auctioneer extracts more revenue in a GSP-1D auction than a VCG-2D auction. A direct conclusion is that in terms of utility (valuation minus prices), advertisers suffer twice from the current most applied scheme: they pay more (due to the GSP rule) and get less from being displayed (from the 1D-auction) than if a 2D-auction with VCG payments were applied.

5 Conclusions

We proposed in this paper an *impression plus click* pricing scheme for online advertisement slots and sponsored searches, which allows to optimally allocate the slots among view- and click-sensitive advertisers. Our scheme presents the advantage of being simple to implement, and efficient. We obtained bounds on the relative efficiency with respect to bids based on clicks or impressions only. In the multiple-slot case, we discussed the VCG implementation of the scheme, but also the GSP one, more likely to be considered since GSP is extensively used for online advertising even if not incentive compatible.

Future works could be in several directions. First, the case where advertisers focus on different (but overlapping) subsets of pages or searches, as highlighted in Section 2.3, is worth studying more deeply, by including it in the decision variables of advertisers. Second, we would like to extend our scheme to the case where CTRs are estimated, in a way to also generalize the work in [7, 8]. Third, the case of budget limits for advertisers in ad campaigns is important to consider, but yields other difficulties since all page visits cannot be treated independently (the budget limits can be reached) [5]. Finally, our results will be strengthened with robustness checks when advertisers have uncertainty about measurements driving bid prices. All this was beyond the scope of the paper, which is centered on the definition of a new auction mechanism, but is worth investigating in a near future.

References

- [1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proc. of ACM conference on Electronic Commerce*, 2006.
- [2] Y. Bachrach, S. Ceppi, I. Kash, P. Key, and M. R. Khani. Mechanism design for mixed ads. In *Proc. of WWW*, 2016.
- [3] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [4] R. DeJarnette. Click-through rate of top 10 search results in Google, 2012.
- [5] S. Dobzinska, R. Lavib, and N. Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.
- [6] 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.
- [7] A. Goel and K. Munagala. Hybrid keyword search auctions. In *Proc. of 18th International World Wide Web Conference*, pages 221–230, 2009.
- [8] S. Goel, S. Lahaie, and S. Vassilvitskii. Contract auctions for sponsored search. In *Proc. of WINE*, 2009.

- [9] T. Groves. Incentives in teams. *Econometrica*, 41(3):617–631, Jul 1973.
- [10] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Proc. of 7th ACM Conference on Electronic Commerce*, pages 218–227, 2006.
- [11] P. Maillé, E. Markakis, M. Naldi, G. Stamoulis, and B. Tuffin. Sponsored search auctions: An overview of research with emphasis on game theoretic aspects. *Electronic Commerce Research Journal*, 12(3):265–300, 2012.
- [12] P. Maillé and B. Tuffin. *Telecommunication Network Economics: From Theory to Applications*. Cambridge University Press, 2014.
- [13] P. Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22, 1989.
- [14] R. Paes Leme and E. Tardos. Sponsored search equilibria for conservative bidders. In *Proc. of the 5th Workshop on Ad Auctions*, Stanford, CA, USA, Jul 2009.
- [15] H.R. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2005.
- [16] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, Mar 1961.