

User Modeling in Position Auctions: Re-Considering the GSP and VCG Mechanisms

Danny Kuminov*
Faculty of Industrial Eng. & Mgmt
Technion

Moshe Tennenholtz†
Microsoft Israel R&D Center &
Faculty of Industrial Eng. & Mgmt
Technion

ABSTRACT

We introduce a study of position auctions, with an explicit modeling of user navigation through ads. We refer to our model as the PPA model, since it is most applicable in the context of pay-per-action position auctions. In this model, which is consistent with other recent approaches to user modeling, a typical user searches sequentially over the ordered list of ads. At each point the user has some probability of performing the action associated with the given ad, some probability of moving to the next ad, and some probability of quitting. In the framework of this model, we re-consider two basic mechanisms: the VCG position auction and the generalized second price (GSP) position auction. We study properties of these mechanisms in the context of the PPA model, and in particular show that the GSP position auction possesses a pure strategy equilibrium, and characterize a set of its equilibria. Our main corollary is that the highest revenue one may obtain in an equilibrium of the GSP position auction matches the revenue obtained in the VCG position auction. This suggests that the VCG position auction is preferable to the ad publisher upon the GSP position auction in the context of the PPA position auctions model. The latter is in sharp distinction to the basic result of the usefulness of GSP over VCG in the standard position auctions model, where user behavior is not modeled explicitly. We also study various possibilities for corruption in the PPA setting, and the relative robustness of the mechanisms against the corresponding manipulations.

1. INTRODUCTION

In this work, we consider the online advertising setting, in which advertisement space on Internet pages is auctioned to advertisers. Currently, three main types of goods are traded in the online advertising market: impressions for brand awareness, clicks for traffic to web sites, and actions for actual sales (or other events that directly create revenue). Those correspond to two different goals of an advertising campaign – brand awareness and increased sales. In the former case, advertisers pay per impression (PPM)

*Email: dannykv@tx.technion.ac.il

†Email: moshet@microsoft.com

Cite as: User Modeling in Position Auctions: Re-Considering the GSP and VCG Mechanisms, Danny Kuminov, Moshe Tennenholtz, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 273–280
Copyright © 2009, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

while in the latter they can either pay per click on their ad (PPC) or pay per action (PPA). The standard PPC model is based on the assumption that an increase in the traffic to a web site can be translated to an increase in sales, by using the conversion rate parameter – the percentage of visitors who perform revenue-generating actions. The pay per action (PPA) model directly links payments to events such as online sales, phone calls, etc. The important distinction is that an “action” needs to be reported by the advertiser, whereas clicks are counted by the ad publisher.

The PPC model studied in the literature [11, 5] assumes that each ad slot has an associated click-through rate (CTR). In addition, it is assumed that each advertiser has a quality multiplier, which is multiplied by the CTR of the slot to which his ad is allocated to determine the number of clicks that his ad actually gets. The slots are sold in an auction, using a pay-per-revenue scheme, in which the advertisers declare their bid-per-click, and are allocated to slots in decreasing order of their bids multiplied by the corresponding quality scores. According to the generalized second price auction (GSP), which is the mechanism mostly used in practice, each advertiser pays an amount that equals the minimal bid he should have submitted in order to retain his current position. There are several inherent problems with the PPC setting, such as the so-called “click fraud” [10]. There are several works that try to tackle this issue from within the existing framework (e.g. [6]), but it seems that a satisfactory solution can be obtained only if we manage to abandon the click-based model and count revenue-generating actions directly. The PPA scheme is a relatively new model, and we are aware of only two studies that deal with it explicitly. [8] adopts the model used for the PPC context, with clicks replaced by actions. [9] designs a mechanism for repeated single-item auction with stochastically distributed agents’ valuations, in which only the winning agent reports his utility. While their mechanism has many desirable properties, it is not immediate how it can be applied to the task of auctioning multiple slots on the same page simultaneously.

In this paper we discuss a new, alternative model for position auctions, in which the user’s navigation through the ads is modeled explicitly. Since this model is especially appealing for the pay-per-action setting, we will refer to it as the PPA position auctions model. Nevertheless, this model is of interest also in the PPC setting. We show that the model describes a setting that is conceptually different from the well-known model of [11, 5] (that is, these models are irreducible one to the other). In the context of the PPA model, we analyze two well-known auction mechanisms, VCG and

GSP, compare their revenue, and arrive to the surprising conclusion that the revenue of the VCG mechanism may be higher, and is higher or equal to that of the GSP mechanism. This is in sharp contrast to the results obtained for the standard position auctions model, where user navigation is not modeled explicitly. Finally, we discuss possible methods of gaming the system by interested parties, and show the relative robustness obtained in the PPA setting.

An explicit modeling of user behavior is advocated in [4]; sequential navigation through the ordered ads is also the topic of [1]. These works however do not present a general equilibrium analysis of GSP, nor compare the revenue in GSP and VCG, as has been done for the standard ad auction setting [11, 5]. Such study, and its surprising ramification is the subject of this paper.

2. THE MODEL

We assume that there are four interested parties in the pay-per-action online advertising setting:

- The Internet users. Specifically, we are interested in users that search the Internet for a web site that supplies a specific good or service, and are willing to commit to a revenue-generating action in return for the good/service. The action may be a monetary transaction, an expression of long-term commitment to a service, etc.
- The set of advertisers. The advertisers are owners of web sites that are interested in attracting users' traffic, with the explicit purpose of increasing the number of revenue-generating actions performed at their site.
- The ad publisher. This is a dedicated Internet service, which is used by the advertisers to strategically place ads on Internet pages that are relevant to them, and is paid on the basis of a per-action commission (when the action is performed by a user following the advertisement).
- The content providers that create the medium in which advertisements are placed. They may coincide with the ad publisher (e.g., the classical Google search advertisement model), or they may be separate entities from which the ad publisher buys ad space (e.g., the Google Ad-Sense campaign). In this work, we use the terms "web page" and "medium" interchangeably.

The focus of this work is on modeling the behavior of Internet users, and building an auction mechanism for allocating ad slots in a given medium to different advertisers, in the pay-per-action setting. We assume that each time a web page is requested by a user, the auction is run anew and ads are placed in the ad slots.

We focus on a specific web page, which contains several slots in which ads can be placed by the ad publisher. Our model assumes that the ad slots are arranged in some linear order (for example, they are placed in a top-down order on the given web page). It is assumed that the user scans the ads in the given order, until he either finds an advertiser with which he commits to a revenue-generating action, or abandons the search. The decision process may include clicking on the ad and checking the web site, performing parallel search for information, etc.; this is transparent to our

model. Once one such appropriate ad is found, no further ads are considered. After an action is performed, the chosen advertiser reports this fact to the publisher and pays him the commission upon which they agreed.¹ Given this model of behavior, we want to define an auction mechanism that allocates ad slots to interested advertisers and determines their per-action payments, given their reported willingness to pay per single action.

The formal definition of our setting is as follows. There are N advertisers that compete for N ad slots.² The ad slots are numbered according to their order on the web page. We refer to the advertisers as the set of *agents*, and use the terms agents and advertisers interchangeably. We assume that a given user considers the advertisers (i.e. their corresponding ads) according to their order on the page, and that his decision is determined as follows:

- Each advertiser has a fixed and commonly known probability p_i , $0 < p_i < 1$ that a user will select its ad, given that he is considering it (and independent of the slot in which it is placed). We will denote the expression $1 - p_i$ by \bar{p}_i .
- In addition, there is a fixed probability $q > 0$ that after considering an ad and deciding against it, the user will continue to the next ad (as opposed to abandoning the search altogether).³
- We denote by $\beta(a, b) = \prod_{k=a}^{b-1} q\bar{p}_k$ the probability that a user will reach ad b , given that he is currently considering ad a .
- Given the above, if agents are numbered according to their slots, the probability of the user to select the ad in position i is $p_i \prod_{j=1}^{i-1} q\bar{p}_j = p_i \beta(1, i)$.

The private value of an agent in the above model is the utility he derives from a single action, which is assumed to be the same regardless of the slot allocated to the agent. We denote the private value of agent i by v_i , $i \in \{1 \dots N\}$; naturally, $\forall i \in \{1 \dots N\} : v_i > 0$.

The agents' bids are interpreted as the maximal price per action they are willing to pay. We denote them by b_i , $i \in \{1 \dots N\}$. Naturally, $\forall i : b_i \geq 0$ by the rules of the auction. For ease of presentation, we define $b_i = 0, v_i = 0$ for all $i > N$.

First, consider the following lemma:⁴

LEMMA 1. *The efficient allocation in the PPA model is the one in which agents are allocated to slots in decreasing order of $\alpha_i v_i$, where $\alpha_i = \frac{p_i}{1 - q\bar{p}_i}$.*

¹The report may be done automatically, if there is a technical way to track user actions, but we do not assume that this mechanism cannot be manipulated to generate false reports. We will discuss this further in section 4.

²This is equivalent to assuming that the number of slots is unlimited, which is a reasonable assumption in the Internet setting.

³All our results hold (with minimal changes) if we allow the parameter q to vary across different advertisers.

⁴For a similar model, this lemma was shown independently by [1]. However, their work focuses on algorithms for computing the efficient allocation when there are more slots than bidders (and therefore multiple allocations satisfy the condition that is implied by this lemma).

PROOF. Note that a necessary condition for an efficient allocation is that switching places between two adjacent ads does not increase the social surplus. Since such switch affects only the revenue of the ads that participate in it (the higher ads are unaffected, and the probability to reach the lower ads also remains the same), this means that:

$$\begin{aligned} p_i v_i + (1 - p_i) q p_{i+1} v_{i+1} &\geq p_{i+1} v_{i+1} + (1 - p_{i+1}) q p_i v_i \\ p_i v_i (1 - q(1 - p_{i+1})) &\geq p_{i+1} v_{i+1} (1 - q(1 - p_i)) \\ \frac{p_i}{1 - q(1 - p_i)} v_i &\geq \frac{p_{i+1}}{1 - q(1 - p_{i+1})} v_{i+1} \end{aligned}$$

This means that in an efficient allocation the agents must be allocated with slots in decreasing order of $\alpha_i v_i$. Since this ordering is unique, then allocating agents according to it must result in the efficient allocation. \square

The intuition behind this lemma is that the users that reach each slot are divided into those that continue to the next slot (a fraction of $q\bar{p}_i$ of the total number of users who reached this slot), and those that exit the process at the given slot (a fraction of $1 - q\bar{p}_i$ of the total number of users who reached this slot) – either because they commit to an action, or because they abandon the search. Therefore, $\frac{p_i v_i}{1 - q\bar{p}_i} = \alpha_i v_i$ is the expected utility *per lost user* that is gained by the advertiser at this slot. Naturally, we put first the advertisers that gain most per user that they take away. From now on, we will assume w.l.o.g. that agents are indexed in decreasing order of $\alpha_i b_i$.

2.1 Comparison with the classical PPC model

The classical PPC model [11, 5], for the setting where there are as many players as slots, can be described as follows:

- There are N advertisers that compete for N ad slots.
- We denote the *clickthrough rate* (CTR) of a slot by $x_i, i \in \{1 \dots N\}$. The CTR is a publicly known property of a slot, which does not depend on the advertiser who is using it. The slots are numbered in decreasing order of CTR: $\forall i : x_i \geq x_{i+1}$.
- Each advertiser has a publicly known and fixed property w_i , which denotes his “quality factor”. It is assumed that the number of clicks advertiser j gets when placed in slot i is $w_j x_i$.
- The private value in this model is the utility that each advertiser derives from a single unit of CTR, which is assumed to be the same regardless of the slot from which it originates. We denote it by $v_i, i \in \{1 \dots N\}$; naturally, $\forall i \in \{1 \dots N\} : v_i > 0$.
- The advertisers’ bids are interpreted as the maximal price per unit of CTR they are willing to pay to the publisher. We denote them by $b_i, i \in \{1 \dots N\}$; w.l.o.g. we assume that $w_i b_i \geq w_{i+1} b_{i+1}$; that is, the advertisers are ordered in decreasing order of the product of bid and quality effect. Naturally, $\forall i : b_i \geq 0$ by the rules of the auction. For ease of presentation, we define $b_i = 0, x_i = 0, v_i = 0$ for all $i > N$.

From now on, we will refer to this model as “the PPC model”, to differentiate it from the PPA model defined in the previous section. First, let us quote the following lemma [5]:

LEMMA 2. *In the PPC model defined above, the efficient allocation is the one in which $\forall 1 \leq i \leq N : w_i v_i \geq w_{i+1} v_{i+1}$; that is, the advertisers are ordered in decreasing order of the product of their valuation and their quality effect.*

Since the above is a well-studied model for the PPC setting (see [2, 6, 7, 3], among others), the first instinct for someone who wishes to develop a model for the PPA setting may be to use the PPC model “as is”, by replacing “clicks” with “actions” in the definition. However, we argue that the PPC model does not capture the inherent advertisers’ competition of the PPA setting. The cornerstone of the PPC model is the assumption that the number of actions that advertiser i gets when placed in slot j is a function of only i and j , and does not depend on the allocation of the other advertisers. This is an arguable assumption even in the PPC setting – if we assume, by introducing the quality parameter w_i , that advertisers can vary in their attractiveness, then how can we claim that the number of advertiser i ’s clicks does not depend on the attractiveness of the advertisers displayed on the same page with him? In the PPA setting, the corresponding assumption makes even less sense. Since we can reasonably assume that a given user will perform at most one action, increasing the attractiveness of the other advertisers must, in general, decrease the amount of actions received by advertiser i (for the same allocation) – otherwise, where do the extra actions come from?

Technically, if there is more than one slot, the PPA model (with parameters N, p_i, q) cannot be reduced to a PPC model (with parameters N, x_i, w_j) – that is, given the values of p_i, q , there is no transformation that will produce x_i, w_j , so that the number of actions each advertiser gets in each possible allocation in the resulting PPC model is the same as in the original PPA model. As an example, consider the following setting: $N = 2, p_a = 0.9, p_b = 0.2, q = 0.5$. When ad a is in the first slot and ad b - in the second, ad a gets 0.9 actions and ad b - 0.01. If the positions are reversed, ad a gets 0.36 actions and ad b - 0.2. Therefore, in the first slot the number of actions that advertiser b gets is $\frac{2}{9}$ of what advertiser a gets (in the same slot), and in the second slot the ratio is $\frac{1}{36}$. This implies that there are no values of w_a, w_b that are consistent with the example – therefore, it cannot be modeled using the PPC model.⁵ This suggests that the PPA and the PPC models are indeed distinct, and therefore we need to analyze the PPA model separately.

3. AUCTION MECHANISMS FOR THE PPA SETTING

Given the PPA model, we now consider two basic auctions for that setting. One is the classical VCG mechanism, known for its truthfulness properties. The other one is the GSP mechanism, which is the mechanism used in practice in the PPC setting, adapted to the PPA setting. Two of the major contributions of the existing work on the PPC setting is in characterizing a well structured set of pure strategy equilibria for the GSP mechanism, and proving using this analysis that the GSP mechanism is revenue-wise preferable to the

⁵There is a more general model for the PPC setting, considered in [2]. It assumes that the number of actions is a general function $CTR(i, j)$ (where i is the advertiser and j is the slot), which is non-increasing in j . However, it can be easily seen that this model cannot be equivalent to a PPA model with $N \geq 3$.

ad publisher compared to the VCG mechanism. In this section we present a similar analysis for the PPA model. Our analysis leads to exposing significant distinctions between the PPA model and the standard PPC model; in particular, we show that the VCG mechanism is revenue-wise preferable to the ad publisher in the PPA model, challenging current practices when transition to a PPA scheme is considered.

3.1 The VCG mechanism

Using lemma 1, we can now define the VCG mechanism for our PPA model:

LEMMA 3. *The VCG mechanism in the PPA model allocates slots to users in decreasing order of $\alpha_i b_i$. Assuming that agents are numbered according to their slots (that is, $\alpha_i b_i \geq \alpha_{i+1} b_{i+1}$), the total payment of user i is*

$$\left(\frac{1}{\bar{p}_i q} - 1\right) \sum_{j=i+1}^N p_j \beta(1, j) b_j$$

PROOF. The definition of the VCG allocation rule follows from Lemma (1). The total payment of advertiser i must be equal to the externality he imposes on the other advertisers. Note that ordering of the other (not- i) advertisers is the same whether i is present or not, therefore the only advertisers harmed by the presence of i are those below him in this ordering. His removal would multiply the number of actions they get by $\frac{1}{(1-p_i)q}$, from which the formula for payments follows immediately. \square

The following corollaries establish the structure of per-action payment and agent utility in the VCG auction for the PPA model.

COROLLARY 1. *The per-action payment of user i in the VCG auction is: $\frac{1}{\alpha_i} \sum_{j=i+1}^N p_j b_j \beta(i+1, j)$.*

PROOF. The per-action payment of user i in the VCG auction is:

$$\begin{aligned} & \frac{1}{p_i \beta(1, i)} \left(\frac{1}{(1-p_i)q} - 1 \right) \sum_{j=i+1}^N p_j \beta(1, j) b_j = \\ & \frac{1-q+qp_i}{p_i \beta(1, i) (1-p_i)q} \sum_{j=i+1}^N p_j \beta(1, j) b_j = \\ & \frac{1-q+qp_i}{p_i \bar{p}_i q} \sum_{j=i+1}^N p_j b_j \prod_{k=i}^{j-1} q \bar{p}_k = \\ & \frac{1}{\alpha_i} \sum_{j=i+1}^N p_j b_j \beta(i+1, j) \end{aligned}$$

\square

COROLLARY 2. *The per-action payment is smaller than the agent's bid.*

PROOF. The per-action payment is:

$$\frac{1}{\alpha_i} \sum_{j=i+1}^N p_j b_j \beta(i+1, j) =$$

$$\frac{1}{\alpha_i} \sum_{j=i+1}^N \alpha_j b_j (1-q(1-p_j)) \beta(i+1, j) =$$

$$\frac{1}{\alpha_i} \sum_{j=i+1}^N [\alpha_j b_j \beta(i+1, j) - q(1-p_j) \alpha_j b_j \beta(i+1, j)] =$$

$$\frac{1}{\alpha_i} \sum_{j=i+1}^N [\alpha_j b_j \beta(i+1, j) - \alpha_j b_j \beta(i+1, j+1)] =$$

$$\frac{1}{\alpha_i} \left[\alpha_{i+1} b_{i+1} + \sum_{j=i+2}^N \alpha_j b_j \beta(i+1, j) - \right.$$

$$\left. \sum_{j=i+1}^N \alpha_j b_j \beta(i+1, j+1) \right] =$$

By shifting indices and adding a zero element in the first sum, we get:

$$\frac{1}{\alpha_i} \left[\alpha_{i+1} b_{i+1} + \sum_{j=i+1}^N \alpha_{j+1} b_{j+1} \beta(i+1, j+1) - \right.$$

$$\left. \sum_{j=i+1}^N \alpha_j b_j \beta(i+1, j+1) \right] =$$

$$\frac{1}{\alpha_i} \left[\alpha_{i+1} b_{i+1} - \sum_{j=i+1}^N (\alpha_j b_j - \alpha_{j+1} b_{j+1}) \beta(i+1, j+1) \right] \leq$$

Since $\forall j : \alpha_j b_j \geq \alpha_{j+1} b_{j+1}$, we have

$$\leq \frac{1}{\alpha_i} [\alpha_{i+1} b_{i+1}] \leq b_i$$

\square

COROLLARY 3. *The utility of agent i in the VCG auction is*

$$\beta(1, i) \left[v_i p_i - \frac{1}{\alpha_i} \sum_{j=i+1}^N p_j b_j \beta(i+1, j) \right]$$

PROOF. The utility of agent i in the VCG auction is

$$\beta(1, i) v_i p_i - \left(\frac{1}{(1-p_i)q} - 1 \right) \sum_{j=i+1}^N \beta(1, j) b_j =$$

$$\beta(1, i) \left[v_i p_i - \frac{1-q+p_i q}{p_i (1-p_i)q} \sum_{j=i+1}^N p_j \beta(i, j) b_j \right] =$$

$$\beta(1, i) \left[v_i p_i - \frac{1}{\alpha_i} \sum_{j=i+1}^N p_j b_j \beta(i+1, j) \right]$$

\square

3.2 The GSP mechanism

Now, let us consider the Generalized Second Price (GSP) auction for our PPA model. In GSP, agents are allocated to slots in decreasing order of $\alpha_i b_i$, and each agent pays the minimal per-action price that he would have had to bid

in order to hold his allocated slot, given the other agent's actions. Therefore, if agents are numbered according to slot order, the per-action payment of agent i is $\frac{\alpha_{i+1}}{\alpha_i} b_{i+1}$.

Under GSP, for a profile of bids \bar{b} to be a pure Nash equilibrium, the following conditions must hold:

- No agent has an incentive to move to a higher slot. That is, for all $i, j < i$:

$$p_i \beta(1, i) \left(v_i - \frac{\alpha_{i+1}}{\alpha_i} b_{i+1} \right) \geq p_i \beta(1, j) \left(v_i - \frac{\alpha_j}{\alpha_i} b_j \right)$$

$$\beta(j, i) \left(v_i - \frac{\alpha_{i+1}}{\alpha_i} b_{i+1} \right) \geq v_i - \frac{\alpha_j}{\alpha_i} b_j$$

$$\alpha_j b_j \geq \alpha_i \left(v_i - \beta(j, i) \left(v_i - \frac{\alpha_{i+1}}{\alpha_i} b_{i+1} \right) \right)$$

$$\alpha_j b_j \geq \alpha_i v_i (1 - \beta(j, i)) + \alpha_{i+1} b_{i+1} \beta(j, i)$$

- No agent has an incentive to move to a lower slot. That is, for all $i, j < i$:

$$p_j \beta(1, j) \left(v_j - \frac{\alpha_{j+1}}{\alpha_j} b_{j+1} \right) \geq$$

$$\geq p_j \prod_{1 \leq k \leq i, k \neq j} q \bar{p}_k \left(v_j - \frac{\alpha_{i+1}}{\alpha_j} b_{i+1} \right)$$

$$v_j - \frac{\alpha_{j+1}}{\alpha_j} b_{j+1} \geq \beta(j+1, i+1) \left(v_j - \frac{\alpha_{i+1}}{\alpha_j} b_{i+1} \right)$$

$$\alpha_{j+1} b_{j+1} \leq \alpha_j \left(v_j - \beta(j+1, i+1) \left(v_j - \frac{\alpha_{i+1}}{\alpha_j} b_{i+1} \right) \right)$$

$$\alpha_{j+1} b_{j+1} \leq \alpha_j v_j (1 - \beta(j+1, i+1)) +$$

$$\alpha_{i+1} b_{i+1} \beta(j+1, i+1)$$

We can summarize the NE conditions as $\forall i, \forall j < i$:

$$\alpha_i v_i (1 - \beta(j, i)) + \alpha_{i+1} b_{i+1} \beta(j, i) \leq \alpha_j b_j$$

$$\alpha_j b_j \leq \alpha_{j-1} v_{j-1} (1 - \beta(j, i)) + \alpha_i b_i \beta(j, i)$$

Similarly to the analysis of the standard PPC model, we will consider a set of bidding profiles that satisfy the conditions $\forall i, \forall j < i$:

$$\alpha_i v_i (1 - \beta(j, i)) + \alpha_i b_i \beta(j, i) \leq \alpha_j b_j$$

$$\alpha_j b_j \leq \alpha_j v_j (1 - \beta(j, i)) + \alpha_i b_i \beta(j, i)$$

We will refer to these profiles as Special Nash Equilibria (SNE), although at this point we have not yet shown that they are a subset of NE.⁶

LEMMA 4 (MONOTONE VALUES). *In an SNE*

$$\alpha_i v_i \geq \alpha_{i+1} v_{i+1}$$

⁶In the study of the standard PPC model, the corresponding equilibria obtained by combining the requirements are called Symmetric Nash Equilibria, and are also termed SNEs. We find this term "symmetric" somewhat misleading and prefer to refer to it as "special Nash equilibria".

PROOF. By re-arranging the inequality that defines the SNE we have: $\forall i \forall j < i : \alpha_i v_i \leq \frac{\alpha_j b_j - \alpha_i b_i \beta(j, i)}{1 - \beta(j, i)} \leq \alpha_j v_j$. \square

LEMMA 5 (NON-NEGATIVE SURPLUS). *In an SNE*

$$v_k \geq \frac{\alpha_{k+1}}{\alpha_k} b_{k+1}$$

PROOF. Let us substitute $i = N + 1, j = k + 1$ into the right-hand inequality from the SNE definition:

$$\alpha_{k+1} b_{k+1} \leq \alpha_{k+1} v_{k+1} (1 - \beta(k+1, N+1))$$

Since $\alpha_k v_k \geq \alpha_{k+1} v_{k+1}$ by the previous lemma:

$$\alpha_{k+1} b_{k+1} \leq \alpha_k v_k (1 - \beta(k+1, N+1))$$

Using $0 < (1 - \beta(k+1, N+1)) < 1$, we have the result. \square

LEMMA 6 (SNE \subset NE). *If a bidding profile is an SNE then it is a NE.*

PROOF. SNE is defined as:

$$\forall i \forall j < i : \alpha_i v_i (1 - \beta(j, i)) + \alpha_i b_i \beta(j, i) \leq \alpha_j b_j$$

$$\alpha_j b_j \leq \alpha_j v_j (1 - \beta(j, i)) + \alpha_i b_i \beta(j, i)$$

Since in an SNE $\alpha_i v_i \geq \alpha_{i+1} v_{i+1}$, and we number the agents so that $\alpha_i b_i \geq \alpha_{i+1} b_{i+1}$, this implies

$$\alpha_i v_i (1 - \beta(j, i)) + \alpha_{i+1} b_{i+1} \beta(j, i) \leq \alpha_j b_j$$

$$\alpha_j b_j \leq \alpha_{j-1} v_{j-1} (1 - \beta(j, i)) + \alpha_i b_i \beta(j, i)$$

which is exactly the requirement for NE. \square

One interesting property of the set of SNE is that it is sufficient to verify the inequalities for one step (up and down) deviations in order to verify that the entire set of inequalities is satisfied:

LEMMA 7 (ONE STEP SOLUTION). *If a set of bids satisfies the SNE conditions for all i and $j = i - 1$, then it satisfies them for all i and $j < i$.*

PROOF. We demonstrate the idea of the proof by example. Suppose that we know that the inequalities hold for $i = 2, j = 1$ and for $i = 3, j = 2$; we want to show that they hold for $i = 3, j = 1$ as well. Specifically, we need to show:

$$\alpha_3 v_3 (1 - \beta(1, 3)) + \alpha_3 b_3 \beta(1, 3) \leq \alpha_1 b_1$$

$$\alpha_1 b_1 \leq \alpha_1 v_1 (1 - \beta(1, 3)) + \alpha_3 b_3 \beta(1, 3)$$

Let us start with the left inequality first. We need to show:

$$\alpha_3 v_3 (1 - \beta(1, 3)) + \alpha_3 b_3 \beta(1, 3) \leq \alpha_1 b_1$$

We know:

$$\alpha_3 v_3 (1 - \beta(2, 3)) + \alpha_3 b_3 \beta(2, 3) \leq \alpha_2 b_2$$

$$\alpha_2 v_2 (1 - \beta(1, 2)) + \alpha_2 b_2 \beta(1, 2) \leq \alpha_1 b_1$$

By substituting $\alpha_2 b_2$ in the second inequality with the LHS of the first one we have:

$$\alpha_2 v_2 (1 - \beta(1, 2)) + \alpha_3 v_3 (1 - \beta(2, 3)) \beta(1, 2) +$$

$$+ \alpha_3 b_3 \beta(2, 3) \beta(1, 2) \leq \alpha_1 b_1$$

Since $\beta(2, 3)\beta(1, 2) = \beta(1, 3)$:

$$\begin{aligned} & \alpha_2 v_2 (1 - \beta(1, 2)) + \alpha_3 v_3 (\beta(1, 2) - \beta(1, 3)) + \\ & \quad + \alpha_3 b_3 \beta(1, 3) \leq \alpha_1 b_1 \\ & (\alpha_2 v_2 - \alpha_3 v_3) (1 - \beta(1, 2)) + \alpha_3 v_3 (1 - \beta(1, 3)) + \\ & \quad + \alpha_3 b_3 \beta(1, 3) \leq \alpha_1 b_1 \end{aligned}$$

Since in SNE $\alpha_2 v_2 \geq \alpha_3 v_3$:

$$\alpha_3 v_3 (1 - \beta(1, 3)) + \alpha_3 b_3 \beta(1, 3) \leq \alpha_1 b_1$$

which is exactly the inequality we set out to prove.

Now, let us prove the right inequality. We need to show:

$$\alpha_1 b_1 \leq \alpha_1 v_1 (1 - \beta(1, 3)) + \alpha_3 b_3 \beta(1, 3)$$

. We know:

$$\alpha_1 b_1 \leq \alpha_1 v_1 (1 - \beta(1, 2)) + \alpha_2 b_2 \beta(1, 2)$$

$$\alpha_2 b_2 \leq \alpha_2 v_2 (1 - \beta(2, 3)) + \alpha_3 b_3 \beta(2, 3)$$

By substituting $\alpha_2 b_2$ in the first inequality with the RHS of the second one we have:

$$\begin{aligned} \alpha_1 b_1 & \leq \alpha_1 v_1 (1 - \beta(1, 2)) + \alpha_2 v_2 (1 - \beta(2, 3)) \beta(1, 2) + \\ & \quad + \alpha_3 b_3 \beta(2, 3) \beta(1, 2) \end{aligned}$$

Since $\beta(2, 3)\beta(1, 2) = \beta(1, 3)$:

$$\alpha_1 b_1 \leq \alpha_1 v_1 (1 - \beta(1, 2)) + \alpha_2 v_2 (\beta(1, 2) - \beta(1, 3)) + \alpha_3 b_3 \beta(1, 3)$$

$$\begin{aligned} \alpha_1 b_1 & \leq \alpha_1 v_1 (1 - \beta(1, 3)) + (\alpha_1 v_1 - \alpha_2 v_2) (\beta(1, 3) - \beta(1, 2)) + \\ & \quad + \alpha_3 b_3 \beta(1, 3) \end{aligned}$$

Since in SNE $\alpha_1 v_1 \geq \alpha_2 v_2$, and $\beta(1, 3) < \beta(1, 2)$ by definition:

$$\alpha_1 b_1 \leq \alpha_1 v_1 (1 - \beta(1, 3)) + \alpha_3 b_3 \beta(1, 3)$$

which is exactly the inequality we set out to prove. \square

Our results imply that we can construct an SNE profile recursively, by starting with player N (the player with the lowest $\alpha_i v_i$), selecting a bid for him so that $\alpha_N b_N \leq \alpha_N v_N (1 - \beta(N, N + 1))$, using his bid to compute the bounds on b_{N-1} , selecting a value for the bid of player $N - 1$, and so on, until all bids are determined. Note that the interval between the upper and lower bounds is always non-empty, and therefore the process is guaranteed to succeed. In fact, since the bounds that define SNE are monotonically increasing in b_i , we can write closed expressions for the lower and upper bounds on the bid of player i in SNE. The one-step SNE conditions are:

$$\alpha_i b_i \geq \alpha_{i+1} v_{i+1} (1 - q\bar{p}_i) + \alpha_{i+1} b_{i+1} q\bar{p}_i$$

$$\alpha_i b_i \leq \alpha_i v_i (1 - q\bar{p}_i) + \alpha_{i+1} b_{i+1} q\bar{p}_i$$

Solving the recursive relations, we have:

$$\begin{aligned} \alpha_i b_i & \geq \alpha_{i+1} v_{i+1} (1 - q\bar{p}_i) + q\bar{p}_i (\alpha_{i+2} v_{i+2} (1 - q\bar{p}_{i+1}) + \\ & \quad + q\bar{p}_{i+1} (\alpha_{i+3} v_{i+3} (1 - q\bar{p}_{i+2}) + \\ & \quad + q\bar{p}_{i+2} (\alpha_{i+4} v_{i+4} (1 - q\bar{p}_{i+3}) + \dots \end{aligned}$$

$$+ q\bar{p}_{N-2} (\alpha_N v_N (1 - q\bar{p}_{N-1}))) \dots)$$

Opening the parenthesis:

$$\begin{aligned} \alpha_i b_i & \geq \sum_{j=i}^{N-1} \alpha_{j+1} v_{j+1} (1 - q\bar{p}_j) \prod_{k=i}^{j-1} q\bar{p}_k = \\ & = \sum_{j=i}^{N-1} p_j \frac{\alpha_{j+1}}{\alpha_j} v_{j+1} \beta(i, j) \end{aligned}$$

Similarly, we get:

$$\alpha_i b_i \leq \sum_{j=i}^N \alpha_j v_j (1 - q\bar{p}_j) \prod_{k=i}^{j-1} q\bar{p}_k = \sum_{j=i}^N p_j v_j \beta(i, j)$$

LEMMA 8. *In the SNE in which the agents' bids match the upper bound, the allocation and the per-action prices are the same as in the dominant-strategy equilibrium of VCG.*

PROOF. The allocation is the efficient allocation both in the dominant strategy equilibrium of VCG and in any SNE. In the SNE in which the agents' bids match the upper bound, agent i 's price per action is:

$$\frac{\alpha_{i+1} b_{i+1}}{\alpha_i} = \frac{1}{\alpha_i} \sum_{j=i+1}^N p_j v_j \beta(i+1, j)$$

Recall from Corollary (1) that this is exactly the VCG price in this setting. \square

The revenue implication: The above shows that in addition to simplifying the decision making for the advertisers, given its truthfulness, the VCG mechanism also brings good revenue to the auctioneer. In fact, one of the major findings of the work on PPC auctions is the advantage for the ad publisher in using GSP upon using VCG. This is obtained by showing that the VCG outcome is obtained in the worst (revenue-minimizing) SNE of the GSP auction is that setting. Surprisingly, our result shows that by considering carefully the PPA model, in this case the preference should be reversed!

4. GAMING THE SYSTEM

Given the surprising lesson on the revenue of the VCG mechanism in the PPA setting, suppose that this mechanism is used for the process of auctioning ad slots. The input of the mechanism includes, in addition to the bids of the advertisers, several parameters that are assumed to be commonly known and fixed; however, in practice, it might be possible for interested parties to strategically influence the values of those parameters in order to manipulate the outcome of the auction. In this section, we would like to discuss several possible ways to do this, and their effect on the outcome and on the incentives that different parties might have for such manipulation. For the purpose of this discussion, we assume that the probabilities p_i and q are computed by the ad publisher, and that he updates them once in a time period according to the action reports by the advertisers and the number of times the relevant page was shown during the previous period.

The first idea that comes to mind is for an advertiser to underreport the actions made by the users, thus decreasing his total payment in the short term. However, this means

that in the next time period, the estimated p_i of that advertiser will decrease and as a result, his α_i will also decrease. A fall in α_i has two effects: First, the advertiser might be moved to a lower slot, losing real customers; Second, even if everything stays the same (he stays in the same slot, submits the same false report and the other advertisers bid the same values), a decrease in α_i means higher per-action price, softening the impact of underreporting. The following proposition quantifies the effect of underreporting in this case:

PROPOSITION 1. *Suppose that advertiser i reports only $1/C$ of his actions (where $C > 1$), and suppose that this misreport has the following effect (compared to truthful reporting): $p'_i = \frac{p_i}{C}$, where p'_i is the advertiser's parameter obtained due to misreporting (while the other advertisers' bids, the allocation of slots, and the number of real actions do not change). Then, the ratio between the advertiser's total payment given misreporting and the total payment he would have paid if he reported truthfully is $1 - \frac{(C-1)qp_i}{C(1-qp_i)}$.*

PROOF. Let us denote by Pr'_i the per-action payment of agent i in case of misreport, and by Pr_i his per-action payment in the case of truthful reporting. Then, using Corollary (1), we have:

$$\begin{aligned} \frac{Pr'_i}{Pr_i} &= \frac{\frac{1-q(1-p'_i)}{p'_i} \sum_{j=i+1}^N p_j b_j \beta(i, j)}{\frac{1-q(1-p_i)}{p_i} \sum_{j=i+1}^N p_j b_j \beta(i, j)} = \\ &= \frac{p_i(1-q+qp'_i)}{p'_i(1-q+qp_i)} = \frac{C(1-q+qp_i/C)}{1-q+qp_i} = \frac{C-Cq+qp_i}{1-q+qp_i} \end{aligned}$$

Therefore, the ratio between the total payment when only $1/C$ of the actions are reported and the total payment in case of truthful report is:

$$\frac{Pr'_i}{Pr_i} \cdot \frac{1}{C} = \frac{1-q+qp_i/C}{1-q+qp_i} = 1 - \frac{(C-1)qp_i}{C(1-q+qp_i)}$$

□

Since in the Internet setting we can assume that $p_i \ll 1-q$ (anecdotal evidence suggests that $q \approx 0.95$, and $p_i \approx 0.001$), this implies that the long-term impact of underreporting is limited. For example, if at some point in time the advertiser starts reporting only half of his actions, then his total payment, once his p_i is updated according to the apparent decrease in the number of actions, will be only slightly smaller than what he would have paid had he reported truthfully. Of course, the above reasoning does not imply that underreporting cannot be profitable at all – after all, an advertiser can always run an ad, report zero actions, pay nothing, close the account with the ad publisher and return under a different identity – but it does mean that it is not as profitable as it seems at first glance.⁷

Another option for an advertiser is to create false visits to the web page that do not result in any actions, thus decreasing the estimated p_i of all advertisers, with the intention of decreasing the per-action price (without affecting the allocation or the reported number of actions). The following proposition quantifies the effect of such manipulation on the per-action price:

⁷As suggested by [8] and others, the false-name-bidding strategy can be countered by requiring a fixed entrance fee from each new advertiser.

PROPOSITION 2. *Suppose that, by generating bogus visits to the web page, an advertiser was able to divide the value of the parameter p_i for all advertisers by $C > 1$ – that is, $\forall i : p'_i = \frac{p_i}{C}$, where p'_i denotes the value of the advertiser's parameter given the manipulation, and p_i is its value without a manipulation.⁸ (while the other advertisers' bids, the allocation of slots, and the number of real actions do not change). Then, the per-action price of the advertiser is at least $1 - \frac{(C-1)qp_i}{C(1-qp_i)}$ of the price he would have paid without the manipulation.*

PROOF. Let us denote by Pr'_i the per-action payment of agent i given the manipulation, and by Pr_i his per-action payment without the manipulation. Then, using Corollary (1), we have:

$$\begin{aligned} \frac{Pr'_i}{Pr_i} &= \frac{\frac{1-q+qp'_i}{p'_i} \sum_{j=i+1}^N p'_j b_j \prod_{k=i}^{j-1} q(1-p'_k)}{\frac{1-q+qp_i}{p_i} \sum_{j=i+1}^N p_j b_j \prod_{k=i}^{j-1} q(1-p_k)} = \\ &= \frac{C-Cq+qp_i}{1-q+qp_i} \frac{\sum_{j=i+1}^N \frac{p_j b_j}{C} \prod_{k=i}^{j-1} q(1-p_k/C)}{\sum_{j=i+1}^N p_j b_j \prod_{k=i}^{j-1} q(1-p_k)} = \\ &= \frac{C-Cq+qp_i}{C-Cq+Cqp_i} \frac{\sum_{j=i+1}^N p_j b_j \prod_{k=i}^{j-1} q(1-p_k/C)}{\sum_{j=i+1}^N p_j b_j \prod_{k=i}^{j-1} q(1-p_k)} \geq \\ &\geq 1 - \frac{(C-1)qp_i}{C(1-q+qp_i)} \end{aligned}$$

□

We can see that although this strategy might cause a slight decrease in the per-action price, it is doubtful whether the gain is worth the effort – the effect is small for the same reasons as in the case of misreporting actions; in addition, in order to cause a non-negligible change in the observed number of page visits, one probably has to perform hundreds of thousands of page requests (while making sure that they appear to come from different computers).

5. CONCLUSION

In this paper we studied a model, which incorporates user modeling into the ad auctions setting; the model is especially appealing for the pay-per-action setting, although can be applicable to the pay-per-click setting as well. Our model fits nicely with the need to incorporate user modeling into ad auctions (see the discussion in [4]), and with the belief that sequential search through the ads is an appropriate assumption (see [1]). We have shown that our model is not reducible to the standard PPC model, and introduced a comparative study of GSP and VCG in the framework of our model. We provided a characterization of pure strategy equilibria for the GSP auction, and proved the revenue-wise usefulness of the VCG mechanism in comparison to the GSP auction, in the framework of the PPA model. This is in sharp distinction to the results obtained for the standard PPC model, where user behavior is not modeled explicitly. We have also illustrated the robustness under manipulation of PPA schemes, using the tools provided in our study.

⁸This is just a rough estimate – the actual effect on p_i is more complex and depends on the method used to compute the parameters. In particular, this manipulation does not necessarily affect all advertisers' parameters in the same way.

6. REFERENCES

- [1] G. Aggarwal, J. Feldman, S. Muthukrishnan, and M. Pal. Sponsored search auctions with markovian users. Working paper, presented at the 4th EC Ad Auctions Workshop, 2008.
- [2] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7, New York, NY, USA, 2006. ACM.
- [3] I. Ashlagi, D. Monderer, and M. Tennenholtz. Mediators in position auctions. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 279–287, New York, NY, USA, 2007. ACM.
- [4] S. Atey and G. Ellison. Position auctions with consumer search. Working paper, presented at the 4th EC Ad Auctions Workshop, 2008.
- [5] 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.
- [6] N. Immorlica, K. Jain, M. Mahdian, and K. Talwar. Click fraud resistant methods for learning click-through rates. In *WINE '05: Proceedings of the First International Workshop on Internet and Network Economics*, pages 34–45, 2005.
- [7] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, New York, NY, USA, 2006. ACM.
- [8] M. Mahdian and K. Tomak. Pay-per-action model for online advertising. In *WINE '07: Proceedings of the Third International Workshop on Internet and Network Economics*, pages 549–557, 2007.
- [9] H. Nazerzadeh, A. Saberi, and R. Vohra. Dynamic cost-per-action mechanisms and applications to online advertising. In *WWW '08: Proceeding of the 17th international conference on World Wide Web*, pages 179–188, New York, NY, USA, 2008. ACM.
- [10] TroubleClicks. Trouble clicks. *The Economist*, November 23 2006.
- [11] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.