

On Expressing Value Externalities in Position Auctions

Florin Constantin
 College of Computing
 Georgia Tech.
 Atlanta GA, USA
 florin@cc.gatech.edu

Malvika Rao
 SEAS
 Harvard University
 Cambridge MA, USA
 malvika@eecs.harvard.edu

Chien-Chung Huang
 Institut für Informatik
 Humboldt-Universität
 12489 Berlin, Germany
 villars@informatik.hu-berlin.de

David C. Parkes
 SEAS
 Harvard University
 Cambridge MA, USA
 parkes@eecs.harvard.edu

Abstract

We introduce a bidding language for expressing negative value externalities in position auctions for online advertising. The *unit-bidder constraints* (UBC) language allows a bidder to condition a bid on its allocated slot and on the slots allocated to other bidders. We introduce a natural extension of the Generalized Second Price (GSP) auction, the *expressive GSP* (eGSP) auction, that induces truthful revelation of constraints for a rich subclass of unit-bidder types, namely *downward-monotonic* UBC. We establish the existence of envy-free Nash equilibrium in eGSP under a further restriction to a subclass of *exclusion constraints*, for which the standard GSP has no pure strategy Nash equilibrium. The equilibrium results are obtained by reduction to equilibrium analysis for reserve price GSP (Even-Dar et al. 2008). In considering the winner determination problem, which is NP-hard, we bound the approximation ratio for social welfare in eGSP and provide parameterized complexity results.

Introduction

A search engine offers positions (slots) for ads adjacent to the organic search results, with slots lower on the page tending to generate fewer clicks. The generalized second-price auction (GSP) is the industry standard for allocating ad slots. In GSP, each ad is associated with a per-click bid indicating the advertiser’s willingness-to-pay for a click. This implies an expected (bid) value for a slot, which we can think of simplistically as the bid times the the clickthrough rate (CTR, i.e. the probability of a user click). A greedy algorithm is used to allocate ads to slots in decreasing order of per-click bid. Whenever an ad in slot j receives a click, the advertiser pays a price equivalent to the bid of the bidder in the “next slot.” This is the smallest per-click bid of the advertiser that would have resulted in it retaining slot j .

For a given slot, the number of clicks that an ad attracts also depends on the *other* ads shown, e.g. via the number of other ads (Reiley, Li, and Lewis 2010), or their relative position (Craswell et al. 2008; Gomes, Immorlica, and Markakis 2009; Jeziorski and Segal 2010). We term this kind of allocative externality a *quantity externality*. We address an orthogonal kind of allocative externality, referred to here as a *value externality*, in which an advertiser’s value, given a

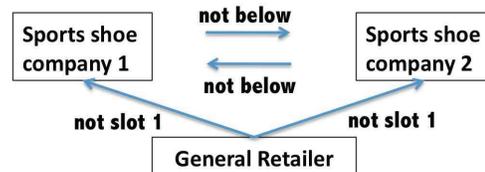


Figure 1: Each sports shoe company’s bid is conditioned on being allocated above the other company. The general retailer’s bid is conditioned on neither shoe company’s ad being in slot one.

click, depends on which other bidders are simultaneously allocated, and where. Whereas quantity externalities are observable to a search engine, value externalities are private to bidders and need to be expressed via a bidding language.

We introduce *unit-bidder constraints* (UBC), enabling a flexible class of languages for expressing *negative value externalities* in position auctions. Each bidder can condition its own bid for a particular slot on another bidder *not* being allocated some other slot. Each bidder can submit multiple such constraints. The constraints associated with a bid are only active *if the bidder is allocated*. See Figure 1. For example, shoe company 1 can say “my bid is only valid if I am allocated above shoe company 2.” In a more general language we introduce *soft-constraints*, with a smaller non-zero bid adopted when a constraint is violated. We refer to the first case, with bid values zero if constraints are violated, as a hard constraint model.

We extend the standard GSP auction to allow bids in the UBC language. This expressive GSP (eGSP) auction greedily allocates bidders to slots. To be eligible for allocation to the next slot, an ad must not be in conflict with the constraints of any bidders already allocated. The bid value for an unallocated ad depends in turn on the allocations already made and the ad’s own constraints. A “next price” payment rule is used, with the payment of an allocated bidder equal to the minimum bid it could have made and still won the same slot given its constraints. The choice of a greedy algorithm is motivated by the need for rapid response time for search engines; moreover, achieving even a reasonable approximation to the optimal allocation in the UBC model is NP-hard.

Although eGSP is not strategyproof, our first result establishes that reporting truthful constraints is a dominant strategy in eGSP for *downward monotonic* UBC, whatever

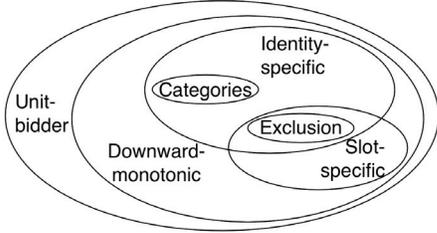


Figure 2: Classes of negative value-externality constraints.

the bid value and whatever the bids of others. While fruitful bid manipulations already exist in standard GSP, our result shows that augmenting the preferences of bidders with downward-monotonic UBC does not introduce new types of fruitful manipulations. Downward-monotonicity insists that a bidder dissatisfied with a particular slot given an allocation to other bidders is also dissatisfied with any lower slot. The downward-monotonic UBC languages include natural languages in which a bidder precludes being below other bidders (identity-specific), or cares about the range of slots it or other bidders are in (slot-specific). See Fig. 2.

We also consider the special case of *exclusion* externalities, in which a bidder insists that it is never allocated simultaneously with another bidder. An exclusion externality is both an identity-specific and a slot-specific externality. Our second result establishes the existence of envy-free Nash equilibrium in eGSP when each bidder is involved in at most one exclusion constraint. The result is obtained by a delicate reduction to the equilibrium in a GSP with bidder-specific reserve prices (Even-Dar et al. 2008). In contrast, there exists no pure-strategy Nash equilibrium in the standard GSP in the same setting. We also establish existence of envy-free equilibria for general degree exclusion constraints.

Turning to algorithmic results, we provide a tight bound on the approximation ratio of the greedy algorithm. For two parametrized classes of constraints, *category* and *local-exclusion*, we identify polynomial time optimal algorithms, assuming fixed parameters. We defer most proofs to the longer version of this paper in the interest of space.

Related Work Position auctions (e.g. GSP) are an active research area (Edelman, Ostrovsky, and Schwarz 2007; Varian 2007), and some of this work focuses on quantity externalities as discussed above. Our class of downward-monotonic UBC languages for expressing value externalities encompasses existing models. Slot-specific constraints generalize the “bid-to-the-top” model of Aggarwal *et al.* (2006), where an advertiser can restrict its bid to appear above some position. The authors describe an easy to implement mechanism and show existence of equilibrium in their special case. UBC can also encode the model of Muthukrishnan (2009), where bidders have bids that depend on the maximum number of ads shown. The latter work proposes a social-welfare maximizing algorithm and critical value pricing scheme, but does not address incentives. Ghosh and Sayedi (2010) consider a model in which an advertiser submits two bids: one for solo placement and another for placement alongside other ads. Revenue and efficiency tradeoffs are examined,

in what is a special case of exclusion UBC with soft constraints. In an incomparable model to ours, Ghosh and Mahdian (2008) earlier considered a setting where an advertiser’s value depends on its quality relative to other ads shown, but irrespective of their location. Other algorithmic work related to externalities includes Krysta *et al.* (2010) for combinatorial auctions and Kash *et al.* (2011) for secondary markets for wireless spectrum.

Preliminaries

Let $N = \{1, \dots, n\}$ denote the bidders in a position auction with m slots. As is standard, we assume $m = n$ (since there is essentially an unlimited number of slots, on multiple pages of search results.) Each bidder i is associated with a per-click value $v_i \geq 0$. We assume that the click-through rate (CTR) falls off from one slot to the next according to discount factor $\delta \in (0, 1)$ and we normalize the first slot’s CTR to 1. Slots $1, 2, \dots, m$ have CTRs $1, \delta, \dots, \delta^{m-1}$.

In our basic model of value externalities, we associate bidder i with constraints C_i , so that his expected value for slot $j \in \{1, \dots, m\}$ is $v_i \delta^{j-1}$ as long as constraints C_i on the allocation are satisfied, and zero otherwise. Later, we also allow “soft constraints” wherein a bidder’s value is v_i or $0 \leq v_i^- \leq v_i$ if the constraints are violated.

Both value and constraints are private to each bidder i , that submits a bid b_i and constraints \hat{C}_i to the seller, perhaps untruthfully. Given reported bids and constraints, the seller would like to solve WDP, i.e. compute an optimal allocation.

Definition 1. Given bids $b = (b_1, \dots, b_n)$ and constraints $\hat{C} = (\hat{C}_1, \dots, \hat{C}_n)$, the winner determination problem WDP is to find a set of winners $W \subseteq N$ and an allocation A (with any $i \in W$ winning slot A_i) solving:

$$\max_{(W,A) \in \mathbb{F}} \sum_{i \in W} b_i \delta^{A_i-1} \quad (1)$$

where \mathbb{F} is the set of feasible solutions (W, A) : $A_i \neq A_j$ for all $i \neq j$ (both in W), and A satisfies \hat{C}_i for every $i \in W$.

Apart from constraints, we assume a standard quasi-linear utility-maximizing bidder model. Bidder i ’s expected utility equals $(v_i \mathbf{1}_{C_i} - p_i) \delta^{A_i-1}$, where p_i is the per-click payment, given that i is allocated slot A_i and $\mathbf{1}_{C_i}$ equals 1 if all constraints in C_i are satisfied and 0 otherwise.

We introduce *unit bidder constraints* (UBC), a natural and expressive constraint model, as argued below. In UBC, any bidder i has a set C_i of $L_i \geq 0$ constraints (encoded as triples)

$$C_i = \{(pos_i, B_\ell, pos_\ell)\}_{\ell=1, \dots, L_i}, \quad (2)$$

where triple $(pos_i, B_\ell, pos_\ell)$ imposes the requirement that if bidder i is allocated to slot pos_i , then bidder B_ℓ ($\neq i$) cannot be allocated to position pos_ℓ .

Different languages would impose restrictions on the specific kinds of UBC constraints enabled. For example, UBC can encode *identity-specific* constraints, where a bidder specifies a set of bidders above which it must be allocated. One can conceptualize this as a directed graph on bidders, with an edge from i to j indicating such an “enemy” of i . With 3 slots, this would be encoded in UBC as

$$C_i = \{(2, j, 1), (3, j, 2), (3, j, 1)\}$$

UBC can also encode *slot-specific* constraints, where a bidder imposes a requirement that its bid is only valid if it is no lower than some slot and one or more other bidders are no higher than bidder-specific slots. For example, if i wants to be in slots 1 or 2 and j should not be in slots 1 or 2, and there are 4 slots, then this can be encoded in UBC as

$$C_i = \{(3, *, *), (4, *, *), (*, j, 1), (*, j, 2)\},$$

where ‘*’ indicates that the entry is instantiated for all valid values. An *exclusion constraint* between i and j is encoded as $C_i = \{(*, j, *)\}$.

We introduce a natural generalization of GSP to UBC.

Expressive GSP Expressive GSP (eGSP) with hard constraints takes reported UBC constraints \hat{C} and bid values b as input, and implements a greedy allocation, collecting next-slot payments. The pseudocode for eGSP is given below. An unallocated bidder i is *eligible* for slot j if this allocation is not precluded by the constraint of some already allocated bidder in slots $1, \dots, j-1$, or by a constraint between i and some already allocated bidder. The allocation rule for eGSP repeatedly allocates the next slot to the unallocated, eligible bidder (if any) with the highest bid price, breaking ties at random. Let $A_i(b, \hat{C}) \in \{1, \dots, m\}$ denote the slot allocated to a winner, with $A_i(b, \hat{C}) = 0$ otherwise.

Let $i_k \in N$ denote the bidder (if any) allocated slot k .

Expressive GSP (eGSP)

Input: bids b_1, \dots, b_n , constraints $\hat{C}_1, \dots, \hat{C}_n$

For slot $k = 1$ to m

Eligible $\leftarrow \{i: \text{allocating } k \text{ to } i \text{ satisfies } \hat{C}_{i_1}, \dots, \hat{C}_{i_{k-1}}, \hat{C}_i\}$
 $i_k \leftarrow \max b_i$ in *Eligible* (if any)

End

The per-click price for bidder i allocated to slot k is

$$p_i(b, \hat{C}) = \min b'_i \text{ s.t. } A_i(b'_i, b_{-i}, \hat{C}) = k, \quad (3)$$

i.e. the smallest bid b'_i (given \hat{C}) for which a bidder is allocated the same slot, with $b_{-i} = (b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$.

Equilibrium concepts As is standard in sponsored search auctions, we study complete information Nash equilibrium:

Definition 2. *Bid profile* (b, \hat{C}) is a Nash equilibrium (NE) in eGSP if $\forall i$ and fixing the bids b_{-i} and reported constraints \hat{C}_{-i} of others, there is no report (b'_i, C'_i) with higher utility for the bidder than (b_i, \hat{C}_i) (given its true v_i and C_i).

The motivation for studying complete information NE in sponsored search is that advertisers can learn each others’ types over time via bidding dynamics. Also of interest in sponsored search auctions are envy-free equilibrium. In the following, let p_i denote the per-click price for bidder i given the bid profile and A_i denote the slot allocated to bidder i .

Definition 3. *Bid profile* (b, \hat{C}) is an envy-free (EF) equilibrium in eGSP if

(1) for every allocated i , $\delta^{A_j-1}(v_i - p_j) \leq \delta^{A_i-1}(v_i - p_i)$ for all bidders $j \neq i$ for which the allocation would satisfy the constraints C_i if i and j switch positions, and

(2) for every unallocated bidder i , $\delta^{A_j-1}(v_i - p_j) \leq 0$, for all bidders j for which the allocation would satisfy C_i if i was allocated j ’s slot (with j going unallocated).

The standard (Edelman, Ostrovsky, and Schwarz 2007) definition of envy-free is not specific about the effect of i receiving a different slot on the rest of the allocation. But this is crucial here because of externalities. The envy-free equilibrium property captures a dynamic stability requirement.¹

Incentives in eGSP

We will establish that eGSP is “semi-truthful”, namely that bidders cannot benefit from misreporting constraints.

Greedy is incompatible with truthfulness. Before continuing, we briefly explain by example why this property would not hold if one was to use a naive application of the standard payment rule used to achieve incentive-compatibility in auctions. For this, say that an allocation algorithm is *monotonic* if, for all \hat{C} , all b_{-i} , all i and all b_i , then $A_i(b'_i, b_{-i}, \hat{C}) \leq A_i(b_i, b_{-i}, \hat{C})$, for all $b'_i \geq b_i$. This insists that a higher bid value can only lead to a higher slot (and thus a lower slot index.) Clearly, the greedy algorithm is monotonic in this sense. Fix reported constraints \hat{C} and let $f_i(b, \hat{C}) = \delta^{A_i(b, \hat{C})-1}$ while i wins, and 0 otherwise. Following Myerson (1981), the standard approach to achieve a truthful auction would charge a winner i an (expected) payment for its allocation to slot $A_i(b, \hat{C})$ of,

$$b_i \cdot f_i(b, \hat{C}) - \int_{w=0}^{b_i} f_i(w, b_{-i}, \hat{C}) dw \quad (4)$$

But we see from the next example that this would not provide truthfulness with respect to *constraints*.

Example 1. Consider 2 slots and bidders 1 and 2, with values 30, 20, where bidder 1’s true constraint precludes bidder 2 from appearing in the top slot when 1 is allocated (but is happy for 2 to appear below 1.) Discount $\delta = 0.9$. Bidder 1’s (expected) payment given this constraint is $30 - [(30 - 20)] = 20$. If bidder 1 did not report this constraint, then it would still win slot 1 but with expected payment of $30 - [(30 - 20) + (20 - 0)(0.9)] = 2$.

The difficulty in achieving truthfulness with this payment rule is that the expected payment (for the same slot) is not independent of the bidder’s report (namely reported constraints), a well-known condition for truthfulness. A bidder can pay less by omitting from its bid any constraints that leave the allocation unchanged but would constrain its allocation for lower bid values. Given the uniqueness of the Myerson payment rule in providing truthfulness in regard to the bid value for fixed constraints, we see that it is impossible to achieve full truthfulness with a greedy allocation method. The classical Vickrey-Clarke-Groves mechanism for achieving truthfulness is undesirable since in our case it requires solving NP-hard optimization problems. We show

¹Consider an allocated bidder i . If (1) is violated then i would like to compete with bidder j for its allocated slot, to drive up j ’s price, and without fear of j retaliating by making i take j ’s slot at j ’s price. Bidder j can always do this by bidding just below i ’s bid price, making i win j ’s slot at i ’s bid price (which was in turn setting j ’s price.) Similarly, consider an unallocated bidder i . If (2) is violated then this bidder would like to compete for j ’s slot and do so without fear of j bidding just below i to make i win the slot.

now however that eGSP achieves truthfulness for constraint reports for a large class within UBC.

Semi-Truthfulness We turn now to the incentive properties of the next-price payment rule in eGSP. We also note that eGSP has the useful property that the price is invariant to bid price b_i while the allocated slot remains unchanged.

Definition 4. A slot auction with UBC constraints is *semi-truthful* if for any reported b_{-i} and \hat{C}_{-i} , of other agents, any v_i and $b_i \leq v_i$ of agent i , it is a dominant strategy for agent i to report its constraints $\hat{C}_i = C_i$ truthfully.

Definition 5. UBC C_i are downward-monotonic (DM) if

$$(pos_i, j, pos_j) \in C_i \Rightarrow (pos_i + 1, j, pos_j) \in C_i \quad (5)$$

Fixing the allocation to other bidders, if i is dissatisfied with slot pos_i , then it is also dissatisfied with any lower slot.

Theorem 1. The eGSP auction is semi-truthful for bidders whose value externalities can be expressed with downward-monotonic UBC constraints.

Proof. Fix any $b_i \leq v_i$. Let k denote the slot allocated to i when reporting true C_i . Conditioned on report \hat{C}_i not changing the allocated slot k , the payment does not change because constraints have no effect on other bidders until a bidder is allocated and so the eligible set is unchanged. Moreover, if i is allocated a slot k then by reporting $\hat{C}_i \neq C_i$ it cannot be allocated a higher slot (for the same bid value) because it is already eligible for slot k and thus all higher slots by DM. Also, i cannot achieve a lower slot by misreport $\hat{C}_i \neq C_i$, fixing its bid value, because any change to preclude i from being eligible for slot k will preclude i from being eligible for all subsequent slots $k' > k$ by DM. Finally, an agent that is unallocated but becomes allocated to slot k' by reporting $\hat{C}_i \neq C_i$ must have a true constraint that is violated upon allocation to slot k' , since the allocation for earlier slots does not change. \square

The easy result below shows natural DM classes of UBC.

Lemma 1. Identity-specific constraints and slot-specific constraints satisfy downward monotonicity.

DM is also necessary for eGSP's semi-truthfulness:

Example 2. Consider 3 slots and 3 bidders with values 60, 40 and 10, and discount $\delta = 0.9$. If bidder 1 is truthful then he wins slot 1 and pays 40 for payoff $60 - 40 = 20$. But by reporting constraint "I do not want slot 1", he wins slot 2 and pays 10 for payoff $(60 - 10)0.9 = 45 > 20$. This constraint is not DM.

Achieving truthfulness for constraints, as we do here, is important since standard GSP is not truthful (for bids only).

Equilibria in eGSP for Exclusion Constraints

We focus now on a subset of DM UBC constraints, and establish a clean separation between the existence of Nash equilibrium in eGSP and its inexistence in standard GSP. We denote an exclusion constraint between i, i' as $i \leftrightarrow i'$.

Theorem 2. Standard GSP may have no pure-strategy NE for bidders with max-degree one exclusion constraints.

Proof. Consider 2 slots and 4 bidders 1, 2, 3, 4 with $1 \leftrightarrow 2, 3 \leftrightarrow 4$ and values $v_1 = v_2 + \epsilon, v_3 = v_4 + \epsilon$ such that $v_1 > v_3$ and $v_1 - v_2 < \delta(v_1 - 0)$. 1 must bid at least v_2 , else 2 can win by bidding v_2 . Similarly 3 must bid at least v_4 . Then $\{1, 3\}$ win slots $\{1, 2\}$. Bidder 1 however prefers bidding below 3, and so this cannot be an NE. \square

Our main theoretical results in regard to the equilibrium properties of eGSP are:

Theorem 3. There exists an envy-free, Nash equilibrium of eGSP under max-degree one exclusion constraints.

Theorem 4. There exists an envy-free equilibrium of eGSP under (general degree) exclusion constraints.

Note that envy-free equilibria and Nash equilibria are incomparable for eGSP. In the standard GSP model however, any envy-free equilibrium is also a Nash equilibrium.

We establish existence by reduction to GSP with bidder-specific reserve prices (rGSP) (Even-Dar et al. 2008). A sketch of the proof is provided below. rGSP operates just as standard GSP except that the price for slot k to bidder i is $\max(r_i, b_{k+1})$, where $b_{k+1} = 0$ if no bidder is allocated in slot $k + 1$. Even-Dar et al. (2008) provide a tâtonnement algorithm to construct an envy-free (and, in their case, Nash) equilibrium for rGSP. They insist that $b_i \geq r_i$ for all bidders (which is achieved through our reduction.)

For simplicity, we adopt in what follows the convention that the bidders are indexed according to the slots allocated.

Definition 6. A bid profile b is an envy-free (EF) equilibrium in rGSP given reserve prices r if $\delta^{j-1}(v_i - \max(p_j, r_i)) \leq \delta^{i-1}(v_i - \max(b_{i+1}, r_i)), \forall j \neq i$.

The utility to every bidder in an envy-free equilibrium of rGSP is at least what it would receive if it could exchange positions with any other bidder. An EF equilibrium of rGSP is also a NE. The reduction identifies a set of *candidates* and reserve prices, such that when an equilibrium is determined for rGSP on the candidates, we can construct a bid profile that is an equilibrium in eGSP. Non candidates will be unallocated in the equilibrium of eGSP. The technical challenge is to establish that the strategic effect of non-candidates on candidates in eGSP is equivalent to the effect of the reserve price on bidders in rGSP. The construction generates an envy-free equilibrium of eGSP that is also a Nash equilibrium for the special case of max-degree one exclusion.

Fleshing this out, assume for simplicity distinct values of bidders. Let E_i denote the *enemies* of i . Namely, the set of bidders with which i has an exclusion constraint. We first determine a *pseudo outcome* (K, X, r, \succ) of eGSP:

- Run eGSP with bids $b = v$ and constraints $\hat{C} = C$. The allocated bidders comprise the set of candidates K .
- Let $X_i \subseteq E_i$ denote bidders that are excluded when i is allocated, i.e. for whom no other enemy was allocated before i . Define $r_i = \max_{j \in X_i} \{v_j\} + \epsilon'$, for a small $\epsilon' > 0$ if $X_i \neq \emptyset$ and $r_i = 0$ otherwise. ϵ' is smaller than the minimum gap between bidder values.
- Define a *priority order* \succ , where $i \succ j$ if i is allocated before j and they share an enemy (necessarily a non-candidate).

If, e.g., $3 \leftrightarrow 2 \leftrightarrow 1$, where values equal IDs, then 3 and 1 are candidates, $X_3 = \{2\}$, $r_3 = 2 + \epsilon'$, $X_1 = \emptyset$, $r_1 = 0$, and $3 \succ 1$. If $4 \leftrightarrow 1, 3 \leftrightarrow 2$ then 4 and 3 are candidates, with $X_4 = \{1\}$, $X_3 = \{2\}$, $r_4 = 1 + \epsilon'$, $r_3 = 2 + \epsilon'$ and no priority order constraints (the case for any max-degree one exclusion as each candidate excludes at most one bidder).

The following observation about the pseudo-outcome is straightforward and stated without proof:

Lemma 2. *In a pseudo-outcome, $v_i > v_j$ and $r_i > v_j$ for candidates i, j with $i \succ j$, and $v_i > r_i$ for all candidates i .*

Consider now running the rGSP on the candidate bidders K and with reserve prices r .

Lemma 3. *In a EF equilibrium of rGSP on the pseudo-outcome, all higher priority $j \succ i$ for any candidate i are allocated to a higher slot than i .*

Denote by b^r the envy-free equilibrium of rGSP (a bid profile for the candidates K). Based on this, we construct an equilibrium bid profile b^* for eGSP as follows:

- All constraints $\hat{C} = C$ are reported truthfully
- Bid $b_i^* = b_i^r$ if $i \in K$ and $b_i^* = v_i$ otherwise.

By the consistent ordering property of Lemma 3, we have

Lemma 4. *Given bids b^* and $C = (C_1, \dots, C_n)$ reported to eGSP, the outcome is identical to that of rGSP on the pseudo-outcome, under a particular tie-breaking rule.*

The proof is by induction on position allocated, from top to bottom. The essential insight is that while the bid ordering corresponding to b^r need not respect the truthful ordering assumed in determining the pseudo-outcome of eGSP, it is sufficient for the bids to respect the priority constraints across bidders that share mutual enemies for the strategic equivalence between bids of non-candidates and reserve prices to hold. In completing the proof for Theorem 3, it is necessary to establish that the Nash constraints hold for all bidders, both candidates and non-candidates. For a non-candidate, we can start with the observation that the price of an allocated enemy is at least the reserve price of the enemy, which is at least the value of the non-candidate by Lemma 2. For a candidate, in considering the Nash constraint to preclude the benefit from a deviation to a higher slot, we inherit this result from the NE of the b^r equilibrium in rGSP. For a downwards deviation, one needs to also argue that the (effective) reserve price (as induced by bids of non-candidates) remains as assumed in the rGSP equilibrium analysis. This part of the reduction needs the assumption of max-degree one. Theorem 4 on EF equilibrium follows directly from the outcome equivalence between rGSP on b^r and eGSP on b^* (Lemma 4).²

²To understand why the max-degree one requirement is important to obtain NE, consider an instance with 3 bidders, values 8, 4 and 6 and exclusion constraints $8 \leftrightarrow 4 \leftrightarrow 6$. Let bidder values equal IDs. We get pseudo-outcome $K = \{8, 6\}$, $X_8 = \{4\}$, $X_6 = \emptyset$, $r_8 = 4$, $r_6 = 0$, and $8 \succ 6$. The corresponding instance of rGSP has values 8, 6 and reserve price 4 and 0. An EF equilibrium is $b^r = (8, 6)$ for $\delta < 1/2$. In this case, we have $8 - 6 > \delta(8 - \max(4, 0))$ and $\delta(6 - 0) > (6 - 6)$. The candidate equilibrium in eGSP is $b^* = (8, 4, 6)$, and indeed, the outcome

Algorithmic considerations

The winner determination problem, to select the bids that maximize total value given constraints, is NP-hard. For exclusion constraints the problem is equivalent to INDEPENDENTSET. This, together with the need for fast algorithms for slot auctions, motivates the greedy algorithm. We first bound the approximation ratio (for social welfare) of the greedy algorithm. The result is stated for any UBC constraints, not only exclusion ones.

Theorem 5. *Consider a directed graph G , with bidders as vertices, in which there is an edge from i to i' if C_i contains at least one constraint $(pos_i, i', pos_{i'})$ with $pos_i > pos_{i'}$: if i is in pos_i then i' cannot be in some higher slot $pos_{i'}$. Let d denote an upper bound on all vertices' in-degrees in G . The greedy algorithm for the WDP problem achieves an $\frac{1-\delta}{1-\delta^{d+1}}$ approximation³ when $\delta < 1$. Formally, letting $b(W, A)$ be the social welfare in the greedy algorithm,*

$$b(W, A) \geq OPT(b, C) \frac{1-\delta}{1-\delta^{d+1}} \quad (6)$$

where $OPT(b, C)$ is the maximum allocation value given bids $b = (b_1, \dots, b_n)$ and constraints $C = (C_1, \dots, C_n)$.

Note that the graph in the theorem is well-defined for any set of UBC constraints. In the special case of exclusion constraints, the in-degree bound is given by the number of exclusion constraints in which an agent i can participate. For identity-specific constraints, it is a bound on the number of other agents for whom agent i is considered an enemy.

The ratio in the theorem is tight (discussion omitted due to space constraints). For exclusion constraints this dependence is intrinsic: INDEPENDENTSET is NP-hard to approximate within a $2^{O(\sqrt{\log d})}/d$ factor (Samorodnitsky and Trevisan 2000). For max-degree one exclusion, we can conclude that the greedy algorithm is optimal.

As positive results for exact winner determination, we identify two parametrized subclasses of constraints with tractable algorithms for WDP. The complexity of these algorithms is polynomial for any fixed parameter value:

Category-specific. This is a special case of the identity-specific model in which every bidder is associated with a category (for example, “sports shoe company” or “general retailer”). Constraints are limited to requiring placement above all bidders in the same category. This cluster structure allows for an efficient solution to the WDP.

Proposition 1. *In the category-specific model, an optimal allocation can be computed in time $O((n \log n + gn)(m^g)(g^g))$ where g denotes the number of categories.*

is equivalent to that of rGSP under these bids, with 8 and 6 allocated slots 1 and 3, at prices 6 and 0 respectively. However b^* is an EF but not a Nash equilibrium in eGSP. Bidder 8 can bid 5 instead, in which case 6 is allocated, 4 is eliminated, leaving 8 to receive slot 2 for price 0. For $\delta > 1/4$ this is better for bidder 8, with $8 - 6 < \delta(8 - 0)$. The effect of non-candidate 4 is no longer strategically equivalent to a reserve price of 4 to bidder 8 when 8 deviates downwards and plays after bidder 6 in eGSP.

³An algorithm is said to achieve a ρ -approximation if the value of the allocation it outputs is within a multiplicative factor of $\rho \leq 1$ of the value of the optimal solution, for all possible instances.

Local-exclusion constraints. Suppose bidders only have exclusion constraints to bidders within some distance w in the bid ranking. w is a locality measure that turns out to be the tree-width of the constraint graph, a standard algorithmic concept. This allows tractable algorithms for constant w .

Proposition 2. *For exclusion constraint locality w , the WDP can be solved in $O(n2^w)$ time and $O(2^w + n)$ space.*

Soft constraints

A natural extension is a *soft constraint* model where, in addition to constraints C_i , a bidder has a pair (v_i, v_i^-) of per-click values, defining its value when no constraints, or at least one constraint in C_i is violated, respectively. The standard hard constraint model has $v_i^- = 0$.

The eGSP auction is generalized as follows: in allocating the next slot, the eligible bidders are those for whom the allocation would not violate a first constraint for an already allocated bidder. The approximation ratio in Theorem 5 continues to hold. For an eligible i , the price adopted is then either (1) b_i or (2) b_i^- , depending on whether or not i 's constraints are still satisfied given the current allocation. If allocated slot k , then the price is the minimal value of b_i or b_i^- , for case (1) or (2) respectively, such that bidder i would still retain the same slot.

Semi-truthfulness no longer holds for soft constraints even with DM constraints:

Example 3. *Consider 4 bidders and 3 slots, with bids $(100, 100c)$, $(70, 70c)$, $(50, 50c)$, and $(30, 30c)$ where $c = 0.5$. Let the discount factor be $1 - \epsilon$, for a small $\epsilon > 0$. If no one has any constraints then eGSP allocates bidders 1, 2, 3 to slots 1, 2, 3. In particular bidder 2 pays 50 for a utility of $70 - 50 = 20$. Now suppose bidder 2 lies and specifies a constraint stating it must be placed above bidder 1. Then eGSP allocates bidders 1, 3, 2 to slots 1, 2, 3. Bidder 2 now pays 30; its utility is $70 - 30 = 40$. Bidder 2 can achieve the same result by misreporting its values instead as $(40, 40c)$.*

We can recover a weaker form of semi-truthfulness. Interestingly, this result does not require downward monotonicity, and extends easily to the earlier hard constraint model.

Theorem 6. *In eGSP with soft-constraints, a bidder i always has a best-response, for any C_i , any (v_i, v_i^-) , and any reports (b_{-i}, \hat{C}_{-i}) , of other bidders, in which i reports C_i truthfully along with some pair (b_i, b_i^-) of bid values.*

The idea of the proof is to establish that for any report with an untruthful constraint set, there exists a report, with bid values b_i, b_i^- and truthful constraint set C_i , such that the slot allocated is unchanged (and thus the price is unchanged), and the subsequent allocation decisions are no worse under the second report than the first report.

Conclusions

We have introduced unit-bidder constraints, an expressive language for negative value externalities in position auctions, and analyzed the strategic properties of an expressive GSP (eGSP) auction. We obtain a “semi-truthfulness” property of eGSP with respect to misreports of downward-monotonic constraints. In this sense, the modified eGSP is

as truthful as the standard GSP and there are no new manipulations. We exhibit a class of such constraints for which Nash equilibria fail to exist in standard GSP, but exist and can be easily constructed in eGSP. A weaker but still useful notion of truthfulness in regard to constraints is established for a generalization of UBC where bidders have a smaller but non-zero bid value for violated constraints.

For future work, it would be interesting to characterize equilibria for more general UBC, thereby enabling revenue and efficiency comparisons to GSP. Turning to complexity results, the “one-enemy” case remains open, where each bidder has a constraint against at most one other bidder. We find this case appealing because it could be achieved through a simple restriction to a bidding language.

Acknowledgments We thank Yiling Chen, Lisa Fleischer, Nicole Immorlica, David Pennock, Ariel Procaccia and the anonymous reviewers for useful feedback. C-C. Huang is supported by an Alexander von Humboldt fellowship.

References

- Aggarwal, G.; Muthukrishnan, S.; and Feldman, J. 2006. Bidding to the top: VCG and equilibria of position-based auctions. In *Workshop on Approximation and Online Algorithms*, 15–28.
- Craswell, N.; Zoeter, O.; Taylor, M.; and Ramsey, B. 2008. An experimental comparison of click position-bias models. In *Proceedings of ACM WSDM '08*, 87–94.
- Edelman, B.; Ostrovsky, M.; and Schwarz, M. 2007. Internet advertising and generalized second price auction: Selling billions of dollars worth in keywords. *Amer. Econ. Review* 97(1):242–259.
- Even-Dar, E.; Feldman, J.; Mansour, Y.; and Muthukrishnan, S. 2008. Position auctions with bidder-specific minimum prices. In *Proc. of 4th WINE*.
- Ghosh, A., and Mahdian, M. 2008. Externalities in online advertising. In *Proc. of World Wide Web Conf. (WWW'08)*, 161–168.
- Ghosh, A., and Sayedi, A. 2010. Expressive auctions for externalities in online advertising. In *Proc. of WWW'10*.
- Gomes, R.; Immorlica, N.; and Markakis, E. 2009. Externalities in keyword auctions: An empirical and theoretical assessment. In *Proc. of the 5th WINE*, 172–183.
- Jeziorski, P., and Segal, I. 2010. What makes them click: Empirical analysis of consumer demand for search advertising. Technical report, Stanford University.
- Kash, I.; Murty, R.; and Parkes, D. C. 2011. Enabling spectrum sharing in secondary market auctions. In *Proc. Workshop on the Economics of Networks, Systems and Computation (NetEcon)*.
- Krysta, P.; Michalak, T.; Sandholm, T.; and Wooldridge, M. 2010. Combinatorial auctions with externalities. In *Proc. of AAMAS*.
- Muthukrishnan, S. 2009. Bidding on configurations in Internet ad auctions. In *Proc. of COCOON*, 1–6.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* 6(1):58–73.
- Reiley, D.; Li, S.; and Lewis, R. 2010. Northern exposure: Field experiment measuring externalities between search adverts. In *Proc. of ACM conf. on electronic commerce (EC'10)*.
- Samorodnitsky, A., and Trevisan, L. 2000. A PCP characterization of NP with optimal amortized query complexity. In *Proc. of ACM symposium on Theory of computing (STOC)*, 191–199.
- Varian, H. 2007. Position auctions. *International Journal of Industrial Organization* 25:1163–1178.