

SPONSORED SEARCH, MARKET EQUILIBRIA, AND THE HUNGARIAN METHOD

PAUL DÜTTING¹ AND MONIKA HENZINGER^{1,2} AND INGMAR WEBER^{1,3}

¹ Ecole Polytechnique Fédérale de Lausanne, Switzerland

E-mail address: {paul.duetting,monika.henzinger,ingmar.weber}@epfl.ch

² University of Vienna, Austria

E-mail address: monika.henzinger@univie.ac.at

³ Yahoo! Research Barcelona, Spain

E-mail address: ingmar@yahoo-inc.com

ABSTRACT. Two-sided matching markets play a prominent role in economic theory. A prime example of such a market is the sponsored search market where n advertisers compete for the assignment of one of k sponsored search results, also known as “slots”, for certain keywords they are interested in. Here, as in other markets of that kind, market equilibria correspond to stable matchings. In this paper, we show how to modify Kuhn’s Hungarian Method (Kuhn, 1955) so that it finds an optimal stable matching between advertisers and advertising slots in settings with generalized linear utilities, per-bidder-item reserve prices, and per-bidder-item maximum prices. The only algorithm for this problem presented so far (Aggarwal et al., 2009) requires the market to be in “general position”. We do not make this assumption.

1. Introduction

Two-sided matching markets play a prominent role in economic theory. A prime example of such a market is the *sponsored search market* [14] where n advertisers (or bidders) compete for the assignment of one of k sponsored search results, also known as “slots”, for certain keywords (or items) they are interested in. Here, as in other markets of that kind, *market equilibria* correspond to *stable matchings*. A stable matching that is preferred by all bidders over all other stable matchings is *bidder optimal*. Mechanisms that compute bidder optimal matchings typically provide the bidders with the incentive to reveal their *true preferences*, i.e., they are *truthful*.

In the most basic model of a two-sided matching market, known as the *stable marriage problem* [9], each bidder has a strict preference ordering over the items and each item has a strict preference ordering over the bidders. In a more general model, see e.g. [16], each bidder has a linear utility function for each item that depends on the price of the item and

1998 ACM Subject Classification: F.2.2 (Nonnumerical Algorithms and Problems).

Key words and phrases: stable matching, envy-free allocation, general auction mechanism, general position. This work was conducted as part of a EURYI scheme award (see <http://www.esf.org/euryi/>).

every item can have a *reserve price*, i.e., a price under which the item cannot be sold to *any* bidder. In the even stronger model that we study here every bidder-item pair can have a *reserve price*, i.e., a price under which the item cannot be sold to this *specific* bidder, and a *maximum price*, i.e., a price above which this bidder does not want to buy this *specific* item. We call this model the *sponsored search market*. An interesting property of this model is that it generalizes standard auction formats such as *VCG* [17, 4, 10] and *GSP* [7].

While the problem of finding a bidder optimal matching in the first two models has been largely solved in the 60s, 70s, and 80s [9, 16, 5, 15], the problem of finding a bidder optimal matching in the sponsored search market has been addressed only recently [2].

The main finding of [2] is that if the market is in “general position”, then (a) there is a unique bidder optimal matching and (b) it can be found in $O(nk^3)$ steps by a truthful mechanism. For a market to be in “general position”, however, any two reserve prices and/or maximum prices must be distinct. In practice, this will rarely be the case and so we typically have to deal with markets that are *not* in general position. The authors of [2] propose to bring such markets into “general position” using random perturbations and/or symbolic tie-breaking. The problem with this approach, however, is that there is no guarantee that a bidder optimal solution of the perturbed market leads to a bidder optimal solution of the original market. In fact, such a solution may not even exist (see Section 3). Additionally, a perturbation-based mechanism may not be truthful.

We improve upon the results of [2] as follows: First, in Section 3, we show how to modify the definition of stability so that a bidder optimal matching is guaranteed to exist for arbitrary markets. Then, in Section 5, 6, and 7, we show how to modify Kuhn’s *Hungarian Method* [13, 8] so that it finds a bidder optimal matching in time $O(nk^3 \log(k))$. Afterwards, in Section 8, we show that with our notion of stability bidder optimality no longer implies truthfulness, unless further restrictions are imposed on the model. Finally, in Section 9, we show how to reduce more general linear utility functions to our setting.¹

Independently of us Ashlagi et al. [3] also improved upon the results of [2] by (a) showing the existence of a unique feasible, envy free, and Pareto efficient solution for *position auctions with budgets* and by (b) providing a truthful mechanism that finds it. The notion of envy-freeness is equivalent to our notion of stability. Their model, however, is a special case of our model as it requires a common preference ordering over the items, it does not incorporate reserve prices, it does not allow the maximum prices to depend on the bidder *and* the item, and it requires the maximum prices to be distinct.

Recently, Kempe et al. [12] presented an efficient algorithm that finds the minimum envy-free prices (if they exist) for a *given* matching.

To summarize our main contributions are: (1) We show how to modify the Hungarian Method so that it finds a bidder optimal solution for *arbitrary* markets, including markets that are *not* in “general position”. (2) We show how different definitions of stability affect the existence of a bidder optimal solution. (3) We show how to reduce more general linear utility functions to the setting that we study in this paper with no loss in performance.

2. Problem Statement

We are given a set I of n bidders and a set J of k items. We use letter i to denote a bidder and letter j to denote an item. For each bidder i and item j we are given a *valuation*

¹These utilities can be used to model that the *click probability* in the *pay-per-click model* has a bidder-dependent component c_i and an item-dependent component c_j . See [1, 7] for details.

$v_{i,j}$, a reserve price $r_{i,j}$, and a maximum price $m_{i,j}$. We assume that the set of items J contains a *dummy item* j_0 for which all bidders have a valuation of zero, a reserve price of zero, and a maximum price of ∞ .²

We want to compute a matching $\mu \subseteq I \times J$ and per-item prices $p = (p_1, \dots, p_k)$. We require that every bidder i appears in exactly one bidder-item pair $(i, j) \in \mu$ and that every non-dummy item $j \neq j_0$ appears in at most one such pair. We allow the dummy item j_0 to appear more than once. We call bidders (items) that are not matched to any non-dummy item (bidder) *unmatched*. We regard the dummy item as unmatched.

We define the *utility* u_i of bidder i to be $u_i = 0$ if bidder i is unmatched and $u_i = u_{i,j}(p_j)$ if bidder i is matched to item j at price p_j . We set $u_{i,j}(p_j) = v_{i,j} - p_j$ if $p_j < m_{i,j}$ and $u_{i,j}(p_j) = -\infty$ if $p_j \geq m_{i,j}$. We say that a matching μ with prices p is *feasible* if (1) $u_i \geq 0$ for all i , (2) $p_{j_0} = 0$ and $p_j \geq 0$ for all $j \neq j_0$, and (3) $r_{i,j} \leq p_j < m_{i,j}$ for all $(i, j) \in \mu$. We say that a feasible matching μ with prices p is *stable* if $u_i \geq u_{i,j}(p_j)$ for all $(i, j) \in I \times J$.³ Finally, we say that a stable matching μ with prices p is *bidder optimal* if $u_i \geq u'_i$ for all i and stable matchings μ' with prices p' .

We say that an algorithm is *truthful* if for every bidder i with utility functions $u_{i,1}(\cdot), \dots, u_{i,k}(\cdot)$ and any two inputs $(u'_{i,j}(\cdot), r_{i,j}, m'_{i,j})$ and $(u''_{i,j}(\cdot), r_{i,j}, m''_{i,j})$ with $u'_{i,j}(\cdot) = u_{i,j}(\cdot)$ for i and all j and $u'_{k,j}(\cdot) = u''_{k,j}(\cdot)$ for $k \neq i$ and all j and matchings μ' with p' and μ'' with p'' we have that $u_{i,j'}(p'_{j'}) \geq u_{i,j''}(p''_{j''})$ where $(i, j) \in \mu$ and $(i, j'') \in \mu''$. This definition formalizes the notion that “lying does not pay off” as follows: Even if bidder i claims that his utility is $u''_{i,j}$ instead of $u_{i,j}$ he will not achieve a higher utility with the prices and the matching computed by the algorithm. Thus, the algorithm “encourages truthfulness”.

3. Motivation

The definition of stability in [2], which we call *relaxed stability* to indicate that every stable solution is also relaxed stable (but not vice versa), requires that for every pair $(i, j) \in I \times J$ either (a) $u_i \geq v_{i,j} - \max(p_j, r_{i,j})$ or (b) $p_j \geq m_{i,j}$. The disadvantage of relaxed stability is that there can be situations where no bidder optimal solution exists if the market is *not* in “general position” (see [2] for a formal definition). Here are two canonical examples:

- *Example 1.* There are three bidders and two items. The valuations and reserve prices are as follows: $v_{1,1} = 1, v_{2,1} = 4, v_{2,2} = 4, v_{3,2} = 1, r_{1,1} = 0, r_{2,1} = r_{2,2} = 2$, and $r_{3,2} = 0$. While $\mu = \{(1, 1), (2, 2)\}$ with $p = (0, 2)$ is “best” for bidder 1, $\mu = \{(2, 1), (3, 2)\}$ with $p = (2, 0)$ is “best” for bidder 3.
- *Example 2.* There are two bidders and one item. The valuations and maximum prices are as follows: $v_{1,1} = 10, v_{2,1} = 10$, and $m_{1,1} = m_{2,1} = 5$. While $\mu = \{(1, 1)\}$ with $p_1 = 5$ is “best” for bidder 1, $\mu = \{(2, 1)\}$ with $p_1 = 5$ is “best” for bidder 2.

In the market of the first example no bidder optimal solution exists as long as there exists a bidder that has the same utility functions and reserve prices for two items and two other bidders that are only interested in one of the items. In the market of the second example no bidder optimal solution exists as long as both bidders have the same maximum price and a non-zero utility at the maximum price. Since these cases are quite general, we conjecture that they occur rather frequently in practice.

²Reserve utilities, or *outside options* o_i , can be incorporated by setting $v_{i,j_0} = o_i$ for all bidders i .

³Since we have $u_i \geq 0$ and $u_{i,j}(p_j) = -\infty$ if $p_j \geq m_{i,j}$, this definition is equivalent to requiring $u_i \geq v_{i,j} - p_j$ for all items j with $p_j < m_{i,j}$.




Figure 1: The left two graphs illustrate Example 1. The right two graphs illustrate Example 2. Bidders are on the left side, items on the right side of the bipartite graph. The numbers next to the bidder indicate her utility, the numbers next to the item indicate its price. The labels along the edge show valuations and reserve prices for the left two graphs and valuations and maximum prices for the right two graphs. With relaxed stability a bidder optimal matching does not exist.

With our notion of stability a bidder optimal solution is guaranteed to exist (e.g. $\mu = \{(2,1)\}$ with $p_1 = p_2 = 2$ in Example 1 and $\mu = \emptyset$ with $p_1 = 5$ in Example 2) for *all* kinds of markets, including markets that are *not* in general position.

4. Preliminaries

We define the *first choice graph* $G_p = (I \cup J, F_p)$ at prices p as follows: There is one node per bidder i , one node per item j , and an edge from i to j if and only if item j gives bidder i the highest utility possible, i.e., $u_{i,j}(p_j) \geq u_{i,j'}(p_{j'})$ for all j' . For $i \in I$ we define $F_p(i) = \{j : \exists (i,j) \in F_p\}$ and similarly $F_p(j) = \{i : \exists (i,j) \in F_p\}$. Analogously, for $T \subseteq I$ we define $F_p(T) = \cup_{i \in T} F_p(i)$ and for $S \subseteq J$ we define $F_p(S) = \cup_{j \in S} F_p(j)$. Note that (1) $p_j < m_{i,j}$ for all $(i,j) \in F_p$ and (2) if the matching μ with prices p is stable then $\mu \subseteq F_p$.

We define the *feasible first choice graph* $\tilde{G}_p = (I \cup J, \tilde{F}_p)$ at prices p as follows: There is one node per bidder i , one node per item j , and an edge from i to j if and only if item j gives bidder i the highest utility possible, i.e., $u_{i,j}(p_j) \geq u_{i,j'}(p_{j'})$ for all j' , and $p_j \geq r_{i,j}$. Note that $\tilde{F}_p \subseteq F_p$. For $i \in I$ we define $\tilde{F}_p(i) = \{j : \exists (i,j) \in \tilde{F}_p\}$ and similarly $\tilde{F}_p(j) = \{i : \exists (i,j) \in \tilde{F}_p\}$. Analogously, for $T \subseteq I$ we define $\tilde{F}_p(T) = \cup_{i \in T} \tilde{F}_p(i)$ and for $S \subseteq J$ we define $\tilde{F}_p(S) = \cup_{j \in S} \tilde{F}_p(j)$. Note that (1) $r_{i,j} \leq p_j < m_{i,j}$ for all $(i,j) \in \tilde{F}_p$ and (2) the matching μ with prices p is stable if and only if $\mu \subseteq \tilde{F}_p$. Also note that the edges in $F_p(i) \setminus \tilde{F}_p(i)$ are all the edges (i,j) with maximum $u_{i,j}(p_j)$ but $p_j < r_{i,j}$.

We define an *alternating path* is a sequence of edges in \tilde{F}_p that alternates between matched and unmatched edges. We require that all but the last item on the path are non-dummy items. The last item can (but does not have to) be the dummy item. A tree in the feasible first choice graph \tilde{G}_p is an *alternating tree* rooted at bidder i if all paths from its root to a leaf are alternating paths that either end with the dummy item, an unmatched item, or a bidder whose feasible first choice items are all contained in the tree. We say that an alternating tree with root i is *maximal* if it is the largest such tree. See Figure 2 for an example.

5. Algorithm

Our algorithm starts with an empty matching and prices all zero. It then matches one bidder after the other by augmenting the current matching along an alternating path. If there is no such path, it repeatedly raises the price of all items in the maximal alternating




Figure 2: The graph on the left is the (feasible) first choice graph. The bidders i_1 to i_6 are on the left. The items j_1 to j_5 are on the right. The dummy item is j_0 . Edges in $\mu \cap \tilde{F}_p$ are thick. Edges in \tilde{F}_p are thin. Edges in $F_p \setminus \tilde{F}_p$ are dashed. The graph on the right is a maximal alternating tree rooted at i_1 .

tree under consideration by the minimum amount (a) to make some item $j \notin F_p(i)$ desirable for some bidder i in the tree, or (b) to make some item $j \in F_p(i) \setminus \tilde{F}_p(i)$ feasible for some bidder i in the tree, or (c) to make some item $j \in \tilde{F}_p(i)$ no longer desirable for some bidder i in the tree. Thus it ensures that eventually an alternating path will exist and the matching can be augmented. Note that a matched bidder i can become unmatched if the price of the item j she is matched to reaches $m_{i,j}$. Case (a) corresponds to δ_{out} , Case (b) corresponds to δ_{res} , and Case (c) corresponds to δ_{max} in the pseudocode below.

Modified Hungarian Method

```

1  set  $p_j := 0$  for all  $j \in J$ ,  $u_i := \max_{j'} v_{i,j'}$  for all  $i \in I$ , and  $\mu := \emptyset$ ,
2  while  $\exists$  unmatched bidder  $i$  do
3    find a maximal alternating tree rooted at bidder  $i$  in  $\tilde{G}_p$ 
4    let  $T$  and  $S$  be the set of bidders and items in this tree
5    while all items  $j \in S$  are matched and  $j_0 \notin S$  do
6      compute  $\delta := \min(\delta_{\text{out}}, \delta_{\text{res}}, \delta_{\text{max}})$  where
7         $\delta_{\text{out}} := \min_{i \in T, j \notin F_p(i)} (u_i + p_j - v_{i,j})$ 4
8         $\delta_{\text{res}} := \min_{i \in T, j \in F_p(i) \setminus \tilde{F}_p(i)} (r_{i,j} - p_j)$ 4
9         $\delta_{\text{max}} := \min_{i \in T, j \in F_p(i)} (m_{i,j} - p_j)$ 
10     update prices, utilities, and matching by setting
11      $p_j := p_j + \delta$  for all  $j \in F_p(T) \setminus \setminus$  leads to a new graph  $\tilde{G}_p$ 
12      $u_i := \max_{j'} (v_{i,j'} - p_{j'})$  for all  $i \in I$ 
13      $\mu := \mu \cap \tilde{F}_p \setminus \setminus$  removes unfeasible edges from  $\mu$ 
14     find a maximal alternating tree rooted at bidder  $i$  in  $\tilde{G}_p$ 
15     let  $T$  and  $S$  be the set of bidders and items in this tree
16   end while
17   augment  $\mu$  along alternating path rooted at  $i$  in  $\tilde{G}_p$ 
18 end while
19 output  $p$ ,  $u$ , and  $\mu$ 

```

⁴We need to define $\min_{i \in T, j \in \emptyset} (\dots) = \infty$ as we might have $F_p(I) = J$ or $F_p(i) \setminus \tilde{F}_p(i) = \emptyset$.

6. Feasibility and Stability

Theorem 6.1. *The Modified HM finds a feasible and stable matching. It can be implemented to run in $O(nk^3 \log(k))$.*

Proof. The matching μ constructed by the Modified HM is a subset of the feasible first choice graph \tilde{G}_p at all times. Hence it suffices to show that after $O(nk^3 \log(k))$ steps all bidders are matched.

The algorithm consists of two nested loops. We analyze the running time in two steps: (1) The time spent in the outer loop *without* the inner loop (ll. 2–4 and 17–18) and (2) the time spent in the inner loop (ll. 5–16). Note that after each execution of the outer while loop the number of matched bidder increases by one. A matched bidder i can only become unmatched if the price of the item j she is matched to reaches $m_{i,j}$. This can happen only once for each pair (i, j) , which implies that each bidder can become at most k times unmatched. Thus, the outer loop is executed at most nk times. Since $|S| \leq k$, it follows that $|T| \leq k$. Thus it is straightforward to implement the outer while loop in time $O(k^2)$.

We call an execution of the inner while loop *special* if (a) right before the start of the execution the outer while loop was executed, (b) in the previous iteration of the inner while loop the maximum price of a pair (i, j) was reached, or (c) the reserve price of a pair (i, j) was reached. As each of these cases can happen at most nk times, there are at most $3nk$ special executions of the inner while loop. Non-special executions increase the number of items in the maximal alternating tree by at least one. Thus there are at most k non-special executions between any two consecutive special executions. We present next a data structure that (1) can be built in time $O(k^2)$ and (2) allows to implement *all* non-special executions of the inner while loop between two consecutive special iterations in time $O(k^2 \log k)$. Thus the total time of the algorithm is $O(nk^3 \log k)$.

Data structure:

- (1) Keep a list of all bidders in T and a bit vector of length n where bit i is set to 1 if bidder i belongs currently to T and to 0 otherwise. Keep a list of all items in S and bit vector of length k , where bit j is set of 1 if item j belongs currently to S and to 0 otherwise. Finally also keep a list and a bit vector of length k representing all items in $F_p(T)$.
- (2) Keep a heap H_{out} and a value δ_{out} , such that H_{out} stores $x_i + p_j - v_{i,j}$ for all pairs (i, j) with $i \in T$ and $j \notin F_p(i)$ and $\delta_{\text{out}} + x_i$ equals u_i for every $i \in T$. Keep a heap H_{res} and a value δ_{res} , such that H_{res} stores $r_{i,j} - y_j$ for all pairs (i, j) with $i \in T$ and $j \in F_p(i) \setminus \tilde{F}_p(i)$ and $\delta_{\text{res}} + y_j$ equals p_j for every $j \in F_p(i) \setminus \tilde{F}_p(i)$. Keep a heap H_{max} and a value δ_{max} , such that H_{max} stores $m_{i,j} - y_j$ for all pairs (i, j) with $i \in T$ and $j \in F_p(i)$ and $\delta_{\text{max}} + y_j$ equals p_j for every $j \in F_p(i)$.
- (3) We also store at each bidder i its current u_i , at each item j its current p_j . Thus given a pair (i, j) we can decide in constant time whether $u_i = v_{i,j} - p_j$, i.e., whether $j \in F_p(i)$. Finally we keep a list of edges in μ .

At the beginning of each special execution of the inner while loop a list of bidders and items currently in T and S are passed in either from the preceding execution of the outer while loop (where T and S are constructed in time $O(k^2)$) or from the previous execution of the inner while loop. Recall that $|S| \leq k$ and thus $|T| \leq k$. Thus we can build the above data structures from scratch in time $O(k^2)$ as follows. To initialize the bit vector for T we use the following approach: At the beginning of the algorithm the vector is once initialized

to 0, taking time $O(n)$. Then at the beginning of all but the first special execution of the inner while loop the bit vector is “cleaned” by setting the bit of all elements of T in the previous iteration to 0 using the list of elements of T of the previous iteration. Then the list of elements currently in T is used to set the appropriate bits to 1. This takes time $O(k)$ per special execution. The bit vector of items in S has only k entries and thus is simply initialized to 0 at the beginning of each special execution. Then the list of elements currently in S is used to set the appropriate bits to 1. Given the list of bidders in T we decide in constant time for each pair (i, j) with $i \in T$ into which heap(s) its appropriate values should be inserted. If $j \in F_p(i)$ we also add j to $F_p(T)$ if it is not already in this set update the bit vector and the list. When we have processed all pairs (i, j) with $i \in T$ we build the three heaps in time linear in their size such that all δ values are 0. Since $|S| = k$ we know that $|T| = k$. Thus, the initialization takes time $O(k^2)$.

To implement each iteration of the inner while loop we first perform a find-min operation on all three heaps to determine δ . Then we remove all heap values that equal δ . Afterwards we update the price of all items in $F_p(T)$ using the list of $F_p(T)$. We also update the utility of all items in T as follows. If $\delta \neq \delta_{\max}$ updating the utilities is just a simple subtraction per bidder. If $\delta = \delta_{\max}$, i.e., p_j becomes $m_{i,j}$ for some pair (i, j) , then updating u_i requires computing $v_{i,j} - p_j$ for all j and potentially removing the edge (i, j) from μ , which in turn might cut a branch of the alternating tree. Thus, in this case we completely rebuild the alternating tree, including S , T , and $F_p(T)$ from scratch. Note however that this can only happen in a special execution of the inner while loop. If $\delta \neq \delta_{\max}$ the elements removed from the heaps tell us which new edges are added to $\tilde{F}_p(T)$ and which new items to add to $F_p(T)$. The new items in $F_p(T)$ gives a set of items from which we start to augment the alternating tree in breadth first manner. For each new item j , we add to $\tilde{F}_p(T)$ the bidder it is matched to as new bidder to S and to $\tilde{F}_p(T)$. For each new bidder i added to $\tilde{F}_p(T)$ we spend time $O(k)$ to determine its adjacent edges in $F_p(i)$ and insert the suitable values for the pairs (i, j) into the three heaps. This process repeats until no new items and no new bidders are added to $F_p(i)$. During this traversal we also update the bit vectors and lists representing T , S , and $F_p(T)$. Let T_{new} be the set of bidders added to T during an execution of the inner while loop and let r be the number of elements removed from the heaps during the execution. Then the above data structures implement the inner while loop in time $O(r * \log k + |T_{\text{new}}| * k)$. Now note that during a sequence of non-special executions of the inner while loop between two consecutive special executions bidders are never removed from T and each (i, j) pair with $i \in T$ is added (and thus also removed) at most once from each heap. Thus the total number of heap removals during *all* such non-special executions is $3k^2$ and the total number of elements added to T is k , giving a total running time of $O(k^2 \log k)$ for *all* such non-special executions. Since there are at most $3nk$ special executions, the total time for all inner while loops is $O(nk^3 \log k)$. ■

7. Bidder Optimality

Theorem 7.1. *The Modified HM finds a bidder optimal matching in $O(nk^3 \log(k))$ steps.*

We say that a (possibly empty) set $S \subseteq J$ is *strictly overdemanding* for prices p wrt $T \subseteq I$ if (i) $\tilde{F}_p(T) \subseteq S$ and (ii) $\forall R \subseteq S$ and $R \neq \emptyset : |\tilde{F}_p(R) \cap T| > |R|$. Using Hall’s Theorem [11] one can show that a feasible and stable matching exists for given prices p if and only if there is no strictly overdemanding set of items S in \tilde{F}_p .

The proof strategy is as follows: In Lemma 7.2 we show that a feasible and stable matching μ with prices p is bidder optimal if we have that $p_j \leq p'_j$ for all items j and all feasible and stable matchings μ' with prices p' . Afterwards, in Lemma 7.3, we establish a lower bound on the price increase of strictly overdemanded items. Finally, in Lemma 7.4 we argue that whenever the Modified HM updates the prices it updates the prices according to Lemma 7.3. This completes the proof.

Lemma 7.2. *If the matching μ with prices p is stable and $p_j \leq p'_j$ for all j and all stable matchings μ' with prices p' , then the matching μ with prices p is bidder optimal.*

Proof. For a contradiction suppose that there exists a feasible and stable matching μ' with prices p' such that $u'_i > u_i$ for some bidder i . Let j be the item that bidder i is matched to in μ and let j' be the item that bidder i is matched to in μ' . Since $p_{j'} \leq p'_{j'}$ and $p'_{j'} < m_{i,j'}$ we have that $u_{i,j'}(p_{j'}) = v_{i,j'} - p_{j'}$. Since the matching μ with prices p is stable we have that $u_i = u_{i,j}(p_j) = v_{i,j} - p_j \geq u_{i,j'}(p_{j'}) = v_{i,j'} - p_{j'}$. It follows that $u'_i = v_{i,j'} - p'_{j'} > u_i = v_{i,j} - p_j \geq v_{i,j'} - p_{j'}$ and, thus, $p'_{j'} < p_{j'}$. This gives a contradiction. ■

Lemma 7.3. *Given $p = (p_1, \dots, p_k)$ let $u_i = \max_j u_{i,j}(p_j)$ for all i . Suppose that $S \subseteq J$ is strictly overdemanded for prices p with respect to $T \subseteq I$ and let $\delta = \min(\delta_{out}, \delta_{res}, \delta_{max})$, where $\delta_{out} = \min_{i \in T, j \notin F_p(i)}(u_i + p_j - v_{i,j})$, $\delta_{res} = \min_{i \in T, j \in F_p(i) \setminus \tilde{F}_p(i)}(r_{i,j} - p_j)$, and $\delta_{max} = \min_{i \in T, j \in F_p(i)}(m_{i,j} - p_j)$. Then, for any stable matching μ' with prices p' with $p'_j \geq p_j$ for all j , we have that $p'_j \geq p_j + \delta$ for all $j \in F_p(T)$.*

Proof. We prove the claim in two steps. In the first step, we show that $p'_j \geq p_j + \delta$ for all $j \in \tilde{F}_p(T)$. In the second step, we show that $p'_j \geq p_j + \delta$ for all $j \in F_p(T) \setminus \tilde{F}_p(T)$.

Step 1. Consider the set of items $A = \{j \in \tilde{F}_p(T) \mid \forall k \in \tilde{F}_p(T) : p'_j - p_j \leq p'_k - p_k\}$ and the set of bidders $B = \tilde{F}_p(A) \cap T$. Assume by contradiction that $\delta' = \min_{j \in \tilde{F}_p(T)}(p'_j - p_j) < \delta$. We show that this implies that $|B| > |A| \geq |\tilde{F}_{p'}(B)|$, which gives a contradiction.

The set of items S is strictly overdemanded for prices p wrt to T and A . Thus, since $A \subseteq S$ and $A \neq \emptyset$, $|B| = |\tilde{F}_p(A) \cap T| > |A|$. Next we show that $A \supseteq \tilde{F}_{p'}(B)$ and, thus, $|A| \geq |\tilde{F}_{p'}(B)|$. It suffices to show that $\tilde{F}_{p'}(i) \setminus A = \emptyset$ for all bidders $i \in B$. For a contradiction suppose that there exists a bidder $i \in B$ and an item $k \in \tilde{F}_{p'}(i) \setminus A$. Recall that we must have (1) $u_{i,k}(p'_k) \geq 0$, (2) $u_{i,k}(p'_k) \geq u_{i,k'}(p'_{k'})$ for all k' , and (3) $p_k \geq r_{i,k}$. Recall also that (1)–(3) imply that $r_{i,k} \leq p'_k < m_{i,k}$ and so $u_{i,k}(p'_k) = v_{i,k} - p'_k$.

We know that there exists $j \in A$ such that $j \in \tilde{F}_p(i)$. Since $j \in A$ we have that $p'_j < p_j + \delta \leq m_{i,j}$ and so $u_{i,j}(p'_j) = v_{i,j} - p'_j$. Thus, since $k \in \tilde{F}_{p'}(i)$, $v_{i,k} - p'_k \geq v_{i,j} - p'_j$. Finally, since $j \in \tilde{F}_p(i)$ and $p_k \leq p'_k < m_{i,k}$, we have that $u_{i,j}(p_j) = v_{i,j} - p_j \geq u_{i,k}(p_k) = v_{i,k} - p_k$.

Case 1: $k \in J \setminus F_p(B)$. Since $\delta \leq \delta_{out} \leq u_i + p_k - v_{i,k}$ and $u_i = v_{i,j} - p_j$ we have that $\delta \leq v_{i,j} - p_j + p_k - v_{i,k}$. Rearranging this gives $v_{i,k} - p_k + \delta \leq v_{i,j} - p_j$. Since $p'_k \geq p_k$ and $p_j > p'_j - \delta$ this implies that $v_{i,k} - p'_k < v_{i,j} - p'_j$. Contradiction!

Case 2: $k \in F_p(B) \setminus \tilde{F}_p(B)$. If $p'_k - p_k \leq p'_j - p_j = \delta'$ then $p'_k \leq p_k + \delta' < p_k + \delta$. Since $\delta \leq \delta_{res} \leq r_{i,k} - p_k$ this implies that $p'_k < r_{i,k}$. Contradiction! Otherwise, $p'_k - p_k > p'_j - p_j$. Since $v_{i,j} - p_j \geq v_{i,k} - p_k$ this implies that $v_{i,j} - p'_j > v_{i,k} - p'_k$. Contradiction!

Case 3: $k \in \tilde{F}_p(B) \setminus A$. Since $j \in A$ and $k \notin A$ we have that $p'_k - p_k > \delta' = p'_j - p_j$. Since $v_{i,j} - p_j \geq v_{i,k} - p_k$ this implies that $v_{i,j} - p'_j > v_{i,k} - p'_k$. Contradiction!

Step 2. Consider an arbitrary item $j \in F_p(T) \setminus \tilde{F}_p(T)$ such that $p'_j - p_j \leq p'_{j'} - p_{j'}$ for all $j' \in F_p(T) \setminus \tilde{F}_p(T)$ and a bidder $i \in T$ such that $j \in F_p(i)$. Assume by contradiction that $\delta' = p'_j - p_j < \delta$. We show that this implies that $\tilde{F}_{p'}(i) = \emptyset$, which gives a contradiction.

First observe that $\delta' < \delta \leq \delta_{\text{res}} \leq r_{i,j} - p_j$ and, thus, $p'_j < p_j + \delta \leq r_{i,j}$, which shows that $j \notin \tilde{F}_{p'}(i)$. Next consider an arbitrary item $k \neq j$. For a contradiction suppose that $k \in \tilde{F}_{p'}(i)$. It follows that $r_{i,k} \leq p'_k < m_{i,k}$ and $u_{i,k}(p'_k) = v_{i,k} - p'_k \geq u_{i,j}(p'_j)$.

Since $p'_j = p_j + \delta' < p_j + \delta \leq m_{i,j}$ we have that $u_{i,j}(p'_j) = v_{i,j} - p'_j$ and so $v_{i,k} - p'_k \geq v_{i,j} - p'_j$. Finally, since $j \in F_p(i)$ and $p_k \leq p'_k < m_{i,k}$, we have that $u_{i,j}(p_j) = v_{i,j} - p_j \geq u_{i,k}(p_k) = v_{i,k} - p_k$.

Case 1: $k \in J \setminus F_p(T)$. Since $\delta \leq \delta_{\text{out}} \leq u_i + p_k - v_{i,k}$ and $u_i = v_{i,j} - p_j$ we have that $\delta \leq v_{i,j} - p_j + p_k - v_{i,k}$. Rearranging this gives $v_{i,k} - p_k + \delta \leq v_{i,j} - p_j$. Since $p'_k \geq p_k$ and $p_j > p'_j - \delta$ this implies that $v_{i,k} - p'_k < v_{i,j} - p'_j$. Contradiction!

Case 2: $k \in F_p(T) \setminus \tilde{F}_p(T)$. If $p'_k - p_k \leq p'_j - p_j = \delta'$ then $p'_k \leq p_k + \delta' < p_k + \delta$. Since $\delta \leq \delta_{\text{res}} \leq r_{i,k} - p_k$ this implies that $p'_k < r_{i,k}$. Contradiction! Otherwise, $p'_k - p_k > p'_j - p_j$. Since $v_{i,j} - p_j \geq v_{i,k} - p_k$ this implies that $v_{i,j} - p'_j > v_{i,k} - p'_k$. Contradiction!

Case 3: $k \in \tilde{F}_p(T)$. From Step 1 we know that $p'_k - p_k \geq \delta > \delta' = p'_j - p_j$. Since $v_{i,j} - p_j \geq v_{i,k} - p_k$ this implies that $v_{i,j} - p'_j > v_{i,k} - p'_k$. Contradiction! ■

Lemma 7.4. *Let p be the prices computed by the Modified HM. Then for any stable matching μ' with prices p' we have that $p_j \leq p'_j$ for all j .*

Proof. We prove the claim by induction over the price updates. Let p^t denote the prices after the t -th price update.

For $t = 0$ the claim follows from the fact that $p^t = 0$ and $p'_j \geq 0$ for all items j and all feasible matchings μ' with prices p' .

For $t > 0$ assume that the claim is true for $t - 1$. Let S be the set of items and let T be the set of bidders considered by the matching mechanism for the t -th price update. We claim that S is strictly overdemanding for prices p^{t-1} wrt to T . This is true because: (1) S and T are defined as the set of items resp. bidders in a *maximal* alternating tree and, thus, there are no edges in $\tilde{F}_{p^{t-1}}$ from bidders in T to items in $J \setminus S$ which shows that $\tilde{F}_{p^{t-1}}(T) \subseteq S$. (2) For all subsets $R \subset S$ and $R \neq \emptyset$ the number of “neighbors” in the alternating tree under consideration is strictly larger than $|R|$ which shows that $|\tilde{F}_{p^{t-1}}(R) \cap T| > |R|$. By the induction hypothesis $p'_j \geq p_j^{t-1}$ for all items $j \in J$ and, thus, Lemma 7.3 shows that $p'_j \geq p_j^{t-1} + \delta$ for all items $j \in F_{p^{t-1}}(t)$. The Modified HM sets $p_j^t = p_j^{t-1} + \delta$ for all items $j \in F_{p^{t-1}}(T)$ and $p_j^t = p_j^{t-1}$ for all items $j \notin F_{p^{t-1}}(T)$ and so $p'_j \geq p_j^t$ for all items $j \in J$. ■

8. Truthfulness

The following example shows that with our notion of stability bidder optimality no longer implies truthfulness, even if (i) there are no reserve prices, i.e., $r_{i,j} = 0$ for all i and j , (ii) maximum prices depend only on the item, i.e., for all i there exists a constant m_i such that $m_{i,j} = m_i$ for all j , and (iii) no two bidders have the same maximum price, i.e., $m_i \neq m_k$ for any two bidders $i \neq k$. More specifically, it shows that a bidder can improve her utility by lying about the valuation of a single item. Since the bidder optimal utilities are uniquely defined, this shows that no mechanism that computes a bidder optimal matching μ with prices p can be truthful. Note that if (i) to (iii) hold and there exists constants

$\alpha_1 \geq \dots \geq \alpha_k$ and v_1, \dots, v_k such that $v_{i,j} = v_i \cdot \alpha_j$ for all i and j , then Ashlagi et al. [3] show the existence of a truthful mechanism.




Figure 3: Bidders are on the left and items are on the right. The numbers next to the bidders indicate their utilities. The numbers next to the items indicate their prices. The labels along the edges show valuations and maximum prices. The graph on the left depicts the bidder optimal matching for the “true” valuations. The graph on the right depicts the bidder optimal matching for the “falsified” valuations. Specifically, in the matching on the right bidder 2 misreports her valuation for item 1. This gives her a strictly higher utility, and shows that lying “pays off”.

9. Generalized Linear Utilities

The following theorem generalizes our results to utilities of the form $u_{i,j}(p_j) = v_{i,j} - c_i \cdot c_j \cdot p_j$ for $p_j < m_{i,j}$ and $u_{i,j}(p_j) = -\infty$ otherwise. This reduction does *not* work if $u_{i,j}(p_j) = v_{i,j} - c_{i,j} \cdot p_j$ for $p_j < m_{i,j}$ and $u_{i,j}(p_j) = -\infty$ otherwise. We prove the existence of a bidder optimal solution for more general utilities in [6].

Theorem 9.1. *The matching $\hat{\mu}$ with prices \hat{p} is bidder optimal for $\hat{v} = (\hat{v}_{i,j})$, $\hat{r} = (\hat{r}_{i,j})$, $\hat{m} = (\hat{m}_{i,j})$ and utilities $u_{i,j}(p_j) = v_{i,j} - c_i \cdot c_j \cdot p_j$ if $p_j < m_{i,j}$ and $u_{i,j}(p_j) = -\infty$ otherwise if and only if the matching μ with prices p , where $\mu = \hat{\mu}$ and $p = (c_j \cdot \hat{p}_j)$, is bidder optimal for $v = (\hat{v}_{i,j}/c_i)$, $r = (c_j \cdot \hat{r}_{i,j})$, $m = (c_j \cdot \hat{m}_{i,j})$ and utilities $u_{i,j}(p_j) = v_{i,j} - p_j$ if $p_j < m_{i,j}$ and $u_{i,j}(p_j) = -\infty$ otherwise.*

Proof. Since $\hat{p}_j < \hat{m}_{i,j}$ if and only if $p < m_{i,j}$ we have that $\hat{u}_{i,j}(\hat{p}_j) = c_i \cdot u_{i,j}(p_j)$. Since $\hat{\mu} = \mu$ this implies that $\hat{u}_i = c_i \cdot u_i$ for all i .

Feasibility. Since $c_i > 0$ for all i we have that $\hat{u}_i \geq 0$ for all i if and only if $u_i = \hat{u}_i/c_i \geq 0$ for all i . Since $c_j > 0$ for all j we have that $\hat{p}_j \geq 0$ for all j if and only if $p_j = c_j \cdot \hat{p}_j \geq 0$ for all j . Since $\mu = \hat{\mu}$ and $r_{i,j} = c_j \cdot \hat{r}_{i,j}$, $p_j = c_j \cdot \hat{p}_j$, and $m_{i,j} = c_j \cdot \hat{m}_{i,j}$ for all i and j we have that $\hat{r}_{i,j} \leq \hat{p}_j < \hat{m}_{i,j}$ for all $(i,j) \in \hat{\mu}$ if and only if $r_{i,j} \leq p_j < m_{i,j}$ for all $(i,j) \in \mu$.

Stability. If $\hat{\mu}$ with \hat{p} is stable then μ with p is stable because $u_i = c_i \cdot \hat{u}_i \geq c_i \cdot \hat{u}_{i,j}(\hat{p}_j) = u_{i,j}(p_j)$ for all i and j . If μ with p is stable then $\hat{\mu}$ with \hat{p} is stable because $\hat{u}_i = u_i/c_i \geq u_{i,j}(p_j)/c_i = \hat{u}_{i,j}(\hat{p}_j)$ for all i and j .

Bidder Optimality. For a contradiction suppose that $\hat{\mu}$ with \hat{p} is bidder optimal but μ with p is not. Then there must be a feasible and stable matching μ' with p' such that $u'_i > u_i$ for at least one bidder i . By transforming μ' with p' into $\hat{\mu}'$ with \hat{p}' we get a feasible and stable matching for which $\hat{u}'_i = c_i \cdot u'_i > c_i \cdot u_i = \hat{u}_i$. Contradiction!

For a contraction suppose that μ with p is bidder optimal but $\hat{\mu}$ with \hat{p} is not. Then there must be a feasible and stable matching $\hat{\mu}'$ with \hat{p}' such that $\hat{u}'_i > \hat{u}_i$ for at least one bidder i . By transforming $\hat{\mu}'$ with \hat{p}' into μ' with p' we get a feasible and stable matching for which $u'_i = \hat{u}'_i/c_i > \hat{u}_i/c_i = u_i$. Contradiction! ■

References

- [1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. *Proceedings of the Conference on Electronic Commerce*, pages 1–7, 2006.
- [2] G. Aggarwal, S. Muthukrishnan, D. Pál, and M. Pál. General auction mechanism for search advertising. *Proceedings of the World Wide Web Conference*, pages 241–250, 2009.
- [3] I. Ashlagi, M. Braverman, A. Hassidim, R. Lavi, and M. Tennenholtz. Position auctions with budgets: Existence and uniqueness. *Working Paper*, 2009.
- [4] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [5] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Political Economy*, 94(4):863–72, 1986.
- [6] P. Dütting, M. Henzinger, and I. Weber. Bidder optimal assignments for general utilities. *Proceedings of the Workshop on Internet and Network Economics*, pages 575–582, 2009.
- [7] 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.
- [8] A. Frank. On Kuhn’s Hungarian Method. *Naval Research Logistics*, 51:2–5, 2004.
- [9] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
- [10] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [11] P. Hall. On representatives of subsets. *London Mathematical Society*, 10:26–30, 1935.
- [12] D. Kempe, A. Mu’alem, and M. Salek. Envy-free allocations for budgeted bidders. *Proceedings of the Workshop on Internet and Network Economics*, pages 537–544, 2009.
- [13] H. W. Kuhn. The Hungarian Method for the assignment problem. *Naval Research Logistics*, 2:83–97, 1955.
- [14] S. Lahaie, D. M. Pennock, A. Saberi, and R. V. Vohra. *Algorithmic Game Theory*, chapter 28, pages 699–716. Cambridge University Press, 2007.
- [15] A. E. Roth and M. Sotomayor. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.
- [16] L. S. Shapley and M. Shubik. The assignment game: The core I. *Game Theory*, 29:111–130, 1972.
- [17] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Finance*, 16:8–27, 1961.