Truthful unit-demand auctions with budgets revisited☆

Monika Henzinger, Veronika Loitzenbauer*

University of Vienna, Faculty of Computer Science
Währinger Straße 29, 1090 Vienna, Austria

Abstract

We consider auctions of indivisible items to unit-demand bidders with budgets. This setting was suggested as an expressive model for single sponsored search auctions. Prior work presented mechanisms that compute bidder-optimal outcomes and are truthful for a *restricted* set of inputs, i.e., inputs in so-called general position. This condition is easily violated. We provide the first mechanism that is truthful in expectation for *all* inputs and achieves for each bidder no worse utility than the bidder-optimal outcome. Additionally we give a complete characterization for which inputs mechanisms that compute bidder-optimal outcomes are truthful.

Keywords: matching market, ascending auction, budget constraint, incentive compatibility, envy-freeness, randomized algorithm

1. Introduction

When a user searches for a specific term in a web search engine, related advertisements are displayed on the search results page. The advertisements are assigned by an auction such that each advertiser receives at most one slot. The advertisers may have different preferences among the slots and budgets that limit the amount of money they can spend. The auctioneer may have a *reserve price* under which she is not interested in selling a slot. This is modeled as follows: The advertisers are called *bidders*, the slots correspond to *items*, and the budgets are modeled as *maximum prices* per bidder and item. Each bidder i can specify a *valuation* $v_{i,j}$ for each item j. A mechanism computes the prices p of the items as well as the assignment μ of the bidders to the items. The preferences of a bidder i are modeled by utility functions $u_{i,j}(p_j)$ such that his utility if he is assigned item j at a price p_j is $u_{i,j}(p_j) = v_{i,j} - p_j$ if the price is lower than his maximum price and $-\infty$ if the price is equal to or higher than his maximum price. A bidder has a utility of zero if he is not assigned any item; thus bidder i only accepts an item j if $u_{i,j}(p_j) \geq 0$. For more details on the expressiveness of quasi-linear utility functions with budgets in sponsored search see [1].

Search engine providers want to satisfy their customers as well as avoid fluctuations in the prices. This corresponds to a bidder-optimal and stable assignment of bidders to items. An outcome (μ, p) is *stable* if a competitive equilibrium is reached. In a *competitive equilibrium* no bidder would prefer a different item or being unmatched to the one he is matched to under the current prices, i.e., every bidder is *envy-free*, and additionally the prices of all unmatched items are equal to their reserve prices. In a *bidder-optimal* outcome each bidder obtains his best utility among all envy-free outcomes. To simplify the bidding for the advertisers

☆NOTICE: this is the authors version of a work that was accepted for publication in Theoretical Computer Science. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in Theoretical Computer Science, 573:(1–15), 2015, DOI 10.1016/j.tcs.2015.01.033

*Corresponding author.
Email address: veronika.loitzenbauer@univie.ac.at (Veronika Loitzenbauer)

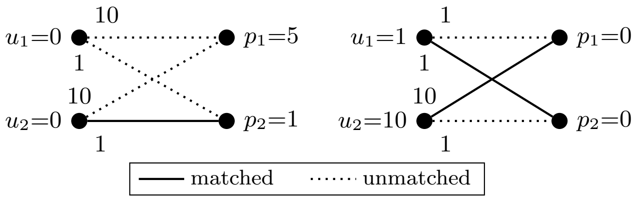

Figure 1: An input for which no bo-mechanism is truthful. The edges are labeled with the valuations the bidders report for the items. The maximum prices of both bidders are 5. The displayed prices are the minimum envy-free prices for the given inputs, respectively. In the left graph the bidders report their true values. In the right graph the first bidder reports a wrong valuation of 1 for the first item. Thus an envy-free outcome with respect to the reported valuations exists already at the initial prices of zero and both bidders obtain a higher utility. Note that the utility gain could be arbitrarily high.

as well as to be able to compute an envy-free outcome with respect to the true values of the bidders, bidders should obtain their best possible utility if they reveal their true preferences to the mechanism. This property is called *truthfulness* or *incentive-compatibility*.

A large body of prior work on this and related problems exists. We summarize below only the most closely related work, see [2] for a more complete overview.

A bidder-optimal outcome exists for all strictly monotonically falling and locally right-continuous utility functions [3] and, thus, for the model used in this work. We say that a mechanism that computes a bidder-optimal outcome for every given input is a *bo-mechanism*. By definition, the utilities of the bidders in a bidder-optimal outcome are unique. Thus one bo-mechanism is incentive-compatible for a given input if and only if all bo-mechanisms are incentive-compatible for that input. Note that we distinguish between incentive-compatibility for a specific input and incentive-compatibility for all inputs.

Without budgets every bo-mechanism is incentive-compatible and its outcome is stable [4, 5]. The inclusion of budgets into the model implies discontinuities in the utility functions, which destroys these desirable properties in general [1, 6, 7]. Aggarwal et al. [1] were among the first who added budget constraints to quasi-linear utility functions. For inputs in *general position*, i.e., certain non-degenerate inputs, Aggarwal et al. provided an incentive-compatible mechanism that computes a bidder-optimal stable outcome in polynomial time.[1] Aggarwal et al. state in [1], with (v, m, r) being the input to the auction:

> *In essence, any auction (v, m, r) can be brought into general position by arbitrarily small (symbolic) perturbations. In practice this assumption is easily removed by using a consistent tie-breaking rule.*

We provide an example that shows that neither a deterministic nor a randomized tie-breaking rule, as suggested above, leads to an incentive-compatible mechanism for *all* inputs. As for the undisturbed input, the gain from lying can be arbitrarily high. Instead, we use further randomization, this time of the prices, and give a mechanism based on randomized tie-breaking and on randomized price extraction that is truthful in expectation. However, as shown in [1, 7], there are degenerate inputs for which no bo-mechanism is incentive-compatible, even if the outcome is stable [3]. Hence, our randomized mechanism cannot be a bo-mechanism. Our mechanism builds upon the results of Dütting et al. in [8], who showed that a modification of the Hungarian Method [9] computes (in polynomial time) a bidder-optimal envy-free outcome for *every* given input, i.e., is a bo-mechanism. If the input is in general position, this mechanism is incentive-compatible and the outcome is a competitive equilibrium [7].

General position is a quite restrictive condition on the input. Intuitively, it forbids that any two maximum prices can be reached simultaneously during an ascending price mechanism. For example, "symmetric inputs",

[1] Chen et al. [6] defined *weak* and *strong* stability, where only for the latter a stable outcome is equal to a competitive equilibrium. The definition of a stable matching in [1] corresponds to weak stability. In their model, the utility of a bidder is set to a negative value if the price of the item strictly exceeds the maximum price of this bidder for the item. On the contrary, in [8] as well as in the definition used in this work, a utility becomes negative as soon as the maximum price is reached. For the utility functions used in [8] and in this work weak and strong stability coincide.

i.e., inputs where two bidders input exactly the same valuations and budgets, violate the general position condition. However, in practice such inputs can easily arise. Consider the example with symmetric bidders in Figure 1. Both bidders have the same budget and prefer the first item over the second. If a bidder-optimal outcome is computed for the true values, the most desirable item is not sold and both bidders have a utility of zero. Furthermore, through lying one of the bidders can ensure that all desirable items are sold. Thus no bo-mechanism can be incentive-compatible. A good outcome in this situation would be to assign the most desirable item with equal probability to each bidder and to charge prices so that the expected utility of each bidder is at least the utility that the bidder could achieve through lying. This is exactly what our algorithm does.

More formally, we improve upon the known results in three ways. The requirement that the input is in general position is a sufficient but not necessary condition for the existence of a truthful bo-mechanism. Specifically, there exist inputs that are not in general position but for which the Modified Hungarian Method (and thus every bo-mechanism) is incentive-compatible. Furthermore, there exists no polynomial-time algorithm that determines whether an input is in general position. Our first contributions are a generalization of the general position condition called *rematch condition* that excludes fewer inputs and a polynomial time algorithm to determine whether an input fulfills the rematch condition. Thus our condition provides new insights on when a bo-mechanism cannot guarantee incentive-compatibility for a given input. We use these insights then in our second and third contributions.

Our second contribution provides the first necessary and sufficient condition for incentive-compatibility of a bo-mechanism for a given input as well as the first general description of the way in which a bidder can lie. Additionally, we give a polynomial-time algorithm for computing the *regret*, i.e., the maximal utility a bidder can gain from lying.

Our third result is the already mentioned randomized mechanism that is incentive-compatible in expectation and extends the previous mechanism of [8]. Contrary to Aggarwal et. al.'s claim, tie-breaking is not sufficient for an incentive-compatible mechanism. However, a combination of randomized tie-breaking with randomized price extraction yields a mechanism that is truthful in expectation. However, when using tie-breaking, envy-freeness is lost. Note that this is unavoidable even for mechanisms that are truthful in expectation: To show truthfulness in expectation we have to compare the outcome of our mechanism for the truthful input with the outcome for the input where one bidder lies. We show that in the latter case the liar will have envy (with regard to his true values). See, for example, the right graph in Figure 1. The first bidder lies such that he is matched to the second item and obtains a higher utility. However, he has envy as he would prefer the first to the second item at the current prices. Thus the truthfulness (in expectation) condition requires that our mechanism computes an outcome that gives the liar at least the (expected) utility of the outcome with envy achievable by a lie, i.e., a utility that is higher than his bidder-optimal utility. In these cases the outcome cannot be envy-free.

As shown by the example in Figure 1, it is not possible to bound the amount of envy that arises after a lie as the valuation that the first bidder has for the first item can be arbitrarily high. Our randomized mechanism would give the first item with probability 0.5 to the first bidder and with probability 0.5 to the second bidder. In both cases the price charged in expectation for the first item would be slightly below the bidders' budgets, while the price of the second item would be zero. Note that the price of the first item cannot be increased up to the minimum envy-free price of 5 as this price equals the budget. Thus the expected utility is at least as high as the utility that a bidder could achieve through lying when a bo-mechanism was used. We can show that this holds in general for our mechanism.

Additionally our randomized mechanism fulfills the following properties: (a) If the maximum prices are public knowledge, our mechanism is ex-post incentive-compatible. (b) Its output is "close" to the output of [8] in the following way: If the input is in general position, an envy-free outcome with minimum prices, i.e., a bidder-optimal outcome, can be obtained by slightly increasing the prices of the outcome of our mechanism. Thus our mechanism matches the known results for these inputs and can additionally guarantee truthfulness in expectation for every input.

In the remaining part of this section we define the relevant terminology and describe the Modified Hungarian Method of [8]. Section 2 is devoted to the rematch condition, Section 3 to the characterization of inputs for which bo-mechanisms are truthful, and Section 4 to our randomized mechanism.

1.1. Definitions

A set I of n bidders and a set J of k items participate in the auction. Each bidder $i \in I$ reports a valuation $v_{i,j} \geq 0$ and a maximum price $m_{i,j} \geq 0$ for each item $j \in J$ to the mechanism. Each item $j \in J$ has a reserve price $r_j \geq 0$ under which the item will not be sold. Let v and m be $n \times k$ matrices with entries $v_{i,j}$ resp. $m_{i,j}$ for $1 \leq i \leq n$ and $1 \leq j \leq k$, and let r be a vector of length k with entries r_j for $1 \leq j \leq k$. Then the input of the mechanism is given by (v, m, r).

The utility u_i of a bidder i is zero if he is not matched to any item and $u_{i,j}(p_j)$ if he is matched to an item j and this item has a price p_j. The function $u_{i,j}(p_j)$ is locally right-continuous and monotonically falling and is defined as

$$u_{i,j}(p_j) = \begin{cases} v_{i,j} - p_j & \text{if } p_j < m_{i,j} \ , \\ -\infty & \text{otherwise} \ . \end{cases} \tag{1}$$

We will also use the continuous and monotonically falling function $u_{i,j}^{-1}(u_i)$ for $u_i \geq 0$. It returns a lower bound on p_j such that $u_{i,j}(p_j) \leq u_i$. For our model it is given by

$$u_{i,j}^{-1}(u_i) = \begin{cases} v_{i,j} - u_i & \text{if } v_{i,j} - u_i < m_{i,j} \ , \\ m_{i,j} & \text{otherwise} \ . \end{cases} \tag{2}$$

A bidder is matched whenever he is assigned to an item and vice versa. To simplify the description of the algorithm, a dummy item $j_0 \in J$ with a reserve price of zero is used to model unmatched bidders. All bidders have a valuation of zero and a maximum price of ∞ for the dummy item. Each bidder can be matched to at most one item and each item except the dummy item can be matched to at most one bidder.

The mechanism outputs a *matching* $\mu \subseteq I \times J$ between bidders and items as well as *prices* $p = (p_1, \ldots, p_k)$ for all items, together denoted as *outcome* (μ, p). We denote by $\mu(i)$ the item a bidder i is matched to, by $\mu^{-1}(j)$ the bidder an item j is matched to, by $\mu(I)$ the set of items a set of bidders I is matched to, and by $\mu^{-1}(J)$ the set of bidders a set of items J is matched to.

For a *feasible* outcome (μ, p) we require $u_{i,\mu(i)}(p_{\mu(i)}) \geq 0$ for all bidders i, $p_{j_0} = 0$ for the dummy item j_0, and $p_j \geq r_j$ for all items j. An *envy-free* outcome (μ, p) is a feasible outcome for which $u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,j}(p_j)$ holds for all $(i, j) \in I \times J$. A *bidder-optimal* outcome (μ, p) is an envy-free outcome with $u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,\mu'(i)}(p'_{\mu'(i)})$ for all bidders i and all envy-free outcomes (μ', p'). A mechanism is a *bo-mechanism* if it computes a bidder-optimal outcome on every input.

Consider for a fixed bidder i the two inputs (v', m', r) and (v'', m'', r) with $v'_{i,j} = v_{i,j}, m'_{i,j} = m_{i,j}$ for bidder i and $v'_{i',j} = v''_{i',j}, m'_{i',j} = m''_{i',j}$ for all $i' \neq i$ and $j \in J$. Let (μ', p') be the outcome for (v', m', r) and (μ'', p'') the outcome for (v'', m'', r). Let $u_{i,\mu'(i)}(p'_{\mu'(i)})$ be denoted by u'_i and $u_{i,\mu''(i)}(p''_{\mu''(i)})$ by u''_i. Then a mechanism is *incentive-compatible* if $u'_i \geq u''_i$ for every fixed bidder i and every two inputs (v', m', r) and (v'', m'', r). If (v'', m'', r) is such that the utility u''_i of bidder i is maximized among all possible inputs (v'', m'', r), then $u''_i - u'_i \geq 0$ is the *regret* of bidder i. That is, a mechanism is incentive-compatible if and only if the regret of each bidder is zero. The mechanism is *incentive-compatible in expectation* if bidder i's expected utility for the input (v', m', r) is at least as large as for (v'', m'', r), i.e., $\mathrm{E}[u'_i] \geq \mathrm{E}[u''_i]$ for all possible inputs (v'', m'', r).

The *first choice graph* $G_p = (I \cup J, E_p)$ at feasible prices p has a node for each bidder and each item and an edge between a bidder i and an item j if bidder i achieves his highest possible utility at prices p when he is matched to j, i.e., $u_{i,j}(p_j) \geq u_{i,j'}(p_{j'})$ for all $j' \in J$. We define the first choice function for a bidder i as $F_p(i) = \{j : \exists(i, j) \in E_p\}$ and for a set of bidders $T \subseteq I$ as $F_p(T) = \cup_{i \in T} F_p(i)$. We define the inverse first choice function for an item j as $F_p^{-1}(j) = \{i : \exists(i, j) \in E_p\}$, and for a set of items $S \subseteq J$ as $F_p^{-1}(S) = \cup_{j \in S} F_p^{-1}(j)$. Note that due to the dummy item, $\max_{j \in J} u_{i,j}(p_j) \geq 0$ holds for all bidders i and all feasible prices. The outcome (μ, p) is envy-free if and only if the matching μ only consists of edges in E_p.

We will compare outcomes and price increases for different inputs several times. To this end, we use matching superscripts to denote the utility functions, the outcome, the first choice graph, and the (inverse) first choice function corresponding to a specific input.

An *alternating path* is a sequence of edges in G_p in which matched and unmatched edges alternate and in which only the last item is allowed to be the dummy item. An *alternating tree* in G_p is a tree rooted at an unmatched bidder i_0 in which all paths from the root to the leaves are alternating paths and every leaf is either an unmatched item, the dummy item, or a bidder whose first choice items are all contained in the tree. We will consider *maximal alternating trees*, i.e., alternating trees that cannot be extended. If a bidder is contained in a maximal alternating tree, then so are all his first choice items.

1.2. Modified Hungarian Method

The Modified Hungarian Method (MHM) was introduced in [8]. The mechanism, restated in Algorithm 1, is initialized with an empty matching and all prices set to their reserve prices (line 1). When a yet unmatched bidder i_0 is considered for matching, first a maximal alternating tree in G_p rooted at bidder i_0 is found (line 3). If this tree contains an alternating path that ends with an unmatched item or with the dummy item, the matching is augmented along this path, i.e., the matched edges along the path become unmatched and vice versa (line 16). Otherwise, the prices of the items in the tree are increased by the smallest possible amount such that either (a) another item becomes a first choice for a bidder in the tree (i.e. $\delta = \delta_{\text{out}}$, line 7) or (b) an item stops being a first choice for a bidder in the tree because it reaches its maximum price ($\delta = \delta_{\max}$, line 8). After a price increase (line 10) the matching μ and the first choice graph are updated. A new maximal alternating tree in G_p rooted at bidder i_0 is found and u_i, the best possible utility for a bidder i under the current prices, is calculated for all bidders in the tree. The prices continue to be increased in the same manner until there is an alternating path in G_p ending with an unmatched item or the dummy item. By augmenting the matching along this path, bidder i_0 is matched and the next unmatched bidder is considered. Note that the mechanism never increases the price of the dummy item.

Algorithm 1: Modified Hungarian Method (MHM)

input : valuations v, maximum prices m, reserve prices r
output : matching μ, prices p

1 $p_j \leftarrow r_j$ for all $j \in J$, $\mu \leftarrow \emptyset$
2 **while** *some bidder* i_0 *is unmatched* **do**
3   construct a maximal alternating tree $\mathcal{T}$ rooted at bidder i_0 in G_p
4   let S and T be the sets of items and bidders in $\mathcal{T}$
5   compute utilities $u_i \leftarrow \max_j u_{i,j}(p_j)$ for all $i \in T$
6   **while** $j_0 \notin S$ *and all items* $j \in S$ *matched* **do**
7         $\delta_{\text{out}} \leftarrow \min_{i \in T, j \in J \setminus S}(u_i - v_{i,j} + p_j)$
8         $\delta_{\max} \leftarrow \min_{i \in T, j \in F_p(i)}(m_{i,j} - p_j)$
9         $\delta \leftarrow \min(\delta_{\text{out}}, \delta_{\max})$
10   update prices $p_j \leftarrow p_j + \delta$ for all $j \in S$
11   update first choice graph G_p
12   update matching $\mu \leftarrow \mu \cap E_p$
13   construct $\mathcal{T}$ rooted at bidder i_0 in G_p
14   let S and T be the sets of items and bidders in $\mathcal{T}$
15   update utilities $u_i \leftarrow \max_j u_{i,j}(p_j)$ for all $i \in T$
16   augment μ along alternating path from i_0 to an unmatched item or j_0, preferably to items j with $p_j > r_j$
17 **return** (μ, p)

The following theorem by [8] shows that the MHM is a polynomial time bo-mechanism. We will additionally use some intermediate results from [8], which are restated below.

Theorem 1 (Dütting et al. [8]). *The outcome of the MHM is feasible, envy-free, and bidder-optimal for the given input. It can be computed in time* $O(n \cdot k^3)$.[2]

[2] In sponsored search the number of items k is usually constant, while the number of bidders n is large.

A set of items $S \subseteq J$ is *strictly overdemanded* in a first choice graph G_p with respect to a set of bidders $T \subseteq I$ if the dummy item j_0 is not in S, all first choice items of bidders in T are contained in S, i.e., $F_p(T) \subseteq S$, and for all non-empty subsets R of S it holds that $|F_p^{-1}(R) \cap T| > |R|$. We say that an item is overdemand in G_p if it is contained in a strictly overdemanded set of items S for some set of bidders T.

Lemma 1 (Dütting et al. [8])**.** *Let $S \subseteq J$ be a set of items that is strictly overdemanded in G_p regarding a set of bidders $T \subseteq I$. Let the price increase δ be equal to $\min(\delta_{out}, \delta_{max})$ with $\delta_{out} := \min_{i \in T, j \in J \setminus S}(u_i - v_{i,j} + p_j)$ and $\delta_{max} := \min_{i \in T, j \in F_p(i)}(m_{i,j} - p_j)$. Then for all envy-free outcomes (μ', p') with $p'_j \geq p_j$ for all j it holds that $p'_j \geq p_j + \delta$ for all items $j \in S$.*

Lemma 2 (Dütting et al. [8])**.** *Let (μ, p) be the outcome of the MHM, and let (μ', p') be any envy-free outcome for the same input. Then $p_j \leq p'_j$ for all items j.*

Lemma 3 (Dütting et al. [8])**.** *Let (μ, p) be an envy-free outcome with $p_j \leq p'_j$ for all items j and all envy-free outcomes (μ', p'). Then (μ, p) is a bidder-optimal outcome.*

1.3. Refined Definition of General Position

The following definition of general position allows more inputs than the definition in [1, 7] and is sufficient for the rematch condition (Lemma 4). We define general position on a bipartite directed multigraph called input graph as follows. The *input graph* for a given input has a node for each bidder and each item and for each bidder-item pair (i, j) a forward edge from i to j with weight $-v_{i,j}$, a backward edge from j to i with weight $v_{i,j}$, and a discontinuity edge from i to j with weight $m_{i,j} - v_{i,j}$. The weight $\omega(P)$ of a walk P in the input graph is the sum of the weights of its edges. Consider all simple paths between two bidders in the input graph that alternate between forward and backward edges and all possibilities to add a discontinuity edge from the last bidder to the last item or to an item not yet in the path. Let $P = (i_1, j_1, i_2, \ldots, j_{s-1}, i_s, j_s)$ and $Q = (i'_1, j'_1, i'_2, \ldots, j'_{t-1}, i'_t, j'_t)$ be two such paths. We call an input in *general position* if for *no* path P with $j_{s-1} = j_s$ there exits a path Q with the same weight and the same starting node as P, i.e., $w(Q) = w(P)$ and $i'_1 = i_1$, and either (a) the discontinuity edge of Q leads back to the previous item, i.e., $j'_{t-1} = j'_t$, and the two paths separate at some bidder $l \in [1, \min(s, t) - 1]$, i.e., $j'_{k-1} = j_{k-1}$ and $i'_k = i_k$ for $2 \leq k \leq l$ and $i'_k \notin P$, $j'_{k-1} \notin P$, $i_k \notin Q$, and $j_{k-1} \notin Q$ for $k > l$, or (b) the discontinuity edge of Q leads to an item not in the path before, i.e., $j'_u \neq j'_t$ for $1 \leq u < t$ and Q is a subpath of P, i.e., $i'_k = i_k$ and $j'_k = j_k$ for $1 \leq k \leq t < s$.

The following example shows that the new definition of general position is less restrictive than the definition in [7].

Example 1. The input consists of three bidders and two items and valuations and maximum prices as follows: $v_{i_1,j_1} = 10$, $v_{i_2,j_2} = 10$, $v_{i_3,j_1} = 10$, $v_{i_3,j_2} = 10$; $m_{i_1,j_1} = 1$, $m_{i_2,j_2} = 2$, $m_{i_3,j_1} = 2$, $m_{i_3,j_2} = 1$. The remaining input values are zero. Let i_1 and i_2 be considered first by the MHM. Then i_1 becomes matched to j_1 and i_2 to j_2. Consider the situation when the MHM tries to match i_3, as shown in the left graph of Figure 2. The items j_1 and j_2 are overdemanded and matched, thus their prices are increased by $\delta_{\max} = m_{i_1,j_1} - p_{j_1} = m_{i_3,j_2} - p_{j_2}$. Item j_1 becomes unmatched but is matched in the next iteration of the algorithm to i_3. The MHM is truthful for this input although it is not in general position as defined in [7] because the two walks $P = (i_3, j_1, i_1, j_1)$ and $Q = (i_3, j_2)$ in the input graph that start with the same bidder, alternate between forward and backward edges, and end with different discontinuity edges have the same weight, $\omega(P) = \omega(Q) = m_{i_1,j_1} - v_{i_1,j_1} = m_{i_3,j_2} - v_{i_3,j_2}$. Since Q is not a subpath of P, the input is in general position following our new definition.

2. Incentive-Compatibility for Non-Degenerate Inputs

The MHM is incentive-compatible if for an input all maximum prices are infinity, i.e., for continuous quasi-linear utility functions. A crucial property in this case is that once an item is matched during the mechanism, it never becomes unmatched. This does not hold for budgets as during the same price increase the maximum prices of many items might be reached, and many might become unmatched. If only one

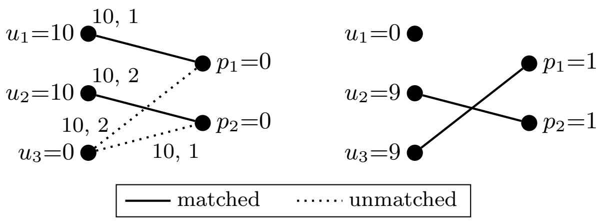

Figure 2: An input not in general position according to former definitions. The MHM is incentive-compatible for this input. In each graph the points on the left side represent the bidders, the ones on the right side the items. The utility bidder i obtains in the current matching is denoted by u_i. The left graph is the first choice graph after the first two bidders were matched and before the third bidder is matched. The edges are labeled with the valuation and the maximum price a bidder reports for the item. The right graph shows the outcome computed by the MHM.

item becomes unmatched during a price increase, then the requirement to preferably match an item j with $p_j > r_j$ (i.e. a previously matched item) guarantees that the unmatched item will be matched in the next iteration as the item is an unmatched leaf at the current maximal alternating tree. The notion of general position of [1, 7] (a refined definition is given in Section 1.3) requires that during *no* price increase two edges of the *first choice graph* G_p disappear because their maximum prices are reached. Of course, if this condition is fulfilled, then it will also never happen that two edges become *unmatched* during the same price increase. Restricting to inputs in general position thus basically avoids the main difficulty introduced by maximum prices. However, even for this purpose, the restriction on the input is stronger than necessary. For example, it could be that two unmatched edges of G_p disappear during the same price increase. This would not unmatch any items and thus not affect the truthfulness of the mechanism. Even if two matched edges of G_p disappear during the same price increase, both could be rematched without another price increase in subsequent iterations if they both lie on the same path to the root of the maximal alternating tree. Our *rematch condition* captures exactly this intuition.

Formally, we say that two nodes x and y are *on the same path* in a maximal alternating tree if x lies on the path from y to the root or vice versa. Otherwise, the two nodes are *on different paths.* A price update is called *problematic*, if (a) two or more items on different paths become unmatched or if (b) one item becomes unmatched and a maximum price for an unmatched edge on the path from this item to the root is reached. Otherwise, a price update is called *unproblematic.* We say that an input fulfills the *rematch condition* if no problematic price update occurs during the execution of the MHM.

The following proof refers to the definition of general position in Section 1.3. It can easily be adapted to the definition in [7].

Lemma 4. *Inputs in general position satisfy the rematch condition.*

Proof. Assume for contradiction that for an input in general position a problematic price update occurs. Let the two maximum prices that are reached during the price increase δ and that cause the problematic price update be m_{i_s,j_s} and $m_{i'_t,j'_t}$. Consider the maximal alternating tree $\mathcal{T}$ with root i_0 from which δ was determined. Since the edges (i_s, j_s) and (i'_t, j'_t) are contained in $\mathcal{T}$, there exists a simple path $\tilde{P}$ from i_0 to i_s and a simple path $\tilde{Q}$ from i_0 to i'_t in $\mathcal{T}$. Let $P = (i_0, j_1, i_2, \ldots, j_{s-1}, i_s, j_s)$ and $Q = (i_0, j'_1, i'_2, \ldots, j'_{t-1}, i'_t, j'_t)$ be the two corresponding paths in the input graph that are constructed by using alternately forward and backward edges along $\tilde{P}$ resp. $\tilde{Q}$ and adding the discontinuity edge for m_{i_s,j_s} resp. $m_{i'_t,j'_t}$.

Let p be the prices before the price increase. Since m_{i_s,j_s} and $m_{i'_t,j'_t}$ are reached during the same price increase δ, we know that $\delta = m_{i_s,j_s} - p_{j_s} = m_{i'_t,j'_t} - p_{j'_t}$, i.e.,

$$m_{i_s,j_s} - m_{i'_t,j'_t} = p_{j_s} - p_{j'_t} \ . \tag{3}$$

Since $\mathcal{T}$ is contained in the first choice graph G_p at prices p and P and Q are constructed from edges in $\mathcal{T}$, $v_{i_u,j_{u-1}} - p_{j_{u-1}} = v_{i_u,j_u} - p_{j_u}$ holds for every bidder i_u, $u \geq 2$, on P and Q. Reformulated, this gives

$$p_{j_u} = v_{i_u,j_u} - v_{i_u,j_{u-1}} + p_{j_{u-1}} \ . \tag{4}$$

Combining (3) and (4) yields

$$
\begin{aligned}
m_{i_s,j_s} - m_{i'_t,j'_t} = &\sum_{u=2}^{s} \left(v_{i_u,j_u} - v_{i_u,j_{u-1}}\right) + p_{j_1} \\
&+ \sum_{w=2}^{t} \left(v_{i'_w,j'_{w-1}} - v_{i'_w,j'_w}\right) - p_{j'_1} \; . \qquad (5)
\end{aligned}
$$

Additionally, we know for the root i_0 that $v_{i_0,j_1} - p_{j_1} = v_{i_0,j'_1} - p_{j'_1}$. Thus, $p_{j_1} = v_{i_0,j_1} - v_{i_0,j'_1} + p_{j'_1}$ and we can eliminate the prices from (5). Reformulation leads to (6) where the left hand side is equal to $\omega(P)$ and the right hand side to $\omega(Q)$, i.e.,

$$
\begin{aligned}
&- v_{i_0,j_1} + \sum_{u=2}^{s} \left(v_{i_u,j_{u-1}} - v_{i_u,j_u}\right) + m_{i_s,j_s} \\
&= -v_{i_0,j'_1} + \sum_{w=2}^{t} \left(v_{i'_w,j'_{w-1}} - v_{i'_w,j'_w}\right) + m_{i'_t,j'_t} \; . \qquad (6)
\end{aligned}
$$

Since a problematic price update occurred, either (a) both (i_s, j_s) and (i'_t, j'_t) were matched and on different paths or (b), w.l.o.g, (i_s, j_s) was matched and (i'_t, j'_t) was unmatched and on the path from j_s to i_0. In case (a), $j_{s-1} = j_s$ and $j'_{t-1} = j'_t$ by the construction of P and Q. "On different paths" is equivalent to $j'_{k-1} = j_{k-1}$ and $i'_k = i_k$ for some integer $l \in [1, \min(s,t) - 1]$ and $2 \leq k \leq l$ and $i'_k \notin P$, $j'_{k-1} \notin P$, $i_k \notin Q$, and $j_{k-1} \notin Q$ for $k > l$, i.e., the two paths could start with the same nodes but deviate at some bidder. In case (b), again $j_{s-1} = j_s$ but $j'_{t-1} \neq j'_t$. Since (i'_t, j'_t) is unmatched and on the path from j_s to i_0, $i'_k = i_k$ and $j'_k = j_k$ for $1 \leq k \leq t < s$. In both cases the equality between $\omega(P)$ and $\omega(Q)$ gives a contradiction to the assumption that the input is in general position. □

We say that an unmatched item j with $p_j > r_j$ is *rematched* if there exists a maximal alternating tree rooted at some unmatched bidder i that contains item j and no other unmatched item j' with $p_{j'} > r_{j'}$. That is, when bidder i is considered by the MHM, item j will be matched without an additional price increase. Thus, j will not remain unmatched.

Lemmata 5 and 6 are analogous to the results for inputs in general position in [7]. Note that Lemma 6 implies that the outcome is a competitive equilibrium.

Lemma 5. *If an item becomes unmatched during an unproblematic price update, it will be rematched.*

Proof. An item j that is matched to a bidder i at some point in the MHM can only become unmatched if $m_{i,j}$ is reached due to a price increase $\delta = \delta_{\max}$. Since the maximal alternating tree $\mathcal{T}$ roots at an unmatched bidder i_0 and each path from the root to a leaf alternates between matched and unmatched edges, every item is nearer to i_0 than the bidder it is matched to. Since the price update is unproblematic, we know that no item on a different path became unmatched and for no unmatched edge on the path from j to i_0 in $\mathcal{T}$ a maximum price is reached. If additionally no matched item on the path from j to i_0 became unmatched, then there exists a path from i_0 to the newly unmatched item j in the maximal alternating tree in the next iteration of the mechanism. Otherwise, some matched item j' on the path from j to i_0 became unmatched. Let i' be the bidder it was matched to. Then there exists a maximal alternating tree rooted at bidder i' that contains the unmatched item j. In both cases j will be rematched. □

Lemma 6. *Let p be the price vector at some step in the MHM. If there are only unproblematic price updates in the mechanism, then every item j with $p_j > r_j$ is matched or rematched during the whole algorithm.*

Proof. At the beginning of the MHM $p_j = r_j$ for all items $j \in J$. Let T and S be the sets of bidders and items in the maximal alternating tree in G_p before some price increase δ. The mechanism only raises prices for items in S. All items in S are already matched because otherwise the mechanism would not increase the prices. Thus, every item in S with a price above its reserve price is matched before the price increase. Since we assume only unproblematic price updates, we know from Lemma 5 that if an item that was matched becomes unmatched, it will be rematched. Thus, every item with a price above its reserve price will be matched or rematched during the whole algorithm. □

Dütting et al. showed in [3] for more general utility functions that a bo-mechanism is incentive-compatible under certain conditions. We restate their result using our terminology and restricted to our setting in Theorem 2. Note that in every envy-free outcome (μ, p) for each matched bidder-item pair $(i, j) \in \mu$ the price of the item j has to be at least $\max_{i' \neq i} u^{-1}_{i',j}(u_{i',\mu(i')}(p_{\mu(i')})))$. Theorem 2 basically states that, for incentive-compatibility of bo-mechanisms for a given input, the bidder-optimal outcome has to be a competitive equilibrium and that in each set $\hat{J}$ of matched items at least one item must be sold at the minimum envy-free price induced by the bidders outside $\mu^{-1}(\hat{J})$. We can show that for any input that fulfills the rematch condition this condition is satisfied.

Theorem 2 (Dütting et al. [3])**.** *Let (μ, p) be an outcome for the input (v, m, r) of a bo-mechanism. Assume that in (μ, p) all items j with $p_j > r_j$ are matched and that for every subset of matched bidders $\hat{I} \subseteq I$ at least one item $j \in \mu(\hat{I})$ has $p_j = \max(r_j, \max_{i \notin \hat{I}} u^{-1}_{i,j}(u_{i,\mu(i)}(p_{\mu(i)})))$. Then every bo-mechanism is incentive-compatible for the input (v, m, r).*

Lemma 7. *Let (v, m, r) be some input for which every item with a price strictly above its reserve price remains matched or is rematched during the execution of the MHM, and let (μ, p) be the outcome of the MHM for this input. Then for every subset of matched bidders $\hat{I} \subseteq I$ at least one item $j \in \mu(\hat{I})$ has $p_j = \max(r_j, \max_{i \notin \hat{I}} u^{-1}_{i,j}(u_{i,\mu(i)}(p_{\mu(i)})))$.*

Proof. Let $\max(r_j, \max_{i \notin \hat{I}} u^{-1}_{i,j}(u_{i,\mu(i)}(p_{\mu(i)})))$ be denoted by $\hat{r_j}$. Assume for contradiction that for some subset of matched bidders $\hat{I} \subseteq I$ it holds that $p_j > \hat{r_j}$ for all $j \in \hat{J} = \mu(\hat{I})$. Consider the iteration in which the last bidder $i \in \hat{I}$ becomes matched to an item $j \in \hat{J}$. Let $p^{(s)}$ be the prices immediately before the matching is augmented to match i to j. We will show that $p^{(s)}_{j'} \leq \hat{r_{j'}}$ for at least one item $j' \in \hat{J}$. We will consider two cases.

Case 1: If there is an unmatched item $j' \in \hat{J}$ before i is matched, then by Lemma 6 either $p^{(s)}_{j'} = r_{j'}$ or j' is rematched in this iteration. In the latter case, j' cannot have been matched to a bidder in $\hat{I}$ because otherwise this bidder would be unmatched and this would not be the last iteration in which a bidder in $\hat{I}$ becomes matched to an item in $\hat{J}$. Thus, j' was matched to some bidder $i' \in I \setminus \hat{I}$. By the definition of $\hat{r_{j'}}$, at most a price of $p_{j'} \leq \hat{r_{j'}}$ is necessary to avoid envy from i' for j'. Since by Lemma 2 the MHM computes minimum envy-free prices, $\hat{r_{j'}} \geq m_{i',j'} = p^{(s)}_{j'}$.

Case 2: If there is no unmatched item in $\hat{J}$, then some item in $j' \in \hat{J}$ is matched to a bidder $i' \in I \setminus \hat{I}$. To be able to match j' to a bidder in $\hat{I}$ in this iteration of the mechanism, bidder i' may not have envy for j' at $p^{(s)}_{j'}$. By the definition of $\hat{r_{j'}}$, we again have $\hat{r_{j'}} \geq p^{(s)}_{j'}$.

In both cases, $p^{(s)}_{j'} \leq \hat{r_{j'}}$ for some item $j' \in \hat{J}$ after bidder i is matched to j. Since after the iteration all bidders in $\hat{I}$ remain matched, no price increase for item j' can be caused by them.

Assume for contradiction that some yet unmatched bidder $i' \in I \setminus \hat{I}$ causes a price increase δ that increases the price of j' strictly above $\hat{r_{j'}}$ and that afterwards all items $j \in \hat{J}$ have a price strictly above $\hat{r_j}$. We denote the prices before the price increase with $p^{(t-1)}$ and the prices after the price increase with $p^{(t)}$. Let $\mathcal{T}$ be the maximal alternating tree from which δ was determined and T the set of bidders in $\mathcal{T}$. Since j' is in $\mathcal{T}$ and all items in $\hat{J}$ are matched to bidders in $\hat{I}$, there is a first choice edge between some bidder $i'' \in I \setminus \hat{I}$ and some item $j'' \in \hat{J}$ before the price increase, but by the definition of $\hat{r_{j''}}$ and $p^{(t)}_{j''} > \hat{r_{j''}}$ there is no such edge after the price increase. Thus, the maximum price $m_{i'',j''}$ was reached during the price increase $\delta = \delta_{\max}$ and $p^{(t)}_{j''} = m_{i'',j''} > \hat{r_{j''}}$. Since $\hat{r_{j''}} \geq u^{-1}_{i'',j''}(u_{i'',\mu(i'')}(p_{\mu(i'')}))$, it holds that $m_{i'',j''} > u^{-1}_{i'',j''}(u_{i'',\mu(i'')}(p_{\mu(i'')})) \geq u^{-1}_{i'',j''}(u_{i'',\mu(i'')}(p^{(t)}_{\mu(i'')}))$. This implies by the definition of $u^{-1}_{i'',j''}$ that $v_{i'',\mu(i'')} - p^{(t)}_{\mu(i'')} > v_{i'',j''} - m_{i'',j''}$. If $\mu(i'') \in F_{p^{(t-1)}}(T)$, then $v_{i'',\mu(i'')} - p^{(t-1)}_{\mu(i'')} - \delta > v_{i'',j''} - p^{(t-1)}_{j''} - \delta$, which gives a contradiction to $j'' \in F_{p^{(t-1)}}(i)$. If $\mu(i'') \notin F_{p^{(t-1)}}(T)$, then $v_{i'',\mu(i'')} - p^{(t-1)}_{\mu(i'')} > v_{i'',j''} - p^{(t-1)}_{j''} - \delta$, which gives a contradiction to $\delta_{\max} \leq \delta_{\text{out}}$. Hence, no such price increase can exist. □

Corollary 8. *The MHM is incentive-compatible for inputs that fulfill the rematch condition.*

3. Calculating the Regret

The rematch condition still does not take into account whether a bidder can actually *profit* from lying. With Theorem 3 we provide the first necessary and sufficient condition for a bo-mechanism to be incentive-compatible for a given input. Additionally, we give an algorithm, Algorithm Regret, that calculates the regret of each bidder.

Theorem 3, below, turns the intuition of how a bidder can lie into an easy-to-follow recipe. The intuition can be described as follows. As we have seen in Section 2, a bidder can only profit from lying if a problematic price update occurs in case he reports truthfully. Lying is helpful if it avoids the problematic price update such that another bidder involved in the problematic price update is matched at lower than envy-free prices. This enables the lying bidder to profit from the reduced demand on the remaining items. A bidder can avoid a problematic price update by reporting lower values than his true ones for the affected item. In Theorem 3 we show that a bidder can simply report maximum prices of zero for all items except the one he wants to be matched to. Intuitively, this does not decrease his utility since these items are either less desirable to him than the item he wants to be matched to or he cannot obtain them anyway due to high prices caused by the demand of the other bidders. As a result, he avoids all problematic price updates he can influence.

Let i be some fixed bidder and let (v, m, r) and (v', m', r) be two inputs with $v_{i',j} = v'_{i',j}$ and $m_{i',j} = m'_{i',j}$ for all bidders $i' \neq i$ and all $j \in J$. Let (μ, p) and (μ', p') be the corresponding outcomes of the MHM and let $u_i = u_{i,\mu(i)}(p_{\mu(i)})$ and $u'_i = u_{i,\mu'(i)}(p'_{\mu'(i)})$.

Theorem 3. *If for some fixed bidder i the utility u'_i bidder i obtains in (μ', p') is strictly higher than u_i, i.e., bidder i has a positive regret and thus no bo-mechanism is incentive-compatible for the input (v, m, r), then bidder i obtains at least the utility u'_i if he reports maximum prices of zero for all items $j' \neq \mu'(i)$ and otherwise reports the same values as in (v, m, r).*

If some fixed bidder i obtains a strictly higher utility than u_i if he reports maximum prices of zero for all items $j' \neq j$ for some item j and all other input values are the same as in the input (v, m, r), then no bo-mechanism is incentive-compatible for the input (v, m, r).

Proof. The second part of the theorem follows from the definition of incentive-compatibility. For the first part assume that some fixed bidder i obtains a strictly higher utility $u'_i > u_i$ in (μ', p'). Let (v', m', r) be such that u'_i is maximized among all possible inputs (v', m', r). Let (v, m'', r) be the input where bidder i reports for all $j \neq \mu'(i)$ a maximum price of $m''_{i,j} = 0$ but $m''_{i,\mu'(i)} = m_{i,\mu'(i)}$ for $\mu'(i)$. We have $u''_{i,\mu'(i)}(p_{\mu'(i)}) = u_{i,\mu'(i)}(p_{\mu'(i)})$ for all $p_{\mu'(i)}$ and $u''_{i,j}(p_j) \leq 0$ for all $j \neq \mu'(i)$ and $p_j \geq 0$. For all other bidders $i' \neq i$ and all items j we have $m''_{i',j} = m_{i',j}$ and thus $u''_{i',j}(p_j) = u_{i',j}(p_j)$ for all p_j.

Let (μ'', p'') be the outcome of the MHM for (v, m'', r). We start by showing $p''_j \leq p'_j$ for all items j. To this end, it is sufficient to show that (μ', p') is envy-free for the input (v, m'', r). Since $u''_{i',j}(p'_j) = u_{i',j}(p'_j)$ for all bidders $i' \neq i$, it remains to show that (μ', p') is envy-free for i and (v, m'', r). As i profits from reporting his values as in the input (v', m', r), we have $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu(i)}(p_{\mu(i)}) \geq 0$. Hence, with $u''_{i,j}(p'_j) \leq 0$ we have that $u_{i,\mu'(i)}(p'_{\mu'(i)}) = u''_{i,\mu'(i)}(p'_{\mu'(i)}) > u''_{i,j}(p'_j)$ holds for all $j \neq \mu'(i)$, i.e., (μ', p') is envy-free for i regarding the input (v, m'', r).

By $p''_{\mu'(i)} \leq p'_{\mu'(i)}$ we have $u''_{i,\mu'(i)}(p''_{\mu'(i)}) \geq u''_{i,\mu'(i)}(p'_{\mu'(i)}) > 0 \geq u''_{i,j}(p''_j)$ for all $j \neq \mu'(i)$. Thus, the envy-freeness of (μ'', p'') for (v, m'', r) implies that bidder i is matched to $\mu'(i)$ in μ'', i.e., we have $\mu''(i) = \mu'(i)$. Hence, the utility $u_{i,\mu''(i)}(p''_{\mu''(i)})$ bidder i obtains with input (v, m'', r) is at least as high as the utility $u'_i = u_{i,\mu'(i)}(p'_{\mu'(i)})$ bidder i obtains by reporting according to (v', m', r). □

It follows that the regret of a bidder is $u'_i - u_i$, where u'_i is the utility bidder i receives if he lies as described in Theorem 3. This implies the correctness of the following algorithm. Theorem 4 follows from Algorithm 2 and the equivalence of incentive-compatibility with a regret of zero for all bidders.

Algorithm 2: Regret

$(\mu, p) \leftarrow$ `MHM`(v, m, r), $u_i \leftarrow u_{i,\mu(i)}(p_{\mu(i)})$, $\mathsf{regret}_i \leftarrow 0$ for all $i \in I$
foreach $(i, j) \in I \times J$ **do**
  $m' \leftarrow m$
  $m'_{i,j'} \leftarrow 0$ for all $j' \neq j$
  compute $(\mu', p') \leftarrow$ `MHM`(v, m', r) and set $u'_i \leftarrow u_{i,\mu'(i)}(p'_{\mu'(i)})$
  $\mathsf{regret}_i \leftarrow \max(\mathsf{regret}_i, u'_i - u_i)$

Theorem 4. *It can be decided in polynomial time whether a bo-mechanism is incentive-compatible for a given input.*

Proof. Algorithm Regret executes the MHM for $n \cdot k + 1$ inputs, each with n bidders and k items. By Theorem 1 this can be done in time $O(n^2 \cdot k^4)$. By definition the MHM is incentive-compatible for a given input if and only if the regret of all bidders is zero. Algorithm Regret calculates the regret of all bidders. The correctness of the algorithm follows from Theorem 3. □

Theorem 5 shows that in order for bidder i to profit from lying there must exist an item j such that bidder i would prefer j at the current prices. Thus bidder i has envy and must have underreported either $v_{i,j}$ or $m_{i,j}$.

Theorem 5. *If the fixed bidder i obtains a higher utility $u'_i > u_i$ in (μ', p'), then bidder i has envy with respect to (v, m, r) in the outcome (μ', p').*

Proof. By the definition of (v', m', r) we have for the fixed bidder i that $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu(i)}(p_{\mu(i)})$. By the envy-freeness of (μ, p) this implies $p'_{\mu'(i)} < p_{\mu'(i)}$. As by Lemma 2 the prices p are the minimum envy-free prices for input (v, m, r), some bidder has envy at prices p'. However, (μ', p') is envy-free with respect to (v', m', r). The two inputs differ only in the values bidder i reported. Thus it has to be bidder i who has envy in the outcome (μ', p'). □

4. Randomization

In this section we describe a mechanism that is truthful in expectation for *all* inputs and achieves for each bidder at least his utility in a bidder-optimal outcome. We first randomize the input such that it is in general position. This requires the assumption that all input values are multiples of some constant $\alpha > 0$, e.g., are integers. Second, we show that the input randomization implies incentive-compatibility of the MHM only with regard to the randomized input. We give an example how a bidder can profit from lying despite the input randomization. Specifically, he has to overreport a maximum price. We show that this is the only way to profit from lying in this case. Using randomized extraction [10] can prevent this type of lying, which yields truthfulness in expectation.

We further prove that the expected utility each bidder obtains is at least the utility in a bidder-optimal outcome as well as at least the utility a bidder could achieve by lying in a bo-mechanism. Finally, we show how a simple rounding of the computed prices can, and in the case of an input in general position will, lead to a bidder-optimal outcome. That is, for inputs in general position our randomized mechanism (in expectation) approximates a bidder-optimal outcome up to an additive factor of α in the prices, and can thus be seen as an extension of the previously known results for the MHM.

As argued in the introduction, envy cannot be avoided in any mechanism that is incentive-compatible (in expectation). As shown in Section 3, a bidder has envy in an outcome in which he profited from lying when a bo-mechanism is used. Any truthful (in expectation) mechanism has to achieve for each bidder at least the same utility as for these outcomes and thus cannot always compute an envy-free outcome. In other words, a truthful (in expectation) mechanism has to somehow break the ties that lead to problematic price updates as defined in Section 2, which will introduce envy in some cases.

4.1. Randomized Input.

We assume that all valuations, maximum prices, and reserve prices in the input (v, m, r) are multiples of some $\alpha > 0$. We apply the following randomization to all maximum prices.

1. Generate $n \cdot k$ random numbers $\Delta_{i,j}$, $1 \le i \le n$, $1 \le j \le k$, uniformly distributed in the interval $(0, \alpha)$.
2. Set $m^{\mathrm{R}}_{i,j} := m_{i,j} - \Delta_{i,j}$ for all $(i, j) \in I \times J$.

We call an input that fulfills our assumption and was randomized in the above way a *randomized input* and denote it with (v, m^{R}, r). Throughout the section, we denote with $(\mu^{\mathrm{R}}, p^{\mathrm{R}})$ the outcome of the MHM for the randomized input, with (μ, p) the outcome of the MHM for the unmodified input, and with (μ', p') any envy-free outcome for the unmodified input.

By reducing the maximum prices, we clearly give up the property of envy-freeness. If the price of an item is increased up to the reduced maximum price of a bidder, this bidder still demands this item regarding his true maximum price. However, the algorithm treats the bidder as if his utility for the item is $-\infty$ and thus will match this item to a different bidder.

Since $\{\Delta_{i,j}\}_{i\in I, j\in J}$ is a negligible set in the interval $(0, \alpha)$, almost surely $\Delta_{i,j} \neq \Delta_{i',j'}$ for any two different bidder-item pairs (i, j) and (i', j'). For a definition of a *negligible set* and *almost surely* see, e.g., [11, pp. 8, p. 54].

Lemma 9. *A randomized input is (almost surely) in general position.*

Proof. Assume that in the input graph for the randomized input two walks $P = (i_a, j_b, i_c, \ldots, i_d, j_e)$ and $Q = (i_a, j_t, i_u, \ldots, i_v, j_w)$ that start with the same bidder, alternate between forward and backward edges, and end with different discontinuity edges have the same weight,

$$-v_{a,b} + v_{b,c} - \cdots + m^{\mathrm{R}}_{d,e} - v_{d,e} = -v_{a,t} + v_{t,u} - \cdots + m^{\mathrm{R}}_{v,w} - v_{v,w} \ . \tag{7}$$

By assumption, all valuations as well as the original unmodified maximum prices are multiples of some $\alpha > 0$ whereas the randomized maximum prices are not multiples of α. Thus, Equation (7) implies

$$\Delta_{d,e} = \Delta_{v,w} \ , \tag{8}$$

which is almost surely not the case. Hence, the input is almost surely in general position. □

Corollary 10. *The MHM is (almost surely) incentive-compatible with respect to every randomized input (v, m^R, r).*

4.2. Counter Example.

The following example, however, shows that the MHM with input randomization is not incentive-compatible with respect to the original input. For some inputs a bidder can improve his utility by reporting a higher maximum price than his true one if he knows that input randomization will be applied. Note that we cannot a-priori bound the regret of this bidder. Lemma 11 shows that this is the only way a bidder can profit from lying.

Example 2. There are three bidders and two items. Each of the bidders likes both items equally and prefers them to being unmatched. All maximum prices are the same and larger than zero. When the input is randomized and the MHM is executed for the randomized input, the prices of the two items are increased up to the lowest of the randomized maximum prices and the two bidders with higher randomized maximum prices are matched to the two items. If one bidder claims he has a higher maximum price, he will always be matched to one of the items. Nevertheless, the prices of the items will be lower than his maximum prices. An instance with concrete values is given in Figure 3.

The example can be adapted to the definitions and the algorithm in [1] to show how a bidder can profit from lying if the lexicographical tie breaking rule proposed in [1] is used instead of a randomization of the input.

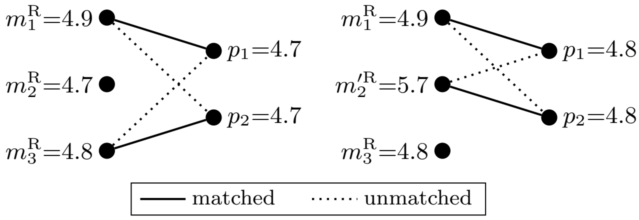

Figure 3: Lying when the input is randomized. All bidders have the same true values for all items, say a valuation of 10 and a maximum price of 5. α is set to one; the chosen randomization is $\Delta_1 = 0.1$, $\Delta_2 = 0.3$, and $\Delta_3 = 0.2$. In each graph the points on the left side represent the bidders, the ones on the right side the items. The left graph displays the final first choice graph for the randomized input if all bidders report truthfully. In the right graph bidder 2 ensures he is matched to one of the items by reporting a higher maximum price. He does not risk a negative utility as long as the other bidders report the same as before.

4.3. Incentive-Compatibility in Expectation.

The following lemma shows that when the MHM is applied to a randomized input, each fixed bidder i can only obtain a higher utility by reporting different values than in the input (v, m, r) if he reports a maximum price strictly higher than $m_{i,j}$ for the item j he is then matched to.

Note that $u^{\mathrm{R}}_{i,j}(p_j) = u_{i,j}(p_j)$ if $p_j < m^{\mathrm{R}}_{i,j}$ and $u^{\mathrm{R}}_{i,j}(p_j) \leq u_{i,j}(p_j)$ for all p_j.

Lemma 11. *Assume that only multiples of some constant $\alpha > 0$ are allowed as input values. Let i be a bidder with utility functions defined by $v_{i,j}$ and $m_{i,j}$ for all items j. Let (v', m', r) and (v'', m'', r) be two inputs with $v'_{i,j} = v_{i,j}$ and $m'_{i,j} = m_{i,j}$ for all items j and $v'_{i',j} = v''_{i',j}$ and $m'_{i',j} = m''_{i',j}$ for all bidders $i' \neq i$ and all items j. Let (v', m'^R, r) and (v'', m''^R, r) be the corresponding randomized inputs for the same $\Delta_{i,j}$ for all bidder-item pairs (i, j) and (μ'^R, p'^R) and (μ''^R, p''^R) the corresponding outcomes for the randomized inputs, respectively. If $m''_{i,\mu''^R(i)} \leq m_{i,\mu''^R(i)}$, $p''^R_{\mu''^R(i)} < m^R_{i,\mu''^R(i)}$, or $p''^R_{\mu''^R(i)} \geq m_{i,\mu''^R(i)}$, then $u_{i,\mu''^R(i)}(p''^R_{\mu''^R(i)}) \leq u_{i,\mu'^R(i)}(p'^R_{\mu'^R(i)})$.*

Proof. By Corollary 10, $u^{\mathrm{R}}_{i,\mu'^{\mathrm{R}}(i)}(p'^{\mathrm{R}}_{\mu'^{\mathrm{R}}(i)}) \geq u^{\mathrm{R}}_{i,\mu''^{\mathrm{R}}(i)}(p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)})$. Assume for contradiction that $u_{i,\mu''^{\mathrm{R}}(i)}(p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)}) > u_{i,\mu'^{\mathrm{R}}(i)}(p'^{\mathrm{R}}_{\mu'^{\mathrm{R}}(i)})$. Then we have $u_{i,\mu''^{\mathrm{R}}(i)}(p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)}) > u_{i,\mu'^{\mathrm{R}}(i)}(p'^{\mathrm{R}}_{\mu'^{\mathrm{R}}(i)}) \geq u^{\mathrm{R}}_{i,\mu'^{\mathrm{R}}(i)}(p'^{\mathrm{R}}_{\mu'^{\mathrm{R}}(i)}) \geq u^{\mathrm{R}}_{i,\mu''^{\mathrm{R}}(i)}(p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)})$. If $p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)} < m^{\mathrm{R}}_{i,\mu''^{\mathrm{R}}(i)}$ or $p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)} \geq m_{i,\mu''^{\mathrm{R}}(i)}$, then $u_{i,\mu''^{\mathrm{R}}(i)}(p''^{\mathrm{R}}_{\mu''(i)})$ is equal to $u^{\mathrm{R}}_{i,\mu''^{\mathrm{R}}(i)}(p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)})$. Thus, we must have $m_{i,\mu''^{\mathrm{R}}(i)} > p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)} \geq m^{\mathrm{R}}_{i,\mu''^{\mathrm{R}}(i)}$. If bidder i has reported a maximum price $m''_{i,\mu''^{\mathrm{R}}(i)}$ at most $m_{i,\mu''^{\mathrm{R}}(i)}$, then $m''^{\mathrm{R}}_{i,\mu''^{\mathrm{R}}(i)} \leq m^{\mathrm{R}}_{i,\mu''^{\mathrm{R}}(i)}$ and $u''^{\mathrm{R}}_{i,\mu''^{\mathrm{R}}(i)}(p''^{\mathrm{R}}_{\mu''^{\mathrm{R}}(i)}) = -\infty$. This gives a contradiction to the feasibility of the outcome $(\mu''^{\mathrm{R}}, p''^{\mathrm{R}})$ for the input $(v'', m''^{\mathrm{R}}, r)$. $\square$

To remove the incentive for a bidder to report a higher maximum price, we use the *Randomized Extraction Scheme* presented in [10]. For an input (v, m, r) in which all values are multiples of some constant $\alpha > 0$, we execute the MHM with input randomization to obtain an outcome $(\mu^{\mathrm{R}}, p^{\mathrm{R}})$. Instead of charging a bidder i the computed price $p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}$ for the item he is matched to, the bidder has to pay $m^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)}$ with a probability $p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}/m^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)}$ and receives the item for free with probability $1 - p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}/m^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)}$. Thus, the expected price $\mathrm{E}[p_{\mu^{\mathrm{R}}(i)}]$ bidder i has to pay is

$$\mathrm{E}[p_{\mu^{\mathrm{R}}(i)}] = m^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)} \cdot \frac{p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}}{m^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)}} + 0 \cdot \left(1 - \frac{p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}}{m^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)}}\right) = p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)} \,.$$

Hence, in expectation all bidders have to pay the same prices and obtain the same utilities as without randomized extraction. Since $m^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)} < m_{i,\mu^{\mathrm{R}}(i)}$, the randomized prices are feasible for all bidders i.

Theorem 6. *Assume that only multiples of some constant $\alpha > 0$ are allowed as input values and that all true input values are multiples of α. Then the MHM combined with input randomization and randomized extraction is incentive-compatible in expectation.*

Proof. If for some item j a bidder i reports a maximum price $m'_{i,j}$ higher than $m_{i,j}$, then $m'_{i,j} \geq m_{i,j} + \alpha$ since only multiples of α are allowed as input. Thus, $m'^{\mathrm{R}}_{i,j} > m_{i,j}$. Hence, if i is matched to j, bidder i either has a positive probability for a utility of $-\infty$ due to the randomized extraction or $p^{\mathrm{R}}_j = 0$. Hence, no bidder has an incentive to report a maximum price higher than $m_{i,j}$ for an item j he could possibly be matched to at a strictly positive price. If a bidder has not reported a higher maximum price than $m_{i,j}$ for the item $\mu^{\mathrm{R}}(i)$ he is matched to in the outcome of the MHM with input randomization or obtains the item at a price of zero, then the algorithm is incentive-compatible by Lemma 11. As the charged prices are only in expectation equal to the prices calculated by the mechanism, the MHM with input randomization and randomized extraction is incentive-compatible in expectation. □

4.4. Utility.

To simplify the notation, we analyze the utility each bidder obtains without randomized extraction; this utility is equal to the expected utility with randomized extraction. Note that the following results hold for every randomization on the input.

In Lemma 13 we capture the following observations. First, the outcome (μ, p) is an envy-free outcome for the input (v, m^{R}, r). Thus, as the MHM computes an envy-free outcome with minimum prices, we have that the computed prices for the randomized input are at most as high as the prices p.

The next observation relates to the envy in the outcome of the MHM for the randomized input. If bidder i has envy for an item j in $(\mu^{\mathrm{R}}, p^{\mathrm{R}})$ then it must be that $m^{\mathrm{R}}_{i,j} \leq p^{\mathrm{R}}_j < m_{i,j}$. Furthermore, in this case, and more generally when $m_{i,j} - \alpha < p^{\mathrm{R}}_j < m_{i,j}$, for every envy-free outcome (μ', p') for the input (v, m, r) the price p'_j will be at least $m_{i,j}$, i.e., bidders can only have envy for items they would not be able to obtain in any envy-free outcome. These results are a formalization of the intuition that the randomization avoids problematic price updates such that some bidders can be matched at lower than envy-free prices to items that would not be rematched otherwise, while other bidders can profit from the reduced demand for the remaining items but have envy. Theorem 7 shows that in terms of utility bidders can only profit from this tie breaking.

Additionally, we can use the characterization in Section 3 of how a bidder can lie to show that the utility of each bidder in the outcome for the randomized input is at least the utility the bidder could obtain from lying if a bo-mechanism was used. This shows (1) that tie breaking is sufficient to achieve the same effect a bidder could generate by lying and (2) that by randomizing the input tie breaking can be done simultaneously for all bidders.

To prove Lemma 13 we first need the following observation.

Lemma 12. *If the MHM is started with an input in which all values are multiples of some constant $\alpha > 0$, then the prices of all items are multiples of α in each step of the mechanism.*

Proof. We denote by $p^{(t)}$ the prices after the t-th price update of the MHM and prove the claim by induction over t.

For $t = 0$ the claim is fulfilled trivially since all prices are initialized with their reserve prices.

Given that the claim holds after the $(t-1)$-st price update, and thus before the t-th price update, we show that the claim holds after the t-th price update. We denote the sets of items and bidders considered by the mechanism for the t-th price update with S and T, respectively. Let δ be $\delta = \min(\delta_{\text{out}}, \delta_{\max})$ with δ_{out} and $\delta_{\max}$ as defined in Lemma 1. The MHM does not change the prices for items $j \notin S$ and increases the prices of all items $j \in S$ by adding δ. Thus, if the prices $p^{(t-1)}$ are multiples of α and the price increase δ is a multiple of α, then also the prices $p^{(t)}$ are multiples of α. Given the prices $p^{(t-1)}$ as well as all input values are multiples of α, $\delta_{\max}$ is obviously a multiple of α. The value of δ_{out} is computed using u_i for $i \in T$ additionally to the input value $v_{i,j}$ and the price $p^{(t-1)}_j$ for some $j \in J \setminus S$. For every bidder i the best possible utility u_i is equal to $u_{i,j'}(p^{(t-1)}_{j'}) = v_{i,j'} - p^{(t-1)}_{j'}$ for some item $j' \in F_{p^{(t-1)}}(i)$. Thus, u_i is a multiple of α for all bidders i. Hence, δ_{out}, and thus δ, is a multiple of α. □

Lemma 13. *(a) For all items j it holds that $p^{R}_j \leq p'_j$.*

(b) For all bidder-item pairs $(i,j) \in \mu^R$ and every item j' with $p^R_{j'} < m^R_{i,j'}$ or $p^R_{j'} \geq m_{i,j'}$, it holds that $u_{i,j}(p^R_j) \geq u_{i,j'}(p^R_{j'})$.

(c) If for some bidder-item pair (i,j) we have $p^R_j > m_{i,j} - \alpha$, then $p'_j \geq m_{i,j}$ and $(i,j) \notin \mu'$.

Proof. (a) By Lemma 2 we have $p_j \leq p'_j$ for all items j. Thus it suffices to show $p^{\mathrm{R}}_j \leq p_j$ for all j.

By Lemma 12 the prices p are multiples of α. For multiples of α the utilities for the randomized input (v, m^{R}, r) are equal to the utilities for the input (v, m, r). Thus the outcome (μ, p) is an envy-free outcome for the input (v, m^{R}, r). By Lemma 2 the MHM computes the envy-free outcome with minimum prices, hence we have $p^{\mathrm{R}}_j \leq p_j$ for all j.

(b) Since μ^{R} is a matching in the first choice graph for the modified maximum prices m^{R} and prices p^{R}, we know $u^{\mathrm{R}}_{i,j}(p^{\mathrm{R}}_j) \geq u^{\mathrm{R}}_{i,j'}(p^{\mathrm{R}}_{j'})$. Since bidder i is matched to item j, $p^{\mathrm{R}}_j < m^{\mathrm{R}}_j$. Thus, $u^{\mathrm{R}}_{i,j}(p^{\mathrm{R}}_j) = u_{i,j}(p^{\mathrm{R}}_j)$, and we have $u_{i,j}(p^{\mathrm{R}}_j) \geq u^{\mathrm{R}}_{i,j'}(p^{\mathrm{R}}_{j'})$.

Case 1: $p^{\mathrm{R}}_{j'} < m^{\mathrm{R}}_{i,j'}$: Since only the maximum prices were modified through the randomization, $u^{\mathrm{R}}_{i,j'}(p^{\mathrm{R}}_{j'}) = u_{i,j'}(p^{\mathrm{R}}_{j'})$ for $p^{\mathrm{R}}_{j'} < m^{\mathrm{R}}_{i,j'}$. Hence, $u_{i,j}(p^{\mathrm{R}}_j) \geq u_{i,j'}(p^{\mathrm{R}}_{j'})$ for this case.

Case 2: $p^{\mathrm{R}}_{j'} \geq m_{i,j'}$: In this case, $u_{i,j'}(p^{\mathrm{R}}_{j'}) = -\infty$. Since $(i,j) \in \mu^{\mathrm{R}}$, we have $u_{i,j}(p^{\mathrm{R}}_j) = u^{\mathrm{R}}_{i,j}(p^{\mathrm{R}}_j) \geq 0$. Consequently, $u_{i,j}(p^{\mathrm{R}}_j) \geq u_{i,j'}(p^{\mathrm{R}}_{j'})$ holds.

(c) By assumption, $m_{i,j}$ is a multiple of α. By (a), $p_j \geq p^{\mathrm{R}}_j$; by Lemma 12, the prices p are multiples of α; thus, $p^{\mathrm{R}}_j > m_{i,j} - \alpha$ implies $p_j \geq m_{i,j}$. By Lemma 2, $p'_j \geq p_j \geq m_{i,j}$. Hence, $u_{i,j}(p'_j) = -\infty$, and bidder i is not matched to item j in μ'. □

Theorem 7. *For all bidders i we define $u^R_i = u_{i,\mu^R(i)}(p^R_{\mu^R(i)})$, $u'_i = u_{i,\mu'(i)}(p'_{\mu'(i)})$, and let $\hat{u}_i$ be the best utility a fixed bidder i can obtain if the MHM is run on the reported input and all other bidders report the same values as in (v,m,r). Then we have (a) $u^R_i \geq u'_i$ and (b) $u^R_i \geq \hat{u}_i$ for all bidders i.*

Proof. (a) For contradiction assume that $u'_i > u^{\mathrm{R}}_i$ for some bidder i. From Lemma 13 (a) we know that $p^{\mathrm{R}}_{\mu'(i)} \leq p'_{\mu'(i)}$. Thus, we have $u_{i,\mu'(i)}(p^{\mathrm{R}}_{\mu'(i)}) \geq u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^{\mathrm{R}}(i)}(p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}) \geq 0$.

Case 1: If $p^{\mathrm{R}}_{\mu'(i)} < m^{\mathrm{R}}_{i,\mu'(i)}$, this contradicts Lemma 13 (b).

Case 2: If $p^{\mathrm{R}}_{\mu'(i)} \geq m^{\mathrm{R}}_{i,\mu'(i)}$, by Lemma 13 (c) $p'_{\mu'(i)} \geq m_{i,\mu'(i)}$, which contradicts $u_{i,\mu'(i)}(p'_{\mu'(i)}) \geq 0$.

(b) Let $(\hat{\mu}, \hat{p})$ be the outcome when some fixed bidder i reports such that his utility is maximized while the other bidders report the same values as in input (v,m,r). By Theorem 3 we can assume that in the reported input $(\hat{v}, \hat{m}, r)$ the maximum prices of bidder i are equal to zero for all items except the one he is matched to in $\hat{\mu}$. Otherwise the input is equal to the input (v,m,r). We compare the outcome $(\mu^{\mathrm{R}}, p^{\mathrm{R}})$ to the outcome if the same randomization as used in (v, m^{R}, r) is applied to $(\hat{v}, \hat{m}, r)$. Let $(\hat{\mu}^{\mathrm{R}}, \hat{p}^{\mathrm{R}})$ be this outcome. Then by (a) we have $\hat{u}_{i,\hat{\mu}^{\mathrm{R}}(i)}(\hat{p}^{\mathrm{R}}_{\hat{\mu}^{\mathrm{R}}(i)}) \geq \hat{u}_{i,\hat{\mu}(i)}(\hat{p}_{\hat{\mu}(i)})$. For $\hat{\mu}(i)$ and bidder i we have $\hat{u}_{i,\hat{\mu}(i)}(p_{\hat{\mu}(i)}) = u_{i,\hat{\mu}(i)}(p_{\hat{\mu}(i)})$ for all $p_{\hat{\mu}(i)}$. Additionally, the item $\hat{\mu}^{\mathrm{R}}(i)$ has to be equal to $\hat{\mu}(i)$ because for the reported input bidder i only obtains a positive utility if he is matched to $\hat{\mu}(i)$. Thus we have $u_{i,\hat{\mu}(i)}(\hat{p}^{\mathrm{R}}_{\hat{\mu}(i)}) \geq u_{i,\hat{\mu}(i)}(\hat{p}_{\hat{\mu}(i)})$.

For bidder i and all j the reported maximum prices $\hat{m}_{i,j}$ are at most $m_{i,j}$. Thus by Lemma 11 $u_{i,\mu^{\mathrm{R}}(i)}(p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}) \geq u_{i,\hat{\mu}^{\mathrm{R}}(i)}(\hat{p}^{\mathrm{R}}_{\hat{\mu}^{\mathrm{R}}(i)}) = u_{i,\hat{\mu}(i)}(\hat{p}^{\mathrm{R}}_{\hat{\mu}(i)})$. Hence, $u_{i,\mu^{\mathrm{R}}(i)}(p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}) \geq u_{i,\hat{\mu}(i)}(\hat{p}_{\hat{\mu}(i)})$, i.e., the utility bidder i obtains with input randomization applied to (v,m,r) is at least the utility he can obtain when the MHM is applied to the reported input and the other bidders report according to (v,m,r). □

4.5. Rounded Prices.

The following theorem shows that for inputs in general position a bidder-optimal outcome can easily be obtained from the outcome for the randomized input by rounding up the prices to the next highest multiples of α. Thus our randomized mechanism extends upon the previously known results for the MHM, as it approximates the outcome of the MHM for inputs in general position and additionally provides incentive-compatibility in expectation for all inputs.

We define the *rounded prices* $\bar{p}^{\mathrm{R}}$ to be the prices we obtain when we round the price p^{R}_j of each item j up the next highest multiple of α.

Theorem 8. *If the input (v, m, r) is in general position, then the outcome $(\mu^R, \bar{p}^R)$ is bidder-optimal.*

We first show that if the prices in the outcome for the randomized input can be rounded up without reaching a maximum price for any matched bidder-item pair, then the outcome with the rounded prices is an envy-free outcome with minimum prices. In such a case we also know that the MHM without randomization is incentive-compatible for the given input.

Lemma 14. *If the outcome $(\mu^R, \bar{p}^R)$ is feasible, then (a) $(\mu^R, \bar{p}^R)$ is bidder-optimal and (b) the MHM is incentive-compatible for the input (v, m, r).*

Proof. (a) Let (μ, p) be the bidder-optimal outcome of the MHM for the input (v, m, r). By Lemma 12, the prices p are multiples of α. By Lemma 13 (a), $p^{\mathrm{R}}_j \leq p_j$ for all j. Thus the rounded prices $\bar{p}^{\mathrm{R}}$ fulfill $\bar{p}^{\mathrm{R}}_j \leq p_j$ for all items j. Since the prices p^{R} were rounded up to multiples of α to reach $\bar{p}^{\mathrm{R}}$, we further know $p^{\mathrm{R}}_j \leq \bar{p}^{\mathrm{R}}_j < p^{\mathrm{R}}_j + \alpha$ for all items j.

In the next step, we show that if we can round the prices p^{R} up to multiples of α without any bidder obtaining a negative utility when charged the new prices $\bar{p}^{\mathrm{R}}$, the outcome $(\mu^{\mathrm{R}}, \bar{p}^{\mathrm{R}})$ is envy-free. Assume by contradiction that some bidder i has envy for item j in the outcome $(\mu^{\mathrm{R}}, \bar{p}^{\mathrm{R}})$, i.e., $u_{i,j}(\bar{p}^{\mathrm{R}}_j) = v_{i,j} - \bar{p}^{\mathrm{R}}_j > u_{i,\mu^{\mathrm{R}}(i)}(\bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}) = v_{i,\mu^{\mathrm{R}}(i)} - \bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)} \geq 0$. As all valuations v and prices $\bar{p}^{\mathrm{R}}$ are multiples of α, it actually holds that $u_{i,j}(\bar{p}^{\mathrm{R}}_j) = v_{i,j} - \bar{p}^{\mathrm{R}}_j \geq v_{i,\mu^{\mathrm{R}}(i)} - \bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)} + \alpha$. Since $u_{i,j}(\bar{p}^{\mathrm{R}}_j) > 0$, it must be that $\bar{p}^{\mathrm{R}}_j < m_{i,j}$; thus, $p^{\mathrm{R}}_j \leq m_{i,j} - \alpha$. Hence, by Lemma 13 (b), $u_{i,\mu^{\mathrm{R}}(i)}(p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}) \geq u_{i,j}(p^{\mathrm{R}}_j)$. We have

$$u_{i,\mu^{\mathrm{R}}(i)}(p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}) \geq u_{i,j}(p^{\mathrm{R}}_j) \geq u_{i,j}(\bar{p}^{\mathrm{R}}_j) \geq u_{i,\mu^{\mathrm{R}}(i)}(\bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}) + \alpha\ . \tag{9}$$

Reformulated, this yields

$$\bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)} \geq p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)} + \alpha\ , \tag{10}$$

which is a contradiction to $\bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)} < p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)} + \alpha$. Thus, no bidder can have envy in $(\mu^{\mathrm{R}}, \bar{p}^{\mathrm{R}})$, and $(\mu^{\mathrm{R}}, \bar{p}^{\mathrm{R}})$ is bidder-optimal for (v, m, r) by Lemmata 2 and 3.

(b) Based on (a) we show that if $(\mu^{\mathrm{R}}, \bar{p}^{\mathrm{R}})$ is feasible, the MHM is incentive-compatible for the original input (v, m, r). Let ${u^{\mathrm{R}}_{i,j}}^{-1}(\cdot)$ be defined as in (2) but with the randomized maximum prices instead of the original ones. The randomized input is in general position; thus, (i) by Lemma 6 in $(\mu^{\mathrm{R}}, p^{\mathrm{R}})$ all items j with $p^{\mathrm{R}}_j > r_j$ are matched, and (ii) by Lemma 7 for every subset of matched bidders $\hat{I} \subseteq I$ and ${\hat{r_j}}^{\mathrm{R}} = \max(r_j, \max_{i \notin \hat{I}} {u^{\mathrm{R}}_{i,j}}^{-1}(u^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)}(p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)})))$ at least one item $j \in \mu^{\mathrm{R}}(\hat{I})$ has $p^{\mathrm{R}}_j = {\hat{r_j}}^{\mathrm{R}}$. Let $\hat{r_j}$ be equal to $\max(r_j, \max_{i \notin \hat{I}} u^{-1}_{i,j}(u_{i,\mu^{\mathrm{R}}(i)}(\bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)})))$. By rounding the prices p^{R} up to the next highest multiples of α, the prices of the items $j \in \mu^{\mathrm{R}}(\hat{I})$ that were equal to ${\hat{r_j}}^{\mathrm{R}}$ become equal to $\hat{r_j}$ because, for all bidders i, $u^{-1}_{i,j}(u_{i,\mu^{\mathrm{R}}(i)}(\bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}))$ is a multiple of α and the difference between $u^{-1}_{i,j}(u_{i,\mu^{\mathrm{R}}(i)}(\bar{p}^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}))$ and ${u^{\mathrm{R}}_{i,j}}^{-1}(u^{\mathrm{R}}_{i,\mu^{\mathrm{R}}(i)}(p^{\mathrm{R}}_{\mu^{\mathrm{R}}(i)}))$ is less than α. The latter follows from the continuity of the functions $u^{-1}_{i,j}(\cdot)$ and ${u^{\mathrm{R}}_{i,j}}^{-1}(\cdot)$. Hence, the properties (i) and (ii) hold for $(\mu^{\mathrm{R}}, \bar{p}^{\mathrm{R}})$.

Consider the following algorithm: 1) Randomize the input. 2) Apply the MHM to the randomized input. 3) Round the computed prices up to the next highest multiples of α to obtain an outcome. 4) If the outcome is feasible, output it. Otherwise, apply the MHM to the original input and output the result. This algorithm always computes a bidder-optimal outcome and outputs $(\mu, \bar{p}^{\mathrm{R}})$ for the input (v, m, r) and the chosen randomization. Since properties (i) and (ii) hold for $(\mu, \bar{p}^{\mathrm{R}})$, by Theorem 2 every bo-mechanism is incentive-compatible for the input (v, m, r). Hence, the MHM is incentive-compatible for (v, m, r). □

For Theorem 8 we further need the following two lemmata.

Lemma 15. *Let T be the set of bidders in a maximal alternating tree in G^R_p. Then $\bar{p}_j - p_j$ is equal for all $j \in F^R_p(T)$.*

Proof. Each bidder i obtains the same utility $u_i = u_{i,j}(p_j) = v_{i,j} - p_j$ from all his first choice items $j \in F_p^{\mathrm{R}}(i)$. Since all valuations are multiples of α, the prices of all the first choice items of bidder i need to have the same difference to their next highest multiple of α. As all bidders $i \in T$ are connected in the maximal alternating tree via their first choice items, for all first choice items j of the bidders in T the value of $\bar{p}_j - p_j$ is the same. □

Lemma 16. *If for some item j with $(i,j) \in \mu^R$ we have $\bar{p}_j^R - p_j^R > 0$, then the last price update for j was due to $\delta_{max}^R = m_{i',j'}^R - p_{j'}$ for some bidder $i' \neq i$ and some item j', where p are the prices before the last price update of j. Furthermore, $\bar{p}_j^R - p_j^R = \Delta_{i',j'} = \bar{p}_{j'}^R - p_{j'}^R$ and $\bar{p}_{j'}^R = m_{i',j'}$. If the outcome $(\mu^R, \bar{p}^R)$ gives bidder i a negative utility and the prices were increased by $\delta_{max} = m_{i',j'} - p_{j'}$ instead of δ_{max}^R, a problematic price update would have occurred.*

Proof. Consider the last price increase δ before the price of item j reached p_j^{R}. Denote with $p^{(t-1)}$ resp. $p^{(t)}$ the prices before resp. after this price increase. Let $\mathcal{T}$ be the maximal alternating tree in $G_{p^{(t-1)}}^{\mathrm{R}}$ from which δ is determined and S and T the sets of items resp. bidders in $\mathcal{T}$. Since the price of j is increased by the price increase δ, we have $j \in S$.

As j is not overdemanded at prices p^{R} (regarding the randomized maximum prices), the price increase δ has to resolve the overdemand for item j. One possibility to resolve the overdemand is to add an unmatched item to the maximal alternating tree by $\delta = \delta_{\mathrm{out}}$. Since the randomized input (v, m^{R}, r) is in general position by Lemma 9, every item with a price strictly higher than its reserve price is already matched by Lemma 6. Thus, after the price increase, the price of the added unmatched item has to be equal to its reserve price, i.e., equal to a multiple of α. The newly added item is a first choice for some bidder in the (new) maximal alternating tree. Thus, if an unmatched item was added to $\mathcal{T}$ by $\delta = \delta_{\mathrm{out}}$, then by Lemma 15 all prices of items in S would be multiples of α at prices $p^{(t)}$, which contradicts $\bar{p}_j^{\mathrm{R}} - \alpha < p_j^{\mathrm{R}} = p_j^{(t)} < \bar{p}_j^{\mathrm{R}}$.

Hence, the overdemand for item j can only be resolved by $\delta = \delta_{\mathrm{max}}^{\mathrm{R}} = m_{i',j'}^{\mathrm{R}} - p_{j'}^{(t-1)}$ for some $i' \in T$, $i' \neq i$, and $j' \in F_{p^{(t-1)}}^{\mathrm{R}}(i')$. If i and j were already matched before the price increase, the overdemand for j can be resolved if the edge (i', j') is either unmatched and on the path from j to the root of $\mathcal{T}$ or matched and on a different path. If i and j were not matched, then the overdemand can be resolved if (i, j) is on the path from j' to the root and i' and j' were matched. In all other cases, i.e., (i', j') is unmatched and not on the path from j to the root or (i', j') is matched and on the path from j to the root, the overdemand for j would not be resolved by the price update $\delta_{\mathrm{max}}^{\mathrm{R}}$.

In each of the possibilities that resolve the overdemand for j, the items j and j' remain connected in the first choice graph. Thus, the last price increase for j also has to be the last price increase for j'.

If rounding up p_j^{R} to the next highest multiple of α leads to a negative utility for i, we have $m_{i,j} - \alpha < p_j^{\mathrm{R}} < m_{i,j}$. Since $m_{i',j'} - m_{i',j'}^{\mathrm{R}} < \alpha$, we have $m_{i',j'} - \alpha < p_{j'}^{(t)} < m_{i',j'}$. Thus, if the prices of all items in $F_{p^{(t-1)}}^{\mathrm{R}}(T)$ were increased by $\delta_{\mathrm{max}} = \min_{i \in T, j \in F_{p^{(t-1)}}^{\mathrm{R}}(i)}(m_{i,j} - p_j^{(t-1)})$, both $m_{i',j'}$ and $m_{i,j}$ would be reached. Each of the possibilities that resolve the overdemand for j would cause a problematic price update if the prices $p^{(t-1)}$ would be increased by δ_{max}. □

Proof of Theorem 8. By Lemma 4 no problematic price update can occur for inputs in general position, independently of the current prices. Hence, by Lemma 16, the outcome $(\mu^{\mathrm{R}}, \bar{p}^{\mathrm{R}})$ is feasible for inputs in general position. Lemma 14 (a) concludes the proof. □

We conjecture that by a tight coupling of the MHM for an input that fulfills the rematch condition with the MHM for the corresponding randomized input Theorem 8 can be extended to inputs that fulfill the rematch condition.

Acknowledgements

This work was funded by the Vienna Science and Technology Fund (WWTF) through project ICT10-002.

References

[1] G. Aggarwal, S. Muthukrishnan, D. Pál, M. Pál, General auction mechanism for search advertising, in: 18th International Conference on World Wide Web, ACM, New York, 2009, pp. 241–250. doi:10.1145/1526709.1526742.

[2] P. Dütting, M. Henzinger, Mechanisms for the marriage and the assignment game, in: 7th International Conference on Algorithms and Complexity, Springer, Heidelberg, 2010, pp. 6–12. doi:10.1007/978-3-642-13073-1_2.

[3] P. Dütting, M. Henzinger, I. Weber, Bidder optimal assignments for general utilities, Theoretical Computer Science 478 (2013) 22–32. doi:10.1016/j.tcs.2013.01.030.

[4] L. S. Shapley, M. Shubik, The assignment game I: The core, International Journal of Game Theory 1 (1971) 111–130. doi:10.1007/BF01753437.

[5] H. B. Leonard, Elicitation of honest preferences for the assignment of individuals to positions, Journal of Political Economy 91 (3) (1983) 461–479. doi:10.1086/261158.

[6] N. Chen, X. Deng, A. Ghosh, Competitive equilibria in matching markets with budgets, SIGecom Exchanges 9 (1) (2010) 5:1–5:5. doi:10.1145/1980534.1980539.

[7] P. Dütting, M. Henzinger, I. Weber, An expressive mechanism for auctions on the web, in: 20th International Conference on World Wide Web, ACM, New York, 2011, pp. 127–136. doi:10.1145/1963405.1963427.

[8] P. Dütting, M. Henzinger, I. Weber, Sponsored search, market equilibria, and the Hungarian method, Information Processing Letters 113 (3) (2013) 67 – 73. doi:10.1016/j.ipl.2012.11.006.

[9] H. W. Kuhn, The Hungarian method for the assignment problem, Naval Research Logistics Quarterly 2 (1955) 83–97. doi:10.1002/nav.3800020109.

[10] S. Bhattacharya, V. Conitzer, K. Munagala, L. Xia, Incentive compatible budget elicitation in multi-unit auctions, in: Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2010, pp. 554–572.

[11] P. Billingsley, Probability and Measure, 2nd Edition, Wiley, New York, 1986.